Measuring Diversity in Populations Employing Cultural Learning in Dynamic Environments

Dara Curran and Colm O'Riordan

National University of Ireland, Galway, Ireland
dara.curran@nuigalway.ie,
colmor@it.nuigalway.ie

Abstract. This paper examines the effect of cultural learning on a population of neural networks. We compare the genotypic and phenotypic diversity of populations employing only population learning and of populations using both population and cultural learning in two types of dynamic environment: one where a single change occurs and one where changes are more frequent. We show that cultural learning is capable of achieving higher fitness levels and maintains a higher level of genotypic and phenotypic diversity.

1 Introduction

A number of learning models may be readily observed from nature and have been the focus of much study in artificial intelligence research. Population learning (i.e. learning which occurs at a population level through genetic material) is typically simulated using genetic algorithms. Life-time learning (i.e. learning which takes place during an organisms's life time through reactions with its environment) can be simulated in a variety of ways, typically employing neural networks or reinforcement learning models.

A relatively new field of study in artificial intelligence is synthetic ethology. The field is based on the premise that language and culture are too complex to be readily analysed in nature and that insight can be gained by simulating its emergence in populations of artificial organisms. While many studies have shown that lexical, syntactical and grammatical structures may spontaneously emerge from populations of artificial organisms, few discuss the impact such structures have on the relative fitness of individuals and of the entire population.

A robust multi–agent system should be able to withstand and adapt to environmental changes. This type of behaviour parallels that of the natural world where species capable of adaptation will have more chance of evolutionary success than ones that are rigid and incapable of such plasticity. At its most basic level, adaptation in nature takes the form of population learning. At a higher level, organisms capable of adapting their behaviour to suit a particular environment during their lifetimes will be more likely to survive in the long term.

The focus of this paper is to attempt to understand the effect of cultural learning on a population of artificial organisms subjected to dynamic environments. This is accomplished by studying its effect on the population's fitness as well as its genotypic and phenotypic diversity. The remainder of this paper is arranged as follows. Section 2 introduces background research, including descriptions of diversity measures and cultural learning techniques that have been employed for this study. Section 3 describes the experimental setup. Section 4 presents the Experiment Results and Section 5 concludes.

2 Background research

2.1 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 artifactual materials. Such information exchange occurs during the lifetime of individuals in a population and can greatly enhance the behaviour of such species. Because these exchanges occur during an individual's lifetime, cultural learning can be considered a subset of lifetime learning.

An approach known as synthetic ethology [10, 17] argues that the study of language is too difficult to perform in real world situations and that more meaningful results could be produced by modeling organisms and their environment in an artificial manner. Artificial intelligence systems can create tightly controlled environments where the behaviour of artificial organisms can be readily observed and modified. Using genetic algorithms, the evolutionary approach inspired by Darwinian evolution, and the computing capacity of neural networks, artificial intelligence researchers have been able to achieve very interesting results.

In particular, experiments conducted by Hutchins and Hazlehurst [8] simulate cultural evolution through the use of a hidden layer within an individual neural network in the population. This in effect, simulates the presence of a Language Acquisition Device (LAD), the physiological component of the brain necessary for language development, the existence of which was first suggested by Chomsky [3]. The hidden layer acts as a verbal input/output layer and performs the task of feature extraction used to distinguish different physical inputs. It is responsible for both the perception and production of signals for the agent.

A number of approaches were considered for the implementation of cultural learning including fixed lexicons [19], indexed memory [16], cultural artifacts [7] and signal–situation tables [10]. The approach chosen was the teacher/ pupil scenario [4, 2] where a number of highly fit agents are selected from the population to act as teachers for the next generation of agents, labelled pupils. 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.

2.2 Diversity

Diversity measures typically quantify the differences between individuals in a population. It is commonly accepted that a population that is capable of maintaining diversity will avoid premature convergence and local maxima.

Diversity measures for populations of neural networks have been the focus of considerable research, focusing mainly on genotypic diversity [18, 14, 1]. Many methods exist for the calculation of genotypic diversity, many based on binary representations. For the purposes of this research however, many schemes are unsuitable due to the nature of the marker-based encoding scheme used to represent each neural network.

Our scheme examines each block of the encoding and compares it to blocks of similar length in other encodings. Each encoding block contains a single node and a number of links emanating from that node. It is therefore intuitive to propose that blocks of similar length (having a similar number of emanating links) are suitable for mutual comparison.

