CO-EVOLUTIONARY LEARNING OF CONTEXTUAL ASYMMETRIC ACTORS

Siang Yew Chong
School of Computer Science
The University of Nottingham, Malaysia Campus
Jalan Broga, Semenyih 43500, Malaysia
Email: Siang-Yew.Chong@nottingham.edu.my

Christopher Hill
The Graduate School
The University of Nottingham, Malaysia Campus
Jalan Broga, Semenyih 43500, Malaysia
Email: Christopher.Hill@nottingham.edu.my

Xin Yao
School of Computer Science
The University of Birmingham
Edgbaston, Birmingham B15 2TT, UK
Email: X.Yao@cs.bham.ac.uk

KEYWORDS

Repeated encounter games, asymmetric payoff, co-evolutionary learning, multi-population, spanish democratic transition.

ABSTRACT

Co-evolutionary learning of the iterated prisoner's dilemma (IPD) has been used to model and simulate interactions, which may not be realistic due to assumptions of a fixed and symmetric payoff matrix for all players. Recently, we proposed to extend the co-evolutionary learning framework for any two-player repeated encounter game to model more realistic behavioral interactions. One issue we studied is to endow players with individual and self-adaptive payoff matrix to model individual variations in their utility expectations of rewards for making certain decisions. Here, we study a different issue involving contextual asymmetric actors. The differences in the utility expectations (payoff matrix) are due to contextual circumstances (external) such as political roles rather than variations in individual preferences (internal). We emphasize the model of interactions among contextually asymmetric actors through a multi-population structure in the co-evolutionary learning framework where different populations representing different actor roles interact. We study how different actor roles modelled by fixed and asymmetric payoff matrices can have an impact to the outcome of co-evolutionary learning. As an illustration, we apply co-evolutionary learning of two contextually asymmetric actors from the spanish democratic transition.

INTRODUCTION

The iterated prisoner's dilemma (IPD) has become the standard metaphor to explain cooperative behaviors among selfish, unrelated individuals (Axelrod, 1984). In its classical form, two players are given two choices, cooperate and defect, and engaged in repeated interactions. The game captures the *dilemma* of cooperative interactions between unrelated, selfish players seeking the highest payoffs through the payoff matrix that defines rewards players received based on the joint choices they made in a behavioral exchange. Although payoffs for mutual cooperation are higher compared to that of mutual defection, a player that exploit its opponents will receive the highest payoff while the opponent receives the lowest payoff.

However, defection is not always the best choice for the IPD. Axelrod (Axelrod, 1980a,b) showed through tournaments of experts-designed strategies that a particular form of cooperative strategy, *tit for tat*, that cooperates in the first round and reciprocates thereafter the choice that the opponent made, can be viable. Other studies (Axelrod, 1984; Nowak and Sigmund, 1992) have explained cooperation in the IPD using the notion of *reciprocal altruism*. i.e., cooperative behaviors arises from repeated encounters that allows cooperation received by an individual to be returned. Further studies have formulated and shown that co-evolutionary models (Axelrod, 1987; Fogel, 1991, 1993; Darwen and Yao, 1995; Julstrom, 1997; Chong et al., 2007) allow cooperative behaviors to be learned through an adaptation process on behavioral representation guided by strategic interactions.

It is noted that all of these studies have assumed that the utility expectation on rewards (payoff matrix) of a strategy for making certain decisions does not change (fixed) and is similar for all strategies (symmetric). These two assumptions may not be realistic if the IPD is used as a model to explain outcomes of real-world interactions due to variations between individuals, which are reflected by them having different utility expectation (payoff matrix) (Johnson et al., 2002). Cooperative outcomes may be due to a different payoff matrix that favors mutual cooperation (i.e., mutualism (Clements and Stephens, 1995)).

