

# Evolutionary and Lifetime Learning in Varying NK Fitness Landscape Changing Environments: An Analysis of both Fitness and Diversity

**Dara Curran\***

Dept. of Computer Science,  
University College Cork  
Email: d.curran@cs.ucc.ie

**Colm O’Riordan**

Dept. of Information Technology  
National University of Ireland, Galway  
Email: colm.oriordan@nuigalway.ie

**Humphrey Sorensen**

Dept. of Computer Science,  
University College Cork  
Email: sorensen@cs.ucc.ie

## Abstract

This paper examines the effects of lifetime learning on populations evolving genetically in a series of changing environments. The analysis of both fitness and diversity of the populations provides an insight into the improved performance provided by lifetime learning. The NK fitness landscape model is employed as the problem task, which has the advantage of being able to generate a variety of fitness landscapes of varying difficulty. Experiments observe the response of populations in an environment where problem difficulty increases and decreases with varying frequency. Results show that lifetime learning is capable of overall higher fitness levels and, in addition, that lifetime learning stimulates the diversity of the population. This increased diversity allows lifetime learning a greater level of recovery and stability than evolutionary learning alone.

## Introduction

A number of researchers have focused on the interaction between learning and evolution, where individuals within a population of artificial organisms are capable of evolving genetically (evolutionary learning) and also of acquiring knowledge during their lifetime (lifetime learning). Hinton and Nowlan (Hinton & 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 (Nolfi & Parisi 1996; Floreano & Mondada 1996; Sasaki & Tokoro 1997; Curran & O’Riordan 2003).

Most of these investigations have been concerned with the effect that learning has on the fitness of the population, finding on the whole that the addition of lifetime learning has a positive effect. Furthermore, studies have shown that both the Baldwin and Hiding effects can be simulated in populations of artificial organisms (Mayley 1997).

However, there is comparatively little research on other effects of learning, in particular with regard to population diversity. It is generally accepted that the maintenance of diversity can prevent an evolutionary algorithm from becoming trapped in local maxima and indeed that such maintenance

can play an important role in allowing populations to better track changes in environment.

The aim of this paper is to examine whether diversity can be maintained by lifetime learning and to analyse how environmental changes of varying difficulty affect the diversity of populations. We employ an abstract model based on Hinton & Nowlan’s work and conduct a number of experiments employing an environment where the problem difficulty is altered in varying degrees. The remainder of this paper is organised as follows. Section 2 describes related research, including the learning models on which this work is based and a discussion on diversity and changing environments in general. Section 3 presents the model employed for these experiments, including a detailed description of the diversity measure. Section 4 outlines the experiment results and Section 5 presents conclusions.

## Related Work

The following sections outline some background material, including learning models, diversity and changing environments.

### Learning Models

Evolutionary learning refers to the mechanism whereby a population of organisms evolves, or learns, by genetic means through a Darwinian process of iterated selection and reproduction of fit individuals (Darwin 1859). 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.

In nature, there are many species that are capable of learning or adapting to environmental changes and novel situations at an individual level. Such learning, known as lifetime learning, is often combined with evolutionary learning further enhancing the population’s fitness through its adaptability. Much research explores the interactions between evolution and learning, showing that the addition of individual lifetime learning can improve a population’s fitness (Nolfi & Parisi 1996; Floreano & Mondada 1996; Curran & O’Riordan 2003).

Hinton and Nowlan employed a genetic algorithm to study the effects of lifetime learning on the performance of genetic evolution (Hinton & Nowlan 1987). Each agent in the model possesses a genome, represented by a string of

---

\*The first author would like to acknowledge the support of the Irish Research Council for Science, Engineering and Technology. Copyright © 2007, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

characters which can be either *1*, *0* or *?*. Each agent is allowed a number of rounds of lifetime learning where, for each *?* in the genotype, they ‘guess’ its value, assigning it either a *1* or a *0*. The *?*s in the genotype represent the agent’s phenotypic adaptability and any guessed values are not retained in the agent’s genome.

The experiments compared the performance of a population endowed with learning to one without. Results showed that the non-learning population was not capable of finding optimal solutions to the problem task. In contrast, once learning was applied, the population converged on the problem solution, showing that individual learning is capable of guiding genetic evolution. A number of researchers have since replicated the Baldwin effect in populations of artificial organisms (Boers, Borst, & Sprinkhuizen-Kuyper 1995; Suzuki & Arita 2003; Arita & Suzuki 2000).

