

The Effect of Learning on Life History Evolution

John A. Bullinaria
School of Computer Science
University of Birmingham
Birmingham, B15 2TT, UK
j.a.bullinaria@cs.bham.ac.uk

ABSTRACT

A series of evolutionary neural network simulations are presented which explore the hypothesis that learning factors can result in the evolution of long periods of parental protection and late onset of maturity. By evolving populations of neural networks to learn quickly to perform well on simple classification tasks, it is shown that better learned performance is obtained if protection from competition is provided during the network's early learning period. 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 individuals not being allowed to reproduce during their protection phase, and the parents suffering increased risk of dying while protecting their offspring.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence – *intelligent agents, multiagent systems*.

General Terms

Algorithms, Measurement, Performance, Experimentation.

Keywords

Artificial Life, Evolution, Learning, Life Histories.

1. INTRODUCTION

The processes of birth, maturation, reproduction and death vary considerably between species, and understanding how the various life histories came about through evolution remains a major research area (e.g., [9, 16]). The idea of trade-offs in life history evolution is now well established (e.g., [15]), with classic examples including the trade-off between reproduction and growth [19], and between nursing and survival [7]. There also appear to exist correlations between other aspects of life histories, for example, fixed ratios within lineages of lifespan to age of maturity [6, 10]. The diversity of species that have existed reflect the different ways the trade-offs can be balanced under different environmental conditions.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GECCO '07, July 7–11, 2007, London, England, United Kingdom.

Copyright 2007 ACM 978-1-59593-697-4/07/0007...\$5.00.

In this paper, one particular trade-off is considered, namely the advantage of parents providing long periods of protection for their offspring, against increased risk of death for the parents, and delayed reproductive activity in the young. In nature, such protection varies enormously from *precocial* species in which the young are born well developed and requiring virtually no protection, to *altricial* species in which the young are born helpless and require long periods of parental care before they are able to survive on their own. It might seem obvious that early maturity is advantageous because it increases the chance that individuals will survive to maturity, and decreases the time wasted not reproducing [8]. However, if parents protect their young, the first advantage will be less relevant, and the second may be balanced by improved performance later in life with increased reproduction overall. In fact, there are many factors involving growth, size, fecundity, environment and so on, known to affect the age at maturity (e.g., [16]). One major advantage of simplified Artificial Life simulations is that one can explore individual factors without all the confounds inherent in empirical measurements of existing biological populations. In this paper, one factor in life history evolution is studied, that appears to have not received much attention before, namely the effect of individuals needing to learn within their own lifetimes, and how that can affect the optimal period of protection parents should offer their young.

This study was prompted by the fact that human infants are particularly altricial, even compared with other primates, requiring extended periods of parental protection and support (e.g., [13]). There are many factors that could affect this (e.g., [12, 14]), but two crucial processes are known to take place in altricial species during the protection stage – the infants are growing, and they are learning. Human infants do need to grow considerably after birth, and parental protection provides obvious survival advantages. Learning is also important for humans, and in other species for which relatively complex behaviour is required, since the environment is unpredictable, and adaptation is needed to cope with rapid growth. Otherwise, innate behaviour would be adequate. The question this paper addresses is: to what extent can learning issues influence the evolution of long protection periods? This is done by simulating the evolution of populations of neural networks that must learn quickly how to perform well on simplified classification tasks. With individuals of all ages competing for survival based on their learned performance, the effect of different protection periods can be explored. Allowing the protection period to evolve shows how the advantages of protection trade-off against the associated

disadvantages so that a particular protection period emerges through evolution.

The next section describes the neural network systems that are evolved. Section 3 then studies the populations that emerge for a selection of fixed protection periods, before the protection period itself is allowed to evolve as discussed in Section 4. The properties of the evolved populations and the associated trade-offs are analyzed in Section 5, and the paper ends with some discussion and conclusions in Section 6.

2. EVOLVING NEURAL NETWORKS

The aim here is to simulate the crucial features of the evolution of most animal populations, with particular emphasis on the aspects of fitness associated with learning. Since neural networks, or brains, are largely responsible for learning in animals, our animal populations will be represented by simple artificial neural networks. A population of individual neural networks (each specified by a set of innate parameters) is therefore maintained, that must learn from a continuous stream of input patterns how to classify future input patterns. Those inputs could, for example, correspond to relevant features of other animals, and the desired output classes could correspond to properties such as being edible or dangerous. The fitness of each individual is defined as how well it classifies the new inputs *before* discovering its correct class and training on it. Requiring the individuals to compete to survive and procreate, according to their relative fitness, should lead to the emergence of populations of increasing fitness. To compete effectively in a population consisting of individuals of all ages, each individual must not only learn how to perform well, but must also be able to learn *quickly* how to achieve that good performance, or at least quickly enough that it can survive after its parents have withdrawn their protection.