Recently, we (Chong and Yao, 2006) proposed to extend the co-evolutionary learning framework for any two-player repeated encounter game that address the restrictive assumptions of having fixed and symmetric payoff

Proceedings 23rd European Conference on Modelling and
Simulation ©ECMS Javier Otamendi, Andrzej Bargiela,
José Luis Montes, Luis Miguel Doncel Pedrera (Editors)
ISBN: 978-0-9553018-8-9 / ISBN: 978-0-9553018-9-6 (CD)

matrix, and allows the modelling of more realistic behavioral interactions. One issue we studied is to endow players with individual and self-adaptive payoff matrix to model individual variations in their utility expectations of rewards for making certain decisions. The result of the co-evolutionary process is that the evolutionary outcome is dependent on the adaptation process of both behaviors (strategy behavioral responses) and utility expectations that determine how behaviors are rewarded (strategy payoff matrices).

In this paper, we study a different issue involving contextual asymmetric actors. The differences in the utility expectations (payoff matrix) are due to contextual circumstances (external) such as political roles rather than variations in individual preferences (internal). The model of interactions among contextually asymmetric actors are achieved using a multi-population structure in the co-evolutionary learning framework. With the single-population framework (Chong and Yao, 2006), strategies undergoing evolutionary (selection and variation) processes in the same population can only represent actors that are contextually similar in roles, despite their variations in individual payoff matrices. With the multi-population framework, the evolution of strategies within a population is guided by strategic interactions with strategies in other populations. This allows the modelling of behavioral exchange between contextually asymmetric actors as different populations represent different actor roles.

We study and show how different actor roles modelled by fixed and asymmetric payoff matrices can have an impact in the co-evolutionary learning of behaviors through empirical studies. As an example, we show that the co-evolution of a population using IPD payoff matrix that interacts with another population using mutualism payoff matrix can lead to mutual cooperation. Mutual cooperation is achieved if rewards for exploitation are sufficiently low. In another example, we consider a real-world political setting through the study of the co-evolutionary learning of two dissimilar actor roles, Franco (dictator) and Juan Carlos (prince), where results suggested that the circumstances of Franco and Juan Carlos learning to mutually cooperate are a result of Juan Carlos reflecting Franco's favorable expectation of a suitable successor.

The rest of the paper is organized as follows. The following section describes the general setting for iterated, two-player, two-choice games to model behavioral interactions. After that, a literature background for the case study on the Spanish Democracy Transition is provided. The following section then describes the co-evolutionary framework used for the experiments in detail. Results of the experiments are then reported with discussions on the observations obtained from the experimental results. Finally, the paper finishes with conclusion with some remarks for future studies.

REPEATED ENCOUNTER GAMES

Repeated encounter games have been used generally to model behavioral interactions. One example is the symmetric, two-player, two-choice game as a model for cooperation (Mesterton-Gibbons and Dugatkin, 1992). Here, each players have the choice of either to cooperate or defect. Depending of the choices that both players have made, each player receives a payoff that is specified by a predefined payoff matrix (Fig. 1). Both players receive R units if both cooperates, or P units if both defects. However, if one player cooperates while the other defects, the cooperator will receive S units while the defector receives T units. The values R, S, T, and P must satisfy the constraints, $R > P$ and $R > (S + T)/2$. Note that for Figure 1, the payoff given in the lower left-hand corner is assigned to the player (row) choosing the move, while that of the upper right-hand corner is assigned to the opponent (column). The game is said to be symmetrical if the payoff matrix is the same for both players (Mesterton-Gibbons and Dugatkin, 1992).

	Cooperate	Defect
Cooperate	R (upper right) / R (lower left)	T (upper right) / S (lower left)
Defect	S (upper right) / T (lower left)	P (upper right) / P (lower left)

Figure 1: The Payoff Matrix Framework of a Two-Player, Two-Choice Game

Games with different characteristics can be produced depending on the specifications for R, S, T, and P. For example, in a simple single iteration game, when $T > R > P > S$, one obtains the prisoner's dilemma game, where the best choice is to defect. However, when $R > T$ and $S > P$, one obtains the mutualism game, where the best choice is to cooperate (Clements and Stephens, 1995). Both of these games can be extended to have more than one iteration in the form of iterated games to model repeated encounters (Johnson et al., 2002).

