

# Mememes in Artificial Life Simulations of 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

The effect that learning has on Life History Evolution has recently been studied using a series of Artificial Life simulations in which populations of competing individuals evolve to learn to perform well on simple abstract tasks. Those simulations assumed that learning was achieved by identifying patterns in sets of training data, i.e. through direct experience. In practice, learning is not only by direct experience, but also by imitation of others. Such imitative information transfer is now often formulated in terms of memes being passed between individuals, and it is clear that this is a substantial part of real learning processes. This paper extends the previous study by incorporating imitation and memes to provide a more complete account of learning as a factor in Life History Evolution.

## Introduction

Computational models based on neural networks that learn from a stream of experience (i.e. representative input-output samples) have provided good accounts of numerous aspects of human behaviour. Extending those models to Artificial Life simulations of evolving populations of competing neural network based individuals can then lead to improved understanding of more general aspects of human development and “life history”, such as the periods of protection that parents offer their young and ages at first reproduction (Bullinaria, 2009). Those simulations elucidated the trade-off between learning quickly and learning well, and showed how evolution can balance the trade-off to result in the emergence of extended periods of parental protection during which learning could be completed slowly and effectively without the impact of fitness based natural selection pressures.

The Bullinaria (2009) Life History Evolution study began by using a simple artificial neural network based system that allowed each individual to learn from a set of training patterns, and then moved on to study non-neural network abstractions of that kind of learning process, that were more computationally efficient for large scale evolutionary simulations. What all those simulations assumed was that the learning was achieved by identifying patterns in relevant training data, i.e. through direct experience. In practice, learning is not purely by direct experience, but also by imitation of learned performance of others. Such information transfer can be formulated in terms of *memes* being passed between individuals (e.g., Brodie, 1996; Blackmore, 1999), and it is clear that this, in its most general form, is a large part of the human learning process, and maybe also of other animal species. It is therefore important to incorporate

imitation and memes into any complete account of learning as a factor in Life History Evolution. As always, there will be trade-offs between the various costs involved (Stearns, 1989, 1992). In many ways, the relevant trade-offs are clear from a theoretical point of view, but the interactions are complex and highly dependent on the associated parameters. It is only by running comprehensive series of simulations that the effect of the various parameter values becomes apparent.

Already Higgs (2000) has simulated the evolution of learning by imitation, but that study didn’t consider how that learning might interact with more traditional neural learning by direct experience, and it is not immediately obvious how best to bring those different forms of learning together. One of the key results of Bullinaria (2009) was that it is possible to abstract out almost all the details of the neural learning, and still be left with a system that resulted in the evolution of the same life history properties. Although it was not the intention at the time, that abstraction process also provides a relatively straightforward way of incorporating imitative learning into the same system. Therefore, the aim of this paper is to introduce a parameterized account of memes and imitation into the approach of Bullinaria (2009), and begin to explore the effect that imitation has on the various life history and human development factors.

In the remainder of this paper, the underlying Artificial Life framework is first described, and then the details are provided about how the direct learning and imitation processes can be modelled efficiently. This is followed by a presentation of the results from a representative series of simulations designed to test and explore many of the key relevant issues. The paper ends with some discussion and conclusions.

## The Artificial Life Framework

The simulation approach involves evolving populations of individuals, each specified by a set of innate parameters, that must learn to perform well on some abstract task. The fitness of each individual at each stage will simply be how well it has so far learned the given task. Forcing the individuals to compete to survive and procreate, according to their relative fitness, results in the emergence of populations of increasing ability. Moreover, 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. This leads to the evolution of





Figure 1: The mean evolved performance levels and protection periods as a function of the ELT  $100/\delta$  when the linear individual performance improvement with age stops with probability  $\rho\delta$  at a random cost in the range  $[0,100]$ , for  $\rho \in \{0, 0.02, 0.05, 0.1\}$ .

riskier learning strategies than over-simplified “generational” approaches that involve weaker selection pressures and do not match real environments so well (Bullinaria, 2007a).

In all the simulations, a fixed population size is maintained (that is consistent with fixed total food resources available to support the population) by replacing the individuals that have died by children of the most fit individuals. Deaths occur by losing a fitness comparison “fight” against other individuals, or randomly due to old age beyond a natural life-span (set here to be around twice the time typically taken to learn the simulated task, namely 30 simulated years). The children are generated by cross-over and mutation from two parents chosen each simulated year by pair-wise fitness comparisons of the eligible individuals. This is 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 gives it a significant chance of falling outside that range. Although these details are clearly over-simplifications of real animal populations, they constitute a manageable approximation of all the key processes, and have proved effective in numerous previous studies (e.g., Bullinaria, 2007a,b, 2009).

