National University of Ireland, Galway
Ollscoil na hÉireann, Gaillimh

Department of Information Technology

technical report NUIG-IT-091204

Robustness in Changing Environments in Populations of Neural Networks Using Cutural Evolution

D. Curran (NUI, Galway)
C. O'Riordan (NUI, Galway)

Robustness in Changing Environments in Populations of Neural Networks Using Cutural Evolution

Dara Curran
Dept. of Information Technology
National University of Ireland, Galway
Email: dara.curran@nuigalway.ie

Colm O'Riordan
Dept. of Information Technology
National University of Ireland, Galway
Email: colmor@geminga.it.nuigalway.ie

***Abstract*— The focus of this paper is to examine the effect of cultural learning on a population of agents in various types of dynamic environment. Cultural learning allows highly fit agents in a population to teach others in order to achieve higher levels of fitness. These agents are placed in environments which may change very frequently, moderately or infrequently. The performance of cultural learning is compared with experiments undertaken using population learning, i.e genetic evolution, of agents in the same environments.**

I. Introduction

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 adaption 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 genetic evolution over generations, also known as 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.

Lifetime learning can take many forms – at its simplest it is a reaction to a particular stimulus and the adjustment of world view that follows the reaction. Thus, very simple organisms are capable of learning to avoid harmful substances and are attracted to other beneficial ones. At its most complex, societies of organisms communicate in order to impart useful knowledge to others in the community. This is the essence of cultural learning, the transmission of information through generations by non–genetic means.

These forces can be modelled in computer systems by employing genetic algorithms to simulate population learning, and neural networks to simulate lifetime learning. Typically, a population of neural networks is randomly generated, and the usual genetic operators such as selection, crossover and mutation are applied at each generation to arrive at higher levels of population fitness. At each generation, the neural network is allowed to learn, usually through an error reducing algorithm such as error back–propagation. To implement cultural learning, highly fit individuals are selected to teach the remaining population through repeated cultural exchanges.

The focus of this paper is to examine the effect of various types of dynamic environment on a population of agents employing population learning alone, and a population employing population and cultural learning. We present a number of experimental results which illustrate the benefit of cultural learning in each type of environment. The remainder of the paper is organised as follows: Section 2 covers some related work, paying particular attention to types of learning previously employed. In section 3 we describe the experimental setup, including the artificial life simulator used for all experiments and the cultural learning framework. In Section 4, we detail the results and in Section 5, conclusions are presented and future directions are outlined.

II. Related Work

A number of learning models can be derived from observation in nature. These have traditionally been classified into two distinct groups: population and life-time 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.

By contrast, there exist species in nature that are capable of learning, or adapting to environmental changes and novel situations at an individual level. Such learning, know as life-time learning, still employs population learning to a degree, but further enhances the population's fitness through its adaptability and resistance to change. Another phenomenon related to life-time learning, first reported by Baldwin[1], occurs when certain behaviour first evolved through life-time learning becomes imprinted onto an individual's genetic material through the evolutionary processes of crossover and mutation. This individual is born with an innate knowledge of such behaviour and, unlike the rest of the populations, does not require time to acquire it through life-time learning. As a result, the individual's fitness will generally be higher than that of the population and the genetic mutation should become more widespread as the individual is repeatedly selected for reproduction.

Research has shown that the addition of life-time learning to a population of agents is capable of achieving much higher levels of population fitness than population learning alone[2], [3]. Furthermore, population learning alone is not well suited to changing environments[4].

A. 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[5], [6] argues that the study of language is too difficult to perform in real world situations and that more meaningful results could be produced by modelling 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[7] 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, whose existence was first suggested by Chomsky[8]. 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[9], [10], indexed memory[11], cultural artifacts[12], [13] and signal–situation tables[5]. The approach chosen was the increasingly popular teacher/pupil scenario[14], [15], [10] 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.

B. Dynamic Environments

Many approaches have been taken to simulate changing environments for multi–agent and artificial life systems[16], [17], [18], [19] 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. Our goal is to evolve a population capable of sustaining repeated iterations of such environmental changes.

III. Experimental Setup

A. Simulator

The experiments outlined in this paper were performed using a previously developed artificial life simulator. The simulator allows populations of neural networks to evolve using a genetic algorithm and each network can also be trained during each generation of an experiment to simulate life–time learning.

The mapping of neural network to genetic code required for the genetic algorithm is achieved using a modified version of marker based encoding. This allows networks to develop any number of nodes and interconnecting links, giving a large number of possible neural network architecture permutations.

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 [20], [21].

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 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.