In (Chong and Yao, 2006), we proposed a generic framework of repeated encounter games, whereby the payoff matrix is different and adaptable for each strategy. The framework is applicable to any repeated encounter game with different payoff matrices, and not just limited to those that satisfy specific constraints such as the IPD or mutualism games.

SPANISH DEMOCRATIC TRANSITION

The Spanish transition to democracy (1975-1982) is a period of significance and interest as it uniquely demonstrates the capacity to transform a right wing dictatorship of Francisco Franco to a socialist government of Felipe

González in a period of seven years, marking it as the only transition of its kind in history (Figure 2). Much has been written in terms of historical discussion but little in the regard of scientific analysis of the actions taken. In particular, the necessity to understand on a more fundamental level the capabilities and constraints associated with individual involvement in democratic transitions is intricately connected to a comprehensive awareness of the form of transitions as a process (Hill, 2007).

1936 – 1939	Spanish Civil War
1938, Jan 5th	Birth of Juan Carlos in Rome
1939, Apr 1st	Franco defeats Republican forces and takes control of Spain
1969	Franco designates Juan Carlos as heir – no real power endowed
1975, Nov 20th	Death of Franco
1975, Nov 22nd	Juan Carlos crowned King of Spain
1982	Socialist Party elected to power

Figure 2: Spanish Democratic Transition Timeline

Throughout the transition period Franco held the dominant position, effecting change and directing the order of progression. The primary objective of the dictator lay in his desire to ensure the survival of his regime after his death and the manner in which he chose to effect this was through the creation of a francoist monarchy. In terms of pure power distribution, Juan Carlos was very much the subservient player even though his interests converged with those of Franco. The desired result would culminate for both men in the nomination of Juan Carlos as King of Spain and the guaranteed continuation and prosperity of the existing regime.

Franco's desire for a francoist monarchy is evident by his actions regarding Juan Carlos. The desire for legitimacy (Gilmour, 1985), as represented by Juan Carlos' bloodline, outweighed any previous evidence of the dictator's animosity towards the young prince's family. Franco called Juan Carlos back from exile in order to establish a dynasty that would sustain his vision for Spain. Franco had allowed Juan Carlos to return to Spain under very specific conditions, with his education and upbringing closely monitored. Juan Carlos was provided scarce opportunity to be anything other than that which Franco required. Franco saw this as the mechanism to ensure continuation of his rule and Juan Carlos believed this as the only manner by which a Borbón (prior royal family) would once again inherit the throne. Franco needed a successor, and Juan Carlos had both the lineage and the ability to create what Franco had been unable to; a *francoist* monarchy. However, Juan Carlos was aware even at this stage that the only way to gain access to the throne was through Franco (Cernuda, 1981) and the dangers inherent in acting against him (Hill, 2007).

As illustrated in Table 1, Juan Carlos (JC) represented

Table 1: Payoff Matrix Scenario for Spanish Democratic Transition

	SCENE1		SCENE2		SCENE3	
	FC	JC	FC	JC	FC	JC
R	1	1	1	1	0	1
S	0	-1	0	-1	-1	-1
T	0	-1	0	-1	1	-1
P	-1	-1	0	-1	1	-1

the ideal choice for Franco (FC) under ideal circumstances (SCENE1) but alternative solutions could be generated based upon a reordering of individual preferences or a more fundamental understanding of relative gains that are made possible simply because of the power Franco had. Deviation on the part of Juan Carlos (Gibson, 1992) and an awareness of alternative choices on the part of Franco would lead to a situation whereby Franco's ideal preference would remain that of Juan Carlos but he would be prepared to accept an alternative candidate if forced to. Juan Carlos on the other hand, would remain powerless to effect change given the emergence of an alternative for Franco (reflected by the reduction of rewards for P in SCENE2 for FC's payoff matrix while no changes are made to JC's payoff matrix) (Hill, 2007).