The Bullinaria (2009) study began with a learning process based on standard fully connected Multi-Layer Perceptron neural networks with one hidden layer, sigmoidal processing units, and training by gradient descent using the cross-entropy error function on simple classification/categorization tasks. The main life history factor explored in that study was the protection of children by their parents until they had reached a certain age, so they could not be killed by competitors before then. That added an implicit cost to the parents in that the more they protected their children, the more likely they were to die themselves through competition. Simulations that evolved the protection period, as well as all the neural learning parameters, established that clear learning advantages and better adult performances were possible if children received longer periods of parental protection, but only if the children were not allowed to reproduce during their period of protection. If procreation was not prevented in that way, the competition to reproduce led to learning strategies that result in worse adult performance. When procreation is prevented

while protected, a compromise protection period evolves that balances the improved learning performance against the reduced period for procreation. It was also shown that the evolved protection period increases with life-span, rather than remaining at a fixed duration determined by the learning task complexity, illustrating the trade-off involved and confirming the importance of learning well.

### Abstracting the Neural Learning Process

An important result of Bullinaria (2009) was that it is possible to approximate the full neural network learning process by a single performance level that varies as a simple parameterized function of age, and still end up with qualitatively the same Life History Evolution results. The simplest stochastic approximation would be to have each individual’s learning performance (i.e. fitness) rise approximately linearly with age from 0 up to 100% in steps drawn randomly each year from the range  $[0, 2\delta]$ . Simulations using different learning rates  $\delta$  then show that the population mean performance falls almost linearly with the Expected Learning Time (ELT), i.e.  $100/\delta$ , and the evolved protection period rises approximately linearly with  $100/\delta$ , but peaks near the point at which individuals start dying of old age. Predictably, the best mean performance is achieved with very high learning rates  $\delta$ , 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  $\delta$  is evolved along with the protection period, it quickly achieves very high levels, and the protection period goes to zero. Of course, with real neural networks one cannot just 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. In that case, the evolutionary process will find the best values for the learning parameters, and having slower learning with longer protection periods does consistently emerge to provide a clear advantage.

A better approximation to the full neural learning process, that has faster learning leading to riskier learning strategies


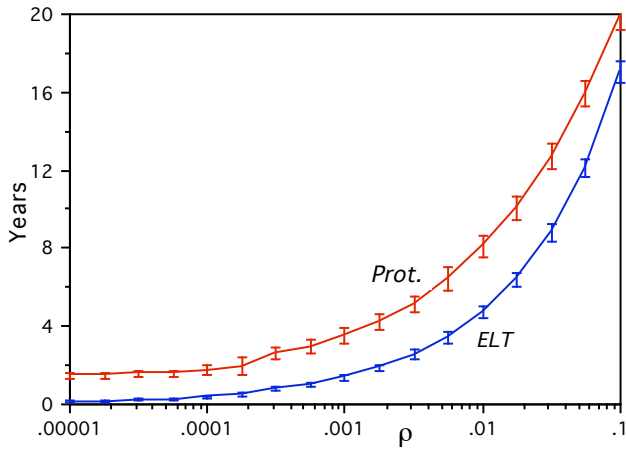


Figure 2: The mean evolved ELT  $100/\delta$  and protection period as a function of learning task difficulty parameter  $\rho$  (left), and the median learning performance as a function of age for the evolved and three other protection periods with  $\rho = 0.04$  (right).

which increasingly lead to persistent poor performance, is achieved by simply having the learning process stop at some random point in the performance range  $[0, 100]$  with a probability  $\rho\delta$  that increases linearly with both the learning rate  $\delta$  and an associated “task difficulty” parameter  $\rho$ . The left graph of Figure 1 shows how the mean performance then depends on the ELT  $100/\delta$  for four representative values of  $\rho$ . The higher  $\rho$  is, the lower the value of  $\delta$  at which significant deviations from the earlier  $\rho = 0$  case arise. The right graph shows that the relation between the evolved protection period and  $100/\delta$  is not much affected by the size of  $\rho$ .