The networks undergo various stages throughout their lifetime. First, the gene codes are decoded to create their neural network structure. Training is then performed using error back–propagation for a given number of iterations (training cycles). Each network is tested to determine its fitness using a fitness function which takes the agent's neural network error into account and the population is ranked using linear based fitness ranking producing fitness values in the range [0.0,1.0]. Roulette wheel selection is employed to generate the intermediate population. Crossover and mutation operators are then applied to create the next generation.

B. Cultural Learning Framework

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 (figure 1).

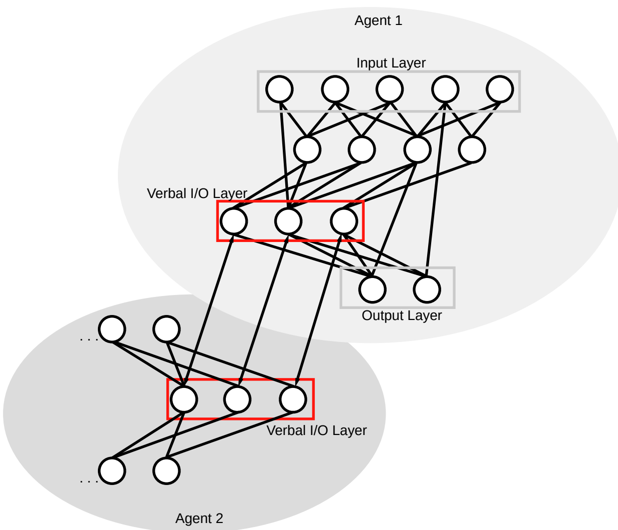

Fig. 1. *Agent Communication Architecture*

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 encounters what it believes to be food or poison bit patterns. The pupil then attempts to emulate its teacher's verbal output using back-propagation. Once the teaching process has been 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.

IV. Experimental Results

The population of agents exist in a world where they are presented with bit patterns representing food and poison items. An agent must learn to distinguish between the two in order to attain high levels of fitness and have a high probability of reproductive opportunities. In this set of experiments, the 5–bit parity problem is used to represent food and poison elements, where food is represented by the value 1 and poison by the value 0. Thus, agents correctly identifying 01001 as poison and 00001 as food will be awarded high levels of fitness.

Experiments were carried out to comprehensively assess the effect of a dynamic environment on agents employing population learning on its own and secondly a population of agents employing both population and cultural learning. A dynamic environment can be generated in a variety of ways, but for the purposes of this experiment we chose to create environments in which changes occur at regular intervals. These changes are implemented by swapping food and poison items in the agent's environment - thus essentially reversing the 5–bit parity problem. Using this approach, three environments are created:

- one in which changes occur often (a quickly changing environment)
- one in which changes occur more slowly (a moderately changing environment)
- one in which changes occur very slowly (a slowly changing environment)

We will discuss the three environments in turn, focusing on population learning alone, and population and cultural learning in combination in each of the experiments.

A. Quickly Changing Environment

The results illustrated in figure 2 show that when the population of agents employs population learning alone, little progress is made. Environmental changes occur every two generations, creating severe oscillations of population fitness. This type of environment poses a serious stability problem to the population, which is unable to retain a steady level of fitness.

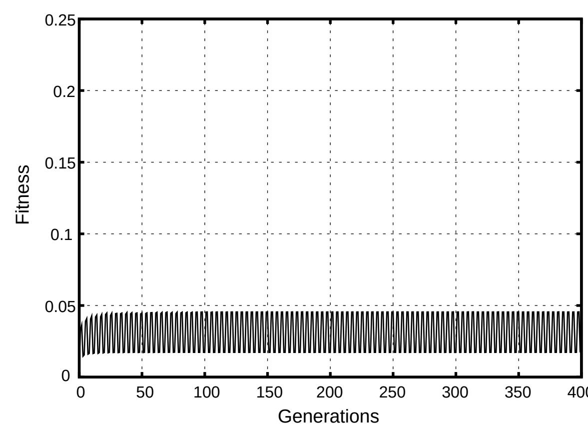

Fig. 2. *Highly Dynamic Environment - Population Learning*

Once cultural learning is introduced, no major gains are made in terms of fitness, but the population has managed to stabilise (figure 3). The difference between the two sets of results is striking: no oscillations of fitness occur once cultural learning is introduced.