Under the scenario listed above, the choice of Juan Carlos was an artificial constraint generated by Franco's desire for legitmacy and monarchy (Gunther, 1992). The nature of this constraint equally depended upon the character and actions of Juan Carlos and therefore endowed Franco with the ability to alter the status quo by reordering his preferences should the actions of Juan Carlos deviate from expectations. If legitimacy and legacy could be achieved without royal blood (Alba, 1978) then the choices available to Franco would change considerably even if those available to Juan Carlos are unaffected (SCENE3).

In analyzing the Spanish transition to democracy, despite Franco's selection of Juan Carlos, we cannot conclude that he constituted the only choice. Franco's ability (outlined by the Law of Succession) to remove Juan Carlos indicates that we must be cautious about any assumptions we make. It became imperative for Juan Carlos to retain Franco's trust and patronage, as this would enable consolidation of his own position and potential for future development. Providing the decisions taken could not be enforced by a third party (e.g., the military) (Powell, 1991), Juan Carlos would be in an advantageous position upon accepting the nomination. The power and influence he gained through his association with Franco would grant him a measure of freedom but the all-important factor became, assuring that he received the nomination at all.

CO-EVOLUTIONARY LEARNING

Strategy Representation

We limit our investigations to only deterministic, memory-one strategy for simplicity. We consider the same direct look-up table strategy representation proposed and studied earlier in (Chong and Yao, 2005), which allows direct behavioral evolution of strategies. Figure 3 illustrates this strategy representation. $m_{ij}, i, j = 1, 2$ specifies the choice to be made, given the inputs i (player's own previous choice) and j (opponent's previous choice). Instead of using pre-game inputs (two for memory-one strategies), the first move is specified independently, m_{fm}. Binary choices of $+1$ and -1 are used to represent cooperate and defect, respectively.

		Opponent's Previous Move	
		$+1$	-1
Player's	$+1$	m_{11}	m_{12}
Previous Move	-1	m_{21}	m_{22}

Figure 3: The Look-Up Table Strategy Representation for the Two-Player Repeated Encounter Game with Two Choices and Memory Length One

A simple mutation operator is used to generate offspring from parents. Mutation replaces the original element with the other possible choice. For example, if mutation occurs at $m_{11} = +1$, then the mutated element m'_{11} takes -1. Each table element has a fixed probability, p_m, of being replaced by the other choice. The value p_m is not optimized, and crossover is not used in any of the experiments as mutation alone is sufficient to introduce variations to strategy representation.

Co-evolutionary Learning Procedure

We describe in detail an example of a multi-population co-evolutionary learning framework to model behavioral interactions of contextually asymmetric actors. The two-population co-evolutionary learning allows the adaptation of strategy behavioral responses. The two-population co-evolutionary learning procedure is described as follows:

1. Generation step, $t = 0$:
   Initialize $N/2$ parent strategies in the first population, $P_i^1, i = 1, 2, ..., N/2$, and in the second population, $P_i^2, i = 1, 2, ..., N/2$, randomly. Each table element, m_{fm} and $m_{ij}, i, j = 1, 2$, is initialized to values, $+1$ or -1, each with equal probability.

2. Generate $N/2$ offspring, $O_i^1, i = 1, 2, ..., N/2$, from $N/2$ parents, $P_i^1, i = 1, 2, ..., N/2$, for the first population and $N/2$ offspring, $O_i^2, i = 1, 2, ..., N/2$, from $N/2$ parents, $P_i^2, i = 1, 2, ..., N/2$, for the second population using a mutation operator with probability $p_m = 0.05$.

3. Each strategy in the first population compete with each strategy in the second population (i.e., full-mixing, where for N strategies in a population, every strategy competes a total of N games).

4. Select the best $N/2$ strategies based on total payoffs of all games played for both populations. Increment generation step, $t = t + 1$.

5. Step 2 to 4 are repeated until termination criterion (i.e., a fixed number of generation) is met.

For the experiments, we use $N = 20$ and repeat the co-evolutionary process for 200 generations, which is sufficiently long to observe an evolutionary outcome (e.g., persistent cooperation). A fixed game length of 150 iterations is used for all game. Experiments are repeated for 30 independent runs.

