

Lifetime Learning as a Factor in Life History Evolution

John A. Bullinaria*
University of Birmingham

Keywords

Artificial life, life history, evolution, learning, neural networks

Abstract An artificial life approach is taken to explore the effect that lifetime learning can have on the evolution of certain life history traits, in particular the periods of protection that parents offer their young, and the age at first reproduction of those young. The study begins by simulating the evolution of simple artificial neural network systems that must learn quickly to perform well on simple classification tasks, and determining if and when extended periods of parental protection emerge. It is concluded that longer periods of parental protection of children do offer clear learning advantages and better adult performance, but only if procreation is not allowed during the protection period. In this case, a compromise protection period evolves that balances the improved learning performance against reduced procreation period. The crucial properties of the neural learning processes are then abstracted out to explore the possibility of studying the effect of learning more generally and with better computational efficiency. Throughout, the implications of these simulations for more realistic scenarios are discussed.

I Introduction

The term “life history” refers to the sequence of events and changes that take place during an organism’s lifetime. This covers issues such as stages of growth, reaching reproductive maturity, production and nurturing of offspring, menopause, and death. There are numerous tradeoffs involved in their timing, with widely different patterns of such events evolving for different species and different environmental conditions (e.g., [17, 27, 30, 31]). Classic examples include the tradeoff between reproduction and growth [35], and between nursing and survival [15]. There also appear to be correlations between various aspects of life history, for example, fixed ratios within lineages of life span to age of maturity [13, 18]. This study takes an artificial life approach to look specifically at one factor that has received relatively little attention in the past, namely, how the ability to learn during an individual’s lifetime affects the optimal period of protection that parents should provide for their children, and how that should affect the children’s age at first reproduction. The general idea that slow growth and extended infant dependency could be associated with learning is not new (e.g., see the discussions in [5, 21, 22, 24]), but the study presented here appears to be the first to simulate explicitly the interaction of learning and life history evolution.

In nature, parental protection varies enormously, from *precocial* species in which the young are born well developed and require virtually no protection, to *altricial* species in which the young are born helpless and require long periods of parental care before they are capable of surviving on their own.

* School of Computer Science, University of Birmingham, Birmingham, B15 2TT, UK. E-mail: j.a.bullinaria@cs.bham.ac.uk

There is also some variation in how the age of first reproduction relates to the period of protection. It might seem obvious that early first reproduction is advantageous because it increases the chance that individuals will survive to reproduce, and decreases the time wasted not reproducing [16]. However, if parents protect their young, the first advantage will be less relevant, and the second may be balanced if avoiding the costs of early reproduction leads to improved performance later in life, with increased reproduction overall. In fact, there are many factors involving growth, size, fecundity, environment, and so on, that are known to affect the age at maturity and first reproduction (e.g., [26, 31]). One major advantage of simplified artificial life simulations is that they render it possible to explore the individual factors without all the confounds inherent in empirical measurements of existing biological populations.

This investigation was prompted by the observation that human infants are particularly altricial, even compared with other primates, and require extended periods of parental protection and support (e.g., [23]). There are numerous factors that could have led to this (e.g., [21, 28]), but there are two particularly important processes that are known to take place during the protection stage of altricial species—the infants are growing, and they are learning. Human infants clearly do need to grow considerably after birth, and consequently survival without parental protection would be virtually impossible. Learning is also crucial for humans, and for other species for which relatively complex behaviors are required, since encoding all the necessary skills genetically is likely to be difficult, and even if that were evolutionarily possible, adaptation would still be needed to cope with their rapid growth processes and the changing and unpredictable nature of their environment. Otherwise, innate behavior would be adequate [7], though, where possible, the costs of learning will tend to result in genetic assimilation of the learned behaviors [3, 36]. This article aims to explore the extent to which such learning processes alone could be responsible for the evolution of long protection periods. Studies of the effect of brain growth and adult brain sizes on life history variables [24], and arguments that the human's unique pattern of early neurological development is responsible for the substantially earlier weaning of humans than of the great apes [22], already lend weight to the idea that neural processes play an important role in determining the details of human life history. Striedter [34] provides a good overview of the factors involved in brain evolution and the differences observed across species.

This study begins by simulating the evolution of populations of learning individuals using artificial neural networks that must learn quickly how to perform well on simplified classification tasks. With individuals of all ages competing to survive and procreate based on their learned performance, the effect of different extended periods of parental protection can be explored. Allowing the protection period to evolve demonstrates how the advantages of protection trade off against any associated disadvantages so that a particular period of protection emerges through evolution. Varying the details of the simulations (in particular, how the typical life span of the simulated species affects the protection period, and how the age of first reproduction might be affected by the parental protection) renders it possible to determine how these various factors interact, and how different patterns of behavior could arise from other species-specific factors. Then, by abstracting out the crucial properties of the neural learning process, it is possible to explore the effect of learning more generally.

The next section describes the approach taken to simulate the evolution of populations of neural network learners. Section 3 then establishes baseline performance results for populations that evolve 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 neural networks are analyzed in Section 5, the robustness of the simulation results is checked in Section 6, and the associated tradeoffs are explored in Section 7. The potential for extending the results to more abstract general learning processes is considered in Section 8. Finally, the article ends in Section 9 with some conclusions and a discussion of the implications for more realistic scenarios.

2 Evolving Neural Network Learners

To maintain as wide a relevance as possible, this study aims to adopt a fairly abstract approach for simulating the crucial features of the evolution of most animal populations, with particular emphasis

on the aspects of fitness associated with lifetime learning. Since it is the brains of animals that are largely responsible for their ability to learn complex behaviors, it is natural to begin by representing the animal populations by simple artificial neural network models. The extent to which it really is necessary to include explicit neural learning, rather than simpler approximations to such learning processes, will be explored later, in Section 8.

The artificial life simulations here will involve populations of individual neural networks, each specified by a set of innate parameters, that must learn from a continuous stream of input patterns how to classify future input patterns. The inputs could, for example, correspond to relevant observable features of plants or other animals, and the desired output classes (categories) could correspond to important properties of them, such as being edible or dangerous. The fitness of each individual will be measured by how well it classifies the new inputs *before* discovering their correct classes and learning from them. By forcing the individuals to compete to survive and procreate, according to their relative fitness, populations of increasing fitness can be expected to emerge. Moreover, in order to compete effectively in a population consisting of individuals of all ages, each individual must not only learn how to perform well, but 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.