The performance plot shows a clear maximum for each value of  $\rho$ , and successful evolutionary processes will result in the emergence of the corresponding optimal learning rates  $\delta$  with their associated non-zero protection periods. The left graph of Figure 2 shows the mean Expected Learning Times  $100/\delta$  and protection periods that actually emerge through evolution as a function of the parameter  $\rho$ . As  $\rho$  increases, the best possible learning time  $100/\delta$  also increases, and the best protection period follows suit. The evolved protection period is always slightly longer than the ELT  $100/\delta$ . This is because of the stochastic nature of the learning process and the fact that the mutations lead to distributions of learning rates and protection periods, and the obvious advantage of protection periods being long enough to accommodate a reasonable number of individuals that are slower than average.

The parameter  $\rho$  is seen to act as an abstract measure of learning difficulty, and can be regarded as an approximate representation of the difficulty the neural network learning algorithm has with its given task. Although this is a rough approximation to reality, it does have the required properties. Relatively easy tasks correspond to low  $\rho$ , are learned quickly, and have short associated protection periods. Harder, or more complex, tasks correspond to higher values of  $\rho$ , take longer to learn, and benefit from longer protection periods. The individual performance levels that emerge in the abstracted learning models were compared directly by Bullinaria (2009) with those arising from the full evolutionary neural network simulations, and a good qualitative correspondence was found for  $\rho = 0.04$ . The right graph of Figure 2 shows the median performance levels as a function of age for this case. The

mean evolved ELT  $100/\delta$  is around 10 years and the mean evolved protection period is around 14 years. As for the full neural simulations, the results arising with evolved protection period (*Ev*) were compared with three fixed protection periods (*1*, *10*, *20*). The linear learning approximation and uniform distribution of residual errors are rough approximations of the real neural learning processes, but the broad pattern of results is found 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 age at onset of “old age”, and of allowing procreation while protected, are also found to be in line with those of the full evolving neural networks.

There certainly remains much scope for more accurate parameterizations for specific real learning processes, as discussed by Bullinaria (2009), but the current set-up will suffice for the preliminary investigation of memes here.

## Incorporating Imitative Learning

The main aim of the abstracted neural learning process was to improve the computational efficiency, and hence allow more detailed Life History factors to be simulated, but it also renders it feasible and fairly straightforward to incorporate learning by imitation into the same performance function.

The basic idea is that it will often be more efficient to imitate the successful behaviour of another individual than it is to learn it from direct experience. One can think of the transmission of behavioral practices or cultural ideas between individuals, and those *memes* will replicate and respond to natural selection pressures in a manner analogous to genes (Dawkins, 1976; Brodie, 1996; Blackmore, 1999). It seems likely that humans have evolved to learn by imitation as well as direct experience across a wide variety of tasks (e.g., Richerson and Boyd, 1992; Offerman and Sonnemans, 1998), though other species appear to imitate to a much lesser extent (e.g., Byrne and Russon, 1998; Blackmore, 1999; Zentall, 2001). There has been considerable recent interest in this idea across a range of disciplines (e.g., Hurley and Chater, 2005;

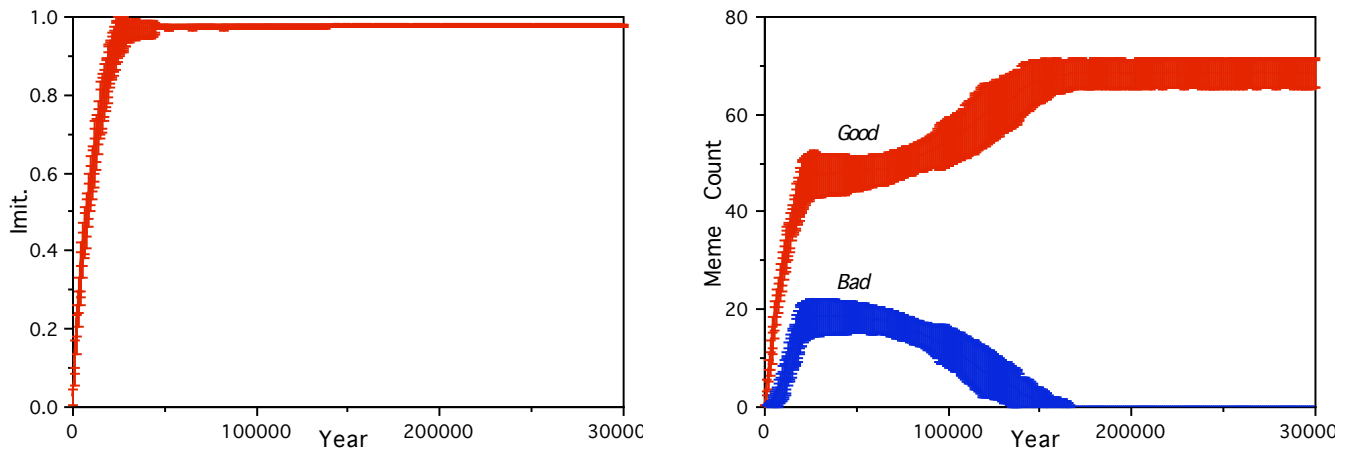


Figure 3: The evolution of imitability (left), and the change in average numbers of good and bad memes known by individuals throughout evolution (right), for 16 runs of the basic imitation-only simulation with limited brain sizes.