There is comparatively little research on phenotypic diversity in evolutionary computation. Typically, phenotypic diversity is measured at the fitness level [5]. However, this measure tends to compress the available diversity information resulting in a coarse grained measure not useful in all situations. The approach adopted in this work is to examine the components of the fitness value of each individual, i.e. an individual's response to each bit-parity stimulus. By comparing the difference between all responses (and not just the aggregate fitness function) a finer grained measure of phenotypic diversity can be obtained.

2.3 Dynamic Environments

Many approaches have been taken to simulate changing environments for multi–agent and artificial life systems[13, 6, 15, 11] focusing on Latent Energy Environments and fitness functions which vary over time. Our approach, while straightforward, has the advantage of clarity: agents are repeatedly presented with a number of bit–patterns representing either food or poison. An agent capable of distinguishing the two by correctly ingesting food and avoiding poison will be rewarded with a high fitness level and reproductive opportunity. At each environmental change all bit–patterns representing food are made to represent poison and vice–versa thus completely reversing the environment. This is partly based on work performed by Nolfi et al[13] who compared the performance of a robotic agent employing genetic evolution (population learning) and that of agents employing back–propagation (life–time learning) in a changing environment.

3 Simulator

The architecture of the artificial life simulator can be seen as a hierarchical structure. At the top-level of the simulator is a command interpreter which allows users to define an experiment's variables including the number of networks, the number of generations to run the experiment, mutation and crossover rates and the actual problem set which the population will be attempting to solve.

The neural network layer takes the variables set using the command interpreter and initialises a given number of neural networks. The layer then performs training and testing of the networks according to the parameters of the experiment. These network memory structures are then passed to the encoding layer which transforms them into genetic code structures for use in the genetic algorithm. The encoding mechanism used for this set of experiments is a modified version of marker based encoding.

Marker based encoding represents neural network elements (nodes and links) in a binary string. Each element is separated by a marker to allow the decoding mechanism to distinguish between the different types of element and therefore deduce interconnections[9, 12].

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 on 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 together with 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

The genetic algorithm layer uses the genetic codes and the data retrieved from the neural network layer's testing of the networks to perform its genetic operators on the population. A new population is produced in the form of genetic codes. These are passed to the decoding layer which transforms each code into a new neural network structure. These structures are then passed up to the neural network layer for a new experiment iteration. Once the required number of generations has been reached, the experiment finishes.

Two-point crossover is employed and weight mutation is employed which takes the weight value and increases/decreases the value according to a random percentage (200%). This approach was found, empirically, to be more successful and was adopted for this set of experiments.

3.1 Simulating Cultural Evolution

In order to perform experiments related to cultural evolution, it was necessary to adapt the existing simulator architecture to allow agents to communicate with one another. This was implemented using an extended version of the approach adopted by Hutchins and Hazlehurst. The last hidden layer of each agent's neural network functions as a verbal input/output layer.

At end of each generation, a percentage of the population's fittest networks are selected and are allowed to become teachers for the next generation. The teaching process takes place as follows: a teacher is stochastically assigned n pupils from the population where $n = \frac{N_{pop}}{N_{teachers}}$, where N_{pop} is the population size and $N_{teachers}$ is the number of teachers. Each pupil follows the teacher in its environment and observes the teacher's verbal output as it interacts with its environment. A teaching cycle occurs when the pupil attempts to emulate its teacher's verbal output using back-propagation. Once the number of required teaching cycles is completed, the teacher networks die and new teachers are selected from the new generation.

Unlike previous implementations, the number of verbal input/output nodes is not fixed and is allowed to evolve with the population, making the system more adaptable to potential changes in environment. In addition, this method does not make any assumptions as to the number of verbal nodes (and thus the complexity of the emerging lexicon) that is required to effectively communicate. It should be noted that neither the parent's nor the pupil's genotype is altered at any time during these cultural exchanges.

4 Experimental Setup

The following set of experiments each employs two populations. One population is allowed to evolve through population learning (by genetic algorithm), while the other employs both population and cultural learning. The problem domain for this set of experiments is the 5-bit parity problem. Each network is exposed to bit patterns and must determine whether the pattern represents an odd or even number. Fitness is assigned according to the mean square error of a network.

