Adaptive Behavior

Increasing Population Diversity Through Cultural Learning

Dara Curran, Colm O'Riordan
Department of Information Technology, National University of Ireland, Galway

A number of learning models are commonly employed in the simulation of social behavior. These include population learning, lifetime learning and cultural learning. Population learning allows populations as a whole to evolve over time, typically through a Darwinian model of natural selection. Lifetime learning allows individuals to acquire knowledge during their lifetimes and cultural learning allows individuals to pass this knowledge to their peers or subsequent generations. This work examines the effects of cultural learning on both the fitness and the diversity of a population of neural network agents. A population employing population learning alone and one employing both population and cultural learning are assigned three benchmark tasks: the 5-bit parity problem, the game of tic-tac-toe and the game of connect-four. Each agent contains a genome which encodes a neural network controller used by the agent to perceive and react to environmental stimuli. Results show that the addition of cultural learning promotes improved fitness and significantly increases both genotypic (the genetic make up of individuals) and phenotypic (the behavior of individuals) diversity in the population.

Keywords cultural learning · diversity · learning models · artificial life

1 Introduction

There has been much research focused on the interaction between learning and evolution, where individuals within a population are capable of evolving genetically (population learning) and also of acquiring knowledge during their lifetime (individual learning). The combination of population and individual learning combines the global search power of genetic evolution with the local optimization of lifetime learning. Hinton and Nowlan (1987) were among the first to show that learning could guide evolution and a number of other researchers have since conducted experiments which support this view (Cangelosi & Parisi, 1996; Curran & O'Riordan, 2003a, 2003b; Floreano & Mondada, 1989; Nolfi & Parisi, 1996; Nolfi, Parisi, & Elman, 1994; Pereira & Costa, 2001; Sasaki & Tokoro, 1997; Watson & Wiles, 2002).

Training data for individual learning can be provided by an external oracle, providing a template by which individuals can be compared and corrected. Alternatively, the population itself can carry out the sharing of information through a cultural process. Culture can be succinctly described as a process of information transfer within a population that occurs without the use of genetic material. Culture can take many forms such as language, signals or artifacts.

A large body of work exists examining both the origins of culture and some of its effects on a population of agents (Billard & Hayes, 1997; Cangelosi, 1999;

Correspondence to: Dara Curran, Department of Information Technology, National University of Ireland, Galway, Ireland.
E-mail: dara.curran@nuigalway.ie
Tel.: +353 91 495041

(2006), Vol 14(4): 315–338.
DOI: 10.1177/1059712306072335
Figure 2 appears in color online: http://adb.sagepub.com

Denaro & Parisi, 1996; Hutchins & Hazlehurst, 1991; MacLennan & Burghardt, 1993; Spector, 1994; Steels, 1996a, 1996b, 1997; Yanco & Stein, 1993). Culture may be formed through a process of imitation, where an individual acquires knowledge from another through repeated observation and mimicry. While the existence of such imitation in nature is still in some dispute, cultural learning can still be found in the animal kingdom (Kawamura, 1963; Whiten & Ham, 1992; Zentall, 2001).

Research examining the effects of cultural learning in populations of agents includes work by Hutchins and Hazlehurst (1991, 1995), Borenstein and Ruppin (2003) and Nolfi et al. (Nolfi & Parisi, 1996; Nolfi et al., 1994). While each model differs slightly in its implementation, the underlying mechanisms are similar. A number of teachers are selected from the population and pupils observe and imitate their actions or utterances. Teachers may be either selected from the current generation (corresponding to the horizontal model of cultural transmission) or the previous generation (corresponding to the vertical model of cultural transmission) (Belew, 1990; Boyd & Richerson, 1985). Such research has shown that the addition of culture to a population of agents is capable of enhancing the population's fitness.

Bull, Holland, and Blackmore (2000) examined the effects of altering the relative speed of cultural, or meme, evolution compared with gene evolution and found that as the speed of cultural evolution increases, genetic evolution degrades considerably, even halting completely. In the words of Blackmore (2000), cultural evolution "takes hold of the leash" by becoming increasingly important in the evolution and adaptation of a population.

This work is concerned with examining the effects of cultural learning on the population and, in particular, the effect of cultural learning on the diversity of the population. It has become clear that cultural learning enhances evolution but it is less clear how exactly the mechanism achieves this. This article is an attempt to bring attention to the largely under-used measure of population diversity as valuable tool in the analysis of the effects of cultural learning. In particular, we wish to examine both genotypic (the genetic differences between individuals in the population) and phenotypic (the behavioral differences between individuals) diversity.

The experiments presented in this work employ two populations. One uses population learning alone, while individuals are allowed to evolve using a genetic algorithm. The second employs both population and cultural learning. Three benchmark problems are used for each experiment: 5-bit parity, the game of tic-tac-toe and the game of connect-four.

The remainder of the article is structured as follows: Section 2 discusses some background material. Section 3 outlines the methods employed for the experiments, including a presentation of the model employed, descriptions of the benchmark learning tasks chosen and an explanation of each diversity measure. Section 4 presents the results obtained. Section 5 provides a discussion on these results and Section 6 summarizes the work.

2 Background

2.1 Learning Models

A number of learning models can be identified from observation in nature. These can roughly be classified into population, life-time and cultural learning.

2.1.1 Population Learning Population learning refers to the process whereby a population of organisms evolves, or learns, by genetic means through a Darwinian process of iterated selection and reproduction of fit individuals. In this model, the learning process is strictly confined to each organism's genetic material: The organism itself does not contribute to its survival through any learning or adaptation process.

2.1.2 Life-Time Learning There exist species in nature that are capable of learning, or adapting to environmental changes and novel situations at an individual level. Such learning, known as life-time learning, is often coupled with population-based learning, further enhancing the population's fitness through its adaptability and resistance to change. Another phenomenon related to life-time learning, first reported by Baldwin (1896), occurs when certain behavior discovered through life-time learning becomes imprinted onto an individual's genetic material through the evolutionary processes of crossover and mutation. To quote Hinton and Nowlan (1987) whose model was the first to demonstrate this effect through simulation, "learning can provide an easy evolutionary path towards co-adapted alleles in environments that have no good evolutionary path for non-learning organisms." Subsequent work has further explored the interactions between evolution and learning and shown that the addition of individual lifetime learning can improve a population's fitness (Curran & O'Riordan, 2003a, 2003b; Floreano & Mondada, 1989; Nolfi & Parisi, 1996; Nolfi et al., 1994; Pereira & Costa, 2001; Sasaki & Tokoro, 1997; Watson & Wiles, 2002).

2.1.3 Cultural Learning Culture can be succinctly described as a process of information transfer within a population that occurs without the use of genetic material. Culture can take many forms such as language, signals or artifacts. Such information exchange occurs during the lifetime of individuals in a population and can greatly enhance the behavior of such species. Because these exchanges occur during an individual's lifetime, cultural learning can be considered a subset of lifetime learning.

A number of approaches have been implemented to simulate cultural learning including fixed lexicons (Cangelosi & Parisi, 1998; Yanco & Stein, 1993), indexed memory (Spector & Luke, 1996), cultural artifacts (Cangelosi & Parisi, 1998; Hutchins & Hazlehurst, 1991) and signal-situation tables (MacLennan & Burghardt, 1993). The approach chosen here was inspired by the teacher/pupil scenario (Billard & Hayes, 1997; Cangelosi & Parisi, 1998; Denaro & Parisi, 1996) where a number of highly fit agents are selected from the population to act as teachers for the next generation. Pupils learn from teachers by observing the teacher's verbal output and attempting to mimic it using their own verbal apparatus. As a result of these interactions, a lexicon of symbols evolves to describe situations within the population's environment.

Experiments conducted by Hutchins and Hazlehurst (1995) simulate cultural evolution through the use of a hidden layer within an individual neural network in the population. The hidden layer acts as a verbal input/output layer and performs the task of feature extraction used to distinguish different physical inputs.

In previous work by Denaro and Parisi (1996), it was suggested that the addition of noise to a teacher's verbal output could enhance a population's ability to retain culturally acquired information. Experiments conducted in our previous work (Curran & O'Riordan, 2004) confirmed that small levels of noise introduced to the communication process improved both the formation of a shared lexicon and agent performance.

2.1.4 Diversity A common view of the evolutionary process is that diversity enhances the performance of a population by providing more opportunities for evolution. A homogeneous population offers no advantage for improvement as the entire population is focused on a particular portion of the search space. By contrast, a diverse population will simultaneously sample a large area of search space, providing the opportunity to locate different, potentially better, solutions. Two types of diversity measure can be identified: genotypic and phenotypic diversity.

Genotypic diversity measures examine each individual's genecode and determines its similarity to all other genecodes in the population. Methods employed in previous research include ancestry and tagging (McPhee & Hopper, 1999), distance measures (Gusfield, 1997; O'Reilly, 1997) and entropy based approaches (Rosca, 1995).

Phenotypic diversity is concerned with the behavior of the population, that is the phenotypical representation of its genotype. Typically this can be achieved by examining each individual's response to stimuli in its environment and comparing it with the responses of all others in the population. Phenotypic diversity measures employed by previous researchers generally apply fitness based approaches (Burke, Gustafson, & Kendall, 2004; McQuesten, 2002).

3 Methods

This section outlines the model employed for this set of experiments (including the encoding scheme used to map neural network structures to genetic code and vice versa), describes the diversity measures developed to examine both genotypic and phenotypic diversity and finally, a presentation of each benchmark learning task.

3.1 Model

A population of agents is used to solve the benchmark tasks described in Section 4. Each agent consists of a neural network controller that allows it to perceive and interact with its environment. The neuro con-

	Start Marker	Node Label	Threshold	Link to Node	Link Weight	Link to Node	Link Weight	End Marker	
...	SM	5	0.8	4	0.83	3	-0.51	EM	...

Figure 1 Marker-based encoding.

troller's architecture and weight values are encoded into the agent's genome and are allowed to genetically evolve over the course of each experiment. Naturally, weight values altered during an agent's lifetime through any cultural process are not reencoded into the genome.

The model allows populations to evolve using population learning alone, or using a combination of population and cultural learning. When cultural learning is applied, teacher agents instruct pupil agents by interacting with their environment (using an approach similar to that outlined in Section 2.2). Agents do not employ lifetime learning: The only form of learning that occurs during their lifetime is cultural learning, which is applied at the start of each agent's life. Thus, an agent's life can be summarized in the following steps:

- Agent is born (neural network structure is formed from inherited genetic encoding)
- Agent is exposed to teaching (where teaching is applied)
- Agent may be selected for reproduction
- Agent may be selected to become a teacher for the next generation

The general algorithm for the experiments presented in this work is as follows:

- Randomly generate initial population
- Repeat:
  - If teaching is applied:
    - For each individual in the population, perform cultural learning for the required number of teaching cycles
    - Calculate population fitness
    - Select 10% of the population to become teachers
  - Otherwise:
    - Calculate population fitness
  - Perform selection, crossover and mutation to generate the next generation

The following subsections outline the encoding scheme used to convert an individual's genetic code to a neural network structure, the processes of crossover and mutation, and the diversity measures and cultural learning implementation employed for this work.

3.1.1 Encoding Scheme One of the most crucial aspects of the model is the translation of genetic codes to neural network structures. Many encoding schemes were considered in preparation of the simulator, prioritizing flexibility, scalability, difficulty and efficiency. These included connectionist encoding (Belew, McInerney, & Schraudolph, 1992), node-based encoding (White & Ligomenides, 1993), graph-based encoding (Pujol & Poli, 1998), layer-based encoding (Mandischer, 1993), marker-based encoding (Moriarty & Miikkulainen, 1995), matrix rewriting (Kitano, 1990; Miller, Todd, & Hedge, 1989), cellular encoding (Gruau, 1994), weight-based encoding (Kolen & Pollack, 1991; Sutton, 1986), architecture encoding (Koza & Rice, 1991) and cascade-correlation architectures (Fahlmann, 1991).

