

A Comparison of Population Learning and Cultural Learning in Artificial Life Societies

Dara Curran and Colm O'Riordan

Department of Information Technology
National University of Ireland, Galway

dara.curran@nuigalway.ie
colmor@it.geminga.nuigalway.ie

Abstract

This paper examines the effect of the addition of cultural learning to a population of agents. Experiments are undertaken using an artificial life simulator capable of simulating population learning (through genetic algorithms) and lifetime learning (through the use of neural networks). To simulate cultural learning, the exchange of information through non-genetic means, a group of highly fit agents is selected at each generation to function as teachers which are assigned a number of pupils to instruct. Cultural exchanges occur through a hidden layer of an agent's neural network known as the verbal layer. Through the use of back-propagation, a pupil agent imitates the teacher's behaviour and overall population fitness is increased. We show that the addition of cultural learning is of great benefit to the population and that in addition, cultural learning causes the population to converge on a fixed lexicon describing its environment.

Introduction

The two primary evolutionary forces in nature can be described as population learning and lifetime learning. The first involves genetic inheritance and its role in the behaviour of living things. Populations of creatures evolve through the processes of genetic recombination and mutation. Over long periods, genetic encodings emerge which produce phenotypic traits suitable for a particular environment, such as webbed feet or enhanced vision, giving individuals competitive advantage over others. As generations progress, useful traits will be passed down to successive populations resulting in increases in overall fitness.

Lifetime learning represents each individual's ability to interact and learn from its environment. Creatures which are capable of correctly interpreting novel situations will be more likely to survive the unpredictable of many habitats. Furthermore, creatures which are capable of memory will recall previous errors in similar situations, greatly reducing the risk of repeated damaging behaviour.

Culture is a means of exchange of information through non-genetic means. Examples of such exchanges are language, symbols and artifacts. Means of cultural exchanges between organisms may be considered a subset of lifetime

learning, known as cultural learning. Generally, these exchanges allow more experienced organisms to impart some domain knowledge to less informed organisms.

These learning types can be simulated in computer systems using genetic algorithms and neural networks respectively. Genetic algorithms represent potential problem solutions as genetic codes which are then evaluated for fitness. Pairs of codes are selected in proportion to their fitness and are combined together to produce offspring. These offspring become part of the next generation and the process is repeated. Genetic algorithms have been shown to be useful in a vast variety of problem domains.

Neural networks are simplified mathematical models of nervous systems, inspired by the neurons and synapses of living creatures. Neural networks function by reading input patterns, feeding these pattern values through a succession of weighted synapses between neurons, and finally displaying output values at specified output neurons. By comparing a network's output pattern and the desired output pattern for a given input, a measure of error can be obtained. This error is then used to alter the weighting value of synapses that connect neurons in the network. One such algorithm for performing this weighting adjustment is known as error back propagation. Through a series of training iterations, the overall error of a network is reduced, improving its performance.

The combination of genetic algorithms and neural networks provides a framework for evolutionary experimentation in popular domains such as language evolution, neural network design optimization and games.

This paper examines the benefits of employing cultural evolution in populations of reactive agents. Experiments are conducted where agents are placed in an environment where they must distinguish between food and poison bit-patterns in order to achieve high levels of fitness. Cultural evolution is introduced using the teacher-pupil model and its performance is compared to that of population learning alone.

We show that the addition of cultural learning is of great benefit to the population and that the population converges to a common lexicon describing its environment.

Copyrighted Material

The remainder of the paper is organised as follows: the next section gives some background information on related work. The third section presents the artificial life simulator employed in the experiments. Each of its components are discussed: genetic algorithm, neural network architecture and the encoding used to generate valid genetic codes from neural network structures. In the fourth section, we detail the results and in the last section, conclusions are presented and future directions are outlined.

Related Work

A number of learning models can be identified from observation in nature. These can roughly be classified into two distinct groups: population and life-time learning. In this paper we consider another form of lifetime learning, cultural learning.

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.

Lifetime Learning

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, known 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 (Baldwin, 1896), 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 (Curran and O'Riordan, 2003c; Curran and O'Riordan, 2003a). Furthermore, population learning alone is not well suited to changing environments (Curran and O'Riordan, 2003b).