Nehaniv and Dautenhahn, 2009). The thinking here is that Artificial Life simulations will be best placed to explore this issue in the context of other Life History traits.

Some interesting preliminary work has already been carried out. Belew (1990) and Best (1999) have introduced imitation based cultural factors into the Hinton and Nowlan (1987) model of learning guiding evolution, but that work is far removed from the neural inspired learning relevant to the life history factors of relevance here. Borenstein and Ruppin (2003) address many of the limitations of those earlier studies, and do incorporate neural learning mechanisms, but they actually prevent cultural evolution by not allowing meme transmission between generations and only allowing innate behaviours to be imitated.

The study of Higgs (2000) comes closest to exploring the life history issues of interest here. That paper considered the evolution of populations of individuals that may invent and imitate memes, and investigated a range of factors that affect how the imitation rates, fitness levels, and number of memes evolve. The key finding was that imitative ability does consistently emerge under a range of conditions, even when some memes have a negative effect on fitness, and/or there is an inherent cost in the ability to imitate. In many ways it is obvious that if there exist memes with a range of positive and negative effects on fitness, then not imitating will leave the fitness at some baseline, whilst imitation will result in a range of fitness levels above and below that baseline. Selection on the basis of fitness will then favour those individuals that have imitated the good memes, and hence favour imitative ability. Moreover, since it favours individuals that have acquired and can pass on those good memes, the good memes will tend to propagate at the expense of the bad memes. Memes acting together (i.e. *memeplexes*), the interplay of genetic and cultural fitness, and the interaction of genetic and mimetic replicators, all complicate this simple picture (e.g., Brodie, 1996; Blackmore, 1999; Best 1999), but these are all things that can be incorporated into future simulations.

The main question this paper aims to address is: how can the Life History Evolution approach of Bullinaria (2009) be extended in a way that enables these issues to be studied in conjunction with direct lifetime learning processes?

### Simulating Memes and Imitation

For the extraction of reliable conclusions from Artificial Life simulations it is important to avoid confounding factors, so to explore general ideas it is usually wise to keep the models much simpler than when the aim is to model particular real life scenarios. Moreover, it is important to parameterize the models (e.g., like introducing the parameter  $\rho$  above) so that they remain relevant to a range of species, tasks, etc. and allow comparisons between them. The aim here is to develop such a parameterized framework that is general enough to cover learning from others in the most general sense, that includes (but is not limited to) simple imitation.

Unfortunately, the details of the Higgs (2000) study do not match with the current aims. In particular, it did not consider the details of any of the processes taking place during the individuals' lifetimes, and it used non-overlapping generations which means a total absence of the competition between individuals of different ages that underlies so many of the issues of interest here. Other factors simply complicate the analysis unnecessarily, such as using Gaussian distributions for the meme fitnesses and mutations, the non-linear relation between learning ability and probability of imitation, and the unbounded number of memes that can be invented. So, instead of following the approach of Higgs (2000), the approach of Bullinaria (2009) will be extended in a minimal computationally efficient manner to include the key concepts of memes and their imitation.

The starting point is to assume that there exist a set of  $M$  memes  $\{m_j : j = 1, \dots, M\}$  and that each individual  $i$  at each stage of its life will have acquired some subset of them to be stored in their brain of size  $B_i$ . There is no need to specify exactly what the memes represent, nor worry about the details of the imitation process. It will also be assumed that all the memes are of equal complexity and imitability, though they may contribute unequally to fitness of the individuals that possess them. To begin with, the individuals' baseline fitness will be 0, and half the memes will be deemed good memes that increase this by 1, and the other half will be bad memes that decrease it by 1. So each individual  $i$  can potentially increase its fitness during its lifetime from 0 up to  $B_i$ .

