
Maturation and the evolution of imitative learning in artificial organisms

Federico Cecconi
Institute of Psychology
National Research Council
Rome 00137, Italy
federico@caio.irmkant.rm.cnr.it

Filippo Menczer
Computer Science & Engr. Dept.
U. C. San Diego
La Jolla, CA 92093-0114, USA
fil@cs.ucsd.edu

Richard K. Belew
Computer Science & Engr. Dept.
U. C. San Diego
La Jolla, CA 92093-0114, USA
rik@cs.ucsd.edu

Author with whom the editors or the publisher should correspond:

Federico Cecconi
Email: federico@caio.irmkant.rm.cnr.it
Phone: +39-6-86090-229
Fax: +39-6-824737

This is an extended version of the paper available as:
Technical Report *CSE 506* University of California, San Diego (1995)

Submitted to *Adaptive Behavior*

Abstract

The traditional explanation of delayed maturation age, as part of an evolved life history, focuses on the increased costs of juvenile mortality due to early maturation. Prior quantitative models of these trade-offs, however, have addressed only morphological phenotypic traits, such as body size. We argue that the development of behavioral skills prior to reproductive maturity also constitutes an advantage of delayed maturation and thus should be included among the factors determining the trade-off for optimal age at maturity. Empirical support for this hypothesis from animal field studies is abundant. This paper provides further evidence drawn from simulation experiments. “Latent Energy Environments” (LEE) are a class of tightly controlled environments in which learning organisms are modeled by neural networks and evolved according to a type of genetic algorithm. An advantage of this artificial world is that it becomes possible to discount all non-behavioral costs of early maturity in order to focus on exclusively behavioral consequences. In spite of large selective costs imposed on parental fitness due to prolonged immaturity, the optimal age at maturity is shown to be significantly delayed when offspring learn from their parents' behavior via imitation.

Key Words: age at maturity, imitative learning, Latent Energy Environments (LEE), genetic algorithms, neural networks, Baldwin effect, genetic assimilation

Short Title: Evolution of Imitative Learning

1 Introduction

To many computational modelers using neural networks as models of learning, it is second-nature to provide as much training experience for their networks as possible, since almost always this leads to increased performance. But when these neural nets are used in conjunction with genetic algorithms (GAs) as models of the *evolution* of learning individuals, it is necessary to consider the problem as it occurs in Nature: Is the increased time required by extended training worth the evolutionary costs of providing it? The fact that many organisms spend their prolonged immaturity as part of family units that can -- potentially -- shape the experience of the learning juvenile in predictable, heritable ways makes the role of the learning period especially important as we attempt to model proto-cultural effects on cognitive development. In this paper we will address restricted versions of these questions, focusing exclusively on imitative types of learning from parent to child.

In the study of the evolution and adaptation of life-history traits in animals, the commonly accepted theory states that any particular trait accomplishes a trade-off between the different selective pressures acting simultaneously upon the phenotypic variants of that trait.¹ One such trait that is central in behavioral and developmental psychology as well as in theoretical biology is the age at which an individual reaches the adult stage. To use a more precise terminology, it is preferable to consider the maturation age, or age at maturity, defined as the onset of the stage when the individual is capable of reproducing.

The selective pressures concerning the evolution of age at maturity considered in theoretical biology are the adaptive costs and benefits associated with anticipating or delaying maturation age. Typically (see, e.g., Stearns 1992, chapter 6), the costs of

¹ For a comprehensive survey of life-history theory, see Stearns (1992).

delayed maturation include: (i) lower population reproductive fitness due to longer generation time; (ii) lower individual reproductive fitness due to decreased probability of reaching the mature stage; and (iii) parenthood cost due to longer immature period requiring parental care. Conversely, the benefits of delayed maturation typically include: (i) higher fecundity of the parent who can grow for a longer time and better endure the reproductive effort; and (ii) lower instantaneous juvenile death rates due to better quality of offspring or parental care.

Ecological field studies attempting to quantify the magnitude of selective pressures toward delayed maturation necessarily focus on easily measurable life-history traits, such as body size and weight. On the other hand, animal and human psychology studies concerned with social cognition emphasize the improvement of offspring phenotypes taking place through cultural learning. For example, Tomasello, Kruger and Ratner (1993) identify three types of cultural learning: imitation, instruction, and collaboration. It would be desirable to be able to quantitatively correlate such behavioral advantages with life-history models of delayed maturation. To apply analytical trade-off models, it is necessary to measure phenotypic traits associated with behavioral development -- a very difficult thing to do in the field. Therefore, behavioral benefits of delayed maturity have been neglected in studies of the evolution of maturation.

