

Relearning and Evolution in Neural Networks

Inman Harvey

School of Cognitive and Computing Sciences

University of Sussex

Brighton BN1 9QH, UK

email: inmanh@cogs.susx.ac.uk

tel: +44 1273 678754

fax: +44 1273 671320

To appear in *Adaptive Behavior*, v.4 no.1, (1996) pp. 79-82

Relearning and Evolution in Neural Networks

Inman Harvey

A recent paper in *Adaptive Behavior* (Nolfi, Elman, & Parisi, 1994) reported simulations of populations of neural networks that evolve (to get fitter at one task) at the population level and may also learn (a different task) at the individual level. One result stated was that average fitness at the evolutionary task is improved when lifetime learning at the different task is introduced. A different explanation will be proposed here for much of the data there presented: that the main results are an artefact of the unconventional evolutionary algorithm used, and can be interpreted rather differently as a form of relearning.

Asexual evolution (mutation only) was used on a population of 100 individuals or animats; the connection weights were genetically specified for a feedforward network for each individual, which transformed sensory inputs of the animat into movements over a grid-like environment on which food had to be found. Mutation of offspring of selected parents perturbed the values of 5 of their weights chosen at random.

The selective pressure used was exceptionally strong. Whereas in population genetics selective differences are typically of the order of 1%, and with conventional genetic algorithms selective pressures are kept low to avoid premature convergence, here the fittest members have 500% more offspring than the average — the top 20 out of 100 each have 5 offspring. In the absence of mutation such selection results in the elite taking over the whole population in just 3 generations (from 1% to 5% to 25% to 100%). With mutation the population consists of near-identical individuals with some weights perturbed around those of the current elite. This premature convergence was mentioned by (Williams & Bounds, 1993) in their analysis of a forerunner of this paper (Parisi, Nolfi, & Cecconi, 1992).

The performance of the elite did not improve when lifetime learning of the second task was introduced, whereas average performance did improve. It seems clear that the effect of lifetime learning was merely to go some way towards restoring performance of networks which had had their weights perturbed (by mutation) away from trained (through evolution) values — a form of relearning. The extreme convergence of the population around the clustered elite members of the previous generation should be borne in mind when reading from (Nolfi et al., 1994), p. 22:

The offspring of a reproducing individual occupy initial positions in weight space that are deviations (due to mutations) from the position occupied by their parent at birth (i.e., prior to learning).

One form of relearning in networks was analysed in (Hinton & Plaut, 1987). In that case a network is first trained by some learning algorithm on a set of input/output pairs; the weights are then perturbed. After retraining on a subset of the original training set, it is found that performance improves also on the balance of the original training set. The present case differs from this, in that the lifetime learning is on a fresh task, rather than on a subset of the original task. Recently just such an effect was predicted and observed in networks (Harvey & Stone, 1995). When good performance on one task is degraded by random perturbations of the weights, then in general training on any unrelated second task can be expected to improve, at least initially, the performance on the first task.




Figure 1: A two-dimensional sketch of weight space.

To briefly summarise the reasons for this, consider the diagram, which represents the weight space of a network in just 2 dimensions. A represents the weights of a network trained on task \mathcal{A} (cf. the evolved elite of the previous generation). B_1 and B_2 are two possible positions for B representing a set of perturbed weights (cf. mutated offspring). Trajectories $B_i \rightarrow C$ represent movement through weight space due to training on a task \mathcal{C} (cf. the lifetime learning task) which is unrelated to \mathcal{A} . It can be seen that trajectory $B_1 \rightarrow C$ passes initially within the circle and hence closer to A in weight space, whereas $B_2 \rightarrow C$ does not. If one assumes that distance from A is locally correlated with degradation in performance on \mathcal{A} , then in the former case performance on \mathcal{A} improves, in the latter case it does not. Any position of B falling on the larger arc PQ rather than the smaller arc PQ should then improve performance. In the two-dimensional weight space illustrated it can be seen that this is more than 50% of the time regardless of the position of C , and in high-dimensional weight spaces this probability increases closer to 100%; if C happens to be within the circle it should always be 100%. Simplifying assumptions that have been used are discussed in (Harvey & Stone, 1995).

There are interesting interactions between learning and evolution such as the Baldwin effect (Hinton & Nowlan, 1987). But it is suggested here that the main effects described in (Nolfi et al., 1994) constitute a form of relearning after weight-perturbation. A further effect is described in that paper, namely an improvement in the rate of individual learning in later generations as compared to earlier. But the data shown there (for instance in Figure 7 for generation 99) effectively gives only a single data point for a converged population of near-identical individuals, rather than the average performance of 100 very different individuals. Until more data is presented it remains a possibility that this single data point is within normal variations, rather than evidence of a reliable effect.

In summary, there are two significant ways in which this new explanation of the effects differs from that given in (Nolfi et al., 1994). Firstly, the effect is not due specifically to some interaction between learning and evolution. If one substituted for the elite member of a population evolved on the food-finding task one individual trained by back-propagation using an external teacher (or any other learning mechanism) on that same task, then one should expect similar responses after weight-perturbations.

Secondly, one should in general expect a *decrease* in food-finding performance in the elite evolved member (or otherwise trained individual) due to learning on the second task¹. The *only* increase in food-finding performance should come from those whose weights have been perturbed away from the values of the peak or elite member, by mutation or any other method. Under this new explanation one would expect this improvement in performance to frequently be only temporary, and for continued training on the second task to then eventually cause decrease in food-finding performance (whenever C in the diagram is outside rather than inside the circle).

References

- Harvey, I., & Stone, J. V. (1995). Unicycling helps your French: Spontaneous recovery of associations by learning unrelated tasks. Tech. rep. CSR 379, School of Cognitive and Computing Sciences, University of Sussex. Submitted for journal publication.
- Hinton, G., & Plaut, D. (1987). Using fast weights to deblur old memories. In *Proceedings of the 7th Annual Conference of the Cognitive Science Society, Seattle*.
- Hinton, G. E., & Nowlan, S. J. (1987). How learning can guide evolution. *Complex Systems*, 1, 495–502.
- Nolfi, S., Elman, J., & Parisi, D. (1994). Learning and evolution in neural networks. *Adaptive Behavior*, 3(1), 5–28.
- Parisi, D., Nolfi, S., & Cecconi, F. (1992). Learning, behavior and evolution. In Varela, F. J., & Bourgine, P. (Eds.), *Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*, pp. 207–216. MIT Press/Bradford Books, Cambridge, MA.
- Williams, B., & Bounds, D. (1993). Learning and evolution in populations of backprop networks. In Deneubourg, J., Bersini, H., Goss, S., Nicolis, G., & Dagonnier, R. (Eds.), *Proceedings of the Second European Conference on Artificial Life*, pp. 1139–1149. Free University of Brussels, Brussels.

¹This may be masked in Figure 4 of the original paper as the best individual after training should usually be a different individual to that best before training.