Hinton & Nowlan’s model was modified by Best to replace individual learning with social, or cultural, learning (Best 1999). A number of model agents are selected from the population according to their fitness to behave as teacher agents. Learner agents are then randomly assigned to teacher agents. For each *?* in a learner’s genome, the learner imitates the corresponding value in the teacher’s genome with some probability. Best found that cultural learning could be used to guide evolution as an alternative to lifetime learning and in addition, that cultural learning resulted in faster convergence on optimal solutions.

The fitness landscape of the original Hinton and Nowlan experiments was a particularly difficult one for populations attempting to evolve genetically, consisting in effect of a ‘needle in the haystack’ problem where no evolutionary feedback is given to partial solutions. An individual with no lifetime learning would have a fitness of 20 unless its genome consisted of exactly twenty 1s, in which case its fitness would be 1 (where 1 and 20 represent highest and lowest fitness respectively). It is highly unlikely that an individual will be randomly generated with such a genome and, since there is no path for the evolutionary process to follow, it is very difficult for evolutionary learning alone to solve this problem.

## Diversity

Diversity measures 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 have been the focus of considerable research, focusing mainly on genotypic diversity (Opitz & Shavlik 1996; Burke, Gustafson, & Kendall 2004; Gusfield 1997; O’Reilly 1997; De Jong 1975; Booker 1985; Collins & Jefferson 1991; Eriksson 2000).

Research has been undertaken to examine the effects of cultural learning on both the phenotypic and genotypic diversity of populations of evolving neural network agents whose problem task is comprised of a number of turn-based games (Curran & O’Riordan 2006). Cultural learning in this context is a form of lifetime learning where agents learn from other members of the population. Results showed that the combination of cultural learning and evolutionary learn-

ing led to improvements in fitness and higher levels of genotypic diversity. However, the complexity of the model employed makes analysis difficult, highlighting the need for a simpler model of both agent and lifetime learning mechanism.

## Changing Environments

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 genetic evolution over generations, also known as evolutionary 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.

Much research has focused on the tracking of changing environments with regard to multi-agent and artificial life systems (Nolfi & Parisi 1996; Grefenstette 1992; Sasaki & Tokoro 1997; Cliff & Miller 1995), focusing on Latent Energy Environments and fitness functions which vary over time. The general goal is to develop the ability of a population to adapt to a change within a reasonable length of time and to guide evolution toward a level of plasticity otherwise difficult to attain.

However, there is comparatively little research on the effects of lifetime learning on the robustness of multi-agent systems facing a changing environment and in particular, with regard to environments of varying difficulty.

## Model

Our model follows the structure of the original Hinton and Nowlan experiments. Each agent comprises a genotype of 20 bits and experiments employ a population of 1000 agents evolved for 400 generations. The experiments employ two types of population: one using evolutionary learning alone (purely genetic evolution) and the other also employing lifetime learning.

Where a population evolves solely by evolutionary learning, agent genomes consist of strings of 20 1s or 0s randomly generated in the initial population with equal probability. Populations employing lifetime learning have genomes containing 1s, 0s or ?s, where the ?s represent the agent’s phenotypic ability to guess either 1 or 0 in a similar manner to the original experiments. In the initial population these are randomly selected with probability 0.25, 0.25 and 0.5 respectively (as in Hinton and Nowlan’s experiments).

Our model employs the NK fitness model as the fitness landscape for the experiments (Kaufmann 1989) as an alternative to the fitness function proposed by Hinton and Nowlan. The NK fitness model alters the ruggedness of a fitness landscape by tuning the level of epistasis present within a chromosome. The model takes two parameters: *N*, the length of the chromosome and *K*, the level of epistasis. The higher the epistasis level, the more rugged the landscape.