The main goal of this paper is to provide quantitative evidence, through a simple computational model, that the benefit of cultural learning may be large enough to provide pressure for delaying maturation. In the simulation model we propose, the advantage of parental care is, in fact, the only benefit of delayed maturation. Lower death rates for a population of individuals may result from improvements of phenotypic behaviors occurring before individuals reach the adult stage. In the model, this improvement is acquired by the offspring through learning by imitation of its parent. Notice that it makes sense to model learning only before maturation; this simplification of life-long learning

retains the important aspect that parents must act as teachers to their immature offspring. While it is not our intention to propose imitation as a universal model for cultural transmission, we use imitation as one possible mechanism by which parents may confer a cultural advantage to their immature offspring.

Our conjecture, that fitness improvements of phenotypes by means of learned behaviors play an important role in the evolution of maturation, is not targeted at any one particular species. Rather, we formulate an idea model to suggest how this biological question could be addressed. We simulate the evolution of age at maturity as a genotypic trait regulating the duration of phenotypic learning by imitation. The model exhibits all of the costs of delayed maturation enumerated above. We hypothesize that the ability to imitate parental behaviors is adaptive, and hence expect that there be selective pressure toward delaying maturation, so that such learning can occur.

In related work, Johnston (1982) has associated learning with immaturity in the study of costs and benefits of phenotypic plasticity as a life-history trait. Cavalli-Sforza and Feldman (1981) have characterized cultural transmission as a diffusion process based on a model of imitation. Another seminal evolutionary account of cultural transmission by imitation has been given by Boyd and Richerson (1985). Tomasello *et al.* (1993) identify imitative learning as the first of the three forms in which cultural learning manifests itself during ontogeny. An approach similar to ours in modeling learning of immature individuals has been taken by Denaro and Parisi (1994); more abstract models of cultural advantage have also been explored (Belew 1990).

The paper is organized as follows. Section 2 contains an overview of the simulation tool used as the framework for our numerical experiments. Section 3 describes how the model and simulations have been implemented. Sections 4 and 5 illustrate and analyse, respectively, the experimental results of simulated experiments. Finally, Section 6 draws some conclusions.

2 LEE model

Simulation experiments described in this paper make use of a recent artificial life model and simulator called Latent Energy Environments (henceforth LEE).² A detailed description of the model and other experiments using it can be found elsewhere (Menczer and Belew 1994, 1995), but its features most relevant to this work are summarized below.

In LEE, environments of measured complexity can be constructed thanks to a careful definition of the energy consumed by organisms and of the work they must do in order to realize it. Energy is always conserved, and behavioral strategies are evolved by populations so as to allow an efficient exploitation of the available energy. Environmental complexity is therefore measured with respect to the difficulty of the survival task.

Organisms in a LEE population live in a shared environment consisting of a rectangular grid with toroidal edge conditions. Each organism behaves according to a neural net mapping sensory information to motor actions (a more detailed description follows). Populations of such organisms are born, reproduce, and die according to the algorithm presented in Figure 1. Organisms in an initial population are randomly constructed and given an initial, random reservoir of energy distributed uniformly in the interval $[0, \alpha]$. A constant, “metabolic” energy cost is paid anytime an individual is active. If an organism's energy level ever reaches zero, it dies. On the other hand, if an organism is consistently capable of realizing more energy than its metabolism demands, it reproduces and gives half of its energy to the offspring anytime its energy passes the fixed reproductive threshold α .

FIGURE 1 APPROXIMATELY HERE

²The most recent release of the LEE software package and documentation is available by URL <http://www-cse.ucsd.edu/users/fil/> or anonymous ftp://cs.ucsd.edu/pub/LEE.

A LEE population evolves according to a “steady-state” GA (De Jong and Sarma 1993). Reproduction is asexual, and the only genetic operator is mutation. Rather than assuming a constant population size, and scaling all fitness values proportionally (as in a conventional GA), the size of the population varies with time, and extinction is possible. “Fitness” is most appropriately measured in terms of net energy intake per unit time; this is directly correlated to the number of offspring per lifetime, as can be easily seen by considering the fixed reproductive threshold (Figure 1). The progression of the adaptive process is measured in terms of time intervals (cycles) that are significantly shorter than generations.

An organism is modeled by its energy reservoir, a sensory-motor system, and a neural network. The sensory system consists of a set of sensors, each mapping a state of the world onto one of the inputs of the network. Two types of sensors were used in these experiments, “contact” and “ambient” sensors. A contact sensor can only sense elements in the cell directly in front of the organism. An ambient sensor is capable of sensing elements anywhere in a region confined and oriented with respect to the organism. These senses are continuous values with the strength of the signal decreasing with distance. In the experiments reported here there are three ambient sensors, two with a range of three moves oriented to the opposite sides of the organism, and one with a range of five moves oriented straight ahead. There is also one contact sensor, oriented ahead. The sensory system is illustrated in Figure 2.