We note the fundamental difference between the multi-population co-evolutionary learning framework and the single-population frameworks studied earlier such as in (Chong and Yao, 2005, 2006). Considering the single-population framework in (Chong and Yao, 2006), strategies undergoing evolutionary processes in the same population can only represent actors that are contextually similar in roles, despite their variations in individual payoff matrices. For the two-population co-evolutionary learning that we study here, strategies undergo evolutionary processes within their own populations (i.e., strategies from the first population will not be mixed with and transferred to the second population, and vice versa). Furthermore, the co-evolution of strategies in the first population will be guided by the results of interactions with strategies from the second population, and vice versa. This allows the modelling of behavioral exchange between contextually asymmetric actors as different populations represent different actor roles.

We also note other forms of mutiple-population co-evolutionary learning framework such as cooperative co-evolution Potter and Jong (2000). However, the motivation for having a multi-population framework in cooperative co-evolution is to have samples of different subcomponents undergoing evolutionary processes that can later be combined to obtain the full solution (i.e., divide-and-conquer). Here, the multi-population framework is applied to a competitive co-evolution approach where the motivation is to have solutions (strategies) competing with one another through interactions (game-play).

RESULTS AND DISCUSSION

Learning Behaviors For Contextually Asymmetric Actors

We start our study by first investigating the impact of having contextually asymmetric actors to the behavioral outcomes of the two-population co-evolutionary learning. For baseline comparison, we first consider the case where the payoff matrices for each strategy are fixed and the same. As an example, we consider both populations with the same IPD payoff matrix. Note that we label an experiment as [First Population Payoff Matrix]-[Second Population Payoff Matrix]. For example, IPD1-IPD1 refers to the experiment where the first and second populations use an IPD payoff matrix as shown in Table 2. We use IPD1-IPD1-P1 and IPD1-IPD1-P2 to refer to the evolutionary outcomes of the first and second population, respectively.

Table 2: Payoff Matrix Labels for Repeated Encounter Games

	IPD1	IPD2	MUT
R	4	4	5
S	0	0	1
T	5	6	4
P	1	1	0

It is noted that a two-population structure is imposed to the co-evolution whereby strategies undergo evolutionary processes in their own respective populations. However, the evolved strategies between the two populations interact with one another through IPD game-plays. As such, we would expect that evolved strategies would learn to play cooperatively through co-evolution starting from random initialization, which is the outcome one would observe from a single-population co-evolutionary learning (Chong and Yao, 2005). Experimental results show that the majority of runs ended with cooperative outcomes (Table 3), which one can further observe from the higher percentages of cooperation being played on average (i.e., more than 50%) (Table 3).

Table 3: Experimental Result for Number of Runs where"$No < 25\%$" Indicates Majority Play is Defection while "$No > 75\%$" is for Cooperation

	$No < 25\%$	$No > 75\%$
IPD1-IPD1-P1	4	24
IPD1-IPD1-P2	4	24
IPD1-MUT-P1	8	22
IPD1-MUT-P2	0	30
IPD2-MUT-P1	12	17
IPD2-MUT-P2	0	29

However, the payoff matrices for both populations need not be the same. We investigate the case of a population with the IPD payoff matrix (IPD1) and the other population with the iterated mutualism payoff matrix (MUT). Although one would expect both populations learning to cooperative (Table 3), the second population using MUT would play more cooperatively as shown by the increase in the number of runs having cooperative outcomes (Table 3) and an increase of cooperation plays that is statistically significant (Table 4). When the first population uses an IPD payoff matrix that further rewards strategies that are tempted to defect and exploit opponents (IPD2), results show that the first population would learn to play significantly less cooperative (Tables 3 and 4) although the second population still plays significantly more cooperative as a result of MUT payoff matrix. Finally, we observe that all populations that learned to cooperate converged to a stable and high cooperative plays compared to the population that did not learn to cooperate, which exhibited fluctuations during co-evolution.

