ANOTHER REASON WHY WE SHOULD LOOK AFTER OUR CHILDREN

JOHN A. BULLINARIA
School of Computer Science, The University of Birmingham
Edgbaston, Birmingham, B15 2TT, UK
j.bullinaria@physics.org

In many ways, it seems obvious that we should look after, feed and protect our children. However, infants of some species are expected to look after themselves from a very early age. In this paper, I shall present a series of simulations that explore the hypothesis that neural network learning issues alone are sufficient to result in the evolution of long protection periods. By evolving populations of neural network systems that must learn to perform simple classification tasks, I show that lengthy protection periods will emerge automatically, despite associated costs to the parents and children.

1. Introduction

Most humans accept that it is part of their role as parents to look after their children until they are old enough to fend for themselves, and it is clear that the children would have an extremely low survival rate if that did not happen. But why have humans evolved to be like that?

Many species are *precocial*, with young that are born well developed and requiring very little parental care. Others are *altricial*, with relatively helpless young requiring periods of parental care before they are able to survive on their own. Human infants are particularly altricial, even compared with other primates, requiring extended periods of parental protection and support (e.g., Lamb, Bornstein & Teti, 2002). For altricial species there are usually two important processes happening during the protection stage – the infants are growing, and they are learning. Human infants do require a lot of growing after birth, and parental protection does provide obvious advantages in term of survival. But why such extended periods compared with other primates? The need to learn will depend on how much innate knowledge the individuals are born with. It is likely that learning is crucial when relatively complex behaviour is required, or when the properties of the environment are variable and each new-born infant needs to learn to adapt accordingly. Individuals will also need to learn to adapt their control processes to compensate for changes caused by their growing (Bullinaria, 2003a). All these processes are clearly applicable to humans. It seems, however, that humans do have excessively long protection periods, and many parents may wonder if they really do need to look after their children for quite so long. Their children might also wonder if they wouldn't be better off "leaving home" and embarking on their reproductive careers at an earlier age.

In this paper, I explore, through a series of simulations, one reason why evolution might have favoured long protection phases in humans. There are actually many possible reasons (e.g., see Sloman & Chappell, 2005), but here I shall focus on the hypothesis that learning issues alone are sufficient to result in the evolution of long protection periods. Moreover, since most human learning takes place in their brains, it is neural network learning that will be studied. I have previously run simulations of the evolution of populations in which learning individuals of all ages compete for survival according to their performance on simplified classification tasks. Not surprisingly, individuals evolved that not only learned how to perform well, but were also able to learn *quickly* how to achieve that good performance. However, it was also observed that the pressure to learn quickly could also have the unfortunate side effect of leading to risky learning strategies that sometimes resulted in very poor performance (Bullinaria, 2007).

In this study, I shall again present results from evolved neural network systems that must learn to perform well on simple classification tasks. One might consider reducing the computational resource requirements of the simulations by using a learning mechanism, or an approximation of learning, that is simpler than an artificial neural network. The problem with attempting this is that the error distributions and associated fitness levels during learning depend in a complex manner on the learning algorithm and its evolved parameters, and these in turn depend in a non-trivial way on the evolutionary pressures and population age distributions which are affected by the protection period we are attempting to study. It is almost impossible to predict what distributions of all these things will emerge across the evolving populations. With so many unknowns and complex interactions, the only reliable way to proceed in the first instance is to run the full evolutionary neural network simulations. Future studies will then be able to safely abstract out the key features for exploration of further issues.

The remainder of this paper will show that evolved neural network systems do exhibit better adult performance if protection from competition is provided during the children's early years. Moreover, if the length of the protection period is allowed to evolve, it does result in the emergence of relatively long protection periods, even if there are other costs involved, such as the children not being allowed to reproduce during their protection phase, and the parents suffering increased risk of dying while protecting their young.