For the simulations, a specific concrete system is clearly needed, and it makes sense to follow one that has already proved instructive and been described in detail elsewhere [2, 4, 5]. Consequently, traditional fully connected Multi-Layer Perceptron neural networks were used, with one hidden layer, sigmoidal processing units, and trained by gradient descent using the cross-entropy error function on simple classification tasks. Since most real-world classification tasks involve learning non-linear classification boundaries in a space of real valued inputs, the set of classification tasks chosen were two dimensional continuous input spaces with particular circular classification boundaries. These proved simple enough to allow extensive simulations, yet recreated the crucial features and difficulties of real world problems. Each individual network was assigned a randomly chosen classification boundary of this type to learn from a stream of randomly drawn data points from the input space that was normalized to a unit square. Then its fitness was its generalization ability, i.e. the average number of correct classifications (e.g., outputs within 0.2 of the binary targets) *before* training on them. The extent to which real parental protection affects the training data available for the infants will be left open by keeping the nature of the training data constant across all cases. As will be seen later, the key learning effect emerges without the need to take this complication into account.

An important feature of human learning of relevance here is the fact that for many skills there are *critical periods* after which further learning becomes more difficult [1, 11]. It has already been demonstrated elsewhere that evolving neural network learning rates that vary during the learning process does lead to improved learning performance, and that the evolved time dependencies are qualitatively similar to human like age dependencies [3, 5]. Such time t dependent learning rates $\eta_L(t)$ can be conveniently approximated by introducing a simple two parameter exponential scale factor $s(t)$:

$$\eta_L(t) = s(t) \eta_L(0) \quad , \quad s(t) = \beta + (1-\beta)e^{-t/\tau}$$

in which the baseline β and the time constant τ are evolved to take on the positive values that result in the best performance. If time dependence proves not to be beneficial, the parameters will evolve so that $s(t) = 1$. The subscript L is introduced because earlier studies [2, 4] have shown empirically that better performance is obtained by evolving separate learning rates η_L and random 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 identical parameters across the whole network. The two variable plasticity parameters β and τ , four learning rates η_L , four initial weight parameters r_L , together with a standard momentum parameter α and weight decay regularization parameter λ , give twelve real valued evolvable innate parameters in all for each network. Since evolving the number of hidden units invariably resulted in the maximum allowed number being used, this was fixed at 20 for all networks, which was more than enough for learning the given tasks.

For the artificial life simulations proposed here, there is a clear need to tie in the neural network learning time-scales with the lifetime and evolutionary time-scales, so a “simulated year of experience” was defined to be 1200 training data samples, and the fitness of each individual was computed at the end of each simulated year as the average over that year. A fixed computationally feasible population size of 200 was maintained throughout, with individuals that die being replaced by children of the most fit individuals. Deaths occurred due to old age, or by losing a fitness comparison “fight” against other individuals. The details of such factors needed to be fixed to encourage evolutionary change and to preserve a reasonably diverse population, but the precise specification proved not to affect the qualitative results a great deal, so convenient “round numbers” were chosen for the parameter values. Old age deaths were taken to occur each year for a random 20% of individuals aged over 30 years, which allowed enough training samples to learn the given tasks, yet prevented the populations becoming dominated by a few very old and very fit individuals. In addition, each year all the unprotected individuals were forced to compete with a random other eligible individual, and would die if their fitness was lower than their competitor. Deaths in this way were limited to 10% of the population each year, so that a reasonable population age distribution was preserved. The children were created by cross-over and mutation from two parents chosen by pair-wise fitness comparisons of the eligible individuals.