The NK fitness model provides some evolutionary feedback to populations evolving genetically and therefore allows a fairer comparison between evolutionary and lifetime

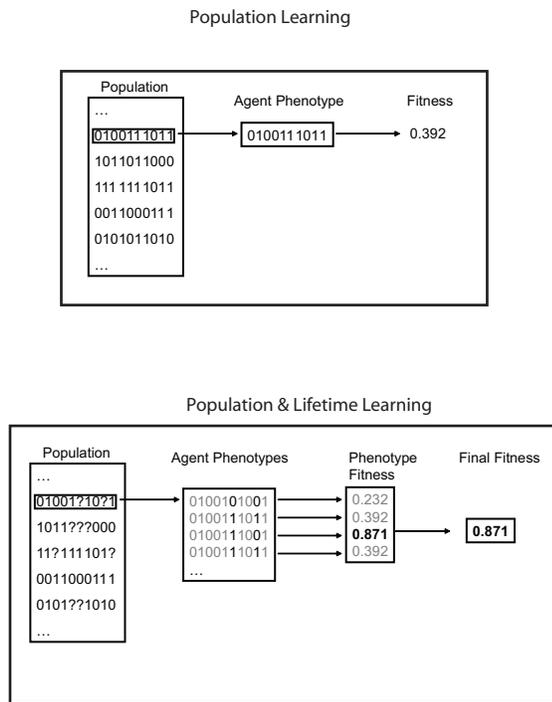


Figure 1: *Fitness evaluation*

learning. Furthermore, the model can be easily configured to provide smoother or more rugged landscapes, making it ideal for this set of experiments.

### Fitness calculation

Populations employing purely genetic evolution (evolutionary learning alone) possess genomes comprised of *1*s or *0*s (no *?*s) and therefore have no opportunity to alter the way they interact with their environment. The fitness of agents in such populations is measured directly from their genomes.

When a population adds lifetime learning to evolutionary learning, its members possess genomes that consist of *1*s, *0*s and *?*s. Individuals are given the opportunity to replace each of the *?*s in their genomes with either *1*s or *0*s. Each agent guesses the value of its *?* with equal probability. The number of guesses an agent makes during its lifetime can be tuned by the experimenter to allow lifetime learning a greater or lesser influence on the evolutionary process. Each guess is evaluated using the NK fitness function and an agent's best guess is taken as its fitness value (see Figure 1).

### Diversity measure

The diversity measure examines the differences between members of a population. The population employing evolutionary learning alone contains genomes comprised of *1*s or *0*s, while the population adding lifetime learning has an additional number of *?*s within its genomes. This difference between the two populations makes a direct genetic comparison difficult. Therefore, a phenotypic measure of diversity may be more suitable. In this model, the phenotype of an agent is defined as its response to its environment.

In the population employing evolutionary learning alone, an agent's response to its environment is fixed and directly influenced by its genome. Therefore, its genotype and phenotype are identical - the agent has no lifetime influence on its behaviour. However, in the population employing lifetime learning, an agent is allowed make a number of guesses which replace the *?*s in its genome with either *1*s or *0*s. Since the agent's fitness is calculated according to the best of these guesses, we chose to select an agent's best guess as its phenotype.

Employing this model, phenotypic diversity measurements of populations employing evolutionary learning alone and populations employing both evolutionary and lifetime learning are directly comparable, as both are measuring differences between strings of values from the same set (containing either *1*s or *0*s).

The diversity is measured in a manner similar to that developed by Eriksson (Eriksson 2000). Diversity is defined as the average of all individual Hamming distances between phenotypes of individuals  $x$  and  $y$ ,  $h(x, y)$ , whose phenotypes are unique within the population. Thus, for a population of size  $n$  containing  $m$  unique phenotypic binary strings, the diversity measure can be calculated as:

$$\frac{2}{n(m-1)} \sum_{x=1}^m \sum_{y=x+1}^m h(x, y)$$

## Experiments

The experiments involve two populations: one employing evolutionary learning alone and one employing both evolutionary and lifetime learning. A parameter termed *learning cycles* controls the number of learning opportunities allowed to each individual in a population employing lifetime learning. Experiments present lifetime learning populations employing 10 learning cycles and populations employing 30 learning cycles. Both fitness and diversity are measured for each experiment. Populations of 1000 agents are allowed to evolve for 200-400 generations. Linear ranking selection is employed, 1-point crossover occurs with 0.6 probability and mutation occurs with 0.02.