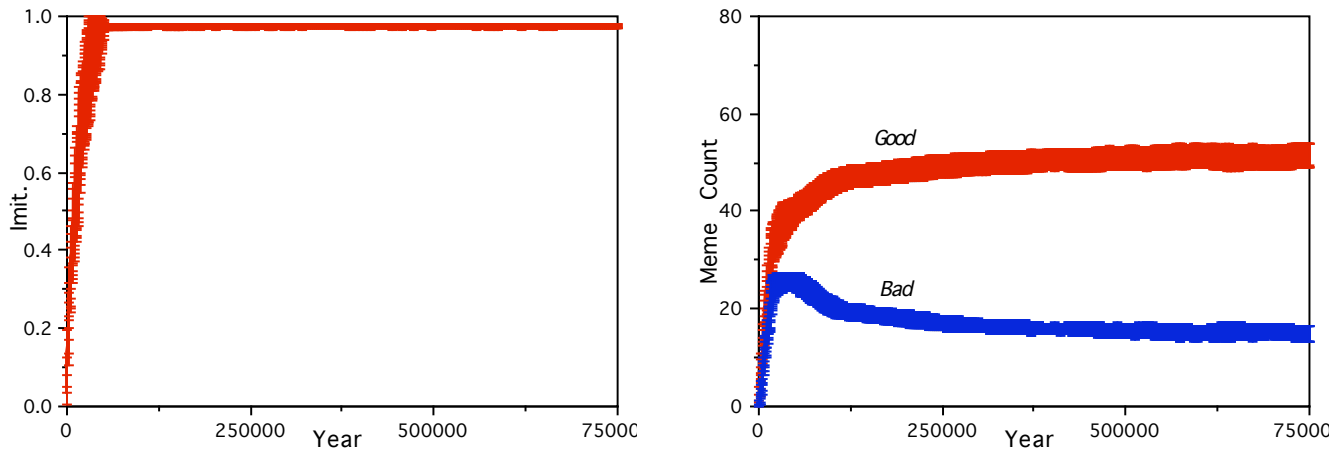


Figure 4: The evolution of imitability (left), and average numbers of good and bad memes known by individuals (right), for 16 runs of imitation-only simulations with limited brain sizes and cultural fitness based imitation selection.

The imitative ability  $\alpha_i$  of all individuals  $i$  in the initial population will be zero, but the mutations and crossovers as described above enable it to evolve from zero up to a maximum of 1 if that proves beneficial. Then during each simulated year, each individual can acquire up to  $\alpha_i \phi B_i$  memes from other individuals, where  $\phi$  is a parameter that specifies the maximum rate at which memes can be copied. To inject memes into the populations with minimal disruption to the imitative process, each year one randomly chosen individual acquires one randomly chosen meme with probability  $r$  if its brain is not already full. Figure 3 shows what happens if  $M=400$ ,  $B_i = 100$ ,  $\phi = 0.1$  and  $r = 0.01$ , with just the imitabilities  $\alpha_i$  allowed to evolve. The tournament based selection of parents, deaths and copied individuals give the good memes an advantage over bad memes, so the number of bad memes rises more slowly than the good memes, and when the number of known memes reaches the level that brains regularly reach full capacity ( $\sim 20,000$  years), the number of bad memes begins to fall and eventually becomes negligible ( $\sim 150,000$  years). There is a clear advantage to acquiring memes throughout, and so the imitability quickly rises to near 1. The behaviour during the lifetime of a typical evolved individual is a simple linear acquisition of memes over the first  $1/\phi = 10$  years, at which point the brain reaches full capacity and maximum performance is achieved. Children are then produced until death due to old age. Most deaths due to competition occur during the meme acquisition period.

There are interesting dependencies on who exactly is imitated to acquire memes. If memes are copied from random individuals, there is still enough selection pressure to eradicate the bad memes, but it takes about twice as long ( $\sim 300,000$  years). If each individual first acquires memes from their own parents, before imitating random others, the number of bad memes disappears more quickly ( $\sim 130,000$  years). If parents are imitated before fitness selected others, the bad memes go even more quickly ( $\sim 120,000$  years). Since parents have already gone through fitness selection to become parents, and are also older and more experienced, they are a better source of memes than other fitness selected individuals. In fact, if individuals *only* copy from their parents, significant numbers of bad memes never build up at any stages of evolution.

Another factor that affects the results is the basing of who to imitate on cultural fitness (Higgs, 2000). In this case, each meme has a cultural fitness that is not correlated with its standard (biological) fitness, and individuals are chosen for imitation according to the total cultural fitness they have acquired. As Figure 4 shows, this allows memes of high cultural fitness to persist in the population, even if they are actually bad memes. This is independent of what contributes to the cultural fitness of those bad memes. Obviously, there are numerous related factors, such as cognitive dissonance (Cooper, 2007) and memes associating into memplexes (Blackmore, 1999), that will increase or decrease this effect to varying degrees, and these are more issues that may be worth attempting to incorporate into future simulations.