Two types of environment were employed for the experiments: an environment with a single dramatic change (at generation 200) and another with a series of regular changes (every 20 generations) during the course of the experiment. The change in environment is implemented by reversing the food and poison representations such that the bit pattern representing food will represents poison and vice-versa.

Each experiment consists of a population of 50 neural networks evolving for 400 generations with crossover and mutation rates set at 0.6 and 0.02 respectively. The population employing cultural learning takes the fittest 10% of each generation as teachers which interact with pupils for five teaching cycles. An additional parameter, cultural mutation, adds noise to each interaction with probability 0.02. The results presented are averaged from 10 independent runs.

5 Experiment Results

The experimental results are divided into two sections. The first examines the relative performance of cultural learning and population learning through analysis of the error values for each population. The second section is concerned with genotypic and phenotypic diversity measures for each population.

5.1 Single Environment Change

The average error values for both populations for the single environment change experiment are presented in figure 1. It is clear from the results that the population employing cultural learning is capable of reducing its error values more successfully than the population using population learning alone. The environment change at generation 200 is clearly marked by a large surge in error values

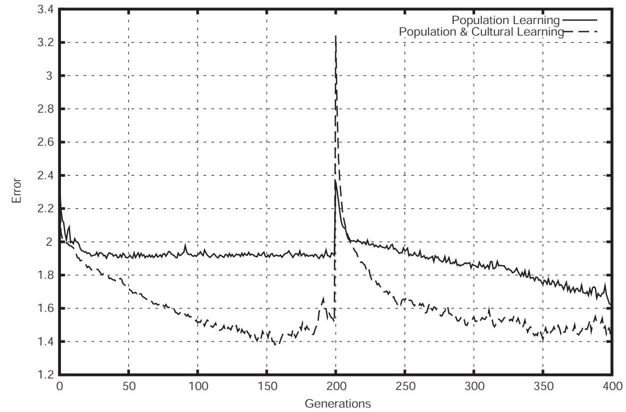

Fig. 1. *Average Error*

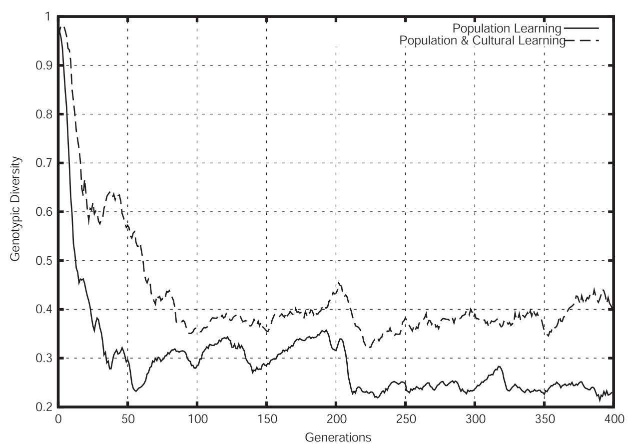

Fig. 2. *Genotypic Diversity Values*

occurring in both populations. However, the sharp increase in error is more evident in the population employing population learning alone, suggesting that cultural learning is softening the environment change.

Figure 2 shows the genotypic diversity for both populations. While both populations have a tendency to reduce diversity as the experiment progresses, the population employing cultural learning is capable of maintaining a higher (and statistically significant) level of diversity throughout the experiment. This trend is reinforced by the results of the phenotypic diversity measure, presented in figure 3. The phenotypic diversity of the population employing population learning alone is considerably lower than that of the population employing cultural learning.

It is clear from these results that in the single change environment, cultural learning is capable of maintaining a high genotypic and phenotypic diversity for its population. This can be correlated to its corresponding superior performance with regard to average error values.

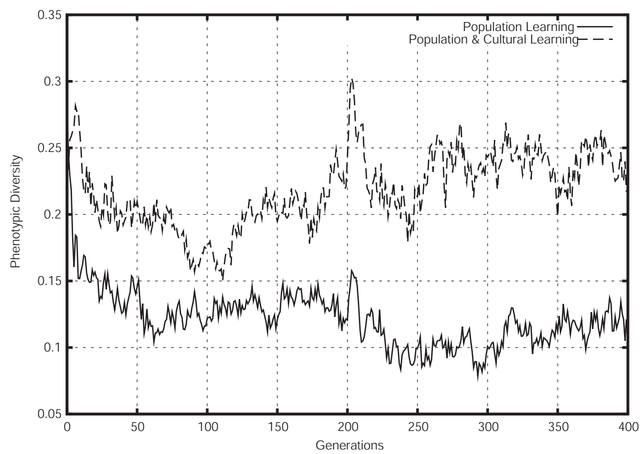

Fig. 3. *Phenotypic Diversity Values*

5.2 Multiple Environment Changes

The multiple environment change experiment presents a considerably more difficult challenge to both populations as the reversals in environment occur very frequently. Figure 4 presents the average error of both populations over the experiment run. Each environment change can be clearly seen as a surge in average error every 20 generations. Clearly both populations experience difficulty in tracking the environmental changes in this experiment.