Table 4: Experimental Results for Average of Cooperation Frequency over 30 Runs in % with a Confidence Interval of 95% in % and t-test where † Indicates a Statistical Significant Difference

	Mean $\pm$ Std Err	t-test
IPD1-IPD1-P1	77.09 ± 11.71	
IPD1-IPD1-P2	78.85 ± 11.88	
IPD1-MUT-P1	70.33 ± 14.87	0.78
IPD1-MUT-P2	95.85 ± 1.10	$-2.82^\dagger$
IPD2-MUT-P1	42.38 ± 16.54	$3.32^\dagger$
IPD2-MUT-P2	94.36 ± 3.27	$-2.43^\dagger$

Although results coincide strongly with our expectation on the outcome of co-evolutionary learning given how payoff matrices are assigned to the two populations, these investigations further highlight the subtle but important differences between a multi-population and single-population co-evolutionary learning framework to model behavioral interactions. In particular, the two-population co-evolution made it possible to study the learning of behavioral responses between two contextually asymmetric actors having different utility expectations to rewards for the same joint actions.

Co-evolutionary Analysis Of The Spanish Democracy Transition

In this section, we study the two-population co-evolutionary learning framework applied to model the Spanish Democratic Transition. Although historical outcome of the interactions between the two actors, Franco (dictator) and Juan Carlos (Prince), is known and that various political science arguments have been given to explain the cooperative outcome leading to the transition, it is not really known why they have chosen to cooperate. Here, we study and offer an analysis to the Spanish Democratic Transition through co-evolutionary learning. We assume that the two actors are rational beings (seeking to maximize payoffs) and capable of learning behavioral responses based only on interactions with one another. In using the two-population co-evolution, the first and second populations will represent samples of behaviors for Franco and Juan Carlos, respectively. We consider the payoff matrices for the populations representing Franco (PFC) and Juan Carlos (PJC) based on possible scenarios identified earlier (Hill, 2007) and listed in Table 1.

Tables 5 and 6 summarize the results of the experi-

Table 5: Spanish Democratic Transitions: Experimental Result for Number of Runs where "$No < 25\%$" Indicates Majority Play is Defection while "$No > 75\%$" is for Cooperation

	$No < 25\%$	$No > 75\%$
SCENE1-PFC	0	30
SCENE1-PJC	0	30
SCENE2-PFC	0	30
SCENE2-PJC	0	30
SCENE3-PFC	30	0
SCENE3-PJC	0	24

Table 6: Spanish Democratic Transitions: Experimental Results for Average of Cooperation Frequency over 30 Runs in % with a Confidence Interval of 95% in % and t-test where † Indicates a Statistical Significant Difference

	Mean $\pm$ Std Err	t-test
SCENE1-PFC	98.16 ± 0.87	
SCENE1-PJC	97.10 ± 1.54	
SCENE2-PFC	97.42 ± 1.16	1.43
SCENE2-PJC	96.90 ± 1.43	0.56
SCENE3-PFC	3.59 ± 1.34	$108^{\dagger}$
SCENE3-PJC	87.70 ± 6.59	$3.02^{\dagger}$

ments for the modelling of Spanish Democratic Transition scenario. For SCENE1 and SCENE2 experiments, both populations evolve to mutual cooperation play starting from random initializations of strategy behaviors. Although results suggest that strategies representing Franco and Juan Carlos learn to cooperate with the political succession (e.g., no statistical differences in the outcomes of cooperative plays), one needs to note the subtle difference between the two scenario that represents different conditions with which mutual cooperation is learned. In both scenario, the payoff matrix of Juan Carlos (Table 1) reflects his expectation that he is not the only available successor, in which case, he is better of cooperating. However, the payoff matrix of Franco for the two scenario differs, i.e., SCENE1 reflects the case that Franco wants Juan Carlos to succeed while SCENE2 reflects the case that Juan Carlos represents the ideal choice but other alternatives are considered (lower P for the case of SCENE2 compared to that of SCENE1).