Proceeding with such simulations obviously requires the choice of a specific concrete learning system and training data set(s), and it makes sense to employ a setup that has already been explored in some detail and proved instructive elsewhere [6, 8–10]. Consequently, standard fully connected multi-layer perceptron neural networks were used, with one hidden layer, sigmoidal processing units, and training by gradient descent weight updates using the cross-entropy error function on simple classification (i.e., categorization) tasks [4]. Since most real-world classification tasks involve learning nonlinear decision boundaries in a space of real-valued inputs, the set of tasks was chosen to have continuous two-dimensional input spaces, conveniently normalized to a unit square, with particular circular classification boundaries. Each input to a neural network was thus represented by a position in the unit square input space, and the network output represented the predicted probability that the input corresponded to a particular class or category. The aim was to learn from past experience (i.e., binary target outputs) the class of each new input, in other words, to discover where the boundaries between classes lay in the input space. This setup was sufficiently simple to allow extensive simulations, yet involved the crucial features and difficulties of real-world problems. It was important that the networks did not evolve to cope only with a particular single data set, so each individual network was assigned a randomly chosen classification boundary of the specified type, and had to learn from a stream of randomly drawn data points from the input space. The individual fitness at each stage could then be defined as the generalization ability, that is, the average number of inputs correctly classified (e.g., defined as having the network outputs—i.e., class probabilities—within 0.2 of the binary targets) *before* training on them.

An additional factor that complicates this learning process for many tasks is that humans and some other animal species have *critical periods* for learning, and outside that period the learning is more difficult [2, 20]. It is not clear how, or even if, that will interact with a period of parental protection. It has certainly been demonstrated in previous computational studies that evolving neural network learning rates that vary during the learning process does lead to improved learning performance, and that the evolved time (t) dependences are qualitatively similar to humanlike age dependences [7, 9, 10]. Those previous studies indicate that, for current purposes, such time-dependent learning rates $\eta_L(t)$ can be conveniently approximated by introducing a simple two-parameter exponential scale factor $s(t)$ to multiply innate initial learning rates $\eta_L(0)$:

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

in which t is the age of the individual in simulated years, the baseline β specifies the ratio of the final to initial learning rates, and the time constant τ determines how quickly the learning rate rises or falls toward the final value. Both β and τ are evolved to take on the positive values that result in the best

performance under the given circumstances. Clearly, if time dependence of this type proves not to be beneficial, the parameters will be able to evolve so that $s(t) = 1$. Note also that the subscript L has been introduced because earlier studies [6, 8, 10] have shown empirically that better performance can be 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. Again, if this proves not to be beneficial in this particular case, the parameters will be able to evolve so that they are equal across any combination of components for which that proves useful.

The four initial weight parameters r_L , four learning rates η_L , and two variable plasticity parameters τ and β , together with a standard momentum parameter α and weight decay regularization parameter λ [4], give a total of 12 real-valued evolvable innate parameters for each network. In principle, the number of hidden units could also be evolved, but doing that invariably results in it rising quickly to whatever limit is placed on it, considerably slowing down the simulations in the process, so this is usually best kept fixed at some appropriate value [10]. For the current study it was fixed at 20 hidden units for all networks, which was found to be more than enough for successful learning of the given tasks.

After specifying the details of the learning individuals, they then needed to be integrated into the evolutionary process. This required the neural network learning time scales to be aligned with the lifetime and evolutionary time scales. This was conveniently done by defining a *simulated year of experience* to be 1,200 training data samples, and computing the fitness of each individual at the end of each simulated year as the average classification performance over that year. This simulated year could then be used as a common unit of time across all the learning, lifetime, and evolutionary processes. A computationally feasible fixed population size of 200 was maintained throughout (consistent with the idea that there are fixed total food resources available to support the population), achieved by replacing the individuals that died with children of the fittest individuals. Deaths occurred by losing a fitness comparison “fight” against other individuals, or randomly due to old age beyond a natural life span. The details of all these evolutionary factors needed to be fixed to encourage evolutionary change and preserve a reasonably diverse population, but fortunately the precise specifications proved not to affect the qualitative results a great deal, so convenient round numbers were chosen for the associated parameter values.

The time scale was chosen so that learning the given tasks typically took a humanlike 10 to 15 years, and the natural life span was defined to be 30 years, beyond which a random 20% of older individuals died each year. This allowed individuals enough training samples to learn their given tasks and have a reasonable period during which they were fit enough to reproduce, yet prevented the populations from becoming dominated by a few very old and very fit individuals. In addition, each of the unprotected individuals was forced to compete each year with another randomly chosen eligible individual, and would die if its fitness proved to be lower than its competitor's. To preserve a reasonable population age distribution, deaths in this way were limited to 10% of the population each year, though for populations with relatively large parental protection periods, this limit was rarely reached. The children were generated by crossover and mutation from two parents chosen each year by pairwise fitness comparisons of the eligible individuals. This was conveniently 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. A litter size of one was chosen, since that appears to be optimal for large primates [25]. Clearly, these details are gross oversimplifications of real biological processes, but they constitute a manageable starting point that includes approximations of all the key processes.

This simple nature-inspired algorithm, tracking an evolving population of aging individuals over time, with the oldest and least fit individuals tending to die and be replaced by children of the fittest individuals at the end of each simulated year, has already been shown to work well in practice [9, 10]. The crucial additional feature here is that the children can be protected by their parents until they reach a certain age, and cannot be killed by competitors before then. This introduces an implicit cost

to the parents in that the more children are protected, the higher the chance they stand of being in the maximum of 10% of the population that die through competition each year, and the higher the chance they have to compete against another adult rather than a less experienced child. A cost that the children can, and will, be forced to bear is that they are prevented from having any children of their own before they leave the protection of their parents. The effect of this cost is something that will be considered in some detail later.

The following sections will present simulation results that 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. First, however, it is important to note some of the simplifications inherent in the definition of the protection process: The parents are assumed to be perfect protectors, whereas in reality their own fitness will normally affect their protection abilities, as will the number of children they are protecting at any given time. Moreover, real parental protection will usually influence the training data available to the infants, but this complication has been ignored and the nature of the training data kept constant across all cases. Obviously, these design choices are oversimplifications, but wherever possible, for each such factor in the simulations, care was taken to make the design decision in such a way that the parental protection effect was least likely to emerge. Then, if it did, there could be some confidence that it really was an important and robust effect. One factor likely to reduce the parental protection relates to the fact that there will usually be more direct costs to parents protecting young (e.g., [14]), but this proved difficult to quantify within the current framework. This complication was therefore ignored for the present purposes, though a better account of the reproductive and protection costs will certainly be required for more realistic models.

3 Simulation Results for Fixed Protection Periods

The natural starting point was to perform the evolutionary neural network simulations described above for a carefully selected range of fixed protection periods to determine if there are any differences between the populations that evolve. The learning time scale was set so that evolved individuals were typically able to learn the given task in 10 to 20 simulated years, and after 30 simulated years they start dying of old age, so protection periods of 1, 10, and 20 years formed a suitable representative sample. The evolution of the initial learning rates $\eta_L(0)$ for these three cases are shown in Figure 1, with means and variances over 10 runs (which proved sufficient to establish statistically significant results). In each case, the pattern of evolved parameters and relatively low variances across runs are similar to those found in earlier studies [9, 10], but subtle differences can be seen between the final parameter values, and the evolutionary process is noticeably slower to settle down for the longer protection periods.