FIGURE 2 APPROXIMATELY HERE

The motor system is made of a pair of “binary” motors, which are controlled by two binary neural net output units. These allow organisms to make one of four possible moves: stay still, turn left or right 90 degrees, or move ahead by one cell. The sensors and motors are connected by a single hidden layer of seven units.

Due to the constant rate of replenishment of environmental resources to be shared and the constant reproductive threshold, fitness is strongly *density dependent*³ (Menczer and Belew 1995). The measured age distribution is stationary for most of the duration of the simulation -- except for the very initial phase in which there is an exponential demographic explosion, quickly taken over by damped oscillations. This model is similar to others' (e.g., Kozłowski and Wiegert 1987), where zero population growth is assumed at any given instant of time. A good measure of fitness is then given by the R_0 statistics, defined as the expected number of offspring per parent. Indeed, over a wide range of LEE environmental conditions, simulations have shown that the number of offspring is highly correlated (corr. coeff. ~ 1) with population size, p . Therefore population size indicates how well the collective behavior of the population is adapted to the environment, and we will use p as the single fitness measure for the evolving population in the remainder of this paper.

3 Experiment setting

The particular task implemented in LEE for the present experiments consists of approaching atoms scattered in the environment with uniform distribution in space and time. This is much like *positive chemotaxis* and represents foraging behaviors common to many biological species (Gruau and Whitley 1993). When an organism reaches a cell containing an atom, it acquires energy E . Atoms appear in the environment with uniform distribution and at a constant rate of r atoms per unit time. We can calculate the carrying capacity of such environments for a population of organisms moving according to a type of “random walk”:⁴

³ The term “density dependence” is used in ecological theory to indicate a condition in which individuals interact almost exclusively by sharing finite resources in a finite area.

⁴ We will use the term “random walk” in the remainder of the paper somewhat unconventionally. Due to the toroidal shape of the LEE world and the uniform distribution of atoms in our experiments, a behavior effectively equivalent to a random walk can be achieved by always moving forward. The equivalence is in terms of encounters with atoms (input is ignored); truly random movement is not a possible behavior in our model because neural nets have neither recursion nor random input.

$$p = \frac{rE}{c} \quad (1)$$

where c is the cost of making a move. Equation (1) is easily derived by setting the population growth to zero (Menczer and Belew 1995). The organisms rapidly evolve to implement this easy, robust, and relatively efficient random behavior, therefore the population size initially oscillates around the value given by Equation (1). Given knowledge of environmental resources we can predict population size, over a wide range of LEE experiments, with very high accuracy. Adaptive behaviors implementing better than random energy-seeking strategies result in population sizes beyond the limit of Equation (1).

To study age at maturity, the life of organisms is divided into two distinct stages separated by an “age of maturity.” A gene, called a_m (age at maturity), represents the age of the transition to maturity in number of cycles. The lower bound for a_m is birth (no immaturity) and the higher bound is death (no maturity). During its immature phase, an organism differs in exactly two respects from its mature form. First, the immature phenotype is not allowed to reproduce. Its energy at birth is conserved until maturity, and its moving costs are entirely transferred to its parent, as will be shown in more detail below. Second, juveniles undergo a process of learning by imitation. The details of this construction are suggested by Figure 3. Parent and child experience identical input stimuli from the world, but the output computed by the parent's neural net is used for determining both organisms' movement. The motor units of the parent's neural net are also used as a training signal for the child's neural net. We think of the offspring as being carried over the shoulders of its parent in this stage (cf. Figure 2). When an organism reaches age a_m it becomes an adult, normal member of the population and stops learning.

FIGURE 3 APPROXIMATELY HERE

Several objections to this construction can be anticipated. First, some may find the notion of “imitation” captured here inappropriately direct and imposed. Certainly the real, biological organism is faced with a much more difficult learning task, for example requiring an appropriate correspondence between the sensory stimuli of *watching* its parent perform a behavior and the proprioceptive and kinesthetic experience of *executing* this behavior itself. We view the biological basis of such identification an extremely important research issue. Our present model elides the issue entirely by assuming complete, perfect identification. Second, why is the immature stage of an organism made to coincide exactly with its learning phase? One motivation behind this modeling choice is the simple observation that in many species learning is in fact contemporaneous with the initial, pre-reproductive phase of life. But our central motivation is that this equivalence is a methodological necessity in order to argue that learning plays a role in the evolution of life-history strategies. By limiting learning to the immature stage, we intrinsically transfer the costs of immaturity to the learning process.