The effect of copying fidelity also needs consideration. This can easily be approximated by having a fraction  $1-f$  of good memes incorrectly copied and thereby transformed into bad memes. As the fidelity  $f$  is reduced from 1, the pattern changes from that like Figure 3 but with increasing times needed to eradicate the bad memes, to something like Figure 4 with persistent levels of bad memes.

Finally, it is important to understand how the results depend on the relation between the total number of memes and the brain capacity. For  $M = 200$ ,  $B_i = 200$  and everything else the same, the simulation results of Figure 3 take on the rather different pattern seen in Figure 5. Now all individuals can acquire all memes, and it proves much more difficult to separate the good from the bad so that selection pressures can act. In this case evolution ends up with only slightly more good memes than bad, and there is little pressure towards high levels of imitability. Interestingly though, the strategy of only imitating ones own parents does manage to prevent the build-up of bad memes in this case too.

A central recurring feature of the Higgs (2000) study was a “mimetic transition” at which there is a dramatic rise in imitative ability and number of memes, and it was shown how numerous factors affected the timing of that transition. In the current framework, that transition virtually always happens right at the start of the evolutionary process.

There is certainly much more to memes and imitation than has been introduced here (e.g., Brodie, 1996; Blackmore,




Figure 5: The evolution of imitability (left), and average numbers of good and bad memes known by individuals (right), for 16 runs of the basic imitation-only simulation with brains large enough to accommodate all known memes.

1999), but the framework as described above already includes all the key ideas necessary to make progress.

### Simulating Direct Learning

Having formulated the key mimetic factors, the direct lifetime learning factors of Bullinaria (2009) can now be reinstated. The natural way to do this in terms of memes is to have  $\delta_i \psi B_i$  random memes learned each year, where  $\delta_i$  is an evolvable learning rate, and  $\psi$  is an intrinsic measure of learning difficulty. The time to learn to brain capacity is then  $1/\delta_i \psi$ , and for  $\psi = 0.01$  the expected learning time matches that of the Bullinaria (2009) simulations. The learning difficulty parameter  $\rho$  that prevents the evolution of unrealistically high learning rates can be implemented easily here by learning a bad meme rather than a good meme with probability  $\rho \delta$ . Then the evolved learning rates balance the trade-off between learning quickly and having too many fitness reducing bad memes, with results equivalent to the full neural network simulations of Bullinaria (2009).

### Life History Simulation Results

The simulations become even more interesting when the imitation and direct learning occur together and interact with life history traits such as protection periods. But, before doing that, there are a few more important details that need to be added to render the simulations reasonably realistic.

First, it is possible for an individual to acquire both good and bad “versions” of the same meme via different routes. The resolution of meme inconsistencies in reality is known to be a complex issue (Cooper, 2007), but a convenient approach to start with here is to have the good and bad memes come in pairs that simply cancel each other out if they occur together. In this way, a bad meme arising from direct learning can be removed if the corresponding good meme is copied from another individual. Similarly, a bad meme arising from poor copying fidelity can be removed by later acquiring the corresponding good meme by direct learning or by copying from a different individual.

Second, in reality, the rate of meme acquisition is unlikely to be as constant as in the processes described above. Instead, more realistic results are produced by a stochastic version, where each usage of the parameters  $\alpha_i$  and  $\delta_i$  are replaced by random numbers from the respective ranges  $[0, 2\alpha_i]$  and  $[0, 2\delta_i]$ , like in the Bullinaria (2009) study.

Figure 6 shows the evolution of the key parameters and resultant meme counts when  $M = 400$ ,  $B_i = 100$ ,  $\phi = 0.1$ ,  $\psi = 0.01$ ,  $r = 0.01$ ,  $f = 0.9$  and  $\rho = 0.001$ . In this case, both copying and direct learning contribute to the learning process, and bad memes are kept to very low levels. The protection period settles to slightly above the typical learning time as in the full neural simulations of Figure 2.

The implementational details obviously affect exactly what emerges from the simulations, and it is those differences that reflect the wide range of life history patterns for the different species that have emerged from biological evolution. Varying the details and parameters allows a systematic exploration of the trade-offs and interactions that lead to specific traits. A few simple examples will now illustrate the kind of factors that can be investigated within this framework.

