A Note on Learning and Evolution in Neural Networks

Brian Carse
Faculty of Engineering
University of the West of England
Coldharbour Lane, Bristol BS16 1QY, England
Brian.Carse@uwe.ac.uk
Tel. +44-117-9656261

Johan Oreland
Faculty of Computer Studies and Mathematics
University of the West of England
Coldharbour Lane, Bristol BS16 1QY, England

Abstract

Interactions between evolution and lifetime learning are of great interest to studies of adaptive behaviour both in the natural world and the field of evolutionary computation. This contribution revisits an earlier discovered observation that the average performance of a population of neural networks which are evolved to solve one task is improved by lifetime learning on a different task. Two existing, and very different, explanations of this phenomenon are summarised and examined. Experimental results are presented that demonstrate that neither of these explanations are sufficient to fully explain the phenomenon. A new explanation, together with experimental justification, is presented which describes the effect in terms of lifetime learning providing a buffer against the potentially deleterious effects of mutation.

1 INTRODUCTION

In 1896 J. Mark Baldwin and others suggested a way in which characteristics acquired during the lifetime of individuals in a population could be indirectly incorporated into the genotype of future generations of these individuals (Baldwin, 1896), (Morgan, 1896). This effect, initially described as “organic selection” and subsequently renamed the “Baldwin effect”, suggests that ontogenetic adaptation by organisms in response to environmental circumstances allows the germ line of such organisms to survive long enough in future generations for the original ontogenetic adjustments to become “genetically assimilated” (Waddington, 1942). After genetic assimilation, a characteristic which previously had to be acquired during the lifetime of an individual becomes innate and can be directly passed on to offspring.

Recently, the interactions between evolution and learning (learning being one particular form of ontogenetic adaptation) have become of interest in the field of artificial evolution. The motivation for this interest appears to be two-pronged: (i) to engender a greater understanding of interplay between biological evolutionary and learning processes through computer simulation and (ii) to incorporate learning into artificial evolutionary algorithms to improve the performance of these algorithms. The artificial neural network paradigm has been widely used in such investigations since neural network parameters can be relatively easily represented for the purposes of (simple) evolution and many efficient learning algorithms are in existence.

A selection of reported research which investigates the interaction between evolution and learning in neural networks includes: (Hinton and Nolan, 1987), (Belew, 1989), (Ackley and Littman, 1991), (Gruau and Whitley, 1993), (French and Messenger, 1994), (Whitley, Scott Gordon and Mathias, 1994), (Carse and Fogarty, 1996), (Mayley, 1996), (Mayley, 1997). In most of these cases the task to solved by evolution and the learning task are the same. What happens if the evolution task and the learning task are different?

Two articles (Nolfi, Elman and Parisi, 1994) and (Harvey, 1996) provide alternative explanations of results reported in simulations in the evolution of neural networks that evolve to perform one task while being subject to lifetime learning on a different task. These simulation results, reported in (Nolfi et al., 1994), (Parisi, Nolfi and Cecconi, 1992), (Parisi and Nolfi, 1996) show that the average evolutionary fitness of the population improves when individual learning takes place alongside evolution compared with evolution alone, whereas peak individual fitness is the same in the two cases. In (Nolfi et al., 1994) these results are interpreted in terms of a "dynamic correlation” between the evolutionary fitness surface and the lifetime learning surface. In (Harvey, 1996), an alternative explanation is offered which suggests that the results may be due to a form of "relearning" of neural network weights which have been perturbed from their evolved values by mutation. Part of this latter explanation relies on the extremely high selection pressure used in the simulation experiments.

We have replicated the simulation environment described in (Nolfi et al., 1994) with good agreement on experimental results. Further experiments have been carried out in order to determine which of the two extant explanations of observed results (dynamic correlation or relearning), if either, is more likely to explain the observed phenomenon. This paper is organised as follows. The next section briefly outlines the simulations described in (Nolfi et al., 1994), their interpretation, and an alternative explanation offered in (Harvey, 1996); for more detail please refer to the original papers. Section 3 describes additional experiments carried out and results obtained in order to investigate these two alternative explanations, and suggests a possible third explanation. Finally, section 4 concludes and suggests areas for further work.