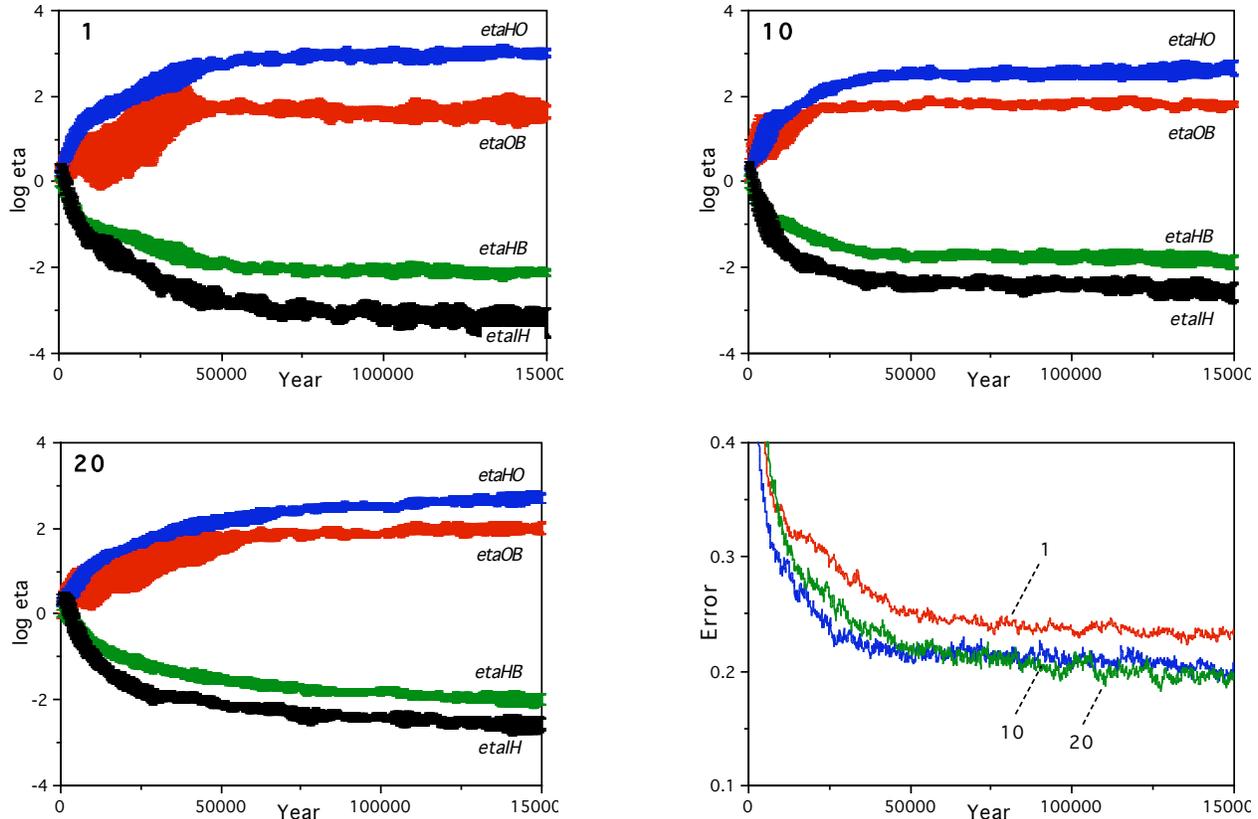


Figure 1: Learning rate evolution for fixed protection periods of 1, 10 and 20 years, and comparison of the corresponding performance error rates.

This was implemented by having each child inherit innate parameters chosen randomly from the corresponding ranges spanned by its two parents, plus a random mutation (from a Gaussian distribution) that gave it a reasonable chance of falling outside that range. All these details are clearly gross over-simplifications of real biological systems, but they do constitute a manageable starting point that includes approximations of all the key processes.

Such an evolutionary process has been shown to work before [5]. The crucial extra feature here is that the children are protected by their parents until they reach a certain age and cannot be killed by competitors before then. (For simplicity, the parents here are assumed to be perfect protectors, though in reality their own fitness will normally affect their protection abilities, as will the number of children they are protecting at any given time.) There is an implicit cost to the parents in that the more children are protected, the higher the chance they stand of being in the 10% of the population that are die through competition each year. (This is another over-simplification, adequate for present purposes, but a better account of the reproductive and protection costs will be required for more realistic models.) A cost the children must bear is that they are prevented from having any children of their own before they leave the protection of their parents. (This may be a reasonable approximation for some species, and a suitable starting point for this general study, but will need more careful consideration in more focused future work.) The

remainder of this paper will explore how the protection period affects the performance of the evolved individuals, and what protection period emerges if it is left free to evolve in the same way as all the other innate parameters.

3. SIMULATION RESULTS FOR FIXED PROTECTION PERIODS

The obvious starting point was to carry out the evolutionary neural network simulations for each of various fixed protection periods and determine if there were any differences between the evolved populations. The learning time-scale was set so that individuals typically learn the given task in 10 to 20 simulated years, and they start dying of old age after 30 years, so it made good sense to begin by looking at protection periods of 1, 10 and 20 years. The evolution of the learning rates for these three cases are shown in Figure 1, with means and variances over six runs (which were sufficient to establish statistically significant results). The evolved parameters and low variances across runs are similar to those found in earlier studies [5]. Subtle differences, however, can be seen between the final parameter values, and the evolutionary process is noticeably slower to settle down for the longer protection periods. The fourth graph in Figure 1 compares the generalization error performance means across populations during evolution for each protection period. It seems that the two longer protection periods *do* have an advantage in this

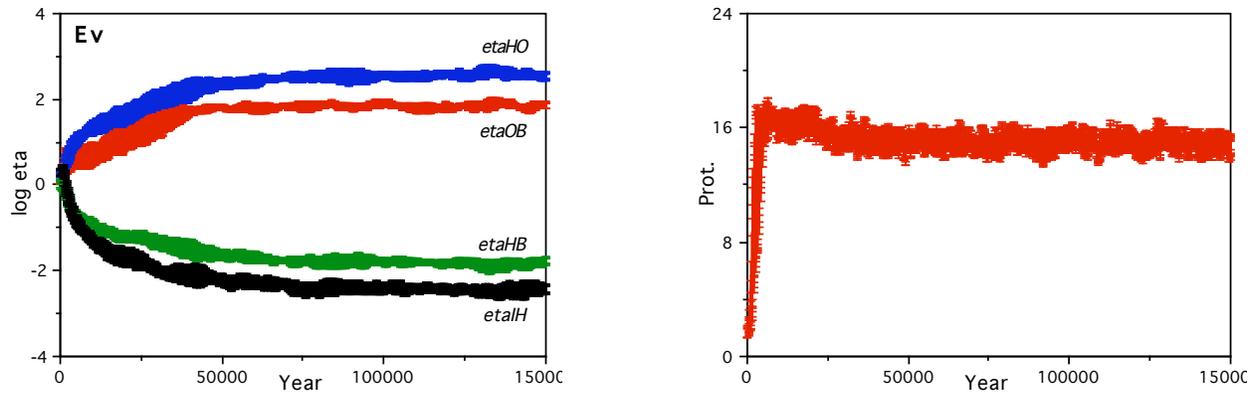


Figure 2: Learning rate evolution when the protection period is allowed to evolve (left), and the evolution of the protection period (right).