Three changing environments are presented. The first sets problem difficulty at its lowest ( $K=0$ ) and increases it every 50 generations (by incrementing  $K$  by 1 every 50 generations). The second begins with problem difficulty at its highest ( $K=19$ ) and decreases it every 50 generations (by decrementing  $K$  by 1 every 50 generations). Both these experiments run for 400 generations. The final environment begins with problem difficult at its lowest ( $K=0$ ) and increases it every 5 generations. Once difficulty reaches its peak, it is decreased every 5 generations. This experiment runs for 200 generations. The results are averaged from 20 independent experiment runs.

### Fitness

The first set of results is illustrated in Figure 2, showing the average fitness of the three populations over 400 generations. The problem difficulty is increased every 50 generations by incrementing the value of  $K$ . Each increase can

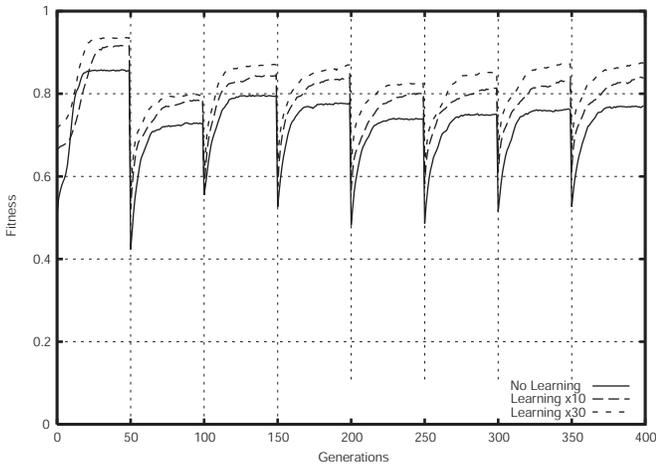


Figure 2: Average fitness, difficulty increases every 50 generations.

be clearly observed by the sudden drop in fitness levels in all three populations. Both populations that employ lifetime learning achieve higher fitness levels than the population employing evolutionary learning alone, though the recovery from each change is similar across all populations.

Figure 3 shows fitness results obtained from the experiment where the problem difficulty decreases every 50 generations (by decrementing the value of K). Again, each change is evidenced by a marked drop in fitness in all three populations and the populations employing lifetime learning exhibit a higher level of fitness than the population using evolutionary learning alone.

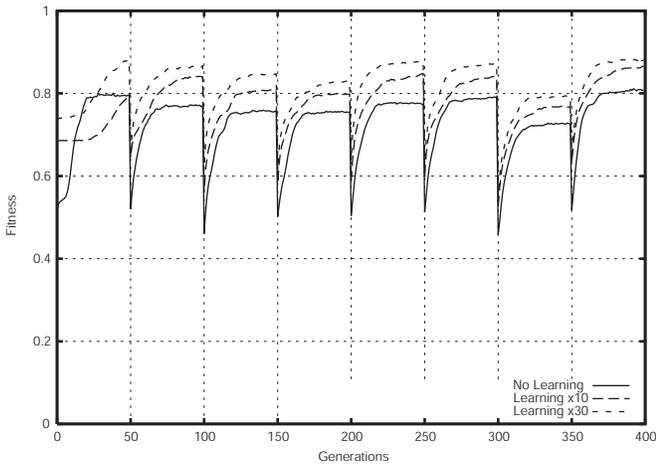


Figure 3: Average fitness, difficulty decreases every 50 generations

The final set of fitness results is illustrated in Figure 4. In this experiment, the problem difficulty is increased every 5 generations and, once K reaches 19, the problem difficulty is decreased every 5 generations. The result is that the problem gradually becomes more difficult in the first half of the

experiment and becomes less so in the second half.

The largest drops in fitness occurring at each change are more pronounced at the beginning and end of the experiment. At the start of the experiment, where the problem difficulty is at its easiest, there is no discernable difference between each population in terms of fitness levels. However, as the experiment progresses, the two lifetime learning populations clearly diverge from the population employing evolutionary learning alone.

Thus, as the problem approaches its peak of difficulty between generations 75 and 125, the populations employing lifetime learning perform significantly better than the population employing evolutionary learning alone. As the problem difficulty then begins to decrease again, the situation reverts to that of the beginning of the experiment.