2 THE ANIMAT SIMULATION AND TWO EXTANT EXPLANATIONS

2.1 SUMMARY OF SIMULATION DETAILS

In (Nolfi et al., 1994) a population of artificial creatures ("O's" or "animats") is simulated. Each individual animat operates by moving around in a two dimensional grid world, consuming food that is randomly scattered and replenished at periodic intervals. The lifetime of each individual lasts for 5000 actions and the evolutionary fitness of the individual is the amount of food consumed during its lifetime. The movement of the animat is controlled by a neural network shown in Figure 1. Inputs to the neural network are the current motor action (two bits: 00=Do_Nothing, 01=Turn_Right, 10=Turn_Left, 11=Move_Forward) and the current sensory input (two real numbers: current distance and angle to the nearest food element). Outputs from the neural network consist of the next motor action (two bits, encoding as before) and the predicted next sensory input comprising two real numbers: predicted next distance and angle to the nearest food element after performing the motor action. Two tasks are defined: an evolutionary task which is to find and consume food, and a lifetime learning task which is to predict the next nearest food position prior to making a move. In (Nolfi et al., 1994), two sets of simulations are carried out and compared. In the first simulation, no lifetime learning is applied: all neural network weights are evolved using a genetic algorithm. In the second simulation, initial network weights are evolved and, in addition, backpropagation is applied to the neural network for modifying network weights from hidden layer nodes to the predicted food position output nodes and also the weights from input layer nodes to hidden layer nodes; the latter weight changes affect the performance of the individual on the evolutionary task during its lifetime. Weight changes (as a result of lifetime learning) are not inherited. It was found that lifetime learning on a different task to that specified for evolution improved mean population fitness on the evolutionary task. The original (Nolfi et al., 1994) explanation and an alternative explanation by Harvey (1996) are summarised next.

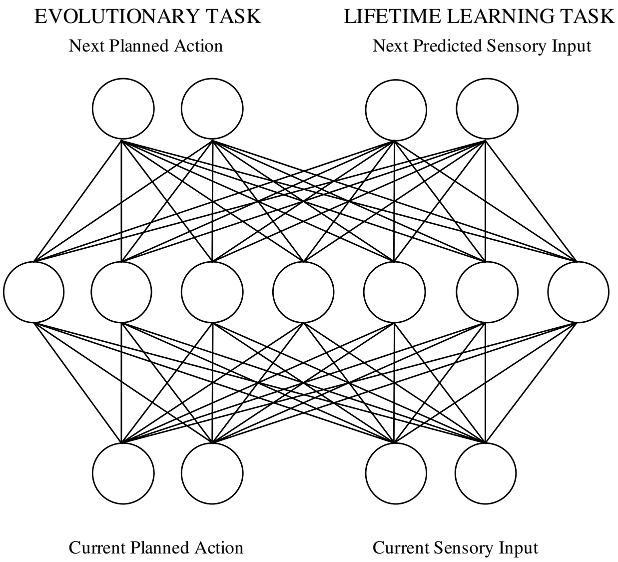

Figure 1: Neural Network used in Experiments (Nolfi et al. 1994)

2.2 EXPLANATION #1-DYNAMIC CORRELATION

Nolfi et al. (1994) introduce the notion of "dynamic correlation" between the evolutionary fitness surface and the lifetime learning surface:

> ".. *evolution selects for initial points in the weight space such that when these points move because of learning, their movement brings them to new locations that correspond to higher levels not only on the learning surface (this is implicit in their learning the learning task) but also on the fitness surface. These points are located in regions of dynamic correlation between the two surfaces.*" (Nolfi et al., 1994)

In a later paper (Parisi and Nolfi, 1996), regarding the same phenomenon, it is stated:

> "*Learning is a mechanism which allows evolution to find out about the regions in weight space surrounding the location of candidates for reproduction...If we assume that at least some of these locations are situated in the surrounding region explored by mutation, we can see that in populations that both evolve and learn, reproductive decisions can be based on knowledge of surrounding regions. In other words, evolution is based on the fitnesses of the currently living individuals but it would be more effective if it could be based also on the fitnesses of the offspring of these individuals. Learning allows evolution 'to look into the future'. Therefore, evolution with learning can be more effective than evolution alone.*"

2.3 EXPLANATION #2 - RELEARNING AFTER MUTATION

In (Harvey, 1996) a totally different explanation for the results reported in (Nolfi et al., 1994) is presented. Harvey explains the improvement in mean fitness but not best fitness as a form of "relearning"; this argument runs briefly as follows. The selection pressure used in these simulations is extremely high, with the elite taking over the population in as few as 3 generations without mutation. With mutation, the population consists of almost identical individuals with some weights perturbed around those of the current elite. The effect of lifetime learning is to bring these weights perturbed by mutation back to nearer their unperturbed values.