The means of generalization error performance across populations during evolution for each protection period are compared in the bottom-right graph of Figure 1. This appears to show that the two longer protection periods do have a clear advantage in this respect. However, such simple population averages are built up from complex age-dependent error distributions, and it is inevitable that the population age distributions will depend considerably on the protection period. In particular, longer protection periods will tend to result in populations with more older, and hence fitter, individuals, so observing improved average population fitness alone is not sufficient to demonstrate that extended protection periods really do result in improved individual performance. To show that they do provide a real evolutionary advantage, the protection period needs to be allowed to evolve alongside all the other parameters.

4 Simulation Results for Evolved Protection Periods

For simulations in which the protection period is allowed to evolve, as another real-valued parameter, the evolution of that period and the associated initial learning rates $\eta_L(0)$ are as shown




Figure 1. Evolution of the average initial learning rates $\eta_L(0)$ and variances over 10 runs for fixed protection periods of 1, 10, and 20 years, and comparison of the corresponding performance error rates, where the subscript refers to the relevant network connections (*IH*, *HO*) and biases (*HB*, *OB*). It is clear that the protection period does affect what evolves.

in Figure 2, again with means and variances over 10 runs. During the early stages of evolution, while all the individuals 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 significant consequences for the evolved learning processes. The effect on the initial learning rates $\eta_L(0)$, as seen in Figures 1 and 2, is barely noticeable, but the effect on the age dependences of the learning rates is very clear. Figure 3 shows how the scale factor time constant τ and baseline β of the evolved learning rate vary with protection period, and the corresponding scale




Figure 2. Evolution of the average initial learning rates $\eta_L(0)$ and variances over 10 runs when the protection period is allowed to evolve (left), and the evolution of the protection period (right). A well-defined optimal protection period of about 15 years emerges.




Figure 3. The effect of different protection periods on the evolved age dependence of the learning rates. Plotted are the evolved baseline β and time scale τ (left), and the resultant learning rate scale factors s (right), for the representative fixed protection periods (1, 10, 20) and evolved period ($Ev \approx 15$). Longer protection periods lead to the evolution of longer learning periods.

factors s as a function of age. Increased protection periods mean there is less urgency to learn, and this results in a lengthening of the period during which significant learning takes place, and a consequent reduction of learning rates needed during that period. This will provide an evolutionary advantage if the extended learning period results in better performance after the period of protection has ended, and could thus provide a reason why extended protection periods should evolve.

As noted above, however, the protection period will also affect the age distribution of the population, and this can also affect its evolution. The averages and variances of the crucial factors are compared for the various protection periods in the four graphs of Figure 4. An increased protection period will reduce the number of deaths per year due to competition from the maximum of 20 (top left). This will result in more individuals living longer and increase the average age of the population (top right), and lead to more individuals living long enough to start dying of old age (top left). 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 increasing protection period (bottom-left), without any introduced direct cost for parents protecting more children. Together these factors result in rather different age distributions for the various populations (bottom right). 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 toward zero.

The various trends observed in Figures 3 and 4 are all monotonic with respect to the protection period, and the evolved-protection-period population results are consistent with what would be expected from their evolved period of around 15 years. The question arises as to what causes the evolved protection period to stabilize at the value it does, rather than at some higher or lower value. It is unlikely that longer protection periods could result in inferior learning, because the evolution can easily adjust the learning process to ignore any final part of the protection period that worsens the performance. However, even if increased protection periods result in better individual performance, there will still be an associated reduction in the number of children per individual, and that will place those individuals at an evolutionary disadvantage, which will tend to drive a reduction in the protection period. There will be a tradeoff between improved individual fitness and reduced potential for procreation. To explore this tradeoff, the next section investigates more carefully the performance of the evolved populations.

5 Analysis of the Evolved Performance

Understanding the evolved performance fully requires going beyond the population averages of the previous two sections. Figure 5 shows the means and variances of the individual performance




Figure 4. The effect of different protection periods on the evolved populations: the average death rates (top left), the average age of individuals (top right), the average number of children per individual (bottom left), and the resultant age distributions within the population (bottom right).

error rates (i.e., inverse fitness) during learning. There do appear to be major reductions in the mean error rates resulting from increased protection periods, with associated delays in reaching those lower error rates, as permitted by the delayed need to compete. However, the variances in the error rates are extremely large, mainly due to long tails at the high end of the distribution of errors, and this renders it difficult to analyze the significance of the mean performance differences. This rather skewed distribution of errors has been observed before, and is a consequence of the evolutionary




Figure 5. The mean individual performance error rates (left) and variances (right) during learning for the evolved populations with different protection periods. Longer protection periods result in slower learning but better final performance.

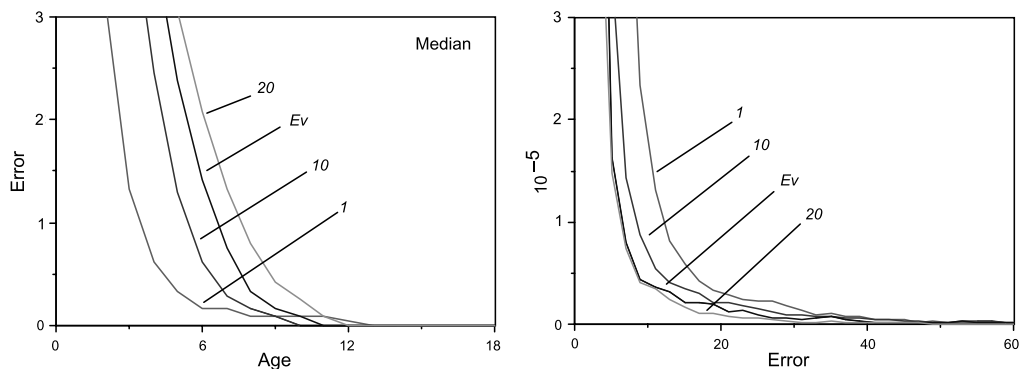


Figure 6. The median performance error rates during learning (left) and the distribution of errors for all individuals between the ages of 50 and 60 (right), for the different protection periods. The median performance is perfect by about 12 years of age, but the population has many large errors persisting into old age.

pressure to learn quickly resulting in individuals adopting risky learning strategies, which lead occasionally to very poor performance [9, 10]. Since extended protection periods reduce the pressure to learn quickly, it is easy to understand the observed performance improvements.

The median learning performance error rates, which are not affected by the long tails of the error distributions, are shown on the left of Figure 6, and indicate essentially perfect performance by age 12 for all protection periods, though with the expected slower initial learning for longer protection periods. The distributions of errors for older individuals (aged between 50 and 60), shown on the right of Figure 6, give an indication of the extent of the small numbers of cases of very large error that persist even for long protection periods. There is certainly a massive peak around zero errors, as would be expected at those ages, but there remain a significant number of very large errors. The upper and lower quartile performance error rates, shown in Figure 7, are consistent with the expectation that longer protection periods will slow the learning, and also confirm that they improve the learned performance at the poorer performance end of the spectrum.