2. Evolving Neural Network Systems

The idea here is to mimic the crucial features of the evolution of most animal populations, but concentrate on the aspects of fitness associated with neural network learning. We therefore take a whole population of individual neural networks, each specified by a set of innate parameters, and expect them to learn from a continuous stream of input patterns how to classify future input patterns. Those inputs could, for example, correspond to specific features of other animals, and the desired output classes could correspond to being dangerous, edible, and such like. Each individual then has a fitness measure determined by how well it classified each new input *before* discovering (somehow) its correct class and learning from it. If the individuals compete to survive and procreate according to their relative fitnesses, we can expect populations of increasing fitness to emerge.

To proceed empirically, we need to concentrate on a specific concrete system, and it makes sense to follow one that has already proved instructive in the past (Bullinaria, 2007). Real-world classification tasks typically involve learning non-linear classification boundaries in a space of real valued inputs. Taking the set of classification tasks corresponding to two dimensional continuous input spaces with circular classification boundaries proves simple enough to allow extensive simulations, yet involves the crucial features and difficulties of real world problems. Each "new-born" neural network is assigned a random classification boundary which it must learn from a stream of randomly drawn data points from the input space, that we can take to be normalized to a unit square. The natural performance measure we shall use is the generalization ability, i.e. the average number of correct outputs (e.g., output neuron activations within 0.2 of their binary targets) before training on them.

We shall take our neural networks to be traditional fully connected Multi-Layer Perceptrons with one hidden layer, sigmoidal processing units, trained by gradient descent using the cross-entropy error function. As previous studies have shown (Bullinaria, 2003b), one gets better performance by evolving separate learning rates η_L and initial weight distributions $[-r_L, +r_L]$ for each of the four distinct network components L (the input to hidden weights IH, the hidden unit biases HB, the hidden to output weights HO, and the output unit biases OB), rather than having the same parameters across the whole network. These, together with the standard momentum parameter α and weight decay regularization parameter λ, result in ten evolvable innate parameters for each network. It is also possible to evolve the number of hidden units, but since the evolution invariably results in the networks using the maximum number we allow, slowing down the simulations considerably, we keep this fixed at 20 for all networks, which is more than enough for learning the given tasks.

For the simulated evolution, we need a single unit of time that covers all aspects of the problem, so we define a "simulated year of experience" to be 1200 training data samples, and compute the fitness of each individual at the end of each year as an average over that year. This number ensures that each individual has its performance sampled a reasonable number of times during its learning phase. Then, using random pair-wise fitness comparisons (a.k.a. tournaments) at the end of each year, we select up to 10% of the least fit individuals to be killed by competitors and removed from the population. In addition, to prevent the populations being dominated by a few very old and very fit individuals, a random 20% of individuals aged over 30 simulated years die of old age each year and are removed from the population.

A fixed population size of 200 is maintained throughout (consistent with the idea that there are fixed total food resources available to support the population), with the removed individuals being replaced by children generated from random pairs of the most fit individuals. Each child inherits innate parameters that are chosen randomly from the corresponding range spanned by its two parents, plus a random mutation (from a Gaussian distribution) that gives it a reasonable chance of falling outside that range. These children are protected by their parents until they reach a certain age and cannot be killed by competitors before then. The direct cost to the children is that they are not allowed to have any children of their own before they leave the protection of their parents. The implicit cost to the parents is that, the more the children are protected, the higher the chance they stand of being in the 10% of the population that are killed each year. In practice, a protected child might receive more from its parents than simple protection (e.g., teaching as well), but throughout these simulations we avoid all potential confounds by always using the set-up that is *least* likely to lead to a positive learning effect.

Various aspects of this basic evolving neural network idea have already been explored in some detail elsewhere (Bullinaria, 2001, 2003a, 2003b, 2007). The two crucial new issues to be investigated here are:

1. How does protecting the children affect the individuals' performance?
2. If the duration of the protection period is free to evolve, what happens?