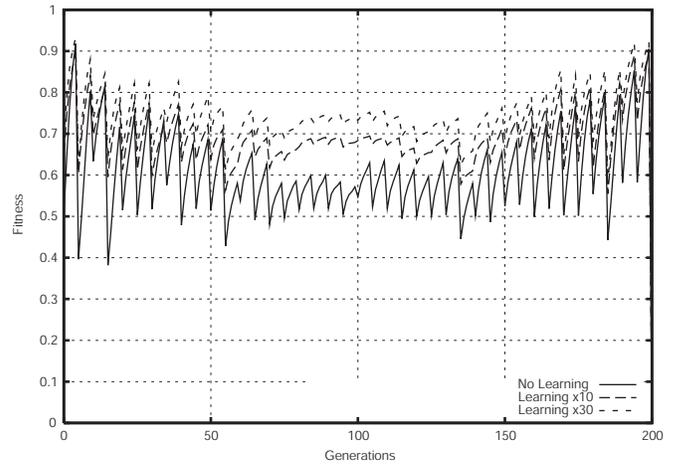


Figure 4: Average fitness, difficulty increases every 5 generations, then decreases every 5 generations

The fact that there is no difference between the populations at the beginning and end of the experiment is as a result of the problem being relatively simple to solve. The addition of lifetime learning does not benefit a population because evolutionary learning is capable of achieving high fitness levels on its own.

The periods of the experiment where the problem is less difficult are marked by much more pronounced fitness drops than the middle of the experiment, where difficulty is at its highest. An evolving population faced with a relatively simple problem will quickly move its individuals into areas of search space where fitness is highest. As a consequence, once a change happens, the population is concentrated in one area of search space and thus is not capable of responding quickly to a sudden change in environment. By contrast, a population in a more difficult environment is unlikely to have its population concentrated in a small area of search space (at least not until it converges) and thus is more likely to respond better to environmental changes.

### Diversity

The diversity measure developed for this set of experiments allows the analysis of the impact of lifetime learning on pop-

ulations employing evolutionary learning. In particular, the study of diversity allows us to determine at what point a population may have converged and to track this as the environment changes.

Figure 5 shows the results obtained from the diversity measure where problem difficulty increased every 50 generations. All populations begin with a high level of diversity and quickly converge to a low level, corresponding to the rise in fitness at the start of the experiment. Each change can be seen as a sharp rise in diversity followed by a new convergence. Thus, at each change, individuals that were previously useful are discarded in favour of a set of diverse individuals in an effort to probe the search space in search of solutions to the new problem.

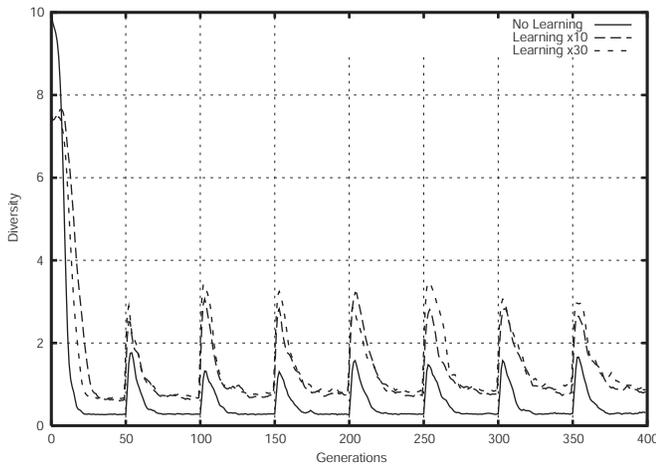


Figure 5: Average diversity, difficulty increases every 50 generations.

The populations employing lifetime learning show a more significant increase in diversity at each environmental change, indicating that the lifetime learning process is generating more diverse populations and thus providing a wider sampling of the search space. This increase in diversity is accompanied by a corresponding increase in fitness in all populations but is more pronounced in those employing lifetime learning.

The second set of results is illustrated in Figure 6, showing diversity for the environment where difficulty decreases every 50 generations. Each population begins with a high level of diversity as before, but, with the exception of the population employing evolutionary learning alone, does not converge as quickly as the previous experiment. This experiment begins with the most difficult problem ( $K=19$ ) and then decreases  $K$  every 50 generations. Thus, at the beginning, a higher level of diversity is maintained in populations employing lifetime learning in response to the initially high problem difficulty. This higher diversity is accompanied by the 'slow start' exhibited by the lifetime learning populations in terms of fitness (see Figure 3).