The scheme chosen is inspired by marker-based encoding, which allows any number of nodes and interconnecting links for each network, giving a large number of possible neural network architecture permutations. Marker-based encoding represents neural network elements (nodes and links) in a sequential list. Each element is separated by a marker to allow the decoding mechanism to distinguish between the different elements and therefore deduce interconnections (Kitano, 1990; Miller et al., 1989).

In this implementation, a marker is given for every node in a network. Following the node marker, the node's details are stored in sequential order in the bit string. This includes the node's label and its threshold value. Immediately following the node's details is another marker which indicates the start of one or more node-weight pairs. Each of these pairs indicates a back connection from the node to other nodes in the network along with the connection's weight value. Once the last connection has been encoded, the scheme places an end marker to indicate the end of the node's encoding.

3.1.2 Crossover As a result of the chosen encoding scheme, crossover may not operate at the bit level as this could result in the generation of invalid gene codes. Therefore, the crossover points are restricted to specific intervals—only whole node or link values may be crossed over.

Two-point crossover is employed in this implementation. Once crossover points are selected, the gene portions are swapped. The connections within each portion remain intact, but it is necessary to adjust the connections on either side of the transferred portion to successfully integrate it into the existing gene code. This is achieved by using node labels for each node in the network. These labels are used to identify individual nodes and to indicate the location of interconnections. Once the portion is inserted, all interconnecting links within the whole gene code are examined. If any links are now pointing to non-existing nodes, they are modified to point to the nearest labeled node.

3.1.3 Mutation The mutation operator introduces additional noise into the genetic algorithm process thereby allowing potentially useful and unexplored regions of problem space to be probed. The mutation operator usually functions by making alterations to the gene code itself, typically by altering specific values randomly selected from the entire gene code. In this implementation, weight mutation is employed. The operator modifies a weight according to a percentage value chosen randomly from the range –200% to +200%. Mutation can alter the value of a start or end marker, thereby introducing structural novelty into the evolutionary process.

3.1.4 Cultural Learning Cultural learning is implemented using an imitation scheme similar to that employed by Hutchins and Hazlehurst (1991, 1995), Borenstein and Ruppin (2003) and Nolfi et al. (Nolfi & Parisi, 1996; Nolfi et al., 1994) all of whom employ a teacher agent's neural network output value as the target for pupils. The fittest agents of the current generation are selected to become teachers that instruct the next generation. The model therefore employs a vertical model of cultural transmission, where information is passed on from one generation to the next (Belew, 1990; Boyd & Richerson, 1985).

Individuals are generated from their genetic code and are immediately exposed to teaching. As an agent encounters stimuli in its environment, it responds both behaviorally (emitting a signal through its output nodes) and verbally (emitting a signal through its verbal nodes—see Figure 2). Pupils iteratively attempt to imitate the teacher's verbal output value using cycles of error back-propagation. These are referred to in the rest of this article as teaching cycles, where one teaching cycle is equivalent to a single exposure to stimulus and a subsequent error back-propagation.

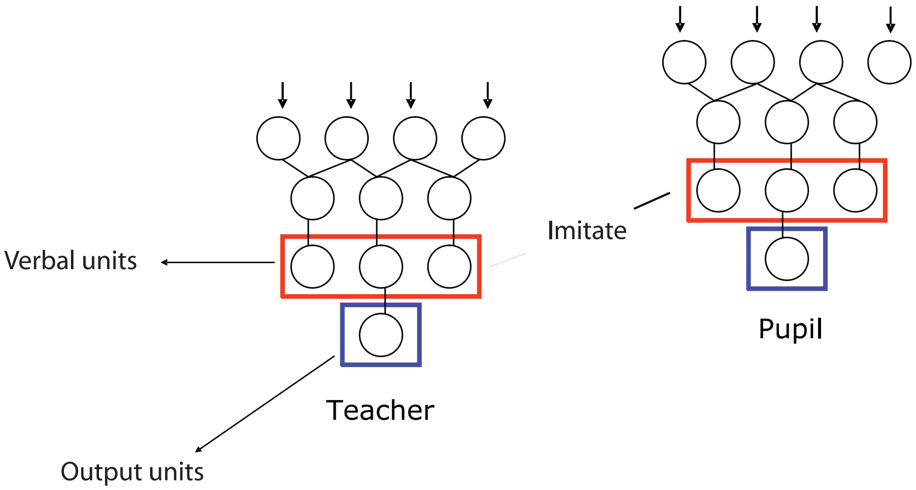

Figure 2 Agent communication architecture.

The model is not restricted to imparting only innate knowledge (knowledge that is genetically inherited and that agents have at their disposal from birth). In other words, teacher agents may impart knowledge that they themselves acquired culturally during their own learning phase.

3.2 Diversity Measures

A common view of the evolutionary process is that diversity enhances the performance of a population by providing more opportunities for evolution. A homogeneous population offers no advantage for improvement as the entire population is focused on a particular portion of the search space. By contrast, a diverse population will simultaneously sample a large area of search space, providing the opportunity to locate good solutions.

Previous work has included the measurement of both genotypic diversity [typically through edit distances (Gusfield, 1997; O'Reilly, 1997) between genomes] and phenotypic diversity [including entropy measurement (Rosca, 1995), crowding (De Jong, 1975) and niching (Mahfoud, 1995a, 1995b)]. In addition, much research has focused on promoting, maintaining or re-introducing diversity into evolving populations of solutions in order to achieve maximum performance. This includes work on mating schemes (Booker, 1985; Collins & Jefferson, 1991, Davidor, 1991; Hillis, 1990; Muhlenbein, 1989; Spiessens & Manderick, 1991) and fitness sharing (Deb & Goldberg, 1989; Goldberg & Richardson, 1987).

The following sections outline the implementation for the genotypic and phenotypic diversity measure employed in this work. The population's diversity is measured at the start of each generation, before any cultural learning takes place.

3.2.1 Genotypic Diversity

The diversity measure employed in this work was attained following the implementation of preliminary cruder measures: node and link diversity.

Preliminary Measures Node diversity simply compares the number of nodes in each agent's neural network with all other agents in the population. Link diversity compares the number of links in each agent's neural network with those present in all other agents in the population. Both measures employ a Euclidian distance measure to compare two individuals and the global diversity measure is an average of all these distances.

To test the measures, an experiment was conducted using the simple 5-bit parity task. A population of 100 agents was presented with 5-bit patterns and attempted to identify whether the number of ones in each pattern was odd or even. Fitness was assigned according to the mean square error of an agent's network. In order to simplify the experiment to examine the effects of these measures, the population employed only population learning to evolve new networks. Both diversity measures were applied to each generation at birth and the experiment was conducted for 400 generations.

Figure 4 shows the averaged results of 20 experiment runs. Both node and link diversity exhibit a similar trend but the two are significantly different to a 95% confidence level (p value < 0.0001, paired t-test).

Final Measure In order to refine the diversity measure, more information than provided by simple raw node and link counts is necessary. To achieve this, we consider the encoding of each agent's neural network in more detail, in an effort to combine node and link information to attain a better representation of the differences between two individuals.

It is important to note that the encoding method imposes two restrictions on a more detailed diversity scheme: The first is that each item in the encoding is a real number, rather than a simple on/off switch. This makes the scheme more suited toward distance measures. Secondly, an encoded genecode is circular in nature, that is, the end of a genecode wraps back to the start. This implementation allows crossover to occur naturally, but gives some restrictions for the purposes of direct comparisons between two genecodes.

The first of these restrictions is easily dealt with. A Euclidian distance measure can easily be found between two items in the gene codes being compared. The second restriction is more complex: Using a traditional, naive approach, two genecodes could be directly compared using the Euclidian measure above. However, such a comparison would yield a high diversity rate for two identical, but rotated, strings. Since the computational time required to search each pair of strings and correct any rotational deficiencies (which may or may not exist) would be impractical, a different method had to be devised.

Our proposed method examines the content of the genotype and breaks each chromosome into meaningful portions, where a meaningful portion is defined as data contained between a start and end marker. In other words, each meaningful portion contains data about a single node and all the links emanating from it to other nodes. Once all meaningful portions have been extracted from the string, any remaining data is kept aside (maintaining its contiguous structure) and labeled as spurious. It is worth noting that this data should be considered dormant rather than purely redundant as a crossover or mutation may reactivate previously isolated spurious data.

Having isolated the meaningful and spurious gene portions from the pair of gene codes being examined, the algorithm proceeds to examine meaningful portions of each genecode. The length of each portion corresponds exactly to the number of links that emanate from each node. Clearly, meaningful portions will not all be of equal length. As a result, an approach must be devised to choose pairs of portions that most merit comparison. For instance, it is not worth comparing an input node, with no incoming links, with an output or hidden node, but it is reasonable to assume that two hidden nodes with an approximately similar number of links are suitable and useful for comparison.

Thus, the algorithm selects pairs of meaningful gene code portions of equal (or as similar as possible) length from the gene codes being examined (Figure 3)

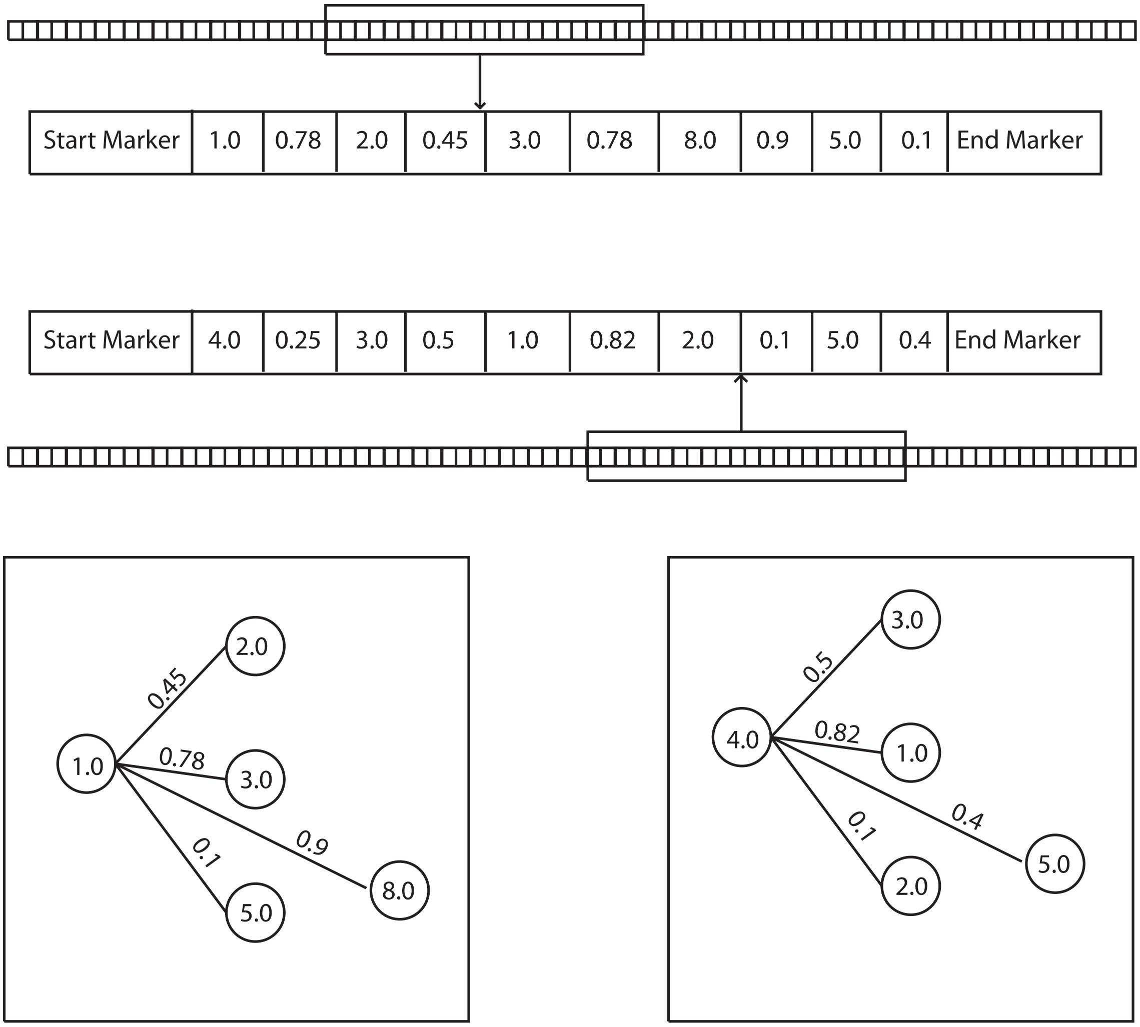

Decoded Gene Portion A

Decoded Gene Portion B

Figure 3 Genotypic diversity measure.