In (Harvey and Stone, 1996) it is demonstrated that such a relearning effect occurs with neural networks which are first trained on one task, A, then trained on a different task, B, to perturb the network weights away from those learned on task A. When the network is trained on task C, unrelated to task A, a transient improvement in performance on task A is observed. A longer lived improvement on performance on task A with training on task C is observed when the network weights learned by the initial training on task A are perturbed randomly, rather than learning on task B. Harvey and Stone suggest that this is a generalisation of a similar spontaneous recovery effect in neural networks described in (Hinton and Sejnowski, 1986).

In (Harvey, 1997) a similar effect to that described in (Nolfi et al., 1994) is reported using genotypes consisting of 50 real numbers. Individuals are evolved to solve an evolutionary task (defined as minimising the distance from the vector, represented by a population individual, to a fixed target vector **E**). Simulations are described using evolution of populations of such vectors with and without learning, the former employing a single application of the delta rule during the lifetime of an individual to move the innate vector towards a target learning vector **L** (distinct from **E**). Each individual vector, **G**, is moved by learning to a new vector $\mathbf{G'} = \mathbf{G} + \delta(\mathbf{L}-\mathbf{G})$, and the fitness of the resulting genotype evaluated as $|\mathbf{G'}-\mathbf{E}|$. Using similar GA parameters to those used in (Nolfi et al., 1994), Harvey reports a very similar result in that population average performance on the evolutionary task is improved when learning is applied compared to evolution only.

In (Harvey, 1997) the effect is analysed geometrically in terms of the movements in vector space brought about by learning using the delta rule, and using these arguments it is shown that the effect can be explained as being due to recovery (through learning) of weight perturbations caused by mutation. This effect, which Harvey calls "Another New Factor", or ANF, is likely to be observed when applying limited learning during evolution on an unrelated task, with relatively large mutation values, and using a real numbered genotype. In (Harvey, 1997) it is suggested that it is possibly the ANF effect, and not the dynamic correlation effect which may explain the results reported in (Nolfi et al., 1994).

3. EXPERIMENTAL STUDY AND DISCUSSION

3.1 THE SIMULATION ENVIRONMENT

As stated earlier, the simulation environment used in (Nolfi et al., 1994) has been recreated. There may be some small differences in detail but the main parameters are the same, namely: the same neural network architecture (see Figure 1); the same 10 x 10 grid world; a population of 100 individuals; 5000 actions in a “lifetime”; a GA strategy of selecting the best 20 individuals per generation, making 5 copies of each, mutating five randomly selected weights within each new individual by adding a random number in the range [-1.0,+1.0] to selected weights to produce the next generation of individuals; periodically reintroducing food every 50 actions; a backpropagation learning rate of 0.2. We have run the GA for only 50 generations since this is sufficient to produce the mean population fitness effect described in (Nolfi et al., 1994).

3.2 BEST AND MEAN POPULATION FITNESS WITH AND WITHOUT LEARNING

Figure 2 shows the average fitness of the fittest individual at each generation versus generation number using a GA only and using a GA with lifetime learning. Figure 3 shows the mean population fitness for the two cases versus generation number. Both these graphs show the average over 20 runs. The main features of these figures match closely with those of the corresponding graphs in (Nolfi et al., 1994) and show that the best fitness is not changed by adding lifetime learning whereas the population mean fitness is improved by adding lifetime learning.

3.3 MEAN OFFSPRING FITNESS BEFORE MUTATION

To investigate the effects of mutation on the mean population fitness, the average fitness of offspring before mutation is applied was measured. This is a new experiment and is intended to determine whether the improved mean populaton fitness for individuals evolved with learning appears before or after mutation of offspring. This is very simple to do, given the way the GA is employed, by taking the average fitness of the 20 fittest individuals (“parents”) from the previous generation. The current generation, before mutation, contains five copies of each of these 20 fittest individuals. Figure 4 shows graphs of the average fitness before mutation without and with lifetime learning versus generation number. These two graphs are virtually identical and the average population fitness before mutation is the same irrespective of whether or not lifetime learning is applied. The benefit of lifetime learning on mean population fitness only appears after mutation has been applied. This suggests that learning is not guiding evolution in the sense of finding better individuals but rather that lifetime learning is diminishing the detrimental effect of mutation on the average population fitness, possibly as a result of relearning of perturbed weights, as suggested in (Harvey, 1997).