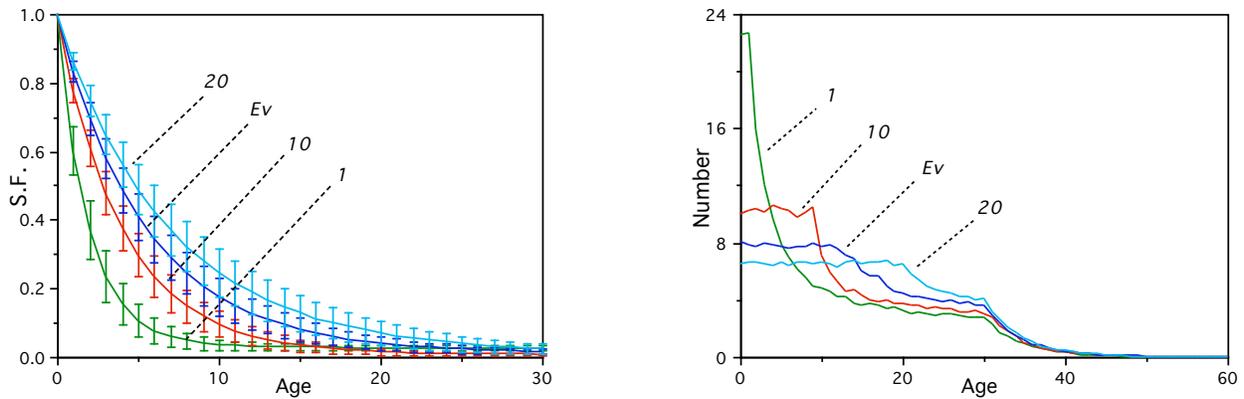


Figure 3: Differences in the age dependent plasticity scale factors (left), and population age distributions (right), for the various protection periods.

respect. However, such population means hide complex age dependent error distributions, and the population age distributions are likely to vary considerably between cases, so this is not sufficient to show that extended protection periods will have an evolutionary advantage. What needs to be done is explicitly simulate the evolution of the protection period along with everything else.

4. SIMULATION RESULTS FOR EVOLVED PROTECTION PERIODS

When the protection period is allowed to evolve, the evolution of that period and the associated learning rates are as shown in Figure 2, again with means and variances over six runs. Early on during the evolution, while the populations are still performing relatively poorly, the protection period rises rapidly to about 17 years, but then falls slightly, settling down at around 15 years.

The length of the protection period, whether it is fixed (at 1, 10, 20 years) or evolved (leading to $Ev \approx 15$ years), has a clear effect on the various age dependencies. On the left of Figure 3 is shown how increased protection periods result in lengthening of the period during which the learning rates remain relatively high. There is clearly no need to rush the learning while being protected. There are also wide variations

in the age distributions of the evolved populations, as shown on the right of Figure 3. Each age distribution is fairly flat during the protection period, and then falls off due to competition until the individuals start dying of old age from the age of 30, at which point there is an exponential fall towards zero. These different patterns of learning and age distributions will obviously affect the other properties of the evolved populations. The averages and variances of the crucial evolved properties are compared in Figure 4, for the various protection periods. As one might 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. As a consequence, more individuals survive to old age, and so the deaths per year due to old age increase slightly. Overall, there is still a net reduction in the number of 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. Finally, it can be seen that the average population performance error rate (i.e. inverse fitness) falls steadily with increasing protection period.

It is reassuring to see that all the observed trends vary monotonically with protection period length, and the evolved protection period population results are consistent with what would be expected from their evolved period of around 15