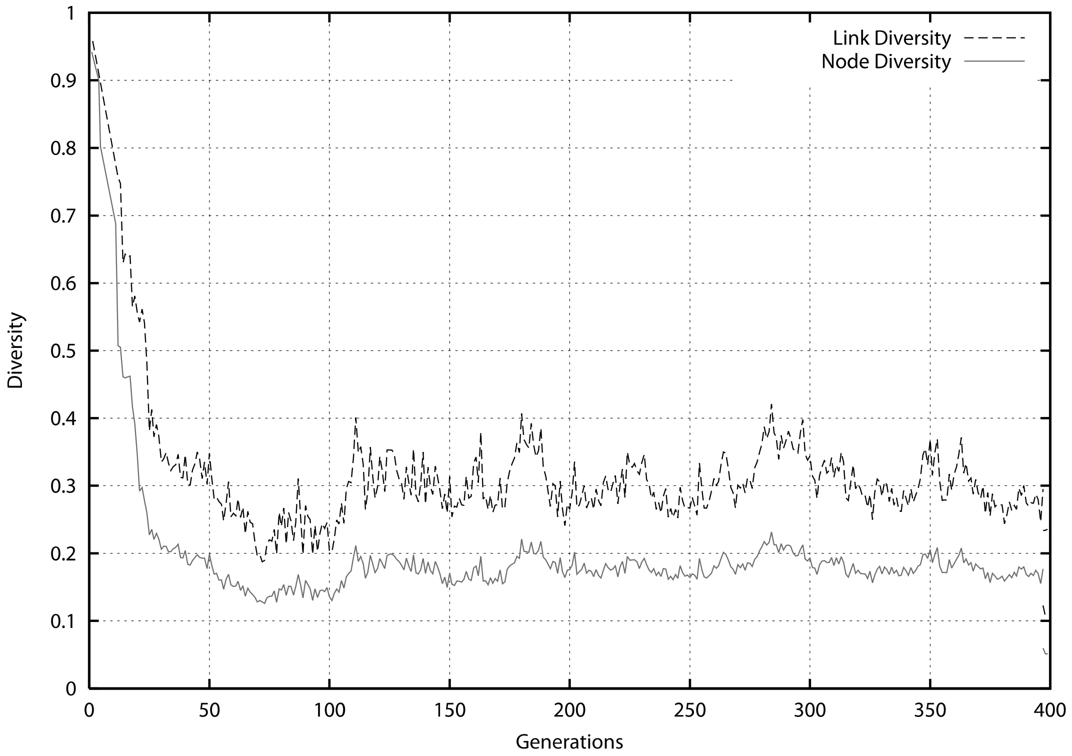

Figure 4 Node and link count diversity measures.

and performs a Euclidian distance measurement for each value to determine their similarity.

$$distance = \sqrt{(gene_a - gene_b)^2}.$$

The distance measures for each pair of portions are averaged together to give a diversity measure for the two full length chromosomes.

Once all pairs of meaningful gene portions have been examined for a given pair of genecodes, the algorithm performs the same comparison task for the spurious data. Since, as previously argued, such data is dormant rather than redundant, it merits comparison as much as active meaningful portions.

Each member of the population is compared with every other member, and the diversity measure is averaged over the entire population, resulting in a global diversity measure:

$$div_{global} = \frac{\sum_{i=1}^{n}(div_{meaningful_i} + div_{spurious_i})}{n}.$$

The new diversity measure takes into account both node and link information present in each neural network encoding. In order to determine how the new diversity measure compares with the cruder measures examined earlier, a new experiment was conducted. Again, a population of 100 agents attempted to solve the 5-bit parity problem using only population learning for 400 generations. Figure 5 shows the results of the new diversity measure along with the cruder measures examined earlier. Interestingly, the new measure falls between the node and link diversity values, showing a very similar trend, with a sharp initial fall followed by stabilization, indicative of convergence.

This result shows that the new measure correlates with cruder node and link counts in terms of its overall trend. In addition, the new measure exhibits considerably less noise than either of the two previous methods. While it is not possible to objectively state that the measure is necessarily more accurate due to its more thorough examination of each genome, it is important for a useful measure to consider as much genetic information as possible.

3.2.2 Phenotypic Diversity The approach employed to measure the phenotypic diversity of the population

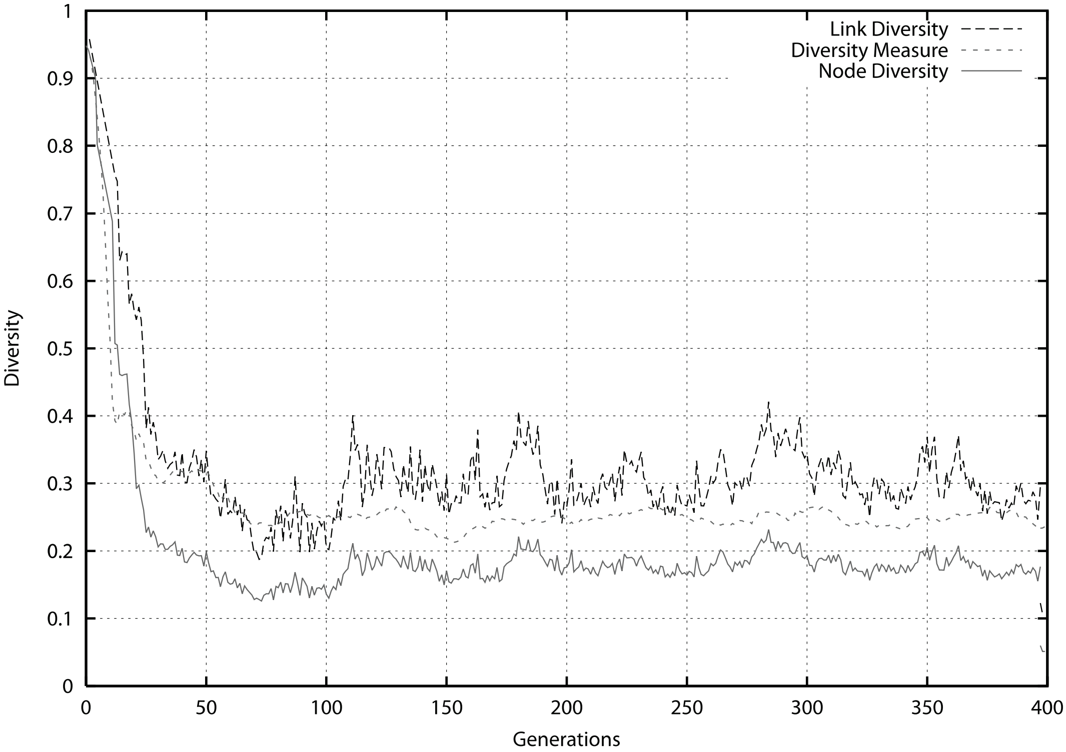

Figure 5 New diversity measure compared with raw node/link count.

is inspired by the work of McQuesten (2002). Specifically, the phenotypic diversity measure examines the actual response of an agent to a given stimulus, rather than relying on pure fitness measures, giving a more meaningful result. Typically, the approach undertaken in this work provides each agent with a number of stimuli to its neural network input layer, records the agent's neural network response and compares all the population's responses using a Euclidian distance measure:

$$distance = \sqrt{(output_a - output_b)^2}.$$

The number of stimuli given to each agent is limited in scope. While it would be possible to examine the response of the entirety of an agent's cognitive capacity, such an approach would be impractical for the more complex problem domains detailed in this work. Instead, the neural network's behavioral response is sampled using a number of fixed stimuli for each member of the population. Thus, the algorithm is comprised of the following steps:

- Take a pair of neural network agents from the population
- Present each neural network with the following patterns:
  - All 0s
  - All 1s
  - Checkerboard 0 (alternating 0s and 1s, beginning with 0)
  - Checkerboard 1 (alternating 0s and 1s, beginning with 1)
- Compare each agent's response using a Euclidian distance measure
- Average the distance measure over the four stimuli

The algorithm is repeated until each agent has been compared with all other agents in the population. The distance measures are averaged to produce a global phenotypic diversity measurement.

3.3 Tasks

The following subsections outline each of the problem tasks assigned to the populations, including the input/output architecture required for each task.

3.3.1 Bit Parity The bit parity problem, also known as the XOR problem, is simple to describe but nontrivial for a neural network to solve. Fundamentally, the bit parity problem requires a neural network to determine whether a given bit string of length *n* contains an odd or even number of 1s (or 0s). For the purposes of this work, 5-bit parity was chosen as the first experiment set, representing a moderately difficult problem for a neural network to solve, as it must distinguish correctly between 32 different bit patterns.

3.3.2 Tic-Tac-Toe Tic-tac-toe, or three in a row, is a very simple two player game played on a 3 × 3 board. Each player is assigned either the X or O symbol and takes turns placing one symbol onto the board at a time. Each player attempts to place three of his/her pieces in a horizontal, vertical or diagonal line of three.

Agents play games against a perfect player, whose minimax engine is modified such that the first move is randomized, allowing agents to play games of some variety. Fitness is assigned according to the length of the game. In other words, agents are rewarded for prolonging the game by forcing the minimax player to a draw. Since it is possible (but difficult) to beat the modified minimax player, a win for the agent is considered equal to a draw.

Each agent's neural network structure contains 18 input nodes, 2 for each board position where 01 is X, 10 is O and 11 is an empty square. Nine output nodes, one corresponding to each board position, are used to indicate the agent's desired move. The node with the strongest response corresponding to a valid move is taken as the agent's choice. The simulator allows agents to evolve any number of hidden layers each with an unrestricted number of nodes, giving maximum flexibility to the evolutionary process.

During the teaching process, a teacher agent plays while the pupil observes. At each move, both the pupil and teacher emit some verbal output in response to the current board position. At every teaching cycle, the pupil's verbal output is corrected with respect to the teacher's using error back-propagation.

3.3.3 Connect-Four The game of connect-four is a two-player game played on a vertical board of 7 × 6 positions into which pieces are slotted from above. Each player is given a number of colored pieces (one color per player) and must attempt to create horizontal, vertical or diagonal piece-lines of length four. Players place one piece per turn into one of the seven slots. The piece then falls down to the lowest free position in the chosen column, creating piles, or towers, of pieces. If a column is full (i.e., when the stack of pieces reaches 7), the slot is no longer available for selection.

As in the tic-tac-toe game, agents play games against a modified perfect player and fitness is assigned as before. At each move, the current board position is taken and the agent's pieces are added iteratively into each slot. At each iteration, the network is shown the board position through 42 input nodes. Unlike the tic-tac-toe game, each neural network has only one output node and the board position which produces the strongest output response is deemed to be the agent's preferred board position.

4 Results

In order to examine the effects of cultural learning, the experiments conducted employ two populations: one using only population learning to evolve, and the other using both population and cultural learning. Populations of 100 agents are evolved for 400 generations. Crossover is set at 0.6, mutation at 0.02 and the number of teaching cycles at five. These values were found empirically to produce the best results.

When cultural learning is applied, teachers are selected from the previous generation. As generation 1 does not yet have teachers available, it cannot perform cultural learning and therefore employs population learning alone. For this reason, results presented for the cultural learning population begin at generation 2, the point where cultural learning is first applied.