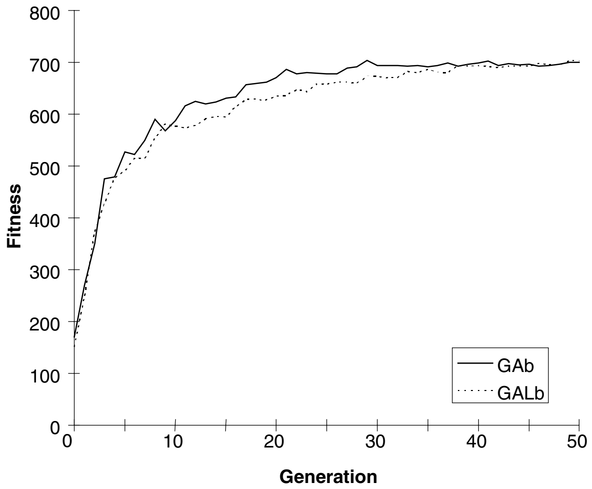

Figure 2: Best Fitness (Food Eaten) versus Generation Number using Nolfi et al's GA Parameters without Learning (GAb) and with Learning (GALb)

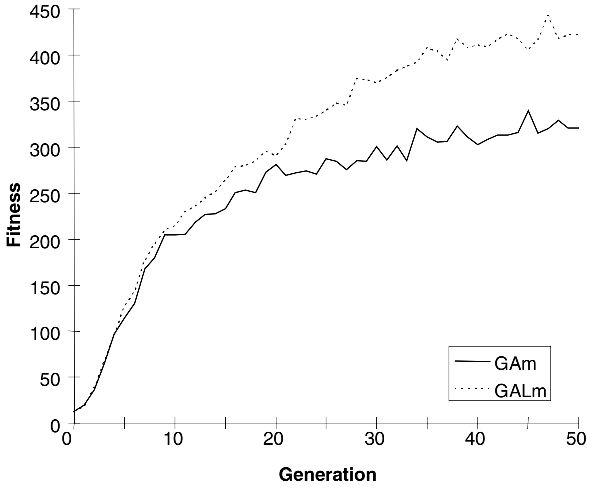

Figure 3: Mean Fitness (Food Eaten) versus Generation Number using Nolfi et al.'s GA Parameters without Learning (GAm) and with Learning (GALm)

3.4 MUTATION OF EVOLVED ELITE INDIVIDUALS

To further investigate this, experiments were carried out on the fittest individuals in the final populations (i.e. at generation 50) of the 20 GA runs with lifetime learning and the 20 GA runs without learning (giving 20 elite individual networks in each case). These experiments evaluated the mean fitnesses of evolved individuals with and without learning (whether or not learning had been applied during their evolution); and also the effects of mutation on the fitness of these individuals.

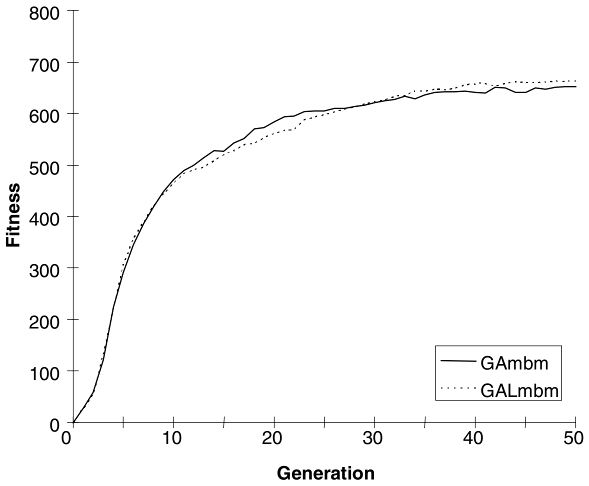

Figure 4: Mean Fitness (Food Eaten) Before Mutation versus Generation Number using Nolfi et al's GA Parameters without Learning (GAmbm) and with Learning (GALmbm)

In the former cases (no mutation), each individual was evaluated only once (on carrying out multiple evaluations, the observed fitness showed little variance). In the cases where mutation was applied, as should be expected with such a random operator, different mutations on the same individual resulted in different individuals with very different fitnesses. Where mutation was applied, 2000 mutant individuals were generated from each parent individual and their resultant fitnesses averaged. The results of these experiments are shown in Table 1.