To reinforce the above understanding, all the evolutionary simulations were repeated without allowing age dependences in the learning rates [12]. In this case, there is no easy mechanism for longer protection periods to lead to slowing and extending of the learning process, as was achieved before by modifying the age-dependent learning rate scale factor as seen in Figure 3. However, evolution of other learning parameter differences still allows longer protection periods to result in significant improvements in the error distribution for old individuals and in lower-quartile learning




Figure 7. The upper- and lower-quartile performance error rates during learning (left and right respectively), for the different protection periods. Longer protection periods massively improve the chances of achieving perfect performance later in life.

performance, with relatively modest slowing of the upper quartile and median learning, as shown in Figure 8. In this case, the evolved protection period is about 16 years, only slightly longer than found before. The differences and improvements afforded by the possibility of age-dependent learning rates confirms the importance of including in the simulations as many as possible of the known features of real learning systems.

6 Robustness of the Evolutionary Simulations

The above simulations and analyses have established that longer protection periods do offer clear learning advantages, and relatively few disadvantages, whether or not age-dependent learning rates are allowed. Evolving the protection period suggests that there is a tradeoff between this learning advantage and the cost of reduced procreation opportunities. However, evolutionary simulations can easily stall in less than optimal configurations, particularly if there is a loss of diversity in the evolving populations, or a poor choice of initial population [10]. A check is therefore required to verify that the evolved protection periods are 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 done by allowing the protection period in the three fixed-period evolved populations to evolve away from its previously fixed values. The results of this are shown on the left of Figure 9, with means and variances across 10 runs. For each case, there is a relatively fast rise or fall to the same evolved period of around 15 years that emerged before. A qualitatively identical pattern of results emerge for populations evolved without age-dependent learning rates [12].

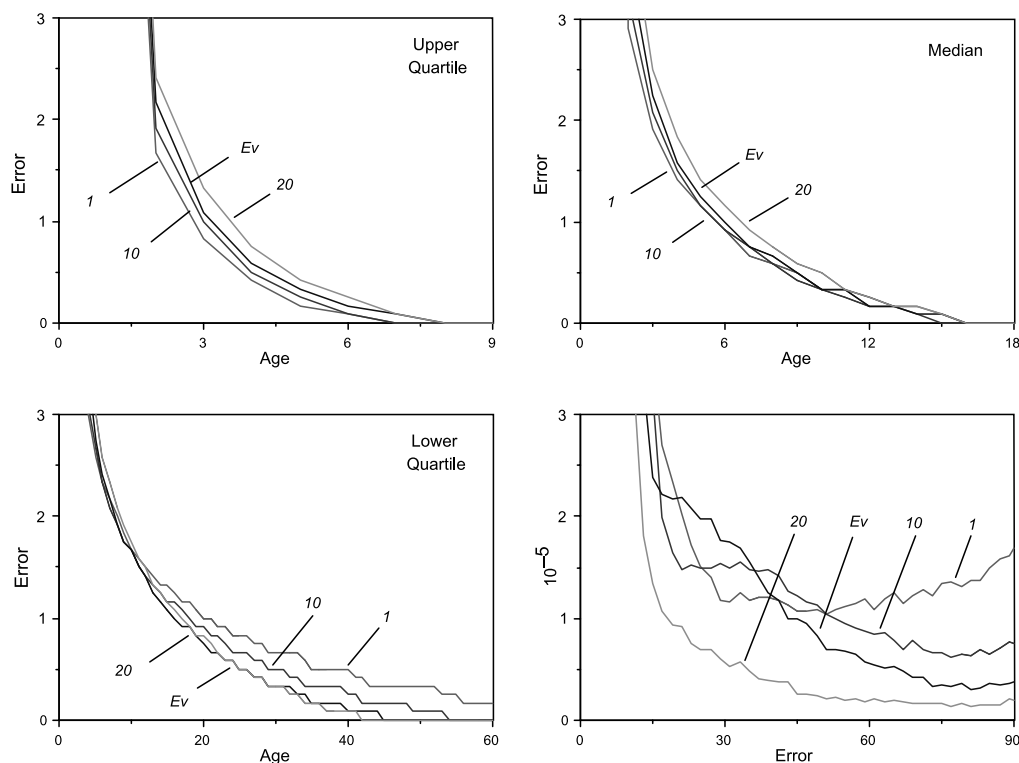


Figure 8. The upper-quartile, median, and lower-quartile performance error rates during learning (top left, top right, bottom left) and the distribution of errors for all individuals between the ages of 50 and 60 (bottom right), for the different protection periods, when age-dependent learning rates are not allowed. Comparison with Figures 6 and 7 reveals a clear reduction in performance levels.




Figure 9. Allowing the protection periods to evolve starting from the fixed-period populations (left), and natural selection working on the combined evolved populations leading to the dominance of the evolved-protection-period individuals (right). These confirm the robustness of the evolved protection period of about 15 years.

An additional test involves combining the 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 the initial constituent populations had already been optimized by evolution, the children inherited characteristics from just their fittest parent, with no further cross-over or mutation allowed. The results of this process are shown on the right of Figure 9, with means and variances across 20 runs. Individuals with virtually no parental protection are wiped out almost immediately, and eventually individuals with the evolved protection period come to dominate the whole population. Again, a qualitatively identical pattern of results emerges for populations evolved without age dependences in the learning rates [12]. It seems that the evolved protection periods are quite robust.

7 Exploration of the Tradeoffs

The simulations described above have demonstrated 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 opportunities, and that in turn places those populations at a serious evolutionary disadvantage. The evolving populations manage this tradeoff, and establish a suitable compromise value for their protection period, appropriate for individuals that start dying of old age after 30 years. Two issues relevant to this tradeoff need further investigation: How does the overall individual life span affect the balance, and how important is the prevention of procreation while being protected?

Repeating the above evolutionary simulations with the onset of old age at different ages, corresponding to different natural life spans, will show how that affects the emergent protection period and associated performance levels. Figure 10 presents the results of doing that. 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. There is also a clear improvement in the individual adult performance levels as the life span becomes longer. These demonstrate that the effect of learning is of sufficient importance that the amount of time devoted to it continues to increase as the life spans become longer, rather than using all the extra time for procreation. Future studies might consider the extent to which such learning-driven extended protection periods will lead to the coevolution of longer growing periods, too. There are known to be correlations between age of maturity and life spans in biological populations, though there are certainly other factors involved there besides learning (e.g., [13, 18, 32, 33]). Interestingly, after controlling for body size, there is also a significant correlation between brain weight and life span among higher primates [1]. It certainly seems plausible that a correlation should exist between the weight of the brain and the protection period that would lead to its optimal learning performance.




Figure 10. The effect of the age of onset of old age, and hence the life span, on the evolved protection period (left), and the corresponding average individual learning performance levels (right). These demonstrate the importance of an extended protection period for learning, relative to a potential increase in procreation period.