Results shown are averages of 20 independent runs and error bars represent 1 standard deviation from the mean.

The following subsections outline results obtained for the population's fitness and diversity measures.

4.1 Fitness

The first set of results examined show both the acquired and innate fitness levels of the two populations. Acquired fitness is measured at the end of an agent's life, after any cultural learning is applied. Innate fitness is measured at the beginning of the agent's lifetime and corresponds to the agent's genotypic performance. Naturally, innate fitness for the population employing population learning alone corresponds to its acquired fitness since no cultural learning occurs.

4.1.1 Bit Parity The first task to be attempted by the two populations is the 5-bit parity problem. Each agent in the population is presented with a 5-bit binary pattern and must determine whether the number of ones is odd or even (and then respond with either a one or a zero through its output node).

Figure 6 shows the average acquired fitness for the two populations over the course of the experiment. It is clear from these results that while population learning is capable of steady (if slow) progress towards an optimal solution, cultural learning greatly enhances the population's performance.

The innate fitness of the two populations is presented in Figure 7 along with the acquired fitness values for the cultural learning population, for comparison purposes. The figure shows the average fitness of the population employing population learning and the fitness values for the cultural learning population before and after teaching is applied.

The results show an interesting trend: While the three curves begin close together, the cultural learning population's innate fitness quickly deteriorates and becomes considerably lower than that of the population-learning population. However, when teaching is applied, the cultural learning population's fitness improves dramatically, easily out-performing population learning.

4.1.2 The Game of Tic-Tac-Toe The second task to be attempted by each population is the game of tic-tac-toe. Agents play a number of games and fitness is calculated according to their performance. Figure 8 shows the average acquired fitness of both populations throughout the experiment run.

It is clear from these results that the population employing cultural learning outperforms the population employing population learning alone from the start of the experiment. While the population-learning population stabilises at around 0.825, the population employing cultural learning achieves fitness values of 0.9 and above.

Figure 9 shows three fitness values: one for the population employing population learning alone, one

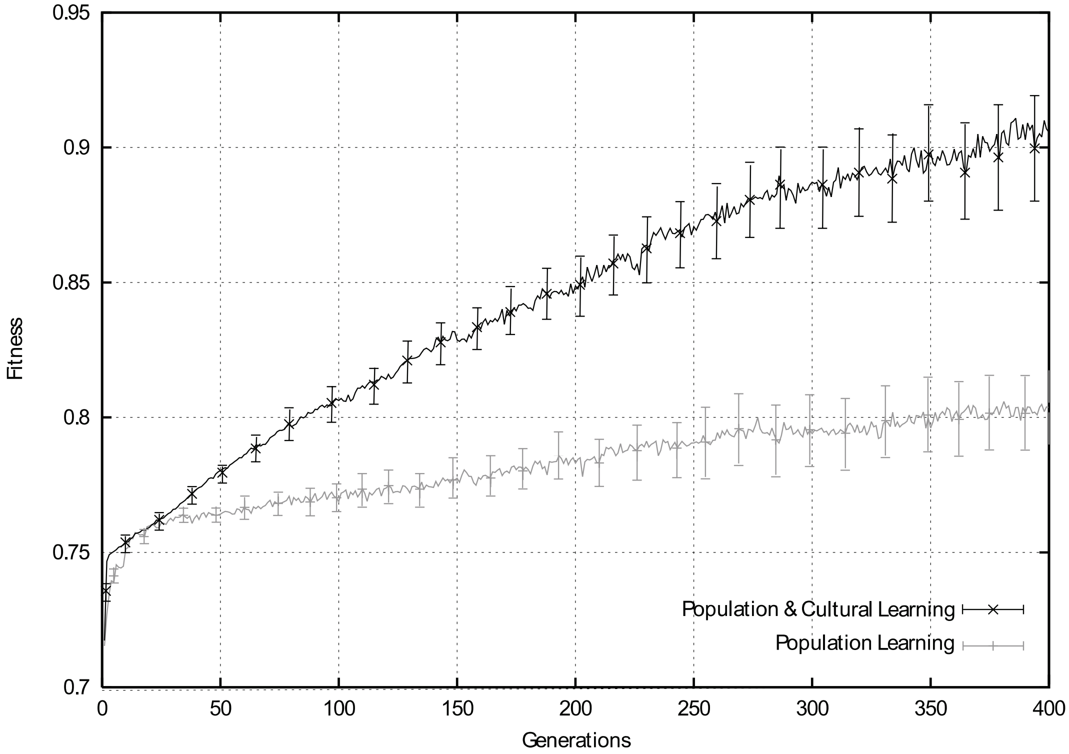

Figure 6 Bit parity average fitness.

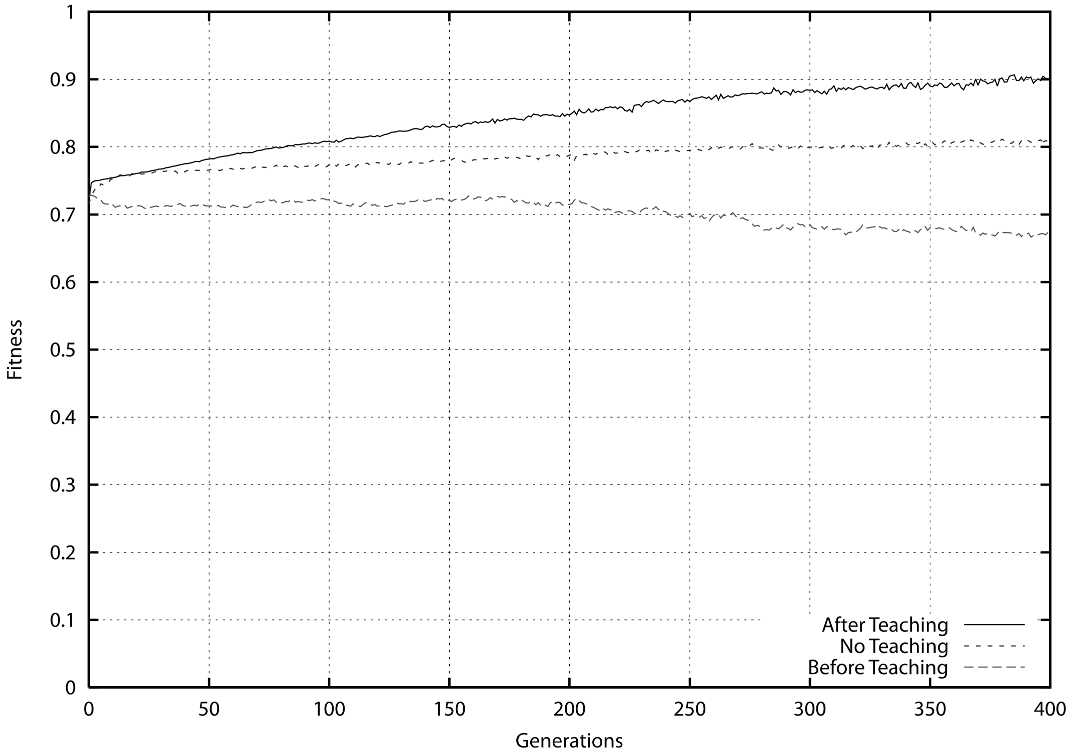

Figure 7 Bit parity average fitness before and after teaching.

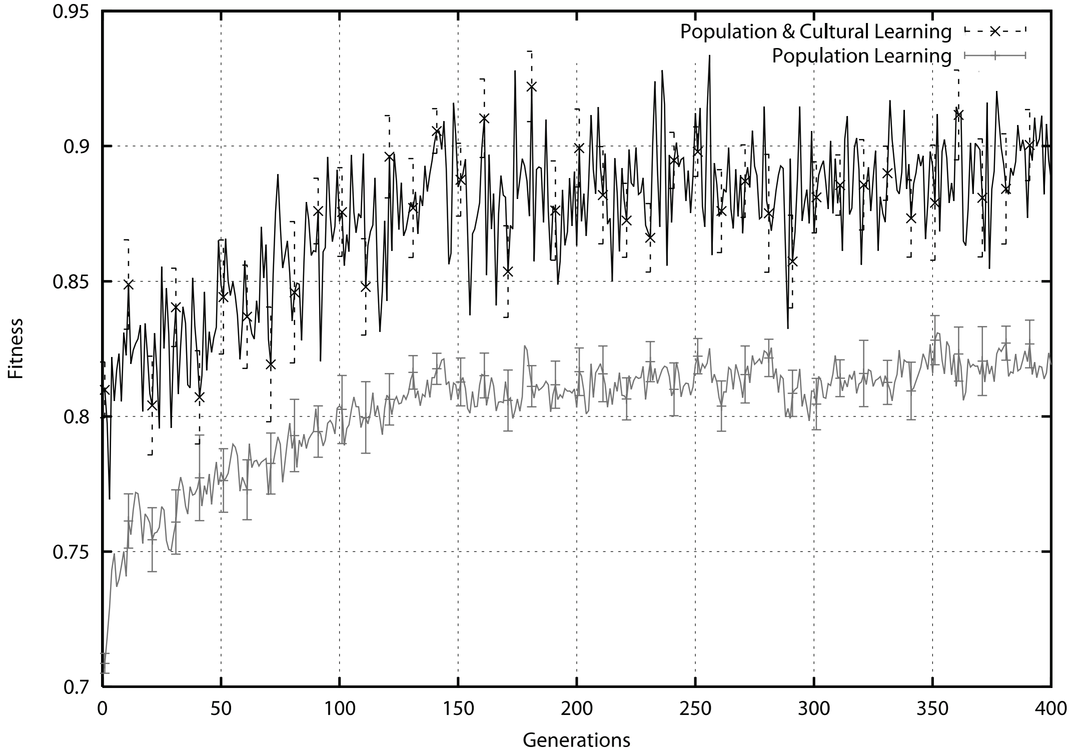

Figure 8 Tic-tac-toe average fitness.

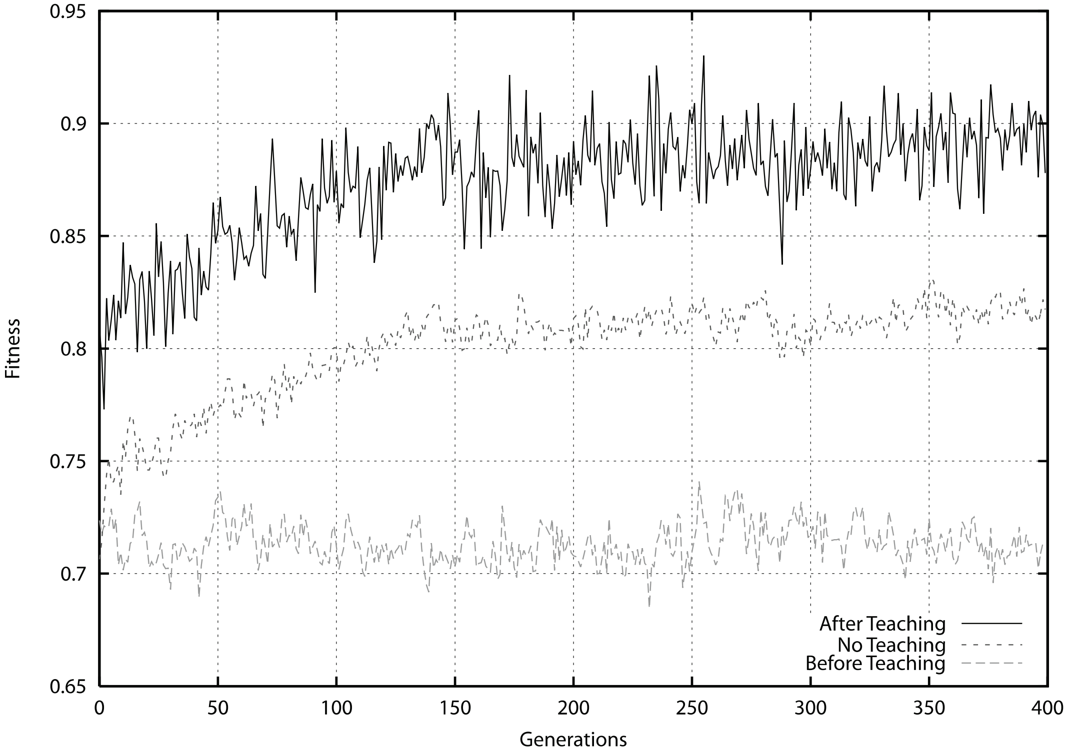

Figure 9 Tic-tac-toe average fitness before and after teaching.

for the cultural learning population prior to teaching (innate fitness) and the last showing the cultural learning population after teaching is applied.