The population employing cultural learning is capable of matching and in some cases improving on the error values acheived by the population employ-

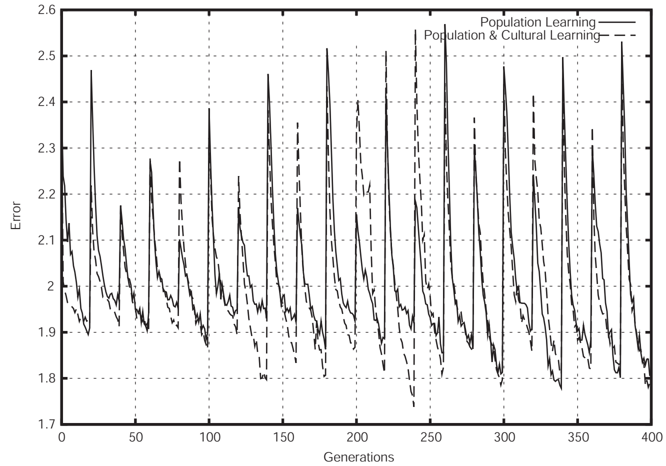

Fig. 4. *Average Error*

ing population learning alone. However, it cannot be said that there is a clear distinction between the two populations.

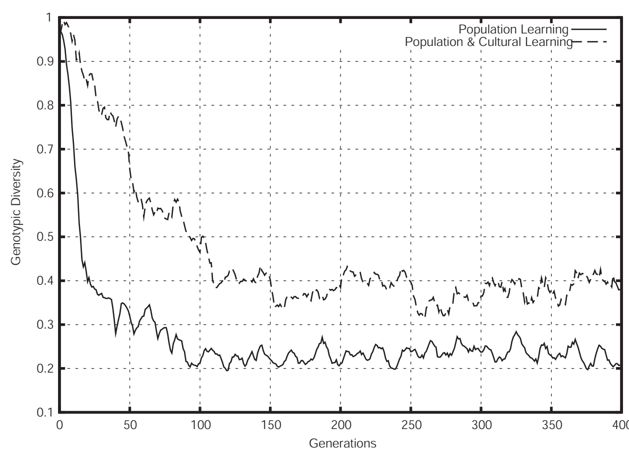

Fig. 5. *Genotypic Diversity Values*

Figure 5 presents the results of the genotypic diversity measure for both populations. The results are similar to those obtained in the previous experiment set, with both populations reducing diversity over the experiment run, but with the population employing cultural learning maintaining a higher level throughout.

Similarly, the phenotypic diversity measure results outlined in figure 6 show that the population employing cultural learning is achieving and maintaining higher levels of phenotypic diversity than that of the population employing population learning alone.

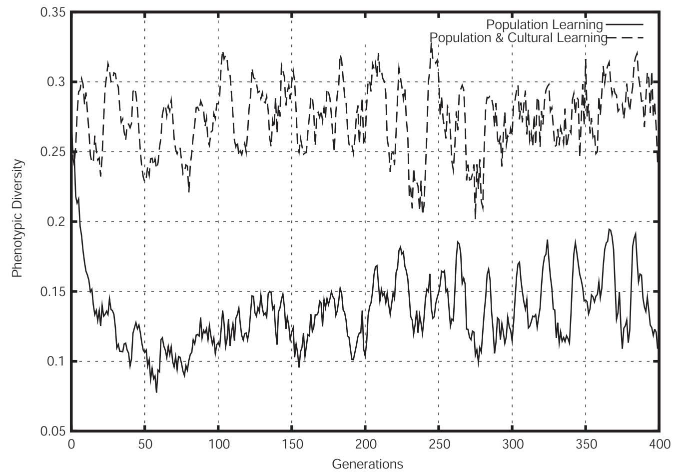

Fig. 6. *Phenotypic Diversity Values*

6 Conclusions

The results presented in this paper suggest that the addition of cultural learning is beneficial to a population subjected to dramatic environmental changes, but is not capable of providing any real advantage in environments where changes occur more frequently. It should be stressed that we do not wish to generalise as to the effects of cultural learning for all problems, rather that this study provides a useful starting point into the analysis of the potential benefits of cultural learning. Diversity measures in particular may allow more detailed analysis into the effects of cultural learning for a variety of problem domains. Future work will focus on more complex problems and environments where changes occur more gradually, rather than simple reversal of problem solutions.