A further set of simulations have demonstrated that if the life span is also allowed to evolve freely within the current setup, it keeps on increasing indefinitely. In biological populations, of course, there are many other factors that act to restrict life spans (e.g., [21, 31]), but simulating such tradeoffs is beyond the scope of this study. Similarly, simulating factors that might result in reproduction stopping early within the typical life span, tendencies for increased reproductive problems at older ages, the effects of grandmothering, and such like [19, 21] are also left for future studies.

In all the above simulations, it has been decreed that procreation is not allowed while being protected, and it has been argued that this is behind the tradeoff that has restricted the protection period. This restriction on procreation is clearly more true of some natural species than others. However, the aim here is not to tie the simulations to particular natural species, but rather to explore the extreme behaviors with a view to understanding the general principles involved. To check the relation between restricted procreation and restricted protection period, the evolutionary simulations were repeated with procreation allowed irrespective of whether the individual concerned was being protected. Allowing procreation from the first fitness comparison point (at age one year) results in the protection period evolving quickly to be safely beyond the normal life span of the individuals, as seen on the left of Figure 11. This means that all individuals are protected all their lives from deaths due to competition, and only die because of old age. Of course, this ignores the fact that most parents will not live long enough to protect their children for that long, but it does show the underlying trend that is likely to persist in more realistic simulations. Rather than simply assuming




Figure 11. Evolution of the protection period for zero non-procreation period (left), and simultaneous evolution of the protection and non-procreation periods (right), if procreation while protected is allowed. In both cases the protection period evolves to be more than the natural life span.

that someone will be around to protect each child for the evolved period, as has been done thus far, future simulations will need to involve much larger populations, which will allow the consideration of explicit families and the patterns of protection they can offer.

If the non-procreation period (or, equivalently, the age of first reproduction) is also allowed to evolve as an independent parameter, that period falls quickly, to finish before the first fitness comparison point at age one, and the protection period again rises to beyond the normal life span, as seen on the right of Figure 11. What both the scenarios of Figure 11 reintroduce is the need to compete at all ages to procreate, and this encourages faster learning again and the return of the unwanted associated side effects, such as the evolution of risky learning strategies that sometimes result in persistent very poor performance at all ages. This can be seen in the return of a shorter learning-rate age-dependence time scale τ and increased mean error rates at age 60, as shown in Figure 12. If the need to compete to procreate is removed (i.e., fitness is not used to choose parents), the protection period again rises to beyond the normal life span, leaving no fitness pressure at all in the evolutionary process, and so no performance improvement takes place.

It seems, then, that there are individual performance advantages that emerge as a result of preventing offspring from reproducing while being protected, but this feature does not evolve in the current simulations. In these simulations, the tradeoff between reproductive opportunity and improved adult performance favors early reproduction, even though allowing the protected children to reproduce will compromise the parents' own reproductive success rate. Interestingly, a significant correlation is found in higher primates between brain weight and age of first female reproduction [1], which makes sense if delayed reproduction really can improve learning performance for complex tasks. Of course, correlation does not imply causation, but such correlations are certainly consistent with the various factors coevolving in these species to provide better overall performances. Clearly, delayed reproduction fails to emerge in the current simulations because there are other important factors that have not been included in them. For example, if the adult performance affects the population size that can be supported by the environment, or affects some absolute ability of individuals to survive (e.g., associated with crucial food gathering skills or competition with other species), then behaviors that increase that performance will be more likely to emerge in simulations that include such factors. There is certainly much scope for future work to explore these issues.

8 Evolution with Abstracted Learning Processes

To render the above artificial life simulations as reliable as possible, the learning task was taken to be something that is likely to be a component of real animal learning, and the learning process was




Figure 12. Comparison of the evolved learning-rate age-dependence time scales τ (left) and mean error rates at age 60 (right), for the case of procreation while protected (PWP) and for the earlier simulations. It is clear that the restricted procreation is crucial for the improved performance.

implemented in the form of a traditional artificial neural network model. Such evolutionary neural network simulations, however, are extremely computationally expensive, and it makes sense to consider now the extent to which the various issues can be explored by abstracting out the key properties of the learning process, and evolving only the life history factors. The problem with attempting this is that the error distributions and associated fitness levels during neural network learning depend in a complex manner on the learning algorithm and its evolved parameters, and these depend in a nontrivial way on the evolutionary pressures and population age distributions, which in turn are affected by the protection periods that the study is attempting to determine. With so many unknowns, and the complexity of their interactions, it is not possible to predict reliably in advance what distributions of all these things will emerge across the evolving populations. However, having run all the above simulations, a good idea of the patterns that emerge is now available, and it becomes feasible to abstract out the key issues and attempt to run the simulations again without the full neural learning processes.

The simplest conceivable approximation to the full learning process would be to have each individual's cost (i.e., abstracted classification error, or inverse fitness) fall linearly with increasing age from 100 down to 0. Figure 13 shows the results of evolving the protection period in this case, for a range of learning rates δ , with means and variances over 20 runs. The deterministic case, which has the individual costs reduced by δ each year down to zero, has similar results to a stochastic version in which the reductions are drawn randomly from the range $[0, 2\delta]$. The population mean cost is approximately linear in the expected learning time (ELT) $100/\delta$ for both the deterministic and stochastic versions. The evolved protection periods also start close to linear in $100/\delta$, but level off around 30 years, the point at which the individuals start dying of old age. Below around 20 years there is a slight, but significant, increase (of 1 to 2 years) in the protection periods for the stochastic case over the deterministic case, which reflects the uncertainty in reaching perfect performance by the expected time ($100/\delta$). Predictably, the best mean performance is achieved with very high learning rates δ , for which all individuals reach perfect performance before their first round of competition to survive or procreate at the end of their first year. Consequently, if the learning rate δ is evolved along with the protection period, it quickly achieves very high levels, and the protection period takes on the associated very low levels as indicated by Figure 13. Of course, with neural networks one cannot keep on increasing the learning rate and expect the learning time to decrease with it. Eventually, at some task-dependent point, the approximation to true gradient descent breaks down, and the learning performance deteriorates. As was observed in the evolving neural network simulations above, evolution can be used to find the best values for the learning parameters, and having slower learning with longer protection periods does provide a clear evolutionary advantage.




Figure 13. Emergent population mean costs (left) and protection periods (right) when the individual costs (i.e., abstracted performance errors) decrease linearly with age at learning rate δ . Deterministic and stochastic versions are shown as a function of the expected learning time $100/\delta$.