The population employing cultural learning performs very differently before and after teaching is applied. Prior to teaching, the cultural learning population's fitness is considerably lower than that of the population learning population. Indeed, the population's genotypic fitness (the fitness measured before any cultural influence is applied) is consistently low and appears to be stable throughout the experiment.

4.1.3 The Game of Connect Four The final task selected for this set of experiments is the game of connect-four. The results illustrated in Figure 10 clearly show the effect of cultural learning on the population. The population employing population learning alone achieves fitness levels of around 0.6 compared with close to 0.8 for the cultural learning population.

Figure 11 shows the fitness values for the population employing population learning alone, the cultural learning population prior to teaching (innate fitness) and the cultural learning population after teaching takes place. Once again, as in previous experiments, cultural learning appears to be selecting individuals for their genetic ability to learn, rather than for their innate ability to solve a particular task. This is illustrated by the fact that the fitness values for the population employing cultural learning are considerably lower prior to teaching than those of the population employing population learning alone.

4.2 Diversity

This section outlines the results obtained for the diversity measures employed throughout these experiments. The section is divided into two subsections. The first, genotypic diversity, examines the differences between individuals at a genetic level. The second, phenotypic diversity, examines the behavioral differences between individuals in the population. Again, results presented are averaged from 20 independent runs.

4.2.1 Genotypic Diversity The genotypic diversity measure examines the genetic makeup of each indi-

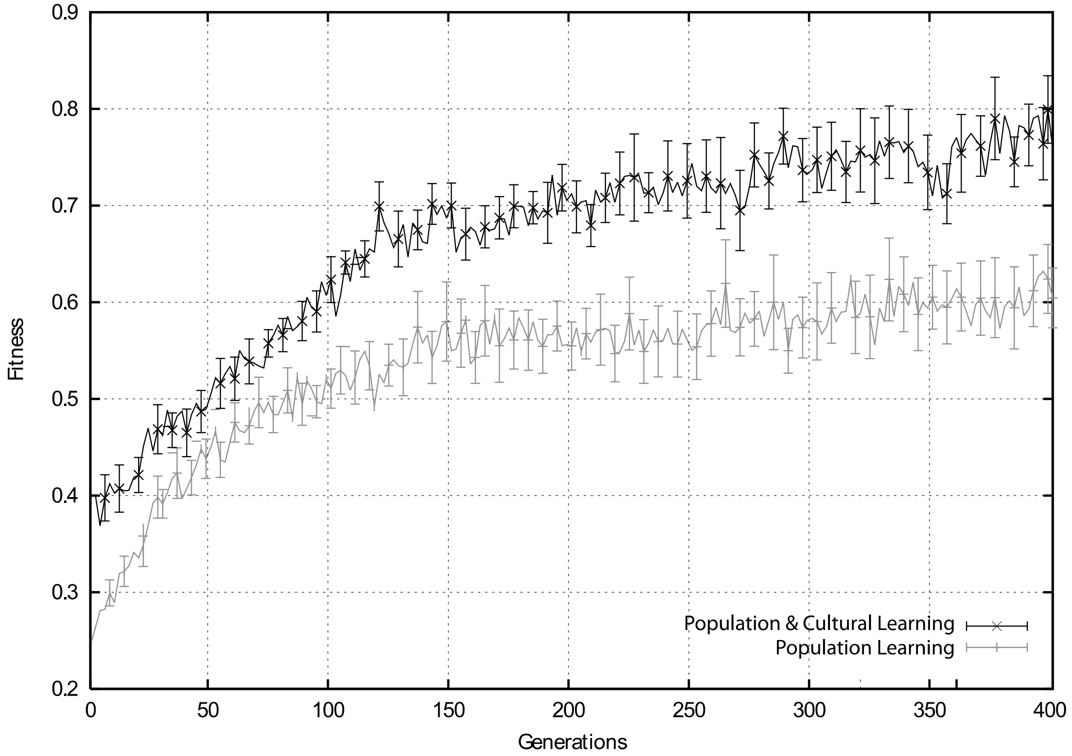

Figure 10 Connect-four average fitness.

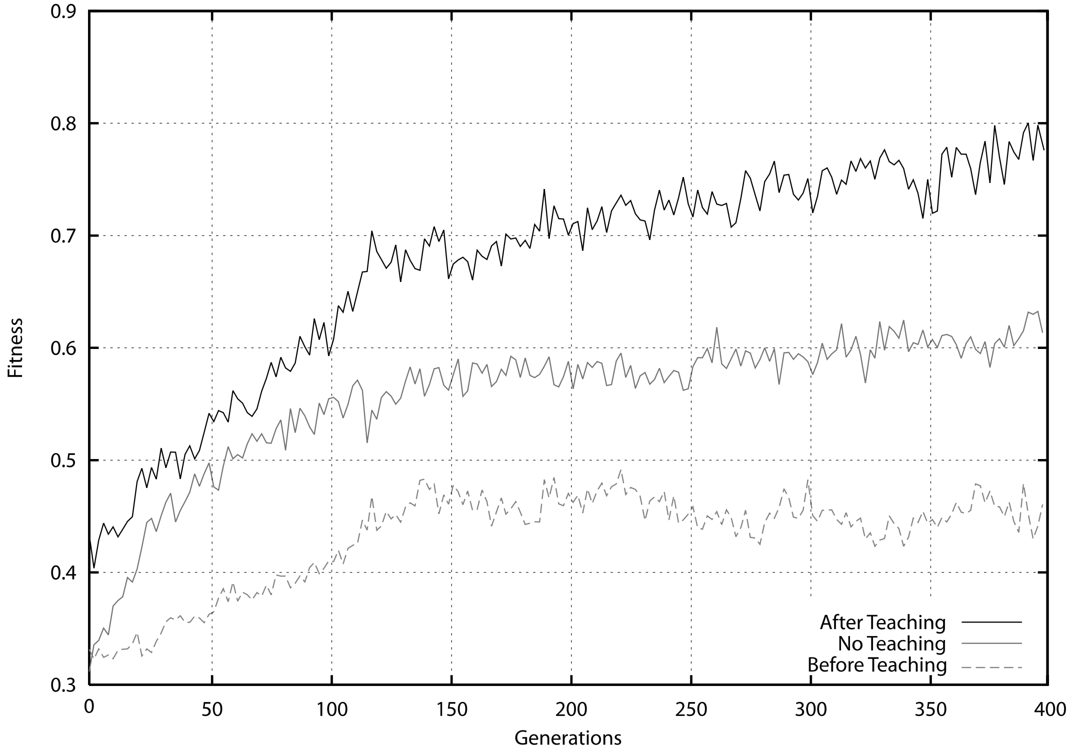

Figure 11 Connect-four average fitness before and after teaching.

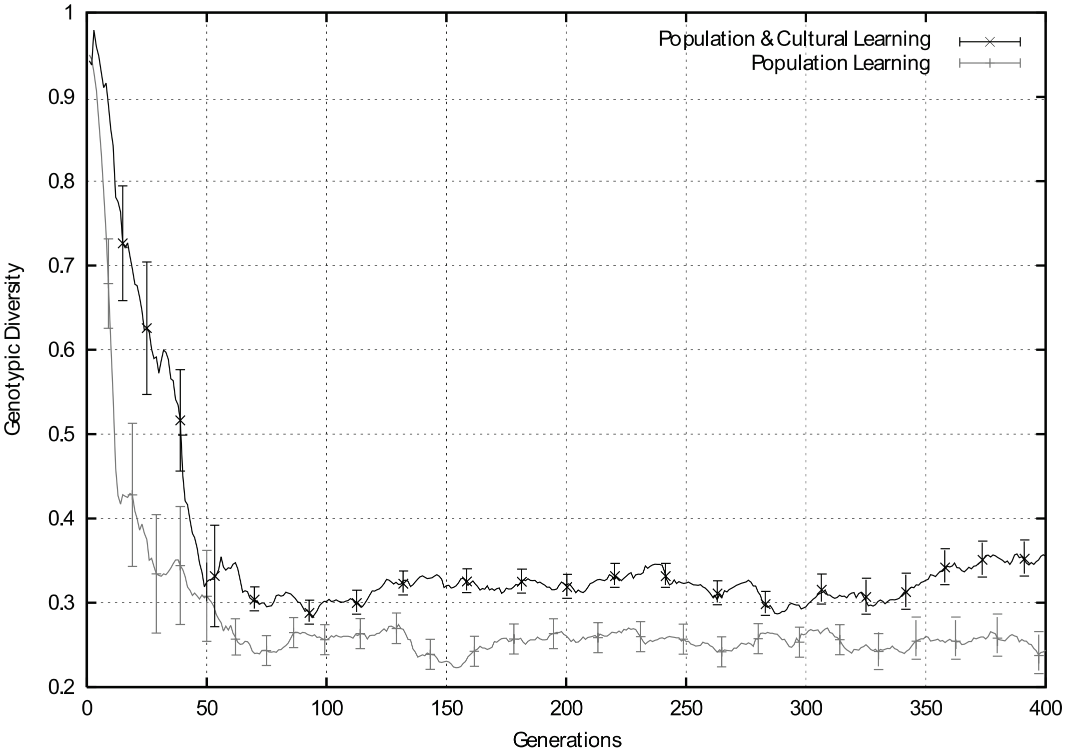

Figure 12 Bit parity genotypic diversity.

vidual agent and compares it with all other agents in the population. The measure is normalized to a value in the range [0, 1], where 1 is highly diverse and 0 is highly similar.

Bit Parity Given the relatively low complexity of the bit parity problem, one would expect that the evolutionary process would quickly converge towards a genotypic architecture suitable for the problem. This is illustrated by Figure 12 which displays the results obtained from the genotypic diversity measure for the population employing population learning and the population employing both population and cultural learning. Both populations quickly converge, showing a sharp drop in genotypic diversity within the first 100 generations, implying that the evolutionary process has selected what it considers to be an optimal agent neural network architecture for the problem.

However, while both populations show a similar overall trend the population employing cultural learning maintains a higher genotypic diversity level throughout the experiment. While the difference is not large, it is significant and shows that cultural learning is contributing to population diversity, something which may help explain its higher level of overall fitness.

Table 1 shows the average, maximum and minimum diversity values for both populations, taken as an average over the whole experiment set. Again, it shows that cultural learning is generating a higher level of average genotypic diversity. The minimum diversity value for the cultural learning population is 0.28, compared with 0.22 for population learning. This implies

Table 1 Bit parity average genotypic diversity.