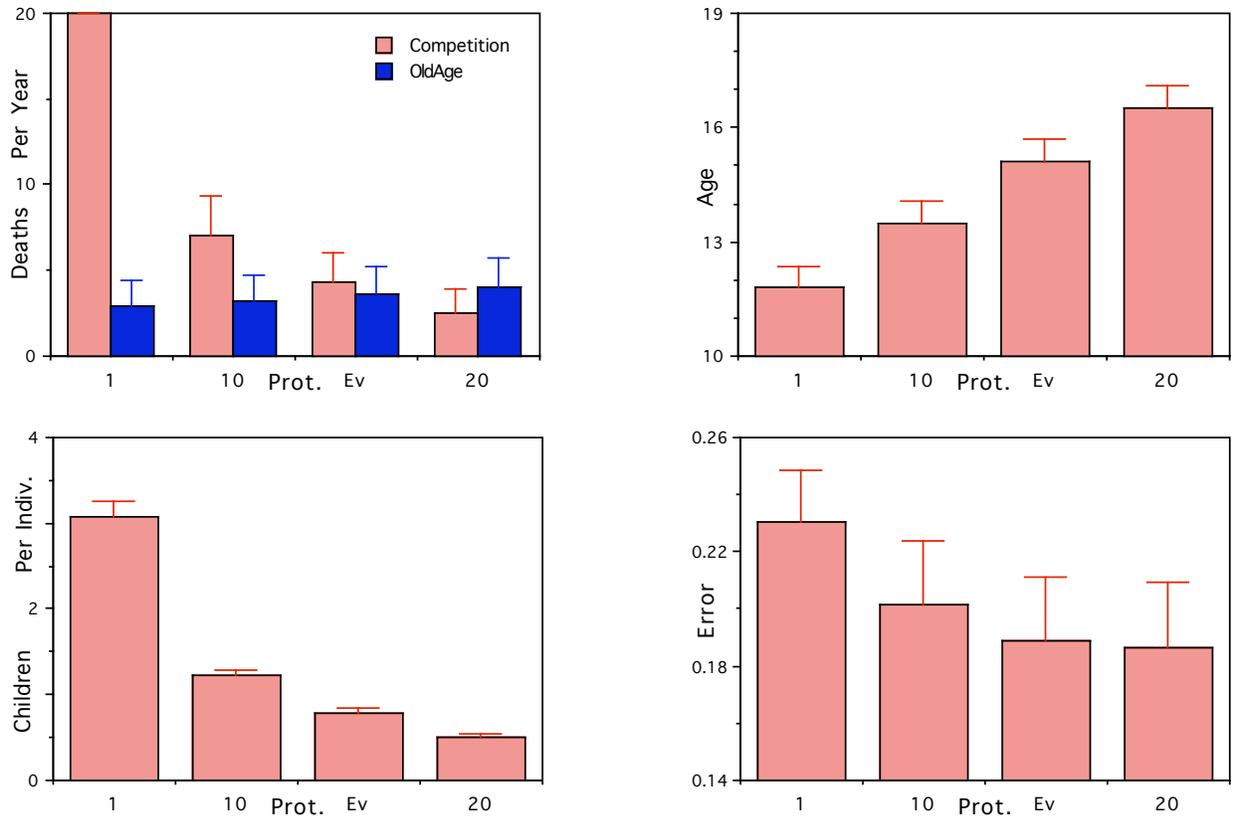


Figure 4: Comparisons of the effects of the various different fixed and evolved protection periods: the competition and old age death rates, the average ages, the average number of children per individual at any given time, and the average performance error rates.

years. The important issue to consider now is why, given that the average population fitness increases with protection period, does the evolved protection withdrawal age not end up higher than 15? Figures 3 and 4 already provide some clues. First, older individuals will obviously have had more time to learn, so they should be fitter by our criteria, and hence the increases in average age of the population will automatically translate into an increased population fitness, even if each individual were no better as a result of the protection period. Moreover, even though there probably are individual fitness advantages too, the reduced number of children per individual that emerge for increased protection periods will place individuals with longer protection periods at an evolutionary disadvantage, and this is likely to result in a decrease in the evolved protection period. Once again there are trade-offs between the relevant factors, and to understand them fully, one needs to look more carefully at the individual fitness profiles. A more detailed analysis of the evolved populations is presented in the next section.

5. ANALYSIS OF THE EVOLVED POPULATIONS

The means and variances of the individual performance error rates (i.e. inverse fitness) during learning are plotted in Figure 5, and there do indeed appear to be significant reductions in the mean errors after protracted protection periods, but with

consequent delays in reaching those lower error rates. However, the error distributions for this type of problem are known to be rather skewed, with the residual mean errors being due to long tails of extremely large errors. Such consequences of risky learning strategies are an unfortunate evolutionary side effect of the pressure to learn quickly [5], that extended protection periods may be expected to help with. The median learning performances shown on the left of Figure 6 are not affected by the long tails of the distributions, and indicate essentially perfect performance by age 12 for all protection periods, though with the expected slower learning for longer protection periods. The error distributions between ages 50 and 60, shown on the right of Figure 6, give an idea of the extent of the small numbers of persistent large errors. As one would expect at that age, there is a massive peak around zero errors, but there remain a significant number of very large errors. As noted above, this is a common problem of evolutionary processes that encourage fast learning [5], and long protection periods do seem to alleviate it. The upper and lower quartile error rates are shown in Figure 7. These confirm the expectation that longer protection periods will slow the learning, and also confirm the improvement in learned performance at the poorer performing end of the spectrum.

Interestingly, allowing evolved age dependent plasticity has a significant effect on what emerges. Without it, rather different median learning performances and error distributions are obtained, as shown in Figure 8. There is now relatively

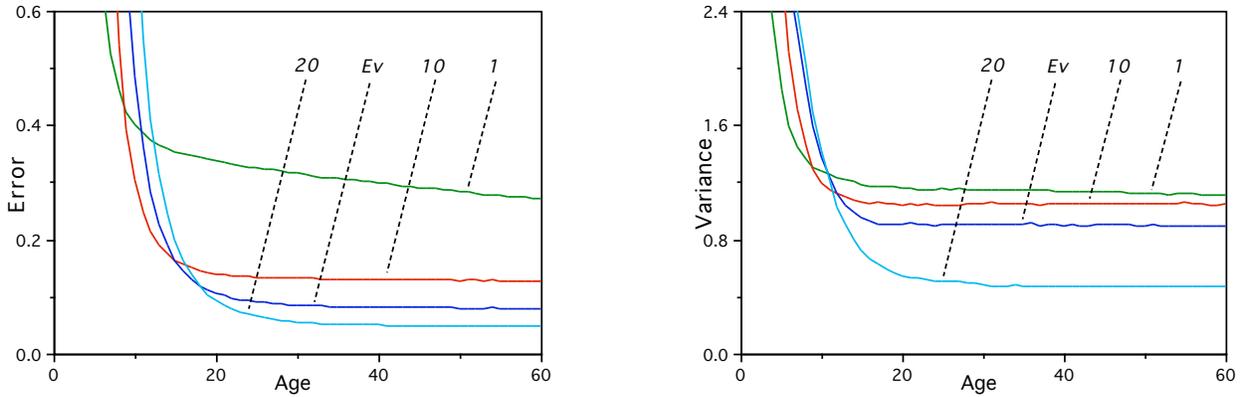


Figure 5: Mean performance error rates (left) and variances (right) during learning for the evolved individuals with the various protection periods.