To provide a better approximation to the observed neural learning process, with faster learning leading to riskier learning strategies that increasingly lead to persistent poor performance, one can consider having the learning process stop at some random point in the performance range $[0, 100]$ with a probability $\rho\delta$ that increases linearly with the learning rate δ and an associated parameter ρ . The top two graphs of Figure 14 show how the stochastic version's performance then depends on δ and $100/\delta$ for four representative values of ρ . The higher ρ is, the lower the value of δ at which significant deviations from the earlier $\rho = 0$ case arise. Eventually, in all $\rho > 0$ cases, the cost begins to increase again with δ until $\rho\delta = 1$, at which point it begins to fall slowly along a ρ -independent curve. The bottom-left graph shows that the relation between the protection period and $100/\delta$ is not much affected by the size of ρ . The bottom-right graph shows how the cost varies with the protection period for the different values of ρ .

The cost plots in Figure 14 show clear minima for each value of ρ , and successful evolutionary processes are likely to result in the emergence of optimal learning rates δ and protection periods that follow similar trends. The top two graphs of Figure 15 show the ELTs $100/\delta$ and protection periods that actually emerge through evolution as a function of the parameter ρ . As ρ increases, the best possible learning time $100/\delta$ also increases, and the best protection period follows suit. In the top-right graph of Figure 14 it can be seen that around $\rho = 0.08$ the global cost minimum shifts from the local minimum at fairly high $100/\delta$ to the point at zero $100/\delta$. The bottom-left graph of Figure 15, showing cost against ρ , indicates that the evolutionary process gets trapped in the local minima at the $\rho = 0.1$ point, but for higher ρ the global optima are found, as is clear from the sudden drops in the top two graphs of Figure 15. The bottom-right graph shows that the mean evolved protection period is always slightly longer than the mean evolved ELT $100/\delta$. This is a

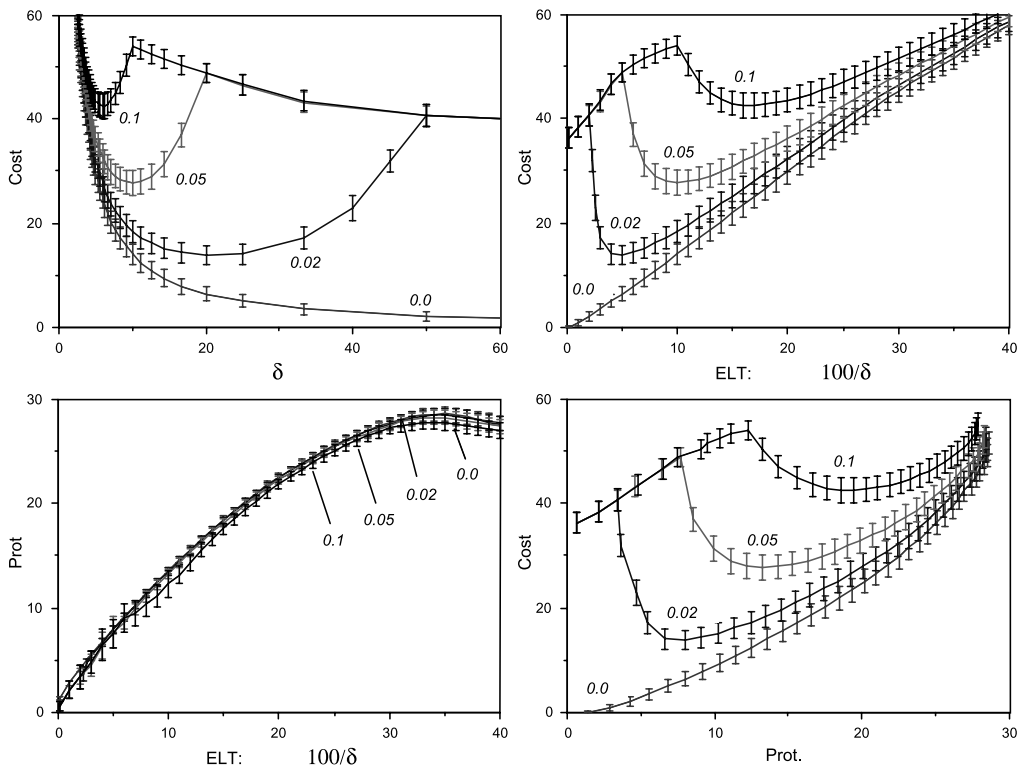


Figure 14. Results when the linear individual cost improvement with age stops with probability $\rho\delta$ at a random cost in the range $[0, 100]$, for $\rho \in \{0, 0.02, 0.05, 0.1\}$: mean population cost as a function of learning rate δ (top left) and as a function of expected learning time $100/\delta$ (top right), protection period as a function of ELT $100/\delta$ (bottom left), and cost as a function of protection period (bottom right). There are clear cost minima that vary with the value of ρ .




Figure 15. Evolutionary simulation results when the linear individual cost improvement with age stops with probability $\rho\delta$ at a random cost in the range $[0, 100]$: mean evolved expected learning time $100/\delta$ as a function of parameter ρ (top left), protection period as a function of ρ (top right), population-mean cost as a function of ρ (bottom left), and protection period as a function of $100/\delta$ (bottom right). The parameter ρ acts as an abstract measure of learning difficulty.

reflection of the fact that mutations always lead to distributions of learning times and protection periods, and that the protection period should always be long enough to accommodate a reasonable number of individuals that are slower than average. For the same reason, the protection periods for the stochastic case are always slightly longer than those for the deterministic case, as previously observed in Figure 13.

The parameter ρ is an abstract measure of learning difficulty; it can be regarded as an approximate representation of the difficulty a neural network learning algorithm has with its given task. It is certainly a rough approximation, but it does have the required properties. Easy tasks will correspond to low ρ , be learned quickly, and have short protection periods. Harder, or more complex, tasks will correspond to higher values of ρ , take longer to learn, and benefit from longer protection periods. The individual performance levels that emerge in the abstracted learning models of Figures 14 and 15 can be compared directly with those from the full evolutionary neural network simulations of Figures 5, 6, and 7. Figure 16 shows the mean, median, upper quartile, and lower quartile costs as a function of age for the abstracted processes with $\rho = 0.04$. In this case, the mean evolved protection period is around 14 years, and the mean evolved ELT $100/\delta$ is around 10 years. As in the full neural simulations, the results for the evolved protection period ($E\nu$) are compared with those for representative fixed protection periods (1, 10, 20). Clearly, the linear learning approximation and flat distribution of residual errors are gross approximations of the real neural learning processes, but the broad pattern of results is seen to be the same: Longer protection periods allow slower learning and result in better adult performance, but not allowing procreation while being protected prevents the evolved protection periods from becoming excessively long. The effects of changing the onset of old age, and of allowing procreation while protected, are shown in Figure 17,