Population	Average diversity	Maximum diversity	Minimum diversity	SD
Population learning	0.279833215	0.9496164	0.2230475	0.00941449
Cultural learning	0.360330611	0.9791067	0.2825165	0.01622283

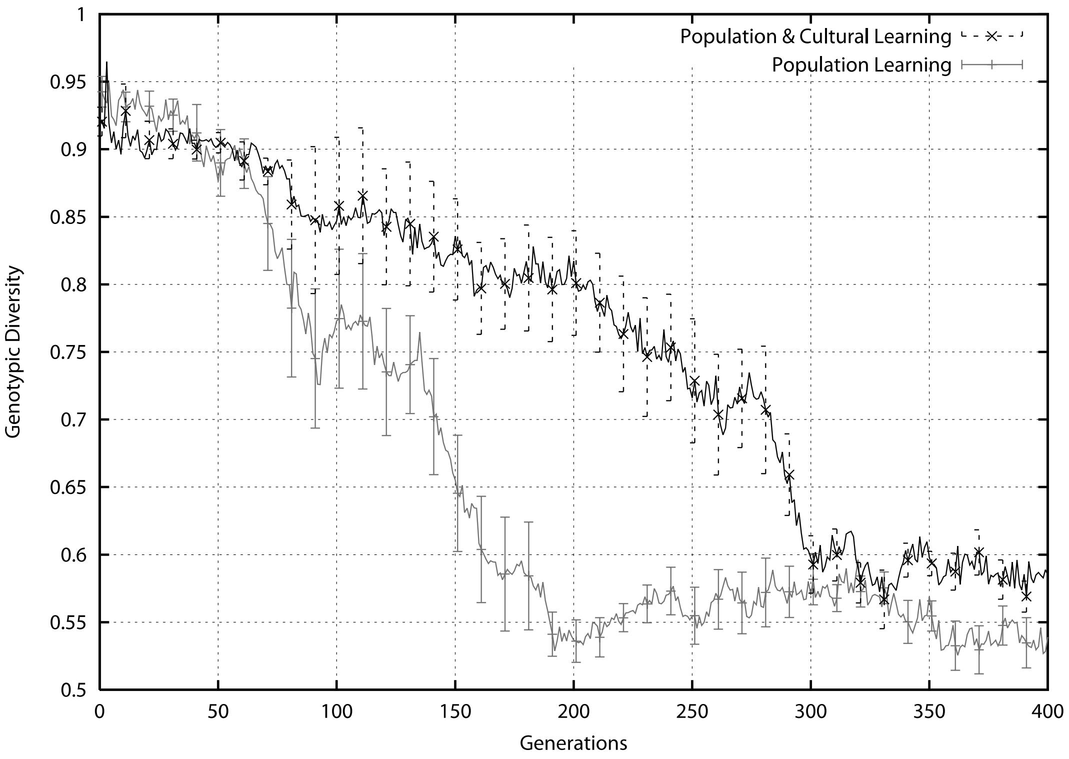

Figure 13 Tic-tac-toe genotypic diversity.

that cultural learning is maintaining a higher level of population diversity throughout the experiment run. The results are statistically significant to a 95% confidence level (p value < 0.0001, paired t-test).

Tic-Tac-Toe Figure 13 shows the results for the genotypic diversity measure throughout the experiment run. The tic-tac-toe experiment follows the pattern found in the bit parity experiment run, that is, both populations begin at a high level of diversity but diverge in behavior after generation 100. In general, diversity is maintained for a considerably longer period in both populations when compared with the results obtained for the parity problem. This is most likely due to the increased complexity of the problem as a more complex network architecture must be found to solve the problem adequately.

It is clear from the population's divergence in diversity following generation 100 that cultural learning is aiding the population in maintaining a higher genotypic diversity than population learning alone. The cultural learning population converges at generation 300, compared with generation 150 for population learning alone.

Cultural learning is capable of maintaining a high level of diversity for the majority of the experiment run, only converging in the last 100 generations, which may explain its superior performance compared with population learning alone. Table 2 shows the average, maximum and minimum genotypic diversity values

Table 2 Tic-tac-toe average genotypic diversity.

Population	Average diversity	Maximum diversity	Minimum diversity	SD
Population learning	0.662021098	0.9496164	0.5255245	0.019489483
Cultural learning	0.75848603	0.9647170	0.5668700	0.013715094

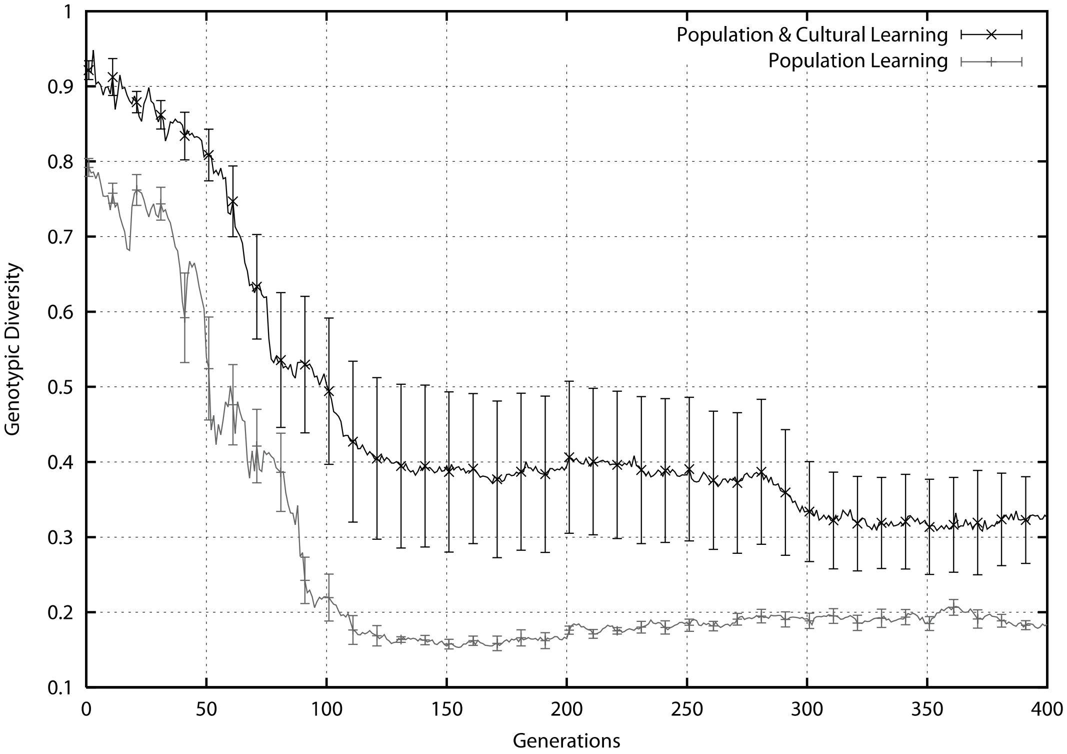

Figure 14 Connect-four genotypic diversity.

taken over the experiment run. While both populations have similar maximum and minimum diversity values, it is clear from the average diversity value that cultural learning is providing consistently higher diversity values than population learning alone. The similar maximum and minimum values are merely indicative of the fact that both populations begin at a similar high level of diversity and eventually converge at a similar low level. The results are statistically significant to a 95% confidence level (*p* value < 0.0001, paired t-test).

Connect-Four Figure 14 shows the results obtained for the genotypic diversity measure for both populations over the course of the experiment. As in the previous experiments, both populations begin the experiment with very high levels of genotypic diversity, consistent with what one would expect from randomly generated neural network populations. By generation 100, the two populations have clearly diverged with respect to their diversity levels. The cultural learning population is maintaining consistently higher levels of genotypic diversity than the population employing population learning alone.

Table 3 shows the average, maximum and minimum values for both populations, taken over the whole experiment run. It is evident from these figures that cultural learning is enhancing the population's average, maximum and minimum diversity values. This result is consistent with the previous experiment sets, although the maximum diversity levels are considerably different. Since this problem is considerably more complex to solve than the previous experiments, it is likely that

Table 3 Connect-four average genotypic diversity.

Population	Average diversity	Maximum diversity	Minimum diversity	SD
Population learning	0.271538822	0.7920000	0.1531017	0.03338738
Cultural learning	0.461263027	0.9482010	0.3080865	0.033095317

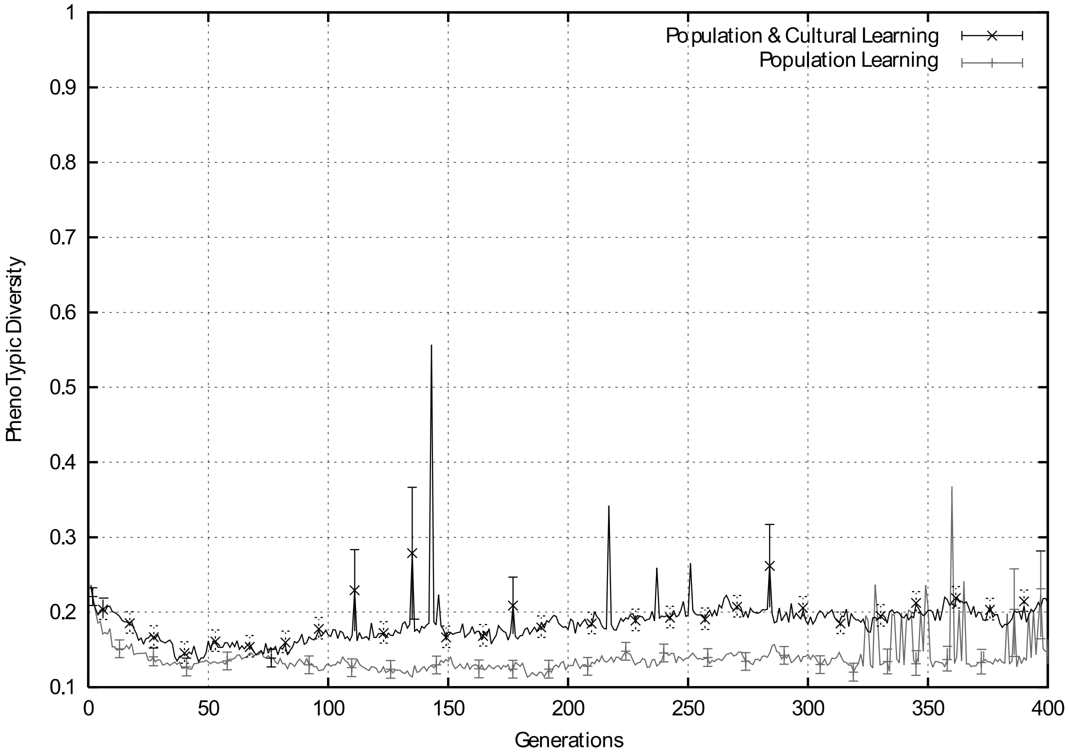

Figure 15 Bit parity phenotypic diversity.

the population employing cultural learning is exhibiting its effects early in the experiment, illustrated by the significantly higher diversity values towards the beginning of the experiment.

4.2.2 Phenotypic Diversity Phenotypic diversity examines the difference in behavior between agents in a population by observing the outputs for a given number of environmental stimuli. The measure gives values in the range [0, 1], where 1 is completely diverse and 0 is completely similar.

Bit Parity Figure 15 shows the results obtained from the phenotypic diversity measure for both populations throughout the experiment run. The populations are effectively equivalent in terms of phenotypic diversity for the first 100 generations, but then begin to diverge, with cultural learning increasing its phenotypic diversity level and maintaining a higher level throughout the rest of the experiment.

Table 4 shows the average, maximum and minimum diversity values for both populations taken over the entire experiment. It is clear from these figures that cultural learning produces a higher level of phenotypic diversity and is capable of higher maximum and minimum diversity values. The results are statistically significant to a 95% confidence level (p value < 0.0001, paired t-test).

Tic-Tac-Toe Figure 16 shows the results obtained for this measure for both populations for the game of tic-

Table 4 Bit parity average phenotypic diversity.