Acknowledgements

This research is funded by the Irish Research Council for Science, Engineering and Technology.

References

1. G. Brown. *Diversity in Neural Network Ensembles.* PhD thesis, University of Birmingham, 2003.

2. A. Cangelosi and D. Parisi. The emergence of a language in an evolving population of neural networks. *Technical Report NSAL–96004, National Research Council, Rome*, 1996.
3. N. Chomsky. On the nature of language. In *Origins and evolution of language and speech*, pages 46–57. Annals of the New York Academy of Science, New York. Vol 280, 1976.
4. D. Denaro and D. Parisi. Cultural evolution in a population of neural networks. In *M.Marinaro and R.Tagliaferri (eds), Neural Nets Wirn-96.New York: Springer*, pages 100–111, 1996.
5. G. Kendall E. K. Burke, S. Gustafson. Diversity in genetic programming: an analysis of measures and correlation with fitness. In *IEEE Trans. Evolutionary Computation*, volume 8(1), pages 47–62, 2004.
6. J. J. Grefenstette. Genetic algorithms for dynamic environments. In R. Maenner and B. Manderick, editors, *Parallel Problem Solving from Nature 2*, pages 137–144, 1992.
7. E. Hutchins and B. Hazlehurst. Learning in the cultural process. In *Artificial Life II, ed. C. Langton et al.*, pages 689–706. MIT Press, 1991.
8. E. Hutchins and B. Hazlehurst. How to invent a lexicon: The development of shared symbols in interaction. In N. Gilbert and R. Conte, editors, *Artificial Societies: The Computer Simulation of Social Life*, pages 157–189. UCL Press: London, 1995.
9. H. Kitano. Designing neural networks using genetic algorithm with graph generation system. In *Complex Systems, 4, 461-476*, 1990.
10. B. MacLennan and G. Burghardt. Synthetic ethology and the evolution of cooperative communication. In *Adaptive Behavior 2(2)*, pages 161–188, 1993.
11. F. Menczer. Changing latent energy environments: A case for the evolution of plasticity. In *Technical Report CS94-336*, 1994.
12. G. F. Miller, P. M. Todd, and S. U. Hedge. Designing neural networks using genetic algorithms. In *Proceedings of the Third International Conference on Genetic Algorithms and Their Applications*, pages 379–384, 1989.
13. S. Nolfi and D. Parisi. Learning to adapt to changing environments in evolving neural networks. In *Technical Report 95-15, Institute of Psychology, National Research Council, Rome, Italy.*, 1995.
14. D. W. Opitz and J. W. Shavlik. Generating accurate and diverse members of a neural-network ensemble. In David S. Touretzky, Michael C. Mozer, and Michael E. Hasselmo, editors, *Advances in Neural Information Processing Systems*, volume 8, pages 535–541. The MIT Press, 1996.
15. T. Sasaki and M. Tokoro. Adaptation toward changing environments: Why darwinian in nature? In Phil Husbands and Inman Harvey, editors, *Fourth European Conference on Artificial Life*, pages 145–153, Cambridge, MA, 1997. MIT Press.
16. Lee Spector. Genetic programming and AI planning systems. In *Proceedings of Twelfth National Conference on Artificial Intelligence*, pages 1329–1334, Seattle, Washington, USA, 1994. AAAI Press/MIT Press.
17. L. Steels. The synthetic modeling of language origins. In *Evolution of Communication*, pages 1–34, 1997.
18. X. Yao Y. Liu and T. Higuchi. Evolutionary ensembles with negative correlation learning. In *IEEE Transactions on Evolutionary Computation*, volume 4(4), pages 380–387, 2000.
19. H. Yanco and L. Stein. 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*, pages 478–485. MIT Press, Cambridge Ma., 1993.