Cultural Learning Culture can be succinctly described as a process of information transfer within a population that occurs without the use of genetic material.

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 (MacLennan and Burghardt, 1993; Steels, 1997; Kirby and Hurford, 1997; Smith, *et al.*) 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 (Hutchins and Hazlehurst, 1995) 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 (Chomsky, 1976). 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 (Yanco and Stein, 1993; Cangelosi and Parisi, 1996), indexed memory (Spector and Luke, 1996), cultural artifacts (Hutchins and Hazlehurst, 1991; Cangelosi, 1999) and signal-situation tables (MacLennan and Burghardt, 1993). The approach chosen was the increasingly popular teacher/pupil scenario (Billard and Hayes, 1997; Denaro and Parisi, 1996; Cangelosi and Parisi, 1996) 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.

Experiment setup

The experiments outlined in the next sections are designed to simulate an environment where agents may encounter food or poison bit patterns. Agents that correctly identify these patterns are awarded higher levels of fitness than agents who

For this set of experiments, food and poison bit patterns are 5-bit patterns representing the 5-bit parity problem, where food is assigned the value 1 and poison the value 0. The next section outlines the artificial life simulator that is employed in these experiments.

Simulator

The architecture of the simulator is based on a hierarchical model (figure 1). Data propagates from the simulator's interface down to the simulator's lowest level. The neural network and genetic algorithm layers generate results which are then fed back up to the simulator's highest level.

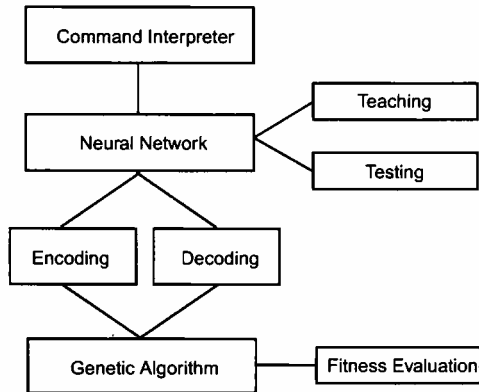


Figure 1: Simulator Architecture

Command Interpreter Layer

The command interpreter is used to receive input from users in order to set all variables used by the simulator. The interpreter supports the use of scripts, allowing users to store parameter information for any number of experiments. A parsing algorithm is used to break down a user's commands and to pass appropriate instructions to the relevant components of the simulator.

Genetic Algorithm Layer

The genetic algorithm layer is responsible for the creation of successive generations of neural networks. The layer is responsible for the simulation of population learning within the simulator. The layer accepts data from the encoding layer as a genetic algorithm data structure. This structure is then manipulated to generate the next generation using the selection, crossover and mutation operators.

Selection The selection process examines each individual in the population and must determine which individuals to place into the intermediate population. The simulator genetic algorithm layer uses the network's error and accuracy as a fitness measure to make this evaluation.

Once the fitness value is computed for each individual in the population, linear based fitness ranking is used to normalize fitness values evenly over the entire population. This

is done to overcome the potential for stagnation across the population gene pool. Traditional proportional fitness assignment may produce a population with very similar fitness values, thereby deteriorating the selection process and possibly resulting in a loss of diversity stemming from a small number of individuals being allowed to reproduce many times. Ranking introduces a uniform scaling across the population's fitness so that fitness values are evenly spread and making the selection process more successful. In addition, fitness ranking provides a simple means of controlling selective pressure, the probability that the best individual is selected compared to the average probability of selection of all individuals (Goldberg, 1989).

The total number of individual networks in the population N_{ind} is sorted such that each individual occupies a ranking position pos where $pos=1$ represents the least fit individual and $pos=N_{ind}$ the most fit individual. For selective pressure SP in $[1.0, 2.0]$, a network's fitness value $Fitness_{net}$ is calculated as:

$$Fitness_{net} = 2 - SP + 2(SP - 1)(pos - 1)/(N_{ind} - 1)$$

Roulette wheel selection is then applied to the population to select individuals for the intermediate population. The number of times that an individual may be copied to the intermediate population is a parameter which can be set using the command interpreter.