Table 1. Results of Experiments on Evolved Elite Individuals

I. Average fitness of individuals evolved with GA only (standard deviations in brackets):	
Ia. Fitness with no learning	674(83)
Ib. Fitness with learning	242(150)
Ic. Fitness after mutation with no learning	351(62)
Id. Fitness after mutation with learning	192(96)
II. Average fitness of Individuals evolved with GA and Learning (standard deviations in brackets)	
IIa. Fitness with no learning	490(137)
IIb. Fitness with learning	646(93)
IIc. Fitness after mutation with no learning	315(66)
IId. Fitness after mutation with learning	406(57)

It should come as no surprise that elite individuals evolved without learning perform more poorly when learning is applied (Table 1: Ia compared to Ib): these individuals already have good neural network weights and any changes to these weights (e.g. by learning a different task as is the case here) is likely to reduce fitness. Individuals evolved with learning (Table 1: IIb compared to IIa) have significantly lower innate fitness compared to their fitness when learning is applied. These individuals do appear to have been "born" in regions of weight space where improvement on the learning task causes improvement on the evolutionary task, i.e. the regions of dynamic correlation suggested by Nolfi et al.

The other results presented in Table 1 offer a further insight into what is going on in relation to the ANF relearning hypothesis in these simulations. In (Harvey, 1996), in relation to Nolfi et al's simulation results, it is stated that:

> "*If one substituted for the elite member of a population evolved on the food-finding task one individual trained by backpropagation using an external teacher (or any other learning mechanism) on that same task, then one should expect similar responses after weight perturbations*"

In Table 1, results are presented (Ic and Id) in which elite members of populations evolved without learning are mutated (i.e. have their weights perturbed), and then evaluated with and without learning. Although these elite individuals have been evolved, they represent good solutions to the animat problem that might have been learned by any other learning mechanism. The effect of applying learning to mutants of these elite individuals appears to actually reduce their fitness, the opposite of that predicted by Harvey's ANF effect, and suggesting that for these individuals learning is not causing a recovery of weights from their mutated (perturbed) values.

In the case of elite individuals evolved with learning, applying the same learning task to their mutants does increase their fitness (Table 1, IIc and IId). We have also carried out some experiments on mutants of elite individuals evolved with learning using a different learning task than that which was used during evolution. This was done simply by having these mutated individuals learn to predict the perceived angle to the nearest food element in an anticlockwise sense, as opposed to the clockwise sense which was used during evolution of the parent individuals (inputs to these individual neural networks were, of course, kept in the clockwise direction). The ANF relearning explanation would suggest that performance of these mutated individuals should improve with learning. We found in our experiments that this was not the case, and training on a task different to that used during evolution/learning reduced the performance of (mutated elite) individuals on the evolutionary (food finding) task.

The magnitude of the decrease in fitness (on the food finding task) caused by mutation was measured for each mutant with and without learning. These results are depicted in Figure 5.

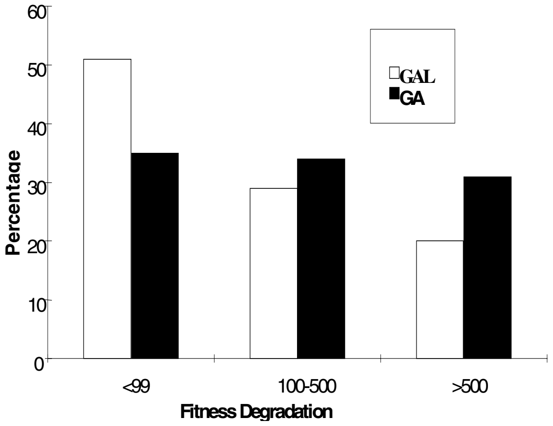

Figure 5: Percentage of Mutants versus Range of Fitness Degradation

This figure shows the percentage of mutants which suffered fitness degradations in the ranges shown. From Figure 5, it appears that elite individuals evolved with learning are somehow more resistant to deleterious mutations (this result could be explained by the ANF effect, but as already stated, this effect does not appear to be in operation in this specific case).

3.5 A NEW EXPLANATION

First, we note that the Nolfi et al. simulations threshold the outputs of the neural network which determine the next movement action. Actual network outputs are in the continuous range [0,1] and these are converted into binary values using a threshold of 0.5 for the purpose of determining the next action. A large number of previous machine learning experiments using neural networks (using either supervised, gradient-based learning or evolutionary learning or combinations of these) have been carried out in simulations which also employ a grid world and thresholding of neural network outputs to determine a discrete action from real-valued outputs. With such thresholding a relatively small change in neural network weights can result in very different animat behaviours. For example, consider the case where the two network outputs in the Nolfi. et al. simulation which determine the next movement are respectively 0.51 and 0.49 (interpreted as 10 in binary i.e. turn left). If minor changes in network weights (either due to backpropagation learning on a second task, or mutation) cause the outputs to change to 0.49 and 0.51 respectively, thresholding then interprets these outputs as 01 in binary, i.e. turn right - a completely different action. The implication of this is that the fitness landscape (in the space of neural network weights) for the evolutionary task is likely to have "plateaux" of high fitness and "plains" of low fitness separated by relatively steep "cliffs". The fitness landscape is also likely to exhibit a high degree of neutrality (Kimura, 1983).