The first question can be conveniently explored by fixing the protection period

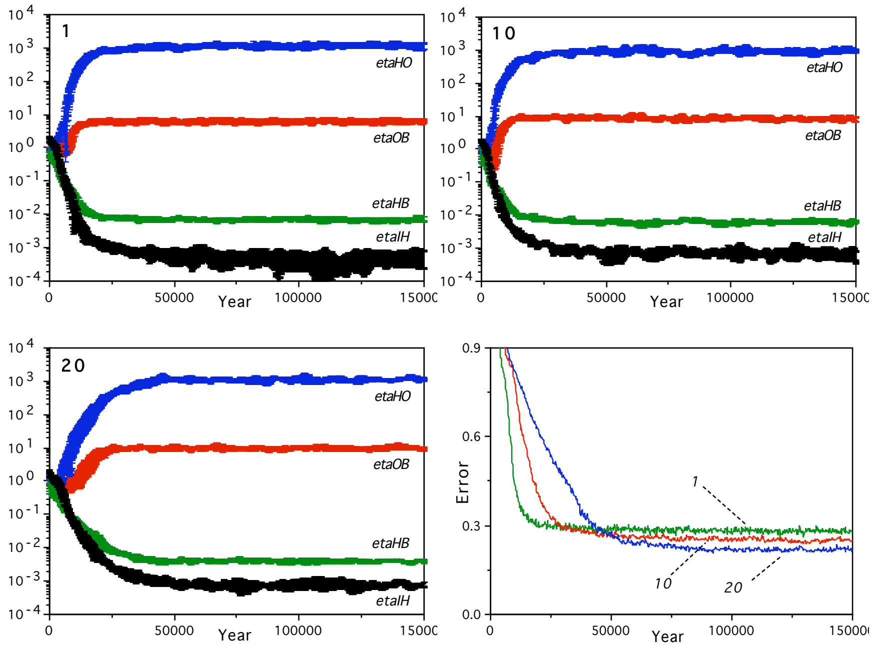

Figure 1: Evolution of the learning rates for protection periods of 1, 10 and 20 years, and comparison of the evolution of the corresponding mean error rates.

at a number of different values by hand, evolving the other innate parameters as before, and comparing the levels of performance that emerge. The second can be answered by running similar simulations with the protection period allowed to be an additional evolvable parameter, and analyzing what happens. The various simulation results are presented in the next three sections.

3. Simulation Results for Different Protection Periods

The natural starting point is to run the evolutionary neural network simulations with a few carefully selected fixed protection periods to determine if that makes any difference to the evolved populations. Since the individuals typically learn their given task in 10 to 20 simulated years, and start dying of old age at age 30, it makes sense to begin by looking at protection periods of 1, 10 and 20 years. Figure 1 shows the evolution of the learning rates for these three cases, with means and variances over six runs. The evolved parameters and low variances across runs are similar to those found in previous studies (Bullinaria, 2007). The evolving parameters in each case have settled down by about 50,000 years, and subtle differences can be seen between the final values.

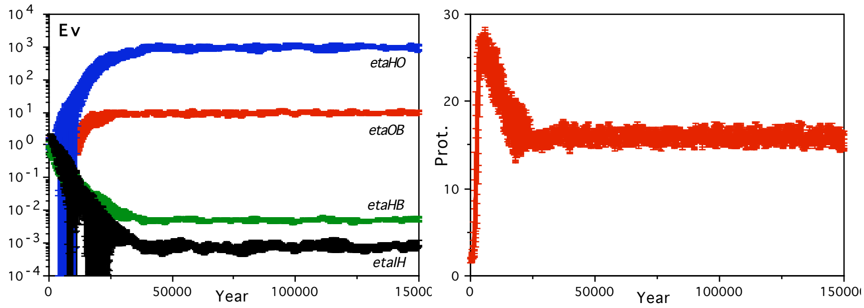

Figure 2: Evolution of the learning rates when the protection period is allowed to evolve, and the evolution of the protection period.

The final panel in Figure 1 compares the generalization performance means across populations during evolution for each protection period. It shows that the evolutionary process is much slower to settle down for the longer protection periods, but longer protection periods *do* appear to have an advantage in terms of final evolved performance. However, these population means hide complex age dependent error distributions, and the population age distributions are unlikely to be the same across the various cases, so to see whether there is a real evolutionary advantage of increased protection periods, we need to simulate their evolution. This is done in the next section.