Results for SCENE1 and SCENE2 are in agreement with the historical outcome of the Spanish Democratic Transitions, suggesting that Franco would have wanted Juan Carlos to succeed. However, what if there are other equally suitable alternatives that are within Franco's expectation. SCENE3 reflects one possibility that the royal lineage of Juan Carlos is not necessarily the important factor in the succession and that Franco has the political power to choose other alternatives. In this case, our experimental results suggest that mutual cooperation is not possible, i.e., that strategies representing Franco would learn to defect (choosing another successor) despite cooperation from strategies representing Juan Carlos.

CONCLUSION

This paper has studied behavioral interactions involving contextually asymmetric actors. The main motivation is to address the issue that real-world interactions involving actor roles that are contextually different. We modelled behavioral interactions of contextual asymmetric actors through a multi-population co-evolutionary learning framework. Unlike the single-population framework where strategies only represent contextually similar actor role, the multi-population framework models behavioral exchange between contextually asymmetric actors with different populations representing different actor roles, with co-evolution guided by strategic interactions with strategies between populations.

We showed that different actor roles modelled by fixed and asymmetric payoff matrices can have an impact in the co-evolutionary learning of behaviors through empirical studies. In one example, we showed that the co-evolution of a population using IPD payoff matrix that interacts with another population using mutualism payoff matrix can lead to mutual cooperation if rewards for exploitation are sufficiently low. In another example, we studied a real-world political setting of the Spanish Democratic Transitions, where results suggested that the circumstances of Franco and Juan Carlos learning to mutually cooperate are a result of Juan Carlos reflecting Franco's favorable expectation of a suitable successor.

This paper presents a preliminary study on behavioral interactions involving contextually asymmetric actors. Future studies would investigate other issues involving more complex interactions in political transitions.

ACKNOWLEDGEMENT

This work was partially support by an EPSRC grant (No. GR/T10671/01).

REFERENCES

Alba, V. (1978). *Transition in Spain: From Franco to Democracy*. Transaction Publishers.

Axelrod, R. (1980a). Effective choice in the prisoner's dilemma. *The Journal of Conflict Resolution*, 24(1):3–25.

Axelrod, R. (1980b). More effective choice in the prisoner's dilemma. *The Journal of Conflict Resolution*, 24(3):379–403.

Axelrod, R. (1984). *The Evolution of Cooperation*. Basic Books, New York.

Axelrod, R. (1987). The evolution of strategies in the iterated prisoner's dilemma. In Davis, L. D., editor, *Genetic Algorithms and Simulated Annealing*, chapter 3, pages 32–41. Morgan Kaufmann, New York.

Cernuda, P. (1981). *Todo un Rey*. Ediciones Temas de Hoy.

Chong, S. Y., Humble, J., Kendall, G., Li, J., and Yao, X. (2007). Learning ipd strategies through co-evolution. In Kendall, G., Yao, X., and Chong, S. Y., editors, *The Iterated Prisoners' Dilemma: 20 Years On*, pages 63–87. World Scientific, Singapore.

Chong, S. Y. and Yao, X. (2005). Behavioral diversity, choices, and noise in the iterated prisoner's dilemma. *IEEE Transactions on Evolutionary Computation*, 9(6):540–551.

Chong, S. Y. and Yao, X. (2006). Self-adaptive payoff matrices in repeated interactions. In *2006 IEEE Symposium on Computational Intelligence and Games (CIG'06)*, pages 103–110, Piscataway, NJ. IEEE Press.

Clements, K. C. and Stephens, D. W. (1995). Testing models of non-kin cooperation: Mutualism and the prisoner's dilemma. *Animal Behaviour*, 50:527–535.

Darwen, P. J. and Yao, X. (1995). On evolving robust strategies for iterated prisoner's dilemma. In *Progress in Evolutionary Computation*, volume 956 of *Lecture Notes in Artificial Intelligence*, pages 276–292. Springer.

Fogel, D. B. (1991). The evolution of intelligent decision making in gaming. *Cybernetics and Systems: An International Journal*, 22:223–236.

Fogel, D. B. (1993). Evolving behaviors in the iterated prisoner's dilemma. *Evolutionary Computation*, 1(1):77–97.