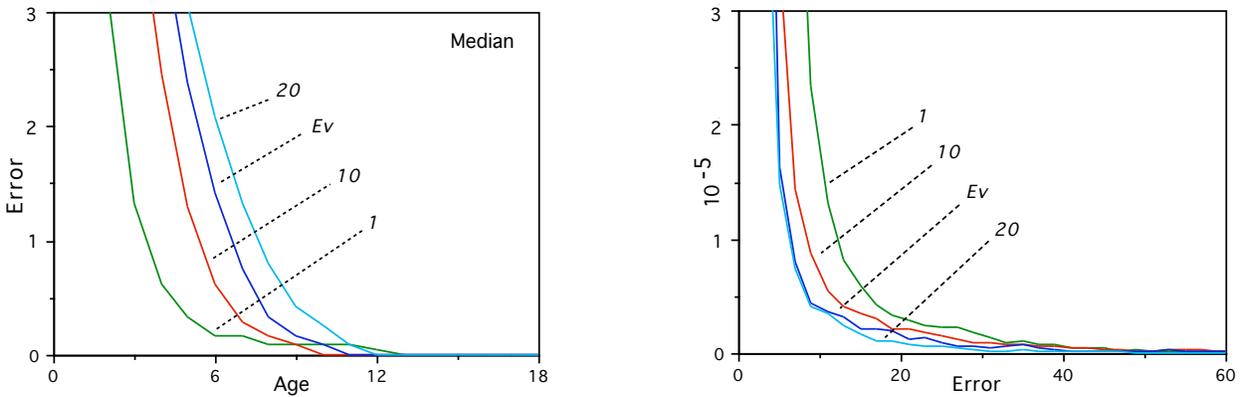


Figure 6: The median learning performances (left) and tails of the error distributions for evolved individuals between the ages of 50 and 60 (right).

little difference between the median performances, but much greater differences in the instances of very large errors. This confirms the importance of including in the simulations as many known features of real learning systems as possible.

6. DISCUSSION AND CONCLUSIONS

The above simulations and analyses have established that longer protection periods do offer clear learning advantages, and relatively few disadvantages. Nevertheless, there is still a need to consider the underlying trade-offs more carefully, in order to understand what prevents even longer evolved protection periods from emerging. First, however, a check is required to verify that the evolved protection period found is not simply some artifact of the chosen evolutionary process, perhaps corresponding to a local optimum of fitness and/or lack of population diversity. This can be conveniently tested by allowing the protection period in the three fixed period evolved populations to evolve away from their previously fixed values. The results of this are shown on the left of Figure 9, with means and variances across six runs. In each case there is a relatively fast rise or fall to the same evolved period of around 15 years that emerged before. A second test involves combining evolved populations from all four cases (the three different fixed period runs, and one evolved period run) into

one big population, and allowing natural selection to take its course. Since each population had already been optimized by evolution, the children inherited characteristics from just their fittest parent, with no further crossover or mutation allowed. The outcome of this process is shown on the right of Figure 9, with means and variances across twelve runs. The individuals with the evolved protection period consistently come to dominate the whole population. Individuals that get virtually no parental protection are wiped out almost immediately.

Overall, the simulations show that, although there are clear learning advantages to having longer protection periods, extending those protection periods into effectively fixed life-spans restricts the available procreation periods and this in turn places those populations at a serious evolutionary disadvantage. The evolving populations have clearly found a suitable trade-off value for their protection period, appropriate for individuals that start dying of old age after 30 years. It is reasonably straightforward to repeat the above evolutionary simulations with the onset of “old age” at different ages, corresponding to different natural “life-spans”, to determine how that affects the emergent protection period and associated performance level. The results of this are plotted in Figure 10. Beyond what might be termed the “natural learning time-scale” of about 10 years, there is a fairly linear relation between the protection period and life-span. It is clear that learning is of

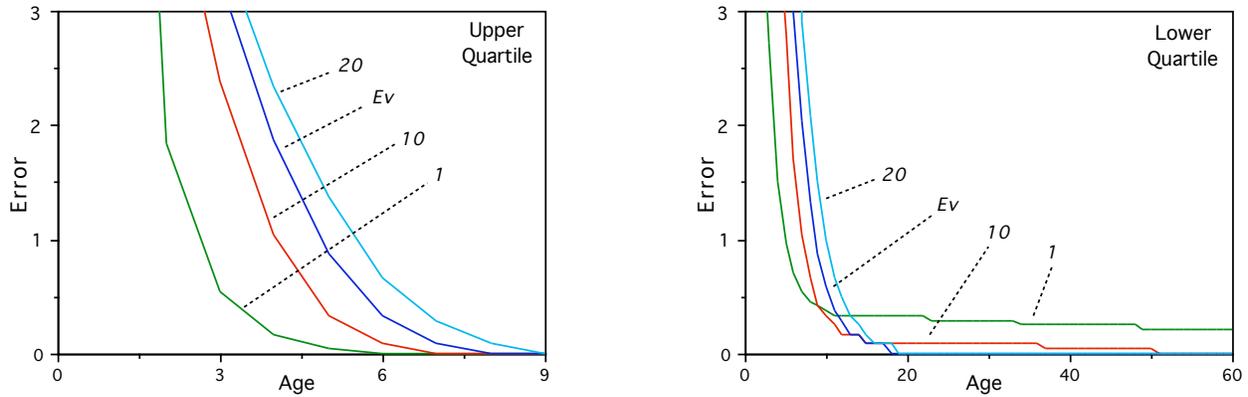


Figure 7: Upper and lower quartile learning performances for the various evolved individuals, shown as error rates as a function of age.