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.

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 on the gene code itself, most typically by altering specific values randomly selected from the entire gene code. In this implementation, weight mutation is employed. The operator takes a weight value and modifies it according to a random percentage in the range $[-200\%, 200\%]$.

Encoding and Decoding Schemes

Before the encoding and decoding layers can perform their respective tasks, it is necessary to arrive at a suitable encoding scheme. Many schemes were considered in preparation of these experiments, prioritising flexibility, scalability, difficulty and efficiency. The scheme chosen is based on Marker Based Encoding which allows any number of nodes and interconnecting links for each network giving a large number of possible neural network 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 (Kitano, 1990; Miller et al., 1989). A gene code produced using this scheme is treated as a circular entity. Thus, the code parsing mechanism reading the end of the gene code will begin reading the start of the gene code once the end is reached until all available information is correctly retrieved.

Encoding and Decoding Layers

To perform genetic algorithm tasks, the neural network structures must be converted into gene codes on which the genetic algorithm will perform its operations. Conversely, once the genetic algorithm has performed its tasks, the genetic code structure must be converted back to a neural network architecture. The encoding and decoding layers must therefore accept both neural network data structures and genetic algorithm data structures to function correctly.

Neural Network Layer

The neural network layer is responsible for all functions carried out by the neural networks in the simulator's population. The neural network layer accepts a population of neural network data structures and performs a number of functions.

Simulating Cultural Evolution

In order to perform experiments related to cultural evolution, it was necessary to adapt the existing simulator architecture developed by Curran and O'Riordan (Curran and O'Riordan, 2003c) 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 2).

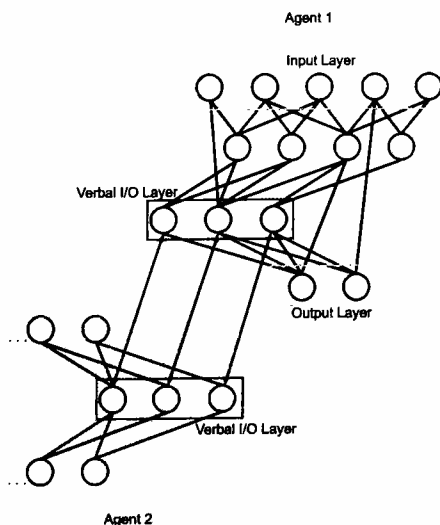


Figure 2: Agent Communication Architecture

At the 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.

Experiments

The purpose of this set of experiments is to identify to what degree the addition of cultural learning aids a population of agents to attain high levels of fitness. We compare the performance of a population of 250 agents endowed with the ability to communicate using the teacher-pupil model to that of a population that uses purely genetic means to advance its fitness levels. The parameters for the cultural learning setting are chosen as follows: the teacher ratio, that is the percentage of the population that is chosen as teachers, is set at 0.1. The number of teaching cycles, the exposure each pupil has to the teachings of a particular teacher is set at 5 cycles. These parameters were chosen following a series of preliminary experiments. The results presented below are averaged from twenty experiment runs over 250 generations.

Results

It is clear from this experiment result illustrated in figure 3 that cultural learning is indeed enhancing the population's fitness and its ability to correctly distinguish between food and poison patterns. Figures 5, 6 and 7 show the evolution of cultural exchanges for one of the food patterns during one of the experiments. The illustration shows the verbal output for the population at generation 1, 125 and 250. These results show that the population's cultural exchanges are converging to specific verbal output patterns and that the population is generating a lexicon of patterns to represent instances of food and poison patterns.

The verbal layer of an agent's neural network is responsible for the cultural exchanges that take place during the agent's lifetime. In order for a cultural exchange to be successful, a teacher's verbal layer must be compatible to that of the pupil. This means that a pupil must have at least the




Figure 3: Cultural Evolution




Figure 5: Verbal Output at Generation 1

same number of verbal nodes in its verbal layer as its teacher. If this does not happen, some information being transmitted from the teacher nodes will be lost.