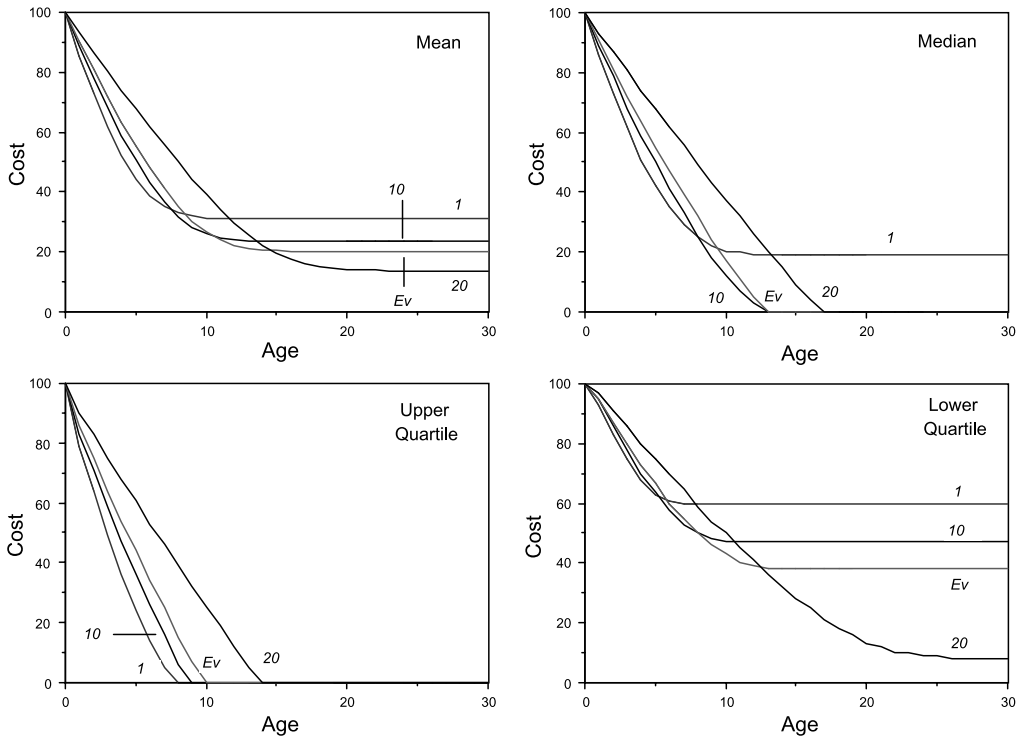


Figure 16. Mean, median, upper quartile, and lower quartile individual performance as a function of age for the abstracted learning processes of Figures 14 and 15 with $\rho = 0.04$. Each plot compares the results for the evolved protection period (Ev) with those for representative fixed protection periods (1, 10, 20). These results are in broad agreement with the full neural simulations of Figures 5, 6, and 7.

and are also found to be in line with those of the full evolving neural network simulations presented in Figures 10 and 11.

9 Conclusions and Discussion

This artificial life study has shown how evolutionary neural network simulations can begin to explore the effects of lifetime learning on life history evolution. It has established that clear learning




Figure 17. Further tests of how well the results on abstracted learning processes match those of the full neural simulations. The effect of life span on protection period (left) for comparison with Figure 10, and evolved periods for the case of procreation while protected (right) for comparison with Figure 11, both show broad agreement.

advantages and better adult performance are possible if children receive longer periods of parental protection, but only if the children are not allowed to reproduce during their period of protection. If procreation is not prevented while being protected, the competition to reproduce leads to learning strategies that result in worse adult performance. When procreation is prevented during protection, a compromise protection period evolves that balances the improved learning performance against the reduced procreation period. It was also found that the evolved protection period increases with life span, rather than remaining at a fixed duration determined by the learning task complexity, illustrating the tradeoff involved and confirming the importance of learning well.

Applying these findings to the understanding of real animal populations is not straightforward. Indirect genetic effects, in which the environmental influence on one individual is affected by a different individual (such as a parent), are known to have important and sometimes nonintuitive consequences for the associated evolutionary processes [37], and there are an enormous number of potentially important, species-dependent factors that have been glossed over in this study. The concept of protection period studied here is certainly a considerably simplified form of the altruistic behaviors known to be common in nature [29]. Moreover, the relative importance of the complexities of social life versus diet and food gathering for the coevolution of brain size, learning, and intelligence remains controversial [21]. Both of these factors are clearly important for primates in general, and humans in particular. However, the simulations presented in this article bypass the distinction, showing that any need for learning complex behavior will result in associated changes to the optimal periods of protection and ages of first reproduction. The physical and biological practicalities associated with the coevolution of the various factors will certainly bias what emerges in biological populations, depending upon the characteristics of the particular species involved. Consequently, mapping the results of the simplified simulations onto particular species is fraught with difficulties. Nevertheless, this article has explored the general principles of lifetime learning effects on life histories, which can be used in the future as the basis for more detailed studies aligned to particular species.

A major problem with most artificial life simulations is the extreme simplifications that are required to render them computationally feasible. If the models are oversimplified, there is a risk that they will become unreliable and misleading. For example, it was seen in this study how important the inclusion of certain details was, such as the differences and improvements that emerge on allowing the evolved neural learning rates to depend on age. To proceed in the face of computational intractability, one can run quite detailed simulations to establish the behavior of particular aspects and how they interact with other aspects, and then attempt to formulate computationally inexpensive approximations of them that may be included in future models that represent the other factors in more detail. Determining how to do that reliably is a major challenge for the field of artificial life in general. If different features of the system under investigation do not interact, then progress in this direction should be straightforward, but usually the assumption of no interaction will be yet another approximation that has to be accommodated.

It was in this spirit that the previous section of this article considered how the neural learning process, which is at the heart of this study, might be approximated to give massive savings in computation time that could then be used to allow other aspects to be modeled in more detail. It was shown how this could be done in such a way that the key results that emerged were qualitatively equivalent to those of the full evolving neural network study. Moreover, there was a simulation time speedup factor on the order of 10,000. It was also seen how the abstract approximated learning process could be parameterized in such a way as to represent learning tasks and processes of differing difficulty, and in that way the simulations could be regarded as extensions of the neural simulations that were tied to a particular simplified task and neural network. Further work would be able to benefit from improvements to the parameterization of the abstract learning processes that might represent known aspects of human/animal learning even better than the neural network models used in this study. Full evolving neural network simulations are still likely to be required to ground the whole process (to relate the abstract performance improvements to the age dependences of real neural learning abilities), so they cannot be avoided completely. For example, it would be virtually impossible to predict the effect of extended protection periods on all the evolved neural network

parameters and the resultant learning performance, without running at least one set of full simulations. But once the key patterns of learning performance have been established and abstracted, the hope is that they can be avoided in future simulations. Checks are required, of course. For example, the correspondence of the abstracted learning results seen in Figure 17 with those in Figures 10 and 11 is evidence that the chosen approximation is at least good enough to explore the effect of varying the life span and allowing procreation while protected.

It seems likely then that future work in this area will, for reasons of computational feasibility, best proceed by pursuing this abstracted learning approach further, with better matching of the learning patterns to real tasks and species. The computational speedups over full neural network simulations will then facilitate explorations over whole ranges of abstract learning parameter values, analogous to the plots against δ and ρ in Figures 14 and 15. As noted at various stages in the above, there certainly remain many related life history issues and refinements that could usefully be incorporated into extensions of the study presented in this article.