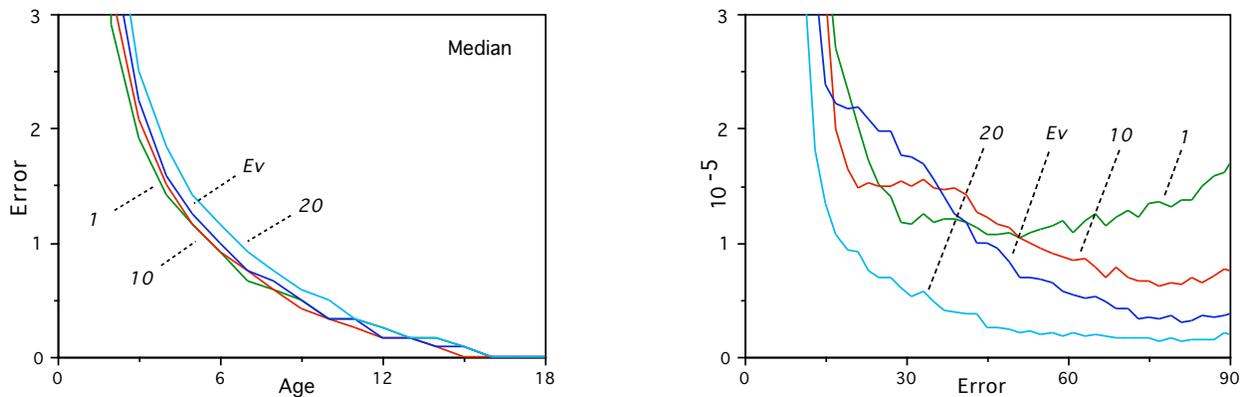


Figure 8: The median learning performances (left) and tails of the error distributions (right) when variable plasticity is not allowed (cf. Figure 6).

sufficient importance that the amount of time devoted to it continues to increase as the life-spans become longer, rather than using all that extra time for reproduction.

In this study, the age of maturity has been tied to the protection period, and the linear relation with life-span appears consistent with biological populations, though in reality there are certainly other factors besides learning involved in determining the age at maturity (e.g., [6, 10, 17, 18]). Figure 10 also shows how the average population performance levels improve with life-span. A further set of simulations have shown that if the life-span is allowed to evolve freely within the current set-up, it keeps on increasing indefinitely. In biological populations, of course, there are many other factors that restrict life-spans (e.g., [12, 16]), but simulating such trade-offs is beyond the scope of this study. As is simulating factors that might result in reproduction stopping a particular ages prior to death, or increased reproductive problems at older ages.

In conclusion, the Artificial Life studies presented in this paper have demonstrated how evolutionary neural network simulations can begin to address learning aspects in life history evolution. As noted throughout, there remain many related issues and refinements that could usefully be incorporated into extensions of this study, such as the

changes to the learning experience that can result from the parental protection (for example due to guided exploration, exploration without risk, “teaching”, and so on), as well as refined patterns of competition and deaths, more realistic costs of reproduction and protection, protection abilities that depend on parental fitness and number of children, and the numerous biological and environmental details that affect the life-span, age of maturity, fecundity and length of the reproductive and growing periods. Further work exploring these factors and the associated trade-offs will hopefully be presented elsewhere in the near future.

7. REFERENCES

- [1] Bailey, J.D., Bruer, B.T., Symons, F. & Lichtman, J.W. (Eds) (2000). *Critical Thinking About Critical Periods*. Baltimore: Brookes.
- [2] 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: LEA.
- [3] Bullinaria, J.A. (2003). From Biological Models to the Evolution of Robot Control Systems. *Philosophical Transactions of the Royal Society of London A*, **361**, 2145-2164.

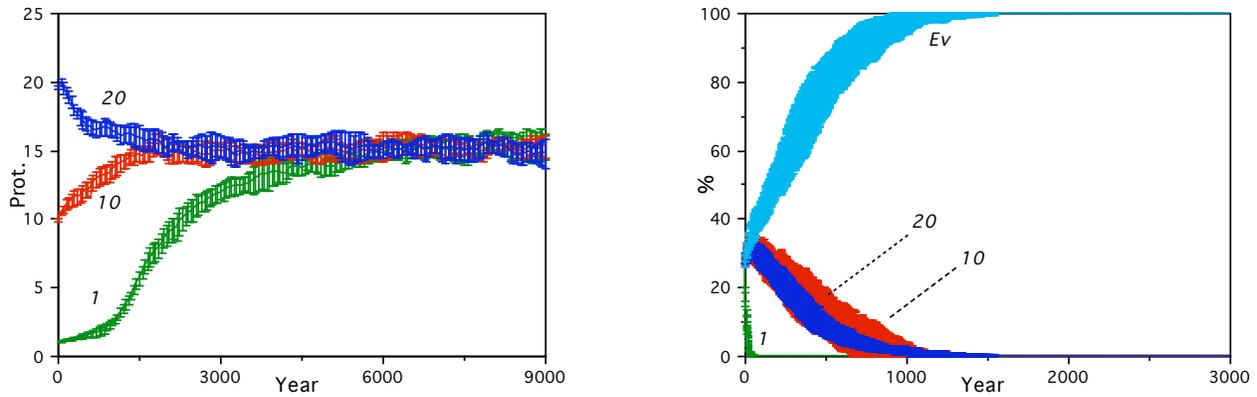


Figure 9: Evolution of protection periods starting from the fixed period populations (left), and competition between the evolved populations (right).