It might be the case that it is this particular feature of the fitness landscape that explains the results obtained in the simulations of Nolfi et al. and our own results. Without learning (and the associated changes of network weights during lifetime) a high fitness individual in the evolving population has no way of "knowing" (through fitness evaluation) if it is close to a cliff leading down to an area in weight space with low fitness and an elitist GA strategy will maintain such individuals in the population. Mutations of such individuals are quite likely to produce low fitness offspring. With learning, a local exploration of the evolutionary fitness landscape is made during an individual's lifetime. Of course this latter observation is not new and represents the first step in the Baldwin Effect, although, in the literature, this effect is usually described in terms of learning improving the chances of genetic hillclimbing to get to higher points in the fitness landscape. However, what we are proposing here is something subtly different. Rather than learning improving the rate of evolution of increasingly fit individuals (certainly not observed in the simulations of Nolfi et al. or our own), learning appears to be endowing individuals with an immunity to the potentially deleterious effects of mutation when they reproduce. In other words, when learning and the associated local search are operating, evolving individuals are likely to move away from cliffs in the fitness landscape (in this case, collections of network weights which produce outputs close to the 0.5 threshold).

As a first step in investigating this, we conducted the following experiments. For the 20 elite individuals evolved without learning, and the 20 elite individuals evolved with learning, offline simulations on the food finding (evolutionary) task were carried out and the distances of both the real-valued outputs of the network from the 0.5 threshold were measured for each action and summed over the complete 5000 action simulation. In the case of individuals evolved with learning, learning was also applied during evaluation. For the 20 individuals evolved without learning the average of these summed distances was 2094 (standard deviation 406); for the 20 individuals evolved with learning, the corresponding figure was 2375 (standard deviation 357). These values are significantly different (95% confidence) suggesting that the individuals evolved without learning are operating nearer the 0.5 threshold.

Of the two network outputs which determine the next animat output, the one which operated nearer the 0.5 threshold was investigated. This is the output most likely to be affected by a deleterious mutation. For each evolved network (with and without learning) the distance of this output from the 0.5 threshold was measured at each step and summed over the 5000 steps of the simulation. Figure 6 shows a scatter plot of this total distance versus the previously measured average fitness degradation after mutation for individuals evolved without learning. Each point represents the (distance, fitness degradation) pair for each of the 20 individuals. The correlation coefficient is -0.89. Figure 7 shows a similar scatter plot for individuals evolved with lifetime learning. The correlation coefficient in this case is -0.6.

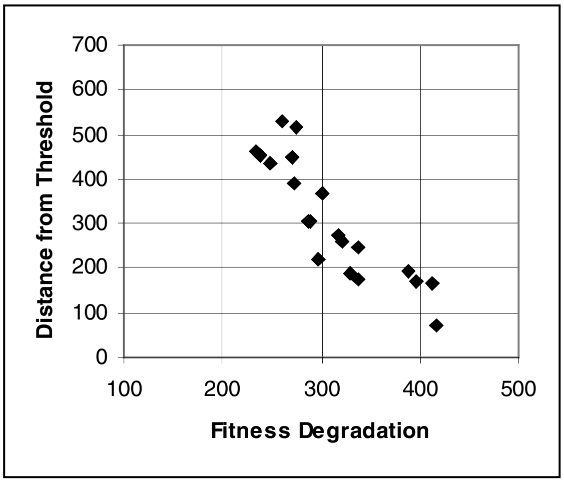

Figure 6: Scatter Plot of Minimum Distance from Threshold versus Fitness Degradation Caused by Mutation for Individuals Evolved without Learning (Correlation Coefficient = -0.89)

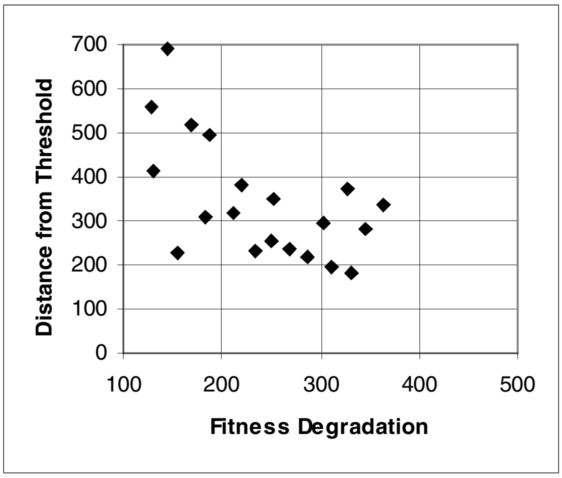

Figure 7: Scatter Plot of Minimum Distance from Threshold versus Fitness Degradation Caused by Mutation for Individuals Evolved with Learning (Correlation Coefficient = -0.60)