4. Allowing the Protection Period to Evolve

If the protection period is allowed to evolve, the evolution of that period and the associated learning rates that emerge are as shown in Figure 2, again with means and variances across six runs. There is higher variance in all the parameters, compared to the fixed protection period runs, until the protection period has settled down after about 40,000 years. Early on in evolution, when the populations are relatively poorly performing, the protection period rises rapidly to about 25 years, but then falls and settles to around 16 years.

Comparison of the averages and variances of the crucial evolved population properties are presented in Figure 3, for both the fixed and evolved protection periods. As one would expect, the number of deaths per year due to competition decreases, from the maximum of 20 per year, as the protection period increases, and this inevitably increases the average age of the population. In turn, more individuals survive to old age, and consequently the deaths per year due to old age increase slightly. Overall, there is still a net reduction in deaths per year, and so, given the fixed population size, the average number of children per individual at any given time decreases with the protection period.

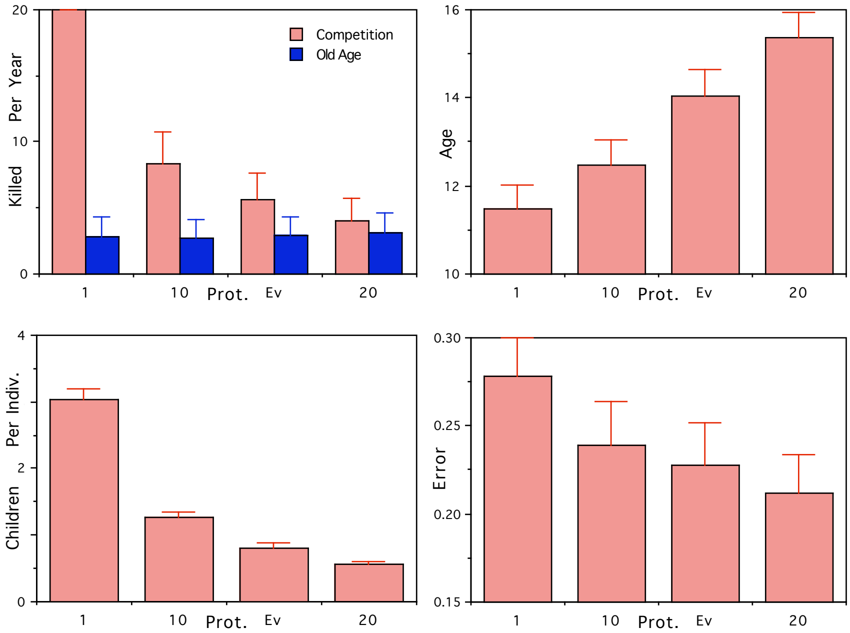

Figure 3: Comparison of the evolved population averages and variances for the various fixed (*1*, *10*, *20*) and evolved (*Ev*) protection periods: deaths per year, ages, children per individual, and performance error rates.

(Note that this is independent of any direct introduced cost of parents protecting more children.) Finally, the average population performance error rate (i.e. inverse fitness) is seen to fall steadily with increasing protection periods. All these trends vary monotonically with protection period, and the evolved protection period population results are consistent with what would be expected from their evolved period of 16 years.

The obvious next question is, given that the average population fitness increases with protection period, why is it that the evolved protection age does not end up higher than 16? Actually, the distributions in Figure 3 already provide us with some clues. First, given that the older individuals will have had more time to learn, they will inevitably be fitter by our criteria, and hence the increases in average age will automatically translate into increased population fitness, even if each individual were no better as a result of the protection period. Moreover, even if there were individual fitness advantages, the reduced number of children per individual for increased protection periods will place individuals with long protection periods at an evolutionary disadvantage, and this will tend to decrease the evolved protection periods. To