Population	Average diversity	Maximum diversity	Minimum diversity	SD
Population learning	0.139159874	0.3615294	0.1133413	0.000521302
Cultural learning	0.185220491	0.5565765	0.1338845	0.000846548

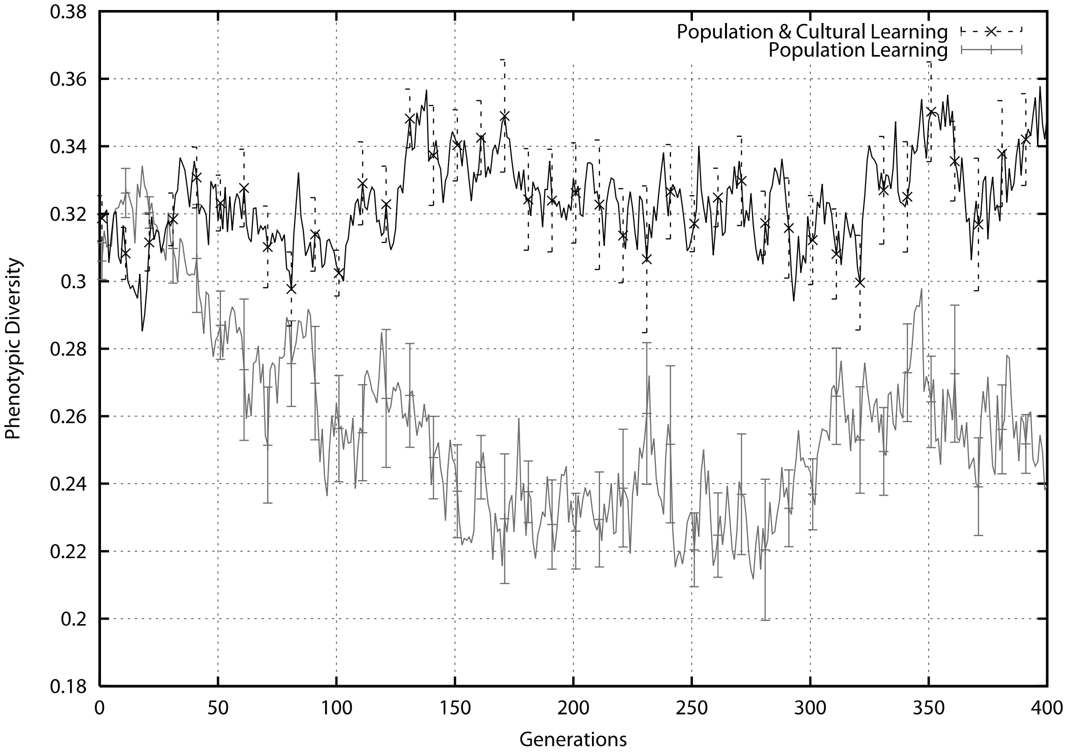

Figure 16 Tic-tac-toe phenotypic diversity.

Table 5 Tic-tac-toe average phenotypic diversity.

Population	Average diversity	Maximum diversity	Minimum diversity	SD
Population learning	0.25763194	0.3341801	0.2117307	0.00074964
Cultural learning	0.324021795	0.3578060	0.2851217	0.00016451

tac-toe. Both populations begin, as expected, with a similar level of phenotypic diversity but diverge after generation 100. Following this, the cultural learning population exhibits a significantly higher phenotypic diversity than does the population employing population learning alone.

Table 5 shows the average, maximum and minimum phenotypic diversity values obtained for the entire experiment run. The similar maximum diversity values for the two populations are indicative of the fact that both begin at a similar level of phenotypic diversity. However, the considerably different minimum diversity measurements obtained by the two populations show that the cultural learning population ends the experiment with a high level of phenotypic diversity. This, coupled with the significantly higher average phenotypic diversity, shows that cultural learning is providing the population with the ability to generate and sustain novel behaviors in its population throughout the experiment run.

Connect-Four Figure 17 shows the results obtained from the phenotypic diversity measure for both populations for the game of connect-four. The populations show a similar trend to the last experiment, as they exhibit very similar levels of phenotypic diversity for the first half of the experiment, only diverging after generation 200. Following this, the cultural learning population shows increased phenotypic diversity, rising until the end of the experiment, while the popula-

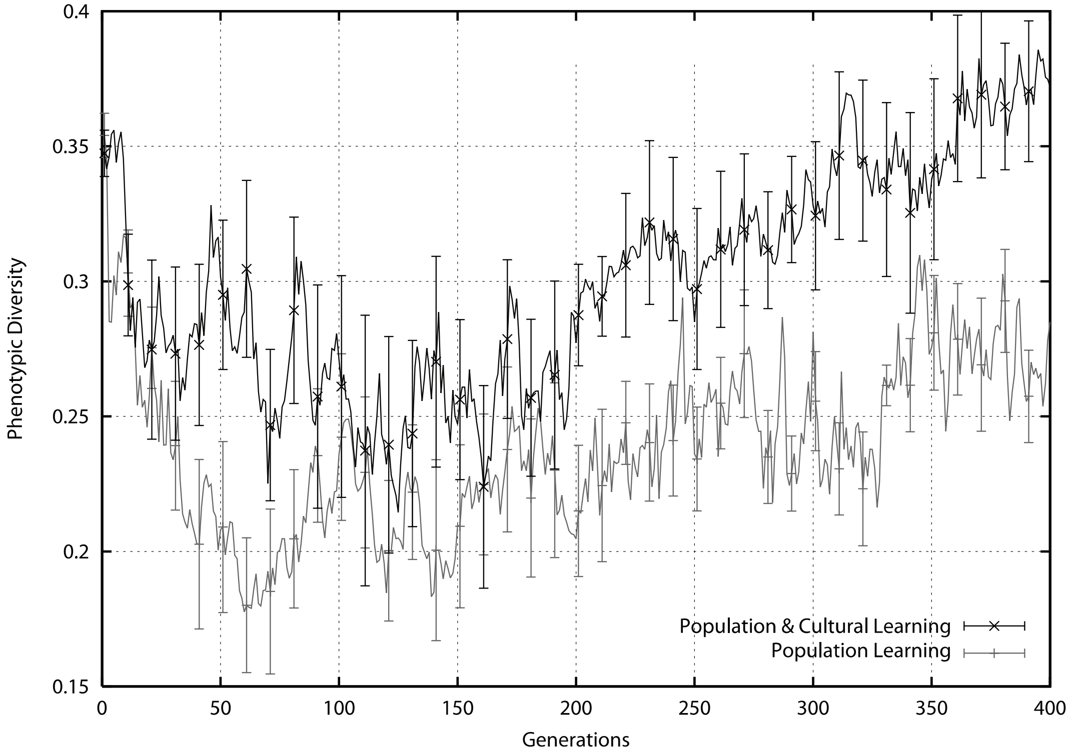

Figure 17 Connect-four phenotypic diversity.

Table 6 Connect-four average phenotypic diversity

Population	Average diversity	Maximum diversity	Minimum diversity	SD
Population learning	0.238738262	0.3540330	0.1776206	0.000935385
Cultural learning	0.301066895	0.3857076	0.2144768	0.00166501

tion employing only population learning shows significantly lower levels.

Table 6 shows the average, maximum and minimum phenotypic diversity levels obtained over the experiment run. It is evident from these results that cultural learning is capable of producing populations with significantly higher levels of phenotypic diversity than population learning alone. The results are statistically significant to a 95% confidence level (p value < 0.0001).

5 Discussion

There are a number of trends that can be observed by examining the results obtained for these three benchmark problems, some of which have already been analyzed in previous research and some which are novel.

Firstly, it is clear that the cultural learning process is producing two distinct effects on the population. One effect is to improve the population's average acquired fitness. The selection process selects individuals in the cultural learning population based on their fitness after teaching is applied and not on the more highly coupled geno-phenotypic mapping employed by population learning. As a result, the cultural process selects individuals based on their phenotypic ability and not on their underlying genetic make-up.

This leads to the second effect, which becomes evident once the innate fitness of the cultural learning population is examined. The innate fitness of the population (measured before any cultural learning is applied) shows significant deterioration when compared with both the acquired fitness and population learning values. Thus, it appears that the population's innate fitness falls significantly during the course of each experiment, regardless of the task.

As the selection process examines individuals at the end of their lifetimes (when cultural learning has already been applied), the underlying genetic affinity to a problem is effectively hidden. However, the evolutionary process must still ensure that individuals in a population are in a position to receive instruction from others in the population. Thus, the genotype of each agent is being selected on the basis of an affinity to learning, rather than a direct affinity to the problem itself.

This result is related to previous work: Nolfi et al. Nolfi and Parisi (1993, 1994) found that their evolved networks were more suited to learning than solving the assigned task, a result that was not the intention of the experimenters. Sasaki and Tokoro (1997) also found that the ability to learn can be more important than the ability to solve a task.

The fact that learning should be important to a population is not in itself surprising. However, the fact that these effects occur with some regularity in such a variety of experiment settings and tasks suggests that it is a more general trend worthy of further analysis. We argue that the results obtained using our diversity measures can provide further insight into this area.

The diversity measures employed examine both genotypic and phenotypic diversity in an attempt to further understand the effects of the cultural process on the population. It is clear that the cultural learning process is producing populations that are both genotypically and phenotypically more diverse than those using only population learning. Increased diversity is generally accepted as a desirable feature of an evolutionary algorithm, so its presence may explain the increase in fitness exhibited by the cultural learning process.

We posit that this increase shows that the selection process is no longer selecting individuals based on their genetic predisposition to solving a particular task. Instead, an individual's genotype is masked by the cultural process, evidenced further by the fact that while the population's acquired fitness may be high, its innate fitness suffers considerably.

The population increasingly relies on the cultural process to produce fit individuals and thus the engine of evolutionary progress is transferred to the culture which is passed down through the generations, resulting in a prolonged retention of genotypic diversity within the population. Phenotypic diversity is also increased as a result, and populations begin to fill with individuals that are both genetically and behaviorally different from one another. Thus, cultural learning's ability to produce fitter populations may be credited to both its ability to transmit information acquired through individual lifetimes and its ability to retain high population diversity.

6 Summary

This work focuses on the effects of the cultural learning process on a population of evolving neural network agents. A population employing population learning alone and a population employing both population and cultural learning are tested using three benchmark tasks: 5-bit parity, the game of tic-tac-toe and the game of connect-four. In addition to studying the effects of cultural learning on the acquired and innate fitness of the populations, we employed both genotypic and phenotypic diversity measures to provide further insight into the cultural learning process.

It is clear from this and previous work that the addition of cultural learning provides an increase in the performance of the population. In addition, the results show that the cultural learning process effectively masks the population's underlying genetic affinity to its environment and that the population comes to increasingly rely on cultural learning for its survival. Furthermore, the results show that cultural learning is accompanied by increased levels of diversity within the population, a contribution which may help explain its improved performance.

Future work will examine the effect of cultural learning on both fitness and diversity in a dynamic environment. In particular, we will focus on the development of a more formal model in the hopes of showing the effects of cultural learning on population diversity in a more general sense. Finally, future work will include the development of more varied and expressive measures of both genotypic and phenotypic diversity.

Acknowledgments

We wish to thank the reviewers for their many constructive comments and suggestions. The first author would also like to acknowledge the support of the Irish Research Council for Science, Engineering and Technology.

References

Baldwin, J. (1896). A new factor in evolution. *American Naturalist, 30*, 441–451.

Belew, R. K. (1990). Evolution, learning and culture: Computational metaphors for adaptive algorithms. *Complex Systems, 4*, 11–49.

Belew, R. K., & Booker, L. B. (Eds.). (1991). *Proceedings of the 4th International Conference on Genetic Algorithms*. San Diego, CA: Morgan Kaufmann.

Belew, R. K., McInerney, J., & Schraudolph, N. N. (1992). Evolving networks: Using the genetic algorithm with connectionist learning. In C. G. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds.), *Artificial life II* (pp. 511–547). Redwood City, CA: Addison-Wesley.

Billard, A., & Hayes, G. (1997). Learning to communicate through imitation in autonomous robots. In *7th International Conference on Artificial Neural Networks* (p. 763–738).

Blackmore, S. (2000). The power of memes. *Scientific American, 284* (4), 52–61.

Booker, L. B. (1985). Improving the performance of genetic algorithms in classifier systems. In *Proceedings of the International Conference on Genetic Algorithms and Their Applications* (pp. 80–92). Pittsburgh, PA.

Borenstein, E., & Ruppin, E. (2003). Enhancing autonomous agents evolution with learning by imitation. *Interdisciplinary Journal of Artificial Intelligence and the Simulation of Behaviour, 1*, 335–348.