Figure 4: Verbal Node Evolution




Figure 6: Verbal Output at Generation 125

The evolution of the number of verbal nodes can be seen in figure 4 showing that the number of verbal nodes in the population has converged. This result is intuitive, because to successfully and consistently communicate with one another, it is necessary that all parties in a cultural exchange have a similar number of verbal outputs. Teachers with a larger number of verbal outputs than pupils would effectively be attempting to communicate information that the pupils cannot perceive, let alone replicate.

Conclusion

These experiments have shown that the addition of cultural learning to a population of agents can greatly improve their overall fitness. This is particularly important when compared to supervised life-time or population learning approaches where problem solutions must be known in advance in order to correctly train populations. By selecting teachers from the populations, we allow the population to




Figure 7: Verbal Output at Generation 250

self-supervise its progress, without the need for problem knowledge or human supervision.

Future work will concentrate on employing the system on problems which would be impossible to solve using traditional life-time learning approaches and will investigate the effect of cultural evolution in changing environments.

Acknowledgements

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

References

- Baldwin, J. (1896). A new factor in evolution. In *American Naturalist* 30, pages 441–451.
- Billard, A. and Hayes, G. (1997). Learning to communicate through imitation in autonomous robots. In *7th International Conference on Artificial Neural Networks*, pages 763–738.
- 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)*, pages 4365–4368, Washington, DC. IEEE Press.
- Cangelosi, A. and Parisi, D. (1996). The emergence of a language in an evolving population of neural networks. *Technical Report NSAL-96004, National Research Council, Rome.*
- Chomsky, N. (1976). 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.
- Curran, D. and O'Riordan, C. (2003a). 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.
- Curran, D. and O'Riordan, C. (2003b). 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.
- Curran, D. and O'Riordan, C. (2003c). On the design of an artificial life simulator. In Palade, V., Howlett, R. J., and Jain, L. C., editors, *Proceedings of the Seventh International Conference on Knowledge-Based Intelligent Information & Engineering Systems (KES 2003)*, University of Oxford, United Kingdom.
- Denaro, D. and Parisi, D. (1996). Cultural evolution in a population of neural networks. In *M. Marinaro and R. Tagliaferri (eds), Neural Nets Wirm-96. New York: Springer*, pages 100–111.
- Goldberg, D. E. (1989). *Genetic Algorithms in Search, Optimization and Machine Learning*. Reading, MA, Addison-Wesley.
- Hutchins, E. and Hazlehurst, B. (1991). Learning in the cultural process. In *Artificial Life II, ed. C. Langton et al.* MIT Press.
- Hutchins, E. and Hazlehurst, B. (1995). How to invent a lexicon: The development of shared symbols in interaction. In Gilbert, N. and Conte, R., editors, *Artificial Societies: The Computer Simulation of Social Life*, pages 157–189. UCL Press: London.
- Kirby, S. and Hurford, J. (1997). Learning, culture and evolution in the origin of linguistic constraints. In *4th European Conference on Artificial Life*, pages 493–502. MIT Press, Cambridge, MA.
- Kitano, H. (1990). Designing neural networks using genetic algorithm with graph generation system. In *Complex Systems*, 4, 461–476.
- MacLennan, B. and Burghardt, G. (1993). Synthetic ethology and the evolution of cooperative communication. In *Adaptive Behavior* 2(2), pages 161–188.
- Miller, G. F., Todd, P. M., and Hedge, S. U. (1989). Designing neural networks using genetic algorithms. In *Proceedings of the Third International Conference on Genetic Algorithms and Their Applications*, pages 379–384.
- Smith, K. (submitted). The cultural evolution of communication in a population of neural networks. *Connection Science*.
- Spector, L. and Luke, S. (1996). Culture enhances the evolvability of cognition. In *Cognitive Science (CogSci) 1996 Conference Proceedings*.
- Steels, L. (1997). The synthetic modeling of language origins. In *Evolution of Communication*, pages 1–34.
- Yanco, H. and Stein, L. (1993). An adaptive communication protocol for cooperating mobile robots.