Changes in difficulty are again evident by drops in fitness every 50 generations and the lifetime learning populations retain an overall higher level of diversity throughout the ex-

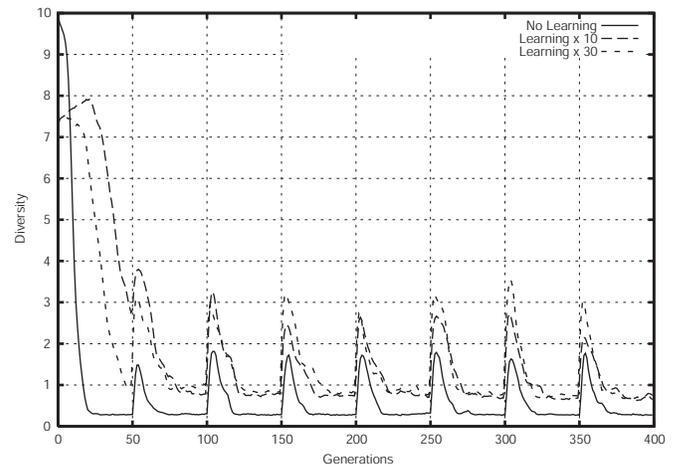


Figure 6: Average diversity, difficulty decreases every 50 generations

periment.

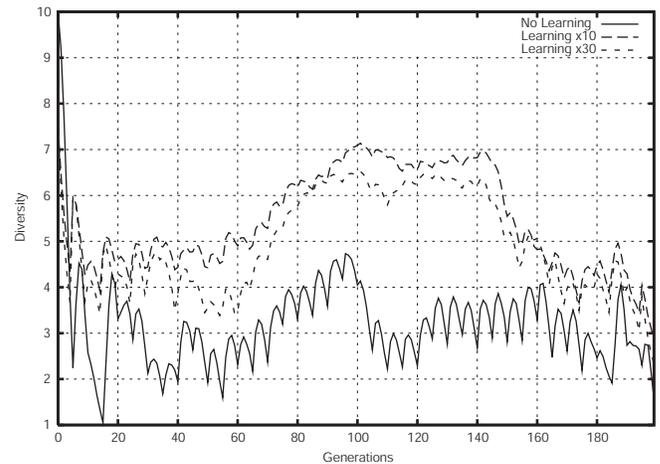


Figure 7: Average diversity, difficulty increases every 5 generations, then decreases every 5 generations

The final set of results is presented in Figure 7, showing the diversity of the three populations in an environment where difficulty is increased every 5 generations and, once maximum difficulty is reached, decreased every 5 generations. The population employing evolutionary learning alone shows an initial sharp fall in diversity at the start of the experiment, followed by an unstable period of changes corresponding to the frequent changes in the environment.

The populations employing lifetime learning show a similar (though much less pronounced) initial drop in diversity but maintain a higher level throughout. Diversity levels then increase dramatically between generations 80 and 140, corresponding to the section of the experiment where problem difficulty is at its highest. The populations respond to the increased difficulty by greatly increasing diversity levels. As a result, changes in environment do not result in severe drops

in corresponding fitness levels (see Figure 4) because the population is spread out across the search space. Thus, lifetime learning, through its stimulation of population diversity, provides some stability in this extreme environment.

## Conclusions

This paper examined the effects of lifetime learning on the fitness and diversity of populations exposed to a series of changing environments. The results corroborate previous work showing that lifetime learning provides a boost to populations in terms of fitness. In addition, the diversity measure developed for these experiments illustrates how lifetime learning may be achieving higher fitness levels. The increased diversity levels exhibited by populations employing both evolutionary and lifetime learning are extremely useful in changing environments. By maintaining a diverse population, those populations incorporating lifetime learning are capable of some stability and quick recovery from environmental changes even in extreme cases.

Future work will examine more complex problems and environments and will analyse the effects of other learning models on the fitness and diversity of evolving populations.