The issue of whether to allow procreation while protected produced interesting results in the direct learning study of Bullinaria (2009). In that case, if procreation was allowed while protected, the protection periods rose so that there were only deaths due to old age and no deaths by competition, and the selection pressure to learn fast to procreate early resulted in higher learning rates that led to poorer adult performance. This no longer happens in the current meme based framework. Since the errors arising from faster learning can now be corrected by copying (or being taught), such fast learning will emerge without a deterioration of the final adult performance. Increased protection periods again remove the worry of early death due to competition, so if some unlucky individuals are slow in correcting their direct learning errors, that is compensated overall by the faster early learning in others. The balance between the two forms of learning, parameterized here by  $\phi$ ,  $\psi$ ,  $f$  and  $\rho$ , will determine exactly what emerges, and the way forward would be to attempt to understand species specific differences in terms of variations in such parameter values.




Figure 6: Evolution of the full imitation and direct learning system with copying fidelity  $f = 0.9$  and  $\rho = 0.001$ : the average imitability  $\alpha$  (top left), learning rate  $\delta$  (top right), protection period (bottom left) and resultant meme counts (bottom right).

The copying fidelity, parameterized by  $f$ , has a particularly large effect on what emerges. If it is raised from the 0.9 of Figure 6 up to 1.0, so that all the copying is exact, evolution results in perfect performance being achieved more quickly and more reliably. One might predict that the evolved direct learning rates  $\delta$  will then decrease to enable more reliable memes for copying, but they actually increase from 12 to 19, because copying can now more effectively correct any direct learning errors. Overall, the evolved protection period can be reduced from 10.0 to 7.6 years to enable a longer procreation period. The trade-offs are such that fidelity differences affect what emerges in different ways depending on the values of the other parameters. This again illustrates the need for a flexible modeling framework to explore such interactions.

If the copying fidelity is very low, a high imitative ability  $\alpha$  never evolves because it introduces too many bad memes into the population, and one ends up with direct learning only, as appears to be the case for most animal species apart from humans. Also, if mechanisms are not available to remove bad memes, interesting changes in imitative ability can arise throughout evolution. For example, Figure 7 shows one such case in which the number of bad memes repeatedly rises to such high levels that the best strategy is to stop copying until all the carriers have died, and then start again.

The brain size is another crucial factor that can be evolved, and in the simulations described above it invariably grows to the maximum allowed. Obviously, for real animals there are significant costs associated with having larger brains, and trading those costs against the improved performance that results from a bigger brain leads to particular brain sizes emerging (e.g., Blackmore, 1999; Striedter, 2005). It actually proves easy to add such costs into the simulations to limit the brain sizes that emerge, but the cost implementations are not yet sophisticated enough that the models can provide reliable testable predictions about particular species.

## Discussion and Conclusions

This paper has made the first steps in introducing imitative learning and memes into Artificial Life simulations of Life History Evolution. The main contribution has been to present a flexible framework which allows a computationally efficient way of parameterizing and exploring any hypotheses in this field. There are certainly numerous simplifications and approximations involved, which have been highlighted throughout, but the basic structures and ideas are in place, and they have already been shown to replicate the key results of earlier approaches and improve upon them.


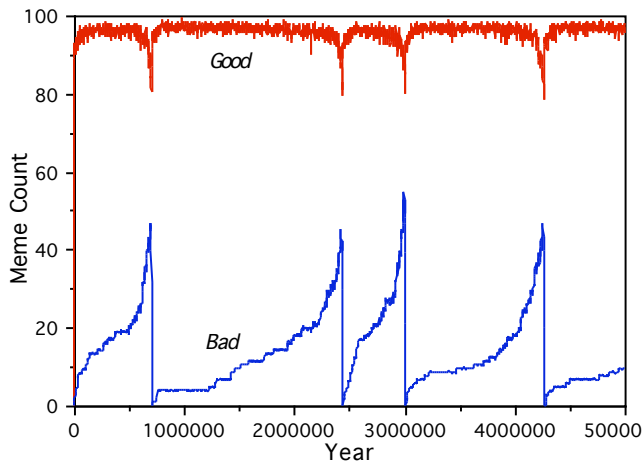


Figure 7: When bad memes (left) are allowed to build up, the evolvable imitative ability (right) can fall quickly to very low values so that the bad memes die out, and then return to the earlier high level until the problem arises again.