It should be expected that, even without learning, the average distance that evolved individuals operate from the 0.5 NN output threshold should increase as evolution proceeds. A further experiment was conducted whereby the average distance of neural network outputs from the 0.5 threshold during individual lifetimes for the whole population was measured with and without learning over 200 generations. The results of this experiment are shown in Figure 8. In both cases (with and without learning), the summed initial distance from the 0.5 threshold at generation 0 is approximately the same. In very early generations this distance falls in both cases to around 200. As evolution proceeds, populations evolved both with and without learning show an increase in distance from the threshold.

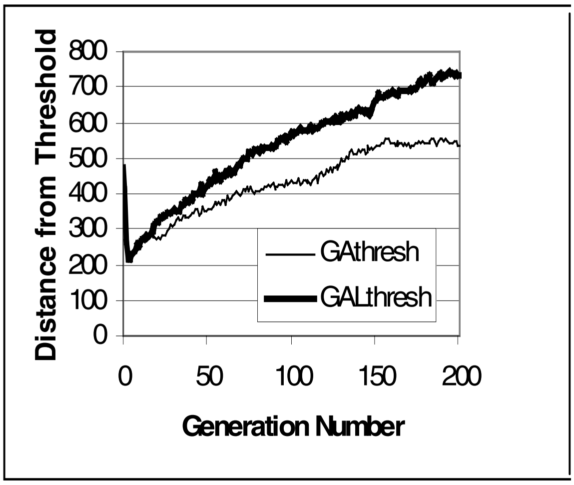

Figure 8: Distance from 0.5 Threshold of Neural Network Action Outputs versus Generation Number for Populations Evolved without Learning (GAThresh) and Populations Evolved with Learning (GALThresh)

However, populations of networks which employ learning move more quickly away from operating near the threshold compared with populations of networks which do not employ learning. In both cases, evolution appears to be reducing the potentially deleterious effects of mutation, but this happens faster if learning takes place.

While these are early results, the observations would seem to support the hypothesis that it is distance from the 0.5 threshold of an individual's network outputs which determines the resilience of the individual to mutation, and that individuals evolved with learning produce outputs which are significantly further away from this threshold than individuals evolved without learning. The latter is due to the fact that local exploration (as a result of lifetime learning) penalises individuals which produce outputs close to the crucial threshold. Highly fit individuals evolved without learning but which produce network outputs close to the threshold are not "put to the test" in this way and are therefore more likely to suffer greater reduction in fitness as a result of mutation. This could explain Nolfi et al.'s original result that, in these and similar simulations, mean population fitness is improved by including lifetime learning on a different task to that which is being addressed by evolution.

4. CONCLUSIONS

The simulations carried out by Nolfi et al. (1994) have been replicated and their published experimental results that lifetime learning on a task different to the evolutionary task improves average population fitness, but not peak population fitness, on the evolutionary task have been confirmed. Additional experimental results have been presented which attempt to investigate two differing explanations of the effect: namely, Nolfi et al.'s "dynamic correlation" explanation and Harvey's "relearning" explanation. While a form of dynamic correlation is undoubtedly occurring in these simulations (individuals evolved with learning perform poorly when learning is removed) this effect does not seem to explain the result that mean population fitness improves when lifetime learning is used. Furthermore, Harvey's relearning effect does not appear to be operating in these simulations (although the effect does undoubtedly occur in other situations).

An alternative explanation of the phenomenon has been tentatively proposed. This explanation relies on the fact that in these, and several other simulation studies, real-valued neural network outputs are thresholded to produce discrete action outputs. The effect of this is likely to be a highly discontinuous fitness landscape with valleys and plateaux separated by steep cliffs. It appears to be that lifetime learning is endowing individuals with greater resistance to the deleterious effects of mutation in such a fitness landscape. Further work is required in order to more fully test this hypothesis and investigate its implications.

References

Ackley D. and Littman M. (1991). Interaction between learning and evolution. In Langton C., Taylor C., Farmer J. and Rasmussen S. (Eds.), Artificial Life II: Proceedings of Santa Fe Conference, February 1990. pp487-509. Reading MA: Addison-Wesley.