Boyd, R., & Richerson, P. (1985). *Culture and the evolutionary process*. University of Chicago Press.

Bull, L., Holland, O., & Blackmore, S. (2000). On meme–gene coevolution. *Artificial Life, 6*, 227–235.

Burke, E. K., Gustafson, S., & Kendall, G. (2004). Diversity in genetic programming: an analysis of measures and correlation with fitness. *IEEE Transactions on Evolutionary Computation 8*, 47–62.

Cangelosi, A. (1999). Evolution of communication using combination of grounded symbols in populations of neural networks. In *Proceedings of IJCNN99 International Joint Conference on Neural Networks* (Vol. 6, pp. 4365–4368). Washington, DC: IEEE Press.

Cangelosi, A., & Parisi, D. (1996). *The emergence of a language in an Evolving population of neural network* (Technical Report NSAL–96004). National Research Council, Rome.

Cangelosi, A., & Parisi, D. (1998). The emergence of a language in an evolving population of neural networks. *Connection Science, 10*, 83–97.

Collins, R. J., & Jefferson, D. R. (1991). Selection in massively parallel genetic algorithms. In R. K. Belew & L. B. Booker (Eds.), *ICGA* (pp. 249–256). San Diego, CA: Morgan Kaufmann.

Curran, D., & O'Riordan, C. (2003a). Artificial life simulation using marker based encoding. In *Proceedings of the 2003 International Conference on Artificial Intelligence (IC-AI 2003)* (Vol. II, pp. 665–668). Las Vegas, NV.

Curran, D., & O'Riordan, C. (2003b). On the design of an artificial life simulator. In V. Palade, R. J. Howlett, & L. C. Jain (Eds.), *Proceedings of the Seventh International Conference on Knowledge-Based Intelligent Information & Engineering Systems (KES 2003)* (pp. 549–555). University of Oxford, UK.

Curran, D., & O'Riordan, C. (2004). The effect of noise on the performance of cultural evolution. In *Proceedings of the 2004 Congress on Evolutionary Computation (CEC 2004).* Portland, OR.

Davidor, Y. (1991). A naturally occurring niche and species phenomenon: The model and first results. In R. K. Belew & L. B. Booker (Eds.), *ICGA* (pp. 257–263). San Diego, CA: Morgan Kaufmann.

Deb, K., & Goldberg, D. E. (1989). An investigation of niche and species formation in genetic function optimization. *In Proceedings of the Third International Conference on Genetic Algorithms* (pp. 42–50). San Francisco, CA: Morgan Kaufmann.

De Jong, K. A. (1975). *Analysis of behavior of a class of genetic adaptive systems*. Unpublished doctoral dissertation, The University of Michigan.

Denaro, D., & Parisi, D. (1996). Cultural evolution in a population of neural networks. In M. Marinaro & R. Tagliaferri (Eds.), *Neural nets wirn-96* (pp. 100–111). New York: Springer.

Fahlmann, S. E. (1991). The recurrent cascade-correlation architecture. In D. S. Touretzky (Ed.), *Advances in neural information processing systems 3 (NIPS'91)* (pp. 190–196). Morgan Kaufmann.

Floreano, D., & Mondada, F. (1989). Evolution of plastic neurocontrollers for situated agents. *Animals to Animats, 4*.

Goldberg, D. E., & Richardson, J. (1987). Genetic algorithms with sharing for multimodal function optimization. In *Proceedings of the Second International Conference on Genetic Algorithms and Their Application* (pp. 41–49). Mahwah, NJ: Erlbaum.

Gruau, F. (1994). *Neural network synthesis using cellular encoding and the genetic algorithm.* Unpublished doctoral dissertation, Centre d'etude nucleaire de Grenoble, Ecole Normale Superieure de Lyon, France.

Gusfield, D. (1997). *Algorithms on strings, trees and sequences: computer science and computational biology.* Cambridge University Press.

Hillis, W. D. (1990). Co-evolving parasites improve simulated evolution as an optimization procedure. In *CNLS '89: Proceedings of the Ninth Annual International Conference of the Center for Nonlinear Studies on Self-Organizing, Collective, and Cooperative Phenomena in Natural and Arti-*

ficial Computing Networks on Emergent Computation (pp. 228–234). Amsterdam, The Netherlands: North-Holland.

Hinton, G. E., & Nowlan, S. J. (1987). How learning guides evolution. *Complex Systems, 1*, 495–502.

Hutchins, E., & Hazlehurst, B. (1991). Learning in the cultural process. In *Artificial life II* (pp. 689–706). Cambridge, MA: MIT Press.

Hutchins, E., & Hazlehurst, B. (1995). How to invent a lexicon: The development of shared symbols in interaction. In N. Gilbert & R. Conte (Eds.), *Artificial societies: The computer simulation of social life* (pp. 157–189). London: UCL Press.

Kawamura, S. (1963). The process of sub-culture propagation among Japanese macaques. *Primates Social Behaviour*, 82–90.

Kitano, H. (1990). Designing neural networks using genetic algorithm with graph generation system. *Complex Systems, 4*, 461–476.

Kolen, J. F., & Pollack, J. B. (1991). Back propagation is sensitive to initial conditions. *Advances in Neural Information Processing Systems, 3*, 860–867.

Koza, J. R., & Rice, J. P. (1991, 8–12). Genetic generation of both the weights and architecture for a neural network. In *International Joint Conference on Neural Networks, IJCNN-91* (Vol. II, pp. 397–404). Seattle, WA: IEEE Computer Society Press.

MacLennan, B., & Burghardt, G. (1993). Synthetic ethology and the evolution of cooperative communication. *Adaptive Behavior, 2*, 161–188.

Mahfoud, S. W. (1995a). A comparison of parallel and sequential niching methods. In L. J. Eshelman (Ed.), *Proceedings of the Sixth International Conference on Genetic Algorithms* (pp. 136–143).

Mahfoud, S. W. (1995b). *Niching methods for genetic algorithms*. Unpublished doctoral dissertation, University of Illinois.

Mandischer, M. (1993). Representation and evolution of neural networks. In R. F. Albrecht, C. R. Reeves, & N. C. Steele (Eds.), *Artificial neural nets and genetic algorithms proceedings of the international conference at Innsbruck, Austria* (pp. 643–649). Wien and New York: Springer.

McPhee, N., & Hopper, N. (1999). Analysis of genetic diversity through population history. In *Proceedings of the Genetic and Evolutionary Computation Conference* (pp. 1112–1120). Morgan Kaufmann.

McQuesten, P. H. (2002). *Cultural enhancement of neuroevolution*. Unpublished doctoral dissertation, University of Texas, Austin.

Miller, G. F., Todd, P. M., & Hedge, S. U. (1989). Designing neural networks using genetic algorithms. In *Proceedings of the Third International Conference on Genetic Algorithms and Their Applications* (pp. 379–384).

Moriarty, D. E., & Miikkulainen, R. (1995). Discovering complex othello strategies through evolutionary neural networks. *Connection Science, 7*, 195–209.

Muhlenbein, H. (1989). Parallel genetic algorithms, population genetics and combinatorial optimization. In *Proceedings of the Third International Conference on Genetic Algorithms* (pp. 416–421). San Francisco, CA: Morgan Kaufmann.

Nolfi, S., & Parisi, D. (1993). Self-selection of input stimuli for improving performance. In *Neural networks in robotics* (pp. 403–418). Kluwer.

Nolfi, S., & Parisi, D. (1994). Desired answers do not correspond to good teaching inputs in ecological neural networks. *Neural processing letters 1* (2), 1–4.

Nolfi, S., & Parisi, D. (1996). Learning to adapt to changing environments in evolving neural networks. *Adaptive Behavior, 5*, 75–97.

Nolfi, S., Parisi, D., & Elman, J. L. (1994). Learning and evolution in neural networks. *Adaptive Behavior, 3*, 5–28.

O'Reilly, U. M. (1997). Using a distance metric on genetic programs to understand genetic operators. In *IEEE International Conference on Systems, Man, and Cybernetics* (Vol. 5, pp. 4092–4097).

Pereira, F. B., & Costa, E. (2001). How learning improves the performance of evolutionary agents: A case study with an information retrieval system for a distributed environment. In *Proceedings of the International Symposium on Adaptive Systems: Evolutionary computation and probabilistic graphical models (ISAS 2001)* (pp. 19–23).

Pujol, J. C. F., & Poli, R. (1998). Efficient evolution of asymmetric recurrent neural networks using a two-dimensional representation. In *Proceedings of the First European Workshop on Genetic Programming (EUROGP)* (pp. 130–141).

Rosca, J. P. (1995). Entropy-driven adaptive representation. In J. P. Rosca (Ed.), *Proceedings of the Workshop on Genetic Programming: From theory to real-world applications* (pp. 23–32).

Sasaki, T., & Tokoro, M. (1997). Adaptation toward changing environments: Why Darwinian in nature? In P. Husbands & I. Harvey (Eds.), *Fourth European Conference on Artificial Life* (pp. 145–153). Cambridge, MA: MIT Press.

Spector, L. (1994). Genetic programming and AI planning systems. In *Proceedings of the Twelfth National Conference on Artificial Intelligence* (pp. 1329–1334). Seattle, WA: AAAI Press/MIT Press.

Spector, L., & Luke, S. (1996). Culture enhances the evolvability of cognition. In *Proceedings of the Eighteenth Annual Conference of the Cognitive Science Society* (pp. 672–677).

Spiessens, P., & Manderick, B. (1991). A massively parallel genetic algorithm: Implementation and first analysis. In R. K. Belew & L. B. Booker (Eds.), *ICGA* (pp. 279–287). San Diego, CA: Morgan Kaufmann.

Steels, L. (1996a). Emergent adaptive lexicons. In P. Maes (Ed.), *Proceedings of the Simulation of Adaptive Behavior Conference*. Cambridge, MA: MIT Press.

Steels, L. (1996b). Self-organising vocabularies. In *Proceedings of Artificial Life V*.

Steels, L. (1997). The synthetic modeling of language origins. In *Evolution of communication* (pp. 1–34).

Sutton, R. S. (1986). Two problems with backpropagation and other steepest-descent learning procedures for networks. In *Proceedings of the 8th Annual Conference of the Cognitive Science Society* (pp. 823–831).

Watson, J., & Wiles, J. (2002). The rise and fall of learning: A neural network model of the genetic assimilation of acquired traits. In *Proceedings of the 2002 Congress on Evolutionary Computation (CEC 2002)* (pp. 600–605).

White, D., & Ligomenides, P. (1993). GANNet: a genetic algorithm for optimizing topology and weights in neural network design. In *Proceedings of the International Workshop on Artificial Neural Networks (IWANN'93)* (pp. 322–327). New York: Springer-Verlag.

Whiten, A., & Ham, R. (1992). On the nature and evolution of imitation in the animal kingdom. In *Advances in the study of behaviour*, 239–283.

Yanco, H., & Stein, L. (1993). An adaptive communication protocol for cooperating mobile robots. In *From animals to animats 2*: *Proceedings of the Second International Conference on Simulation of Adaptive Behavior* (pp. 478–485). Cambridge, MA: MIT Press.

Zentall, T. (2001). Imitation in animals: evidence, function and mechanisms. *Cybernetics and Systems, 32*, 53–96.

About the Authors

Dara Curran is a researcher at the Department of Information Technology, National University of Ireland, Galway. His primary research interests are artificial life, language evolution and cultural learning in artificial societies and multi-agent systems.

Colm O' Riordan lectures at the Department of Information Technology, National University of Ireland, Galway. He teaches, among other things, agent-based systems and information retrieval. His main research interests are in the fields of agent-based systems, artificial life and information retrieval. His current research focuses on cooperation and coordination in artificial life societies and multi-agent systems. *Address:* Department of Information Technology, National University of Ireland, Galway, Ireland.
E-mail: colmor@geminga.nuigalway.ie