Gibson, I. (1992). *Fire in the Blood*. BBC Books.

Gilmour, D. (1985). *Transformation of Spain: From Franco to the Constitutional Monarchy*. Quartet he Books Limited.

Gunther, R. (1992). Spain: The very model of the modern elite settlement. In Higley, J. and Gunther, R., editors, *Elites and Democratic Consolidation in Latin America and Southern Europe*, pages 38–80. Cambridge University Press, Cambridge.

Hill, C. (2007). *The Role of Elites in the Spanish Transition to Democracy (1975-1981) Motors of Change*. The Edwin Mellen Press.

Johnson, D. D. P., Stopka, P., and Bell, J. (2002). Individual variation evades the prisoner's dilemma. *BMC Evolutionary Biology*, 2(15).

Julstrom, B. A. (1997). Effects of contest length and noise on reciprocal altruism, cooperation, and payoffs in the iterated prisoner's dilemma. In *Proc. 7th International Conf. on Genetic Algorithms (ICGA'97)*, pages 386–392, San Francisco, CA. Morgan Kauffman.

Mesterton-Gibbons, M. and Dugatkin, L. A. (1992). Cooperation among unrelated individuals: Evolutionary factors. *The Quarterly Review of Biology*, 67(3):267–281.

Nowak, M. A. and Sigmund, K. (1992). Tit for tat in heterogeneous populations. *Nature*, 355:250–253.

Potter, M. A. and Jong, K. A. D. (2000). Cooperative coevolution: An architecture for evolving coadapted subcomponents. *Evolutionary Computation*, 8(1):1–29.

Powell, C. T. (1991). *El Piloto del cambio. El rey, la Monarquía y la transición a la democracia*. Editorial Planeta.

AUTHOR BIOGRAPHIES

SIANG YEW CHONG is an Assistant Professor with the School of Computer Science, University of Nottingham, Malaysia Campus, a member of the Automated Scheduling, Optimization and Planning (ASAP) Research Group, University of Nottingham, UK, and an Honorary Research Fellow with the School of Computer Science, University of Birmingham, UK. He was a Research Associate with the Centre of Excellence for Research in Computational Intelligence and Applications (CERCIA), University of Birmingham, UK, in 2007. His major research interests include evolutionary computation, neural networks, and evolutionary game theory. Dr. Chong received the Ph.D. degree in Computer Science from the University of Birmingham, UK, in 2007, and was awarded the 2009 IEEE Computational Intelligence Society Outstanding Ph.D. Dissertation Award for his work on theoretical frameworks for performance analysis in co-evolutionary learning.

CHRISTOPHER HILL is the Director of the Graduate School, University of Nottingham, Malaysia Campus. He was a Teaching Fellow with the School of Politics and International Relations, University of Nottingham, UK, from 2005 to 2007 and a member of the Centre for the Study of European Governance during his time there as both Ph.D. candidate and member of staff. His major research interests include the role of elite actors in transition processes, iterated prisoners dilemma, game theory and its implication for the structure-agency debate. Dr. Hill received his Ph.D. degree in Political Science from the University of Nottingham, UK, in 2006, and was nominated for the 2006 Lord Bryce Prize – Comparative and International Relations Best Dissertation Prize as sponsored by Political Studies Association.

XIN YAO is a Professor (Chair) of Computer Science and the Director of the Centre of Excellence for Research in Computational Intelligence and Applications (CERCIA) in the School of Computer Science, University of Birmingham, UK. He is an IEEE Fellow, an IEEE Computational Intelligence Society Distinguished Lecturer, the former Editor-in-Chief (2003-08) of the IEEE Transactions on Evolutionary Computation, an associate editor or editorial board member of twelve other journals, and the Editor of the World Scientific Book Series on Advances in Natural Computation. He has given more than 50 invited keynote and plenary speeches at conferences and workshops worldwide. His major research interests include evolutionary computation and neural network ensembles. He has more than 300 refereed publications. Prof. Yao won the 2001 IEEE Donald G. Fink Prize Paper Award for his work on evolutionary artificial neural networks.