B. Moderately Changing Environment

In a moderately dynamic environment, environmental changes occur less frequently allowing agents to evolve and increase their fitness between changes. Figure 4 illustrates the

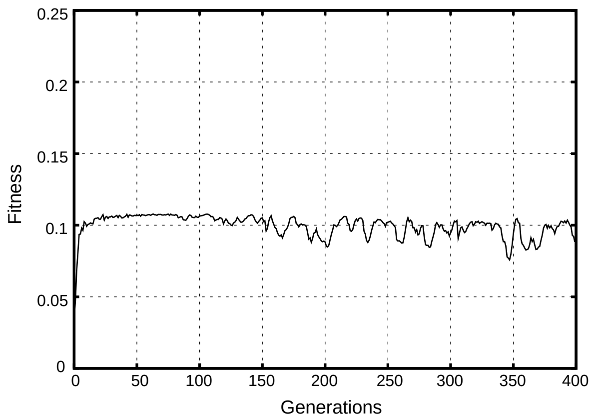

Fig. 3. *Highly Dynamic Environment - Population and Cultural Learning*

results of the moderately dynamic environment on agents employing population learning alone. The environmental changes every 20 generations are clearly depicted in the results as a sudden drop in fitness level, followed by a small recovery. Once the environment reverts to its original state 20 generations later, the population's fitness quickly rises to its original level. This suggests that the population has evolved to tackle the first environment, but is making little progress once the environment is inverted (i.e. poison becomes food and vice–versa), implying that population learning is not swift enough to track such changes.

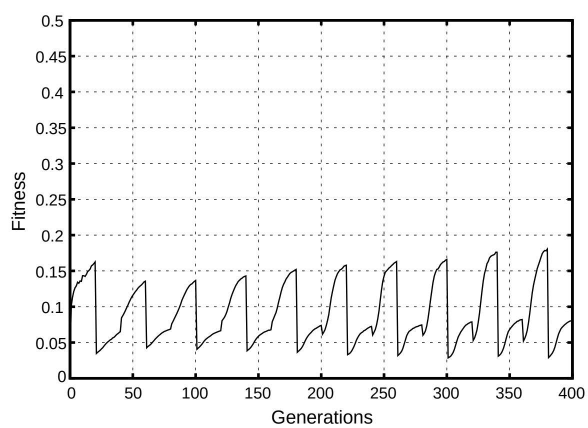

Fig. 4. *Moderately Dynamic Environment - Population Learning*

Figure 5 shows the results of applying cultural learning to this environment. Following a relatively stable period, a series of sharp oscillations occur following generation 60. The environmental changes are even more clear in this set of results, with each change producing a clear drop in fitness. However, unlike the previous result, there is a fast recovery following each drop, showing that the population is capable of recovering from each change. It is clear that the population is capable of evolving successfully in both types of environment and that the addition of cultural learning means that agents are capable of thriving in both environments.

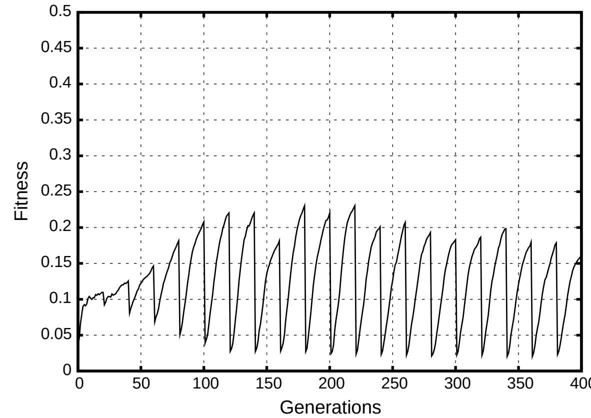

Fig. 5. *Moderately Dynamic Environment - Population and Cultural Learning*

C. Slowly Changing Environment

In this last set of experiments, environmental changes occur much more slowly, every 100 generations. In this type of environment, agent populations have the opportunity to evolve relatively undisturbed for a prolonged period, resulting in higher levels of fitness. Figure 6 shows the results for population learning alone. Each environmental change is again clearly demarcated by a sudden plunge in population fitness. It is clear from these results that the population has evolved to successfully inhabit the first environment, but that once the environment changes to the inverted version, the population is incapable of sustained recovery.