Baldwin J.M. (1896). A New factor in evolution, American Naturalist *30*, 536-553.

Belew R.K. (1989). When both individuals and populations search: adding simple learning to the genetic algorithm. In Schaffer J.D. (Ed.*)*, Proceedings of the Third International Conference on Genetic Algorithms. pp34-41, San Matteo, CA. Morgan Kauffmann,

Belew R.K. and Mitchell M. (1996) (Eds.). Adaptive individuals in evolving populations. Addison-Wesley, Reading MA.

Cecconi F. and Parisi D. (1991). Evolving organisms that can reach for objects. In Meyer J. and Wilson S. (Eds.) Proceedings of 1st International Conference on Simulation of Adaptive Behaviour pp391-399. MIT Press, Cambridge MA.

Carse B. and Fogarty T.C. (1996) 'Tackling the "curse of dimensionality" of radial basis function neural networks using a genetic algorithm'.In Voigt H-M, Ebeling W., Rechenberg I and Schwefel H-P(eds.) Proceedings of the Fourth Conference on Parallel Problem Solving from Nature, Springer-Verlag Lecture Notes on Computer Science, vol.1141, pp710-719.

French R.M. and Messinger A. (1994). Genes, phenes and the Baldwin effect: learning and evolution in a simulated population. In Brookes R.A. and Maes P. (Eds.*)*, Artificial Life IV: Proceedings of the Fourth International Workshop on Synthesis and Simulation of Living Systems. pp277-282, Cambridge MA: MIT Press.

Gruau F. and Whitley D. (1993). The cellular development of networks: the interaction of learning and evolution. Research Report No. 93-04, Lyon: Ecole Normale Supérieure, Laboratoire de l'Informatique du Parallélisme.

Harvey I. (1996). Relearning and evolution in neural networks, Adaptive Behaviour, *4(1)*, 81-84.

Harvey I. and Stone J.V. (1996). Unicycling helps your French: Spontaneous recovery of associations by learning unrelated tasks. Neural Computation, *8*, 697-704.

Harvey I. (1997). Is there another new factor in evolution, Evolutionary Computation, *4(3)*, 313-329.

Hinton G.E. and Nowlan S.J. (1987). How learning can guide evolution. Complex Systems, *1*, 495-502.

Hinton G.E. and Sejnowski T. (1986). Learning and relearning in Boltzmann machines. In Rumelhart D., McClelland J. and the PDP Research Group (Eds.), Parallel Distributed Processing: Explorations in the Microstructure of Cognition. Volume 1: Foundations, pp282-317. MIT Press/Bradford Books, Cambridge, MA.

Kimura M. (1983). The Neutral Theory of Molecular Evolution. Cambridge University Press.

Mayley G. (1996). The Evolutionary cost of learning. In Maes P., Mataric M., Meyer J., Pollack J. and Wilson S. (Eds.). From Animals to Animats 4: Proceedings of the 4th International Conference on Simulation of Adaptive Behaviour, pp458-467. MIT Press.

Mayley G. (1997). Landscapes, learning costs and genetic assimilation. Evolution, Learning and Instinct: A Special Edition of Evolutionary Computation *4(3)*, 213-234.

Morgan C. Lloyd (1896). On modification and variation. Science IV *(99),* November 20, 1896. pp733-740

Nolfi S., Elman J. and Parisi D. (1994). Learning and evolution in neural networks, Adaptive Behaviour, *3*, 5-28.

Parisi D., Nolfi S. and Cecconi F. (1992). Learning, behaviour and evolution. In F.J. Varela & P. Bourgine (Eds.), Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life. Cambridge MA: MIT Press/Bradford Books.

Parisi D. and Nolfi S. (1996). The Influence of learning on evolution. In R.K. Belew & M. Mitchell (Eds.), Adaptive Individuals in Evolving Populations, SFI Studies in the Sciences of Complexity, Vol. XXVI, Addison-Wesley

Waddington C.H. (1942) Canalization of development and the inheritance of acquired characters. Nature 3811, November 14, 1942, pp563-565.

Whitley D., Scott Gordon V. and Mathias K. (1994). Lamarckian evolution, the Baldwin effect and function optimisation. In Davidor Y., Schwefel H.P. and Männer R. (Eds.), Proceedings of the Third Conference on Parallel Problem Solving from Nature, Springer-Verlag Lecture Notes on Computer Science, vol.866, pp6-15, Springer Verlag.