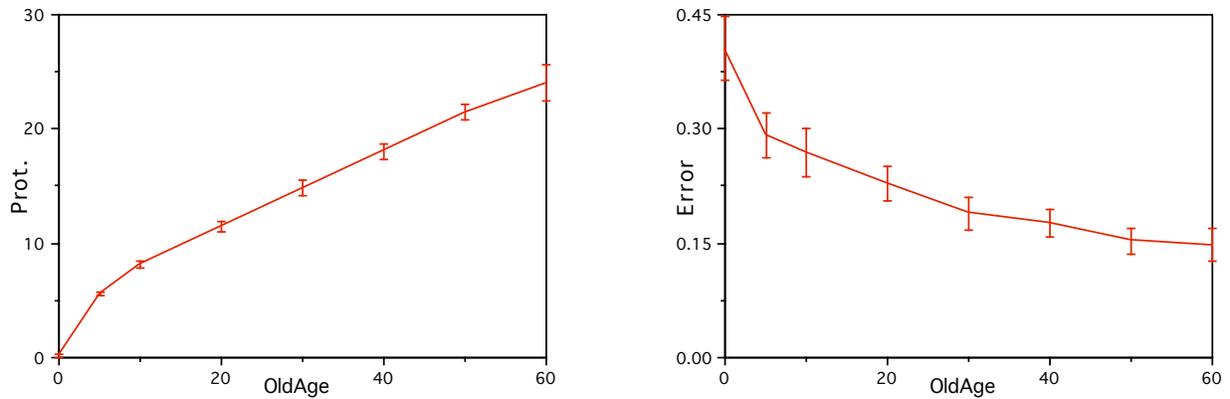


Figure 10: Effect of age of onset of "old age" and life-span on the evolved protection period (left), and average population performance (right).

- [4] Bullinaria, J.A. (2003). Evolving Efficient Learning Algorithms for Binary Mappings. *Neural Networks*, **16**, 793-800.
- [5] Bullinaria, J.A. (2005). Evolved Age Dependent Plasticity Improves Neural Network Performance. In: *Proceedings of the Fifth International Conference on Hybrid Intelligent Systems (HIS 2005)*, 79-84. Piscataway, NJ: IEEE.
- [6] Charnov, E.L. & Berrigan, D. (1990). Dimensionless Numbers and Life History Evolution: Age at Maturity Versus the Adult Lifespan. *Evolutionary Ecology*, **4**, 273-275.
- [7] Clutton-Brock, T.H., Guinness, F.E. & Albon, S.D. (1983). The Cost of Reproduction to Red Deer Hinds. *Journal of Animal Ecology*, **52**, 367-383.
- [8] Cole, L.C. (1954). The Population Consequences of Life History Phenomena. *Quarterly Review of Biology*, **29**, 103-137.
- [9] Gadgil, M. & Bossert, W. (1970). Life History Consequences of Natural Selection. *American Naturalist*, **104**, 1-24.
- [10] Harvey, P.H. & Zammuto, R.M. (1985). Patterns of Mortality and Age at First Reproduction in Natural Populations of Mammals. *Nature*, **315**, 319-320.
- [11] Julesz, B. & Kovacs, I. (Eds) (1995). *Maturation Windows and Adult Cortical Plasticity*. Reading, MA: Addison-Wesley.
- [12] Kaplan, H., Hill, K., Lancaster, J. & Hurtado, A.M. (2000). A Theory of Human Life History Evolution: Diet, Intelligence, and Longevity. *Evolutionary Anthropology*, **9**, 156-185.
- [13] Lamb, M.E., Bornstein, M.H. & Teti, D.M. (2002). *Development in Infancy: An Introduction*. Mahwah, NJ: Lawrence Erlbaum Associates.
- [14] Sloman, A. & Chappell, J. (2005). The Altricial-Precocial Spectrum for Robots. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 1187-1193. IJCAI.
- [15] Stearns, S.C. (1989). Trade-offs in Life History Evolution. *Functional Ecology*, **3**, 259-268.
- [16] Stearns, S.C. (1992). *The Evolution of Life Histories*. Oxford, UK: Oxford University Press.
- [17] Stearns, S.C. & Crandall, R.E. (1981). Quantitative Predictions of Delayed Maturity. *Evolution*, **35**, 455-463.
- [18] Stearns, S.C. & Koella, J. (1986). The Evolution of Phenotypic Plasticity in Life-History Traits: Predictions for Norms of Reaction for Age- and Size-At-Maturity. *Evolution*, **40**, 893-913.
- [19] Warner, R.R. (1984). Deferred Reproduction as a Response to Sexual Selection in Coral Reef Fish: A Test of the Life Historical Consequences. *Evolution*, **38**, 148-162.