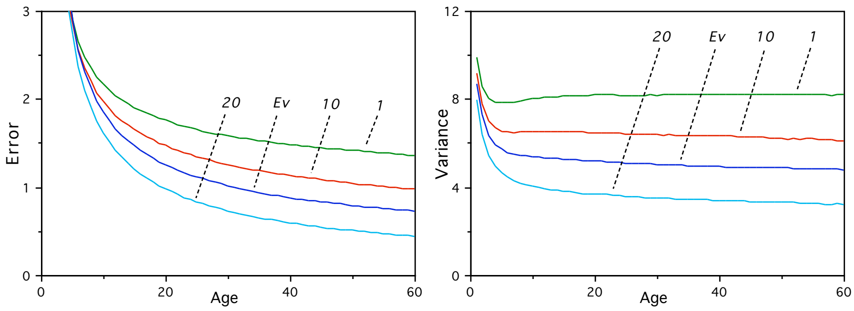

Figure 4: Mean errors and variances during learning for evolved individuals, with evolved protection period (*Ev*) and the three fixed protection periods (*1*, *10*, *20*).

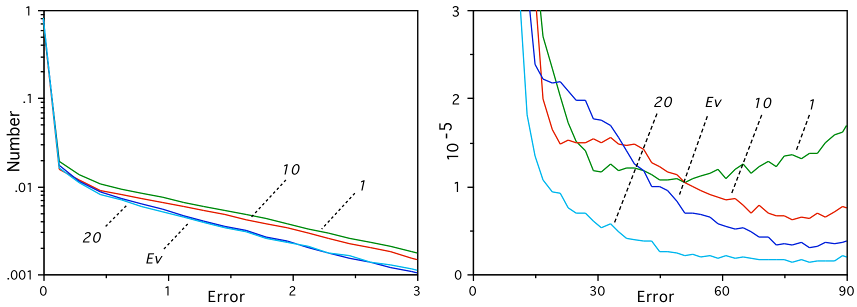

Figure 5: The peaks and tails of the error distributions for evolved individuals aged between 50 and 60, for each of the four protection period cases (*Ev*, *1*, *10*, *20*).

understand the advantages and disadvantages to individuals, and explore the detailed effects of such trade-offs, we need to look more carefully at the individual fitness profiles. This will be done in the next section.

5. Analysis of the Evolved Performance

The means and variances of the individual error rates (i.e. inverse fitness) during learning are shown in Figure 4, and there do indeed appear to be significant advantages for protracted protection periods. However, the error distributions for this type of problem are known to be rather skewed, with the residual mean errors due largely to instances from the long tails of very large errors (Bullinaria, 2007). This is clear in the peaks and tails of the error distribution for individuals aged between 50 and 60 years shown in Figure 5. There is a massive peak around zero errors, as one would expect at that age, but there also remain significant numbers of very large errors. This is a common feature of evolutionary processes that encourage fast learning (Bullinaria, 2007), and

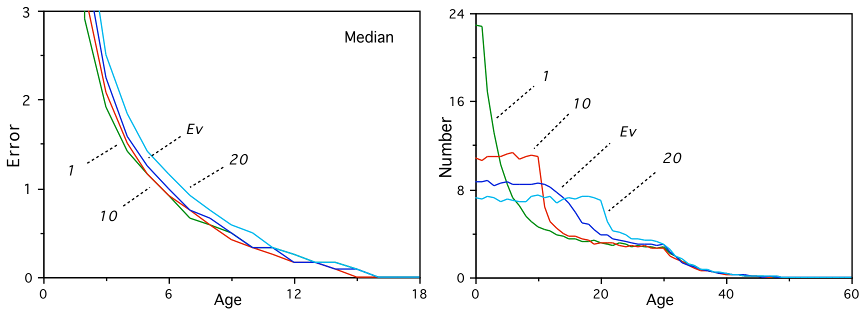

Figure 6: The median error rates during learning and the age distributions of the evolved populations, for each of the four protection period cases (*Ev*, *1*, *10*, *20*).