Even this simplified framework can be used to investigate an enormous number of interactions and trade-offs. This paper has only presented results from a small selection of simulations to illustrate the kinds of issues that can be explored. Experiments studying further issues will be reported in a longer paper elsewhere. The simulation results so far are in line with existing intuitions, which instills confidence that they can now be taken further with some reliability to explore issues for which our intuitions are not so clear and controversy remains.

There are numerous aspects of the current set-up that could be improved further without too much effort. One would be the refinement of the parameterization of direct learning, and the relation of that to different types of animal learning. Some preliminary attempts involving more parameters and different distributions of good and bad memes have shown that they do indeed re-balance the trade-offs slightly, but no fundamentally different behaviours have yet emerged. Specific details of the mechanisms for removing bad memes tend to have a more dramatic effect on the results, as Figure 7 shows. Building in associations between good and bad memes and simulating the creation of memplexes (Blackmore, 1999), and introducing related mechanisms for the resolution of cognitive dissonance (Shultz and Lepper, 1996; Cooper, 2007), are obvious avenues for future enhancement of the framework in that direction, but it is not clear what fundamentally new results might emerge from that. More challenging future work will involve the incorporation into the existing framework of more realistic additional indirect performance costs related to biological factors (such as the cost of running a larger brain, or of providing parental protection, or of allowing copying, or of teaching), and better distinction between types of learned behaviour and related factors such as ease of copying.

## References

Belew, R.K. (1990). Evolution, learning, and culture: Computational metaphors for adaptive algorithms. *Complex Systems*, 4:11-49.  
 Best, M.L. (1999). How culture can guide evolution: An inquiry into gene/meme enhancement and opposition. *Adaptive Behavior*, 7: 289-306.

Blackmore, S. (1999). *The Meme Machine*. Oxford, UK: OUP.  
 Borenstein, E. and Ruppin, E. (2003). Enhancing autonomous agents evolution with learning by imitation. *AISB Journal*, 1:335-347.  
 Brodie, R. (1996). *Virus of the Mind: The New Science of the Meme*. Seattle, WA: Integral Press.  
 Bullinaria, J.A. (2007a). Using evolution to improve neural network learning: Pitfalls and solutions. *Neural Computing and Applications*, 16:209-226.  
 Bullinaria, J.A. (2007b). Understanding the emergence of modularity in neural systems. *Cognitive Science*, 31:673-695.  
 Bullinaria, J.A. (2009). Lifetime learning as a factor in Life History Evolution. *Artificial Life*, 15:389-409.  
 Byrne, R.W. and Russon, A.E. (1998). Learning by imitation: A hierarchical approach. *Behavioral and Brain Sciences*, 21:667-684.  
 Cooper, J. (2007). *Cognitive Dissonance: 50 Years of a Classic Theory*. London, UK: Sage.  
 Dawkins, R. (1976). *The Selfish Gene*. Oxford, UK: OUP.  
 Higgs, P.G. (2000). The mimetic transition: A simulation study of the evolution of learning by imitation. *Proceedings of the Royal Society B: Biological Sciences*, 267:1355-1361.  
 Hinton, G.E. and Nowlan, S.J. (1987). How learning can guide evolution. *Complex Systems*, 1:495-502.  
 Hurley, S. and Chater, N., editors (2005). *Perspectives on Imitation: From Neuroscience to Social Science*. Cambridge, MA: MIT Press.  
 Nehaniv, C.L. and Dautenhahn, K., editors (2009). *Imitation and Social Learning in Robots, Humans and Animals*. Cambridge, UK: CUP.  
 Offerman, T. and Sonnemans, J. (1998). Learning by experience and learning by imitating successful others. *Journal of Economic Behavior and Organization*, 34:559-575.  
 Richerson, P.J. and Boyd, R. (1992). Cultural inheritance and evolutionary ecology. In: E.A. Smith and B. Winterhalder, editors, *Evolutionary Ecology and Human Behaviour*, 61-91. Chicago, IL: Aldine de Gruyter.  
 Shultz, T.R. and Lepper, M.R. (1996). Cognitive dissonance reduction as constraint satisfaction. *Psychological Review*, 103: 219-240.  
 Stearns, S.C. (1989). Trade-offs in life history evolution. *Functional Ecology*, 3:259-268.  
 Stearns, S.C. (1992). *The Evolution of Life Histories*. Oxford, UK: OUP.  
 Striedter, G.F. (2005). *Principles of Brain Evolution*. Sunderland, MA: Sinauer.  
 Zentall, T. (2001). Imitation in animals: Evidence, function, and mechanisms. *Cybernetics and Systems*, 32:53-96.