The most obvious set of extensions are those that relate directly to the learning process, such as the changes to the learning experience that can result from parental protection—for example, due to guided exploration, exploration without risk, mimicking of parental behavior, parental instruction, and so on. It will probably actually be easier to incorporate these factors into an abstracted learning process than it would be to adjust a full neural learning process, so this should be a particularly promising avenue to explore further. Then there are more realistic accounts of the costs of reproduction and protection, how the protection abilities depend on the parental fitness and number of children, and how these affect the patterns of competition and deaths. Similarly, the effect of numerous other biological and environmental details that influence the life span, age of maturity, fecundity, and length of the reproductive and growing periods could be incorporated into the models. Individual age-related factors, such as biological deterioration affecting the ability to compete or reproduce at older ages, may also have important consequences under certain circumstances. There is also the need to pay more attention to the evolutionary pressures and consequences that arise from the introduction of competition with, and coevolution with, other species. For example, the performance differences seen in Figure 12 are likely to have different consequences if the competition is between species, rather than purely within one species. Parameterized representations of such additional factors will allow further explorations of the associated tradeoffs, better understanding of how these factors interact with learning, and how different patterns arise for different species. Various aspects of these extensions will hopefully be presented elsewhere in the near future.

Acknowledgments

Preliminary versions of many of the key ideas in this article were first presented at the 2007 Neural Computation and Psychology Workshop (NCPW10) [12] and the 2007 Genetic and Evolutionary Computation Conference (GECCO) [11].

References

1. Allman, J., McLaughlin, T., & Hakeem, A. (1993). Brain weight and life-span in primate species. *Proceedings of the National Academy of Sciences of the USA*, *90*, 118–122.
2. Bailey, J. D., Bruer, B. T., Symons, F., & Lichtman, J. W. (Eds.) (2000). *Critical thinking about critical periods*. Baltimore, MD: Brookes.
3. Baldwin, J. M. (1896). A new factor in evolution. *The American Naturalist*, *30*, 441–451.
4. Bishop, C. M. (1995). *Neural networks for pattern recognition*. Oxford, UK: Oxford University Press.
5. Bogin, B. (1997). Evolutionary hypotheses for human childhood. *Yearbook of Physical Anthropology*, *40*, 63–89.
6. Bullinaria, J. A. (2001). Simulating the evolution of modular neural systems. In *Proceedings of the Twenty-Third Annual Conference of the Cognitive Science Society* (pp. 146–151). Mahwah, NJ: Lawrence Erlbaum Associates.

7. Bullinaria, J. A. (2003). From biological models to the evolution of robot control systems. *Philosophical Transactions of the Royal Society A*, *361*, 2145–2164.
8. Bullinaria, J. A. (2003). Evolving efficient learning algorithms for binary mappings. *Neural Networks*, *16*, 793–800.
9. 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)* (pp. 79–84). Piscataway, NJ: IEEE.
10. Bullinaria, J. A. (2007). Using evolution to improve neural network learning: Pitfalls and solutions. *Neural Computing & Applications*, *16*, 209–226.
11. Bullinaria, J. A. (2007). The effect of learning on life history evolution. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2007)* (pp. 222–229). New York: ACM.
12. Bullinaria, J. A. (2008). Another reason why we should look after our children. In R. M. French & E. Thomas (Eds.), *From associations to rules: Connectionist models of behavior and cognition* (pp. 41–52). Singapore: World Scientific.
13. Charnov, E. L., & Berrigan, D. (1990). Dimensionless numbers and life history evolution: Age at maturity versus the adult lifespan. *Evolutionary Ecology*, *4*, 273–275.
14. Clutton-Brock, T. H., Albon, S. D., & Guinness F. E. (1989). Fitness costs of gestation and lactation in wild mammals. *Nature*, *337*, 260–262.
15. 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.
16. Cole, L. C. (1954). The population consequences of life history phenomena. *Quarterly Review of Biology*, *29*, 103–137.
17. Gadgil, M., & Bossert, W. (1970). Life history consequences of natural selection. *American Naturalist*, *104*, 1–24.
18. Harvey, P. H., & Zammuto, R. M. (1985). Patterns of mortality and age at first reproduction in natural populations of mammals. *Nature*, *315*, 319–320.
19. Hawkes, K., O'Connell, J. F., Blurton Jones, N. G., Alvarez, H., & Charnov, E. L. (1998). Grandmothering, menopause, and the evolution of human life histories. *Proceedings of the National Academy of Sciences of the USA*, *95*, 1336–1339.
20. Julesz, B., & Kovacs, I. (Eds.) (1995). *Maturation windows and adult cortical plasticity*. Reading, MA: Addison-Wesley.
21. 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.
22. Kennedy, G. E. (2005). From the ape's dilemma to the weanling's dilemma: Early weaning in its evolutionary context. *Journal of Human Evolution*, *48*, 123–145.
23. Lamb, M. E., Bornstein, M. H., & Teti, D. M. (2002). *Development in infancy: An introduction*. Mahwah, NJ: Lawrence Erlbaum Associates.
24. Leigh, S. R. (2004). Brain growth, life history and cognition in primate and human evolution. *American Journal of Primatology*, *62*, 139–164.
25. Leutenegger, W. (1979). Evolution of litter size in primates. *American Naturalist*, *114*, 525–531.
26. Marrow, P., Stevenson, I. R., McNamara, J. M., Houston, A. I., & Clutton-Brock, T. H. (1996). State-dependent life history evolution in Soay sheep: Dynamic modelling of reproductive scheduling. *Philosophical Transactions of the Royal Society B*, *351*, 17–32.
27. Roff, D. A. (2002). *Life history evolution*. Sunderland, MA: Sinauer Associates.
28. Sloman, A., & Chappell, J. (2005). The altricial-precocial spectrum for robots. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 1187–1193.
29. Sober, E., & Wilson, D. S. (1999). *Unto others: The evolution and psychology of unselfish behavior*. Boston, MA: Harvard University Press.
30. Stearns, S. C. (1989). Trade-offs in life history evolution. *Functional Ecology*, *3*, 259–268.
31. Stearns, S. C. (1992). *The evolution of life histories*. Oxford, UK: Oxford University Press.

32. Stearns, S. C., & Crandall, R. E. (1981). Quantitative predictions of delayed maturity. *Evolution*, *35*, 455–463.
33. 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.
34. Striedter, G. F. (2005). *Principles of brain evolution*. Sunderland, MA: Sinauer Associates.
35. 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.
36. Weber, B. H., & Depew, D. J. (2003). *Evolution and learning: The Baldwin effect reconsidered*. Cambridge, MA: MIT Press.
37. Wolf, J. B., Brodie, E. D., Cheverud, J. M., Moore, A. J., & Wade, M. J. (1998). Evolutionary consequences of indirect genetic effects. *Trends in Ecology and Evolution*, *13*, 64–69.