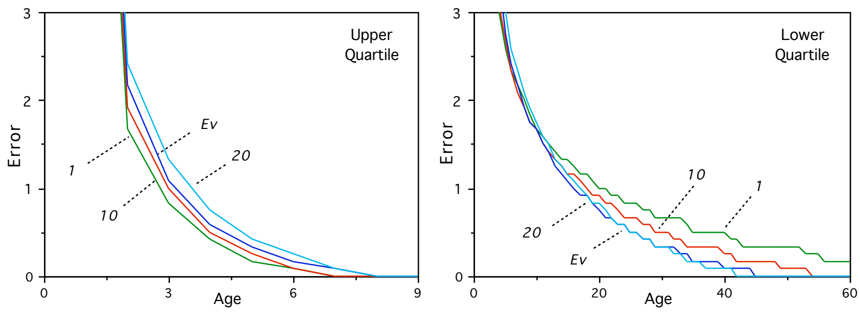

Figure 7: The upper and lower quartile error rates during learning for evolved individuals, for each of the four protection period cases (*Ev*, *1*, *10*, *20*).

longer protection periods, which limit the need for fast learning at early ages, seem to alleviate the problem.

One can get a better idea of the population performances, that is not skewed by a few instances of very poor performance, by looking at the medians rather than the means. The median error rates during learning, shown on the left of Figure 6, are in accordance with the expectation of learning the task essentially perfectly by a certain age, but there is surprisingly little difference in median performance across the four protection periods. There is at most two years difference in learning across the cases, despite the twenty years range of protection periods, and the wide variations in the age distributions of the evolved populations shown on the right of Figure 6. Each age distribution is fairly flat during the protection period, then falls off due to competition till the individuals start dying of old age from the age of 30, at which point there is an exponential fall to zero. The upper and lower quartile error rates are shown in Figure 7. The faster learning quartiles are still remarkably similar to each other,

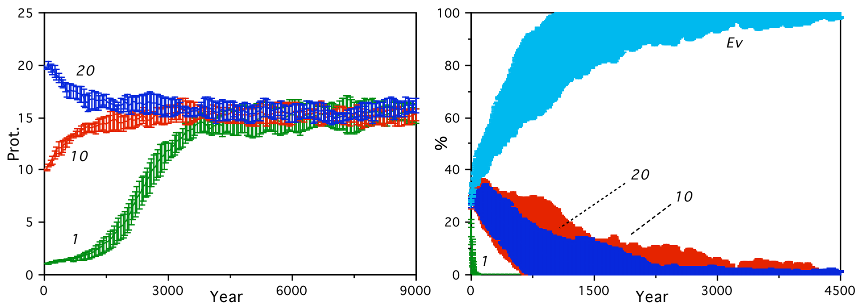

Figure 8: Evolutionary improvement of protection periods (left) and competition between the evolved populations from the four protection period cases (right).

with just a slight increase in learning speed resulting from shorter protection periods. However, there are clear differences in the slower learning quartiles, with large improvements seen for the longer protection periods, as already evident in Figures 4 and 5.

6. Discussion and Conclusions

The above simulations have established that longer protection periods *do* offer a clear learning advantage, and relatively little disadvantage. So we now need to return to the question of what it is that prevents the evolved protection period from becoming even longer. But first, we need to make sure that the evolved period we found is not simply an artifact of the chosen evolutionary process. This can be checked by freeing the protection period in each of the three fixed period evolved populations, and allowing them to evolve further. The results of this are shown on the left of Figure 8, with means and variances across six runs. In each case there is either a rise or fall to the same evolved period of around 16 years that emerged before.

A second check involves combining the evolved populations from the four cases into one big population, and then allowing natural selection to take its course. Since each case has already been optimized by evolution, no further crossover and mutation was allowed. The outcome is shown on the right of Figure 8, with means and variances across twelve runs. There is quite a large variance across runs, but individuals with the evolved protection period consistently come to dominate the whole population. Individuals with virtually no protection are wiped out almost immediately.