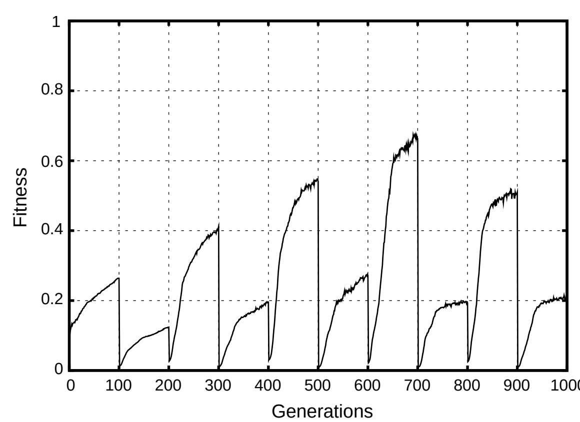

Fig. 6. *Slow Dynamic Environment - Population Learning*

The final set of results, shown in figure 7, once again displays the advantage of cultural learning. Environmental changes are still clearly evident in the graph, but at each change, the population recovers swiftly. While this population's fitness does not rise to the higher levels of the previous experiment, agents employing cultural learning seem by far more flexible in the face of altering environments.

V. Conclusions

It is clear from the experimental results that the addition of cultural learning to populations of agents allows greater

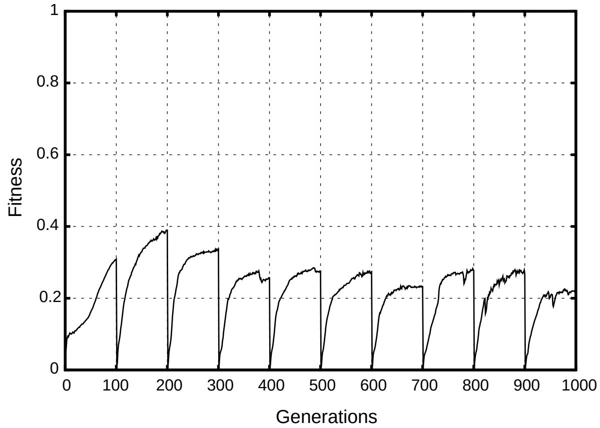

Fig. 7. *Slow Dynamic Environment - Population and Cultural Learning*

stability and recovery to take place. This holds true for all three types of dynamic environment presented and suggests that further work should be carried out examining other types of dynamism and more complex problems.

References

[1] J.M. Baldwin. A new factor in evolution. In *American Naturalist 30*, pages 441–451, 1896.

[2] D. Curran and C. O'Riordan. On the design of an artificial life simulator. In V. Palade, R. J. Howlett, and L. C. Jain, editors, *Proceedings of the Seventh International Conference on Knowledge-Based Intelligent Information & Engineering Systems (KES 2003)*, University of Oxford, United Kingdom, 2003.

[3] D. Curran and C. O'Riordan. Artificial life simulation using marker based encoding. In *Proceedings of the 2003 International Conference on Artificial Intelligence (IC-AI 2003)*, volume II, pages 665–668, Las Vegas, Nevada, USA, 2003.

[4] D. Curran and C. O'Riordan. Lifetime learning in multi-agent systems: Robustness in changing environments. In *Proceedings of the 14th Irish Conference on Artificial Intelligence and Cognitive Science (AICS 2003)*, pages 46–50, Dublin, Ireland, 2003.

[5] B. MacLennan and G. Burghardt. Synthetic ethology and the evolution of cooperative communication. In *Adaptive Behavior 22* , pages 161–188, 1993.

[6] L. Steels. The synthetic modeling of language origins. In *Evolution of Communication*, pages 1–34, 1997.

[7] 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.

[8] 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.

[9] H. Yanco and L. Stein. An adaptive communication protocol for cooperating mobile robots, 1993.

[10] C. Angelo and D. Parisi. The emergence of a language in an evolving population of neural networks. *Technical Report NSAL–96004, National Research Council, Rome*, 1996.

[11] L. Spector and S. Luke. Culture enhances the evolvability of cognition. In *Cognitive Science (CogSci) 1996 Conference Proceedings*, 1996.

[12] E. Hutchins and B. Hazlehurst. Learning in the cultural process. In *Artificial Life II, ed. C. Langton et al.* MIT Press, 1991.

[13] A. Cangelosi. 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)*, pages 4365–4368, Washington, DC, 1999. IEEE Press.

[14] A. Billard and G. Hayes. Learning to communicate through imitation in autonomous robots. In *7th International Conference on Artificial Neural Networks*, pages 763–738, 1997.

[15] 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.

[16] S. Nolfi and D. Parisi. Learning to adapt to changing environments in evolving neural networks.

[17] 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.

[18] T. Sasaki and M. Tokoro. Adaptation toward changing environments: Why darwinian in nature?, 1997.

[19] F. Menczer. Changing latent energy environments: A case for the evolution of plasticity, 1994.

[20] H. Kitano. Designing neural networks using genetic algorithm with graph generation system. In *Complex Systems, 4, 461-476*, 1990.

[21] 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.