## References

- Arita, T., and Suzuki, R. 2000. Interactions between learning and evolution: The outstanding strategy generated by the baldwin effect. In *Proceedings of Artificial Life VII*, 196–205. MIT Press.
- Best, M. L. 1999. How culture can guide evolution: An inquiry into gene/meme enhancement and opposition. *Adaptive Behavior* 7(3/4):289–306.
- Boers, E.; Borst, M.; and Sprinkhuizen-Kuyper, I. 1995. Evolving Artificial Neural Networks using the “Baldwin Effect”. Technical Report TR 95-14.
- Booker, L. B. 1985. Improving the performance of genetic algorithms in classifier systems. In *Proc. of the International Conference on Genetic Algorithms and Their Applications*, 80–92.
- Burke, E. K.; Gustafson, S.; and Kendall, G. 2004. Diversity in genetic programming: an analysis of measures and correlation with fitness. In *IEEE Trans. Evolutionary Computation*, volume 8(1), 47–62.
- Cliff, D., and Miller, G. F. 1995. Tracking the red queen: Measurements of adaptive progress in co-evolutionary simulations. In Moran, F.; Moreno, A.; Guervos, J. J. M.; and Chacon, P., eds., *ECAL*, volume 929 of *Lecture Notes in Computer Science*, 200–218. Springer.
- Collins, R. J., and Jefferson, D. R. 1991. Selection in massively parallel genetic algorithms. In *ICGA*, 249–256.
- Curran, D., and O’Riordan, C. 2003. On the design of an artificial life simulator. In Palade, V.; Howlett, R. J.; and Jain, L. C., eds., *Proceedings of the Seventh International Conference on Knowledge-Based Intelligent Information & Engineering Systems (KES 2003)*, 549–555.
- Curran, D., and O’Riordan, C. 2006. Increasing population diversity through cultural learning. *Adaptive Behavior* 14(4).
- Darwin, C. 1859. *The Origin of Species: By Means of Natural Selection or the Preservation of Favoured Races in the Struggle for Life*. Bantam Classics.
- De Jong, K. A. 1975. *Analysis of Behavior of a Class of Genetic Adaptive Systems*. Ph.D. Dissertation, The University of Michigan.
- Eriksson, R. I. 2000. An initial analysis of the ability of learning to maintain diversity during incremental evolution. In Freitas, A. A.; Hart, W.; Krasnogor, N.; and Smith, J., eds., *Data Mining with Evolutionary Algorithms*, 120–124.
- Floreano, D., and Mondada, F. 1996. Evolution of plastic neurocontrollers for situated agents. *Animals to Animats* 4.
- Grefenstette, J. J. 1992. Genetic algorithms for dynamic environments. In Maenner, R., and Manderick, B., eds., *Parallel Problem Solving from Nature 2*, 137–144.
- Gusfield, D. 1997. *Algorithms on strings, trees and sequences: computer science and computational biology*. Cambridge Univ. Press. Gusfield.
- Hinton, G. E., and Nowlan, S. J. 1987. How learning guides evolution. *Complex Systems* 1:495–502.
- Kaufmann, S. A. 1989. Adaptation on rugged fitness landscapes. *Lectures in the Sciences of Complexity* 1:527–618.
- Mayley, G. 1997. Guiding or hiding: Explorations into the effects of learning on the rate of evolution. In *Proceedings of the Fourth European Conference on Artificial Life*. MIT Press.
- Nolfi, S., and Parisi, D. 1996. Learning to adapt to changing environments in evolving neural networks. *Adaptive Behavior* 5(1):75–97.
- Opitz, D. W., and Shavlik, J. W. 1996. Generating accurate and diverse members of a neural-network ensemble. In Touretzky, D. S.; Mozer, M. C.; and Hasselmo, M. E., eds., *Advances in Neural Information Processing Systems*, volume 8, 535–541. The MIT Press.
- 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*, volume 5, 4092–4097.
- Sasaki, T., and Tokoro, M. 1997. Adaptation toward changing environments: Why darwinian in nature? In Husbands, P., and Harvey, I., eds., *Fourth European Conference on Artificial Life*, 145–153. Cambridge, MA: MIT Press.
- Suzuki, R., and Arita, T. 2003. The baldwin effect revisited: Three steps characterized by the quantitative evolution of phenotypic plasticity. In Banzhaf, W.; Christaller, T.; Dittrich, P.; Kim, J. T.; and Ziegler, J., eds., *ECAL*, volume 2801 of *Lecture Notes in Computer Science*, 395–404. Springer.