The natural conclusion from the results in Figure 8 is that, although there are clear learning advantages to having longer protection periods, the best periods from an evolutionary point of view are shorter than they could be. This can be

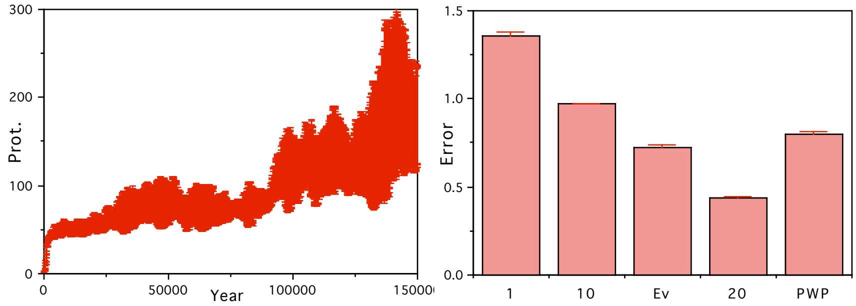

Figure 9: Evolution of the protection period when procreation while protected (PWP) is allowed (left), and comparison of the mean error rates at age 60 (right).

understood in terms of the number of children per individual seen in Figure 3. Because the individuals effectively have fixed life-spans, extended periods of protection will use up a significant proportion of the potential procreation period and thus put those populations at a serious evolutionary disadvantage. The evolutionary simulations are able to establish a suitable trade-off value for the protection period, that balances the improved performance against the loss of reproductive opportunities.

The obvious check of this conclusion is to repeat the whole evolutionary process with procreation allowed while being protected. As seen on the left of Figure 9, the protection period then evolves to be way beyond the normal life-span of the individuals, so that there are no deaths at all due to competition, only due to old age. What this scenario has re-introduced, however, is the need to compete at all ages to procreate, and this encourages faster learning again. Of course, that brings back with it the unwanted associated side effects, such as the use of risky learning strategies that sometimes result in persistent very poor performance at all ages. This can be seen in the increased mean error rates at age 60 shown on the right of Figure 9. One could imagine allowing individuals to procreate randomly without having to compete to do so, and that would remove the pressure to learn quickly, but that would leave no evolutionary pressure to improve fitness at all, and the individual performances would end up even worse. It seems then, that there is a real advantage to preventing offspring from reproducing while being protected, and that goes beyond enhancing the parents' own reproductive success rate.

In conclusion, the results presented in this paper have shown how evolutionary neural network simulations can begin to address aspects of human behaviour such as the protection of offspring. Of course, there are many related issues that remain to be taken into account in future work. First, more attention could be paid to the changes in the learning experience that might result from the parental protection, for example due to guided exploration, exploration without risk, “teaching”, and so on. The costs to parents of protecting their children should also be accounted for more carefully, particularly for situations where each parent has to protect many children for long periods. We also need to consider the evolutionary pressures and consequences that would arise due to the introduction of competition with, and co-evolution with, other species. The evolved protection periods are also likely to interact with changes to the natural life-span of the species, and that needs to be explored. There are also many other “life history” factors which may evolve (e.g., Stearns, 1992; Roff, 2002), with associated trade-offs, and these could also usefully be incorporated into improved future models.

References

Bullinaria, J.A. (2001). Simulating the Evolution of Modular Neural Systems. In: *Proceedings of the Twenty-Third Annual Conference of the Cognitive Science Society*, 146-151. Mahwah, NJ: Lawrence Erlbaum Associates.

Bullinaria, J.A. (2003a). From Biological Models to the Evolution of Robot Control Systems. *Philosophical Transactions of the Royal Society of London* A, **361**, 2145-2164.

Bullinaria, J.A. (2003b). Evolving Efficient Learning Algorithms for Binary Mappings. *Neural Networks*, **16**, 793-800.

Bullinaria, J.A. (2007). Using Evolution to Improve Neural Network Learning: Pitfalls and Solutions. *Neural Computing & Applications*, **16**, 209-226.

Lamb, M.E., Bornstein, M.H. & Teti, D.M. (2002). *Development in Infancy: An Introduction*. Mahwah, NJ: Lawrence Erlbaum Associates.

Roff, D.A. (2002). *Life History Evolution*. Sunderland, MA: Sinauer Associates.

Sloman, A. & Chappell, J. (2005). The Altricial-Precocial Spectrum for Robots. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 1187-1193. IJCAI.

Stearns, S.C. (1992). *The Evolution of Life Histories*. Oxford, UK: Oxford University Press.