With these caveats, let us review the way in which the three costs of delayed maturation (cf. Section 1) are modeled. First, longer generation time is implicit in the LEE model: longer immature stages correspond to shorter times in which offspring can be generated, and thus to lower R_0 fitness. Second, increasing the immature stage decreases the probability of reaching maturity, due to the cumulative probability of death of the parent. When a parent dies, its “orphan” immature offspring are dropped into the environment before becoming adults. An orphan cannot move until it reaches its mature age, but may run out of energy and die before becoming adult. The probability of survival to mature age is higher the closer to maturation the offspring is when the parent dies. The third cost of the immature stage is associated with parental care: a parent pays an additional energy toll equal to the cost of a move for each offspring it carries on its shoulders. Energy is still conserved, since immature offspring neither incur living costs nor receive energy from their parents, as long as they are carried on their shoulders.

Learning is the only benefit of delayed maturity in our model. Thus if a delay of maturation age is observed in simulations, it can only be attributed to the advantage provided by learned behaviors. Note that this benefit is conferred upon the offspring and not the parent. In evolutionary ecology, similar benefits for the offspring are mainly attributed to morphologic development; for example, offspring size has been considered by Kozlowski and Wiegert (1987). The benefit modeled here is less explicit because mediated through the behavior of the offspring: juvenile mortality can be decreased thanks to the experience accumulated by offspring during their immature life stage, by way of parental imitation.

An offspring inherits its parent's genotype at reproduction. The genotype is composed of the a_m gene and a description of the neural network's connections. Genotypes evolve through selection and mutation. a_m is mutated (with probability 0.1) by a random additive deviate uniformly distributed in the interval $[-a_m, +a_m]$, so that it remains constrained to positive values. Genotypic connections are represented, as a sequence of real numbers, by the weights between the four inputs, seven hidden units, and two outputs. They are mutated (with probability 0.15) by random additive deviates uniformly distributed over the range $[-2.5, +2.5]$.

At birth, an organism's phenotype is constructed by copying the connection weights from its genotype. Learning is modeled via forced motor imitation of the parent, given that both parent and its offspring receive the same sensory input. The actions encoded in the parent's outputs are used as teaching signals for the offspring on its shoulders. Standard error back-propagation (Rumelhart, Hinton and Williams 1986) is then used for altering the phenotypic connection weights of the offspring. The learning rate parameter is set to a relatively high value (0.8). The reason is that the learning neural networks of immature organisms in this model are given a relatively short training experience (the duration of the immaturity stage), and previous research has shown that the more volatile learning

associated with high learning rates is effective in the context of evolving populations (Belew, McInerney and Schraudolph 1991). It is important to stress that since in this simulation weight changes only affect phenotypes, learning is strictly non-Lamarckian -- a requirement of the model in order for our results to transfer to natural systems.

4 Results

In this section we outline the main results of a series of experiments aimed at evaluating the potential role of imitation learning in the evolution of age at maturity. Simulations were run for 150,000 cycles (as defined above). The first simulations are controls designed to facilitate the interpretation of subsequent experiments.

In Figure 4 population size as a function of time (measured in cycles) is plotted for two single runs in which a_m is held at the constant values 0 (no immaturity) and 100 cycles. The first case ($a_m = 0$) provides a baseline with neither the costs nor advantages of an immature period. The observable increase in population corresponds to an improvement in the approaching behaviors of the organisms due to the evolution of their network weights. The second case ($a_m = 100$) gives us a measure of the magnitude of the cost of delayed maturation: this is large enough to always drive the population to extinction in less than 25,000 cycles. Thus we expect strong selective pressure against delayed maturation. For comparison, the (analytically derived) expected size of a population of random walkers is also shown in the figure; cf. Equation (1).

FIGURE 4 APPROXIMATELY HERE

In order to study the evolution of life-history strategies, however, the age of maturity must itself be assumed to be an evolved trait, the ability to learn via imitation being an independent variable. Our next experiments therefore add the a_m gene to the genome of the evolving organisms. The population is initialized with a_m uniformly distributed over the interval $[0, 100]$. Figure 5 shows the population size for two simulations, one in

which imitative learning is enabled during the immature phase and one in which it is not. In both cases, after the initial stochastic fluctuations, the populations are able to evolve individuals with behaviors significantly more adaptive than random ($p < 0.01$). However, with imitative learning, the evolution of good approaching behaviors is accelerated by approximately 75 cycles.

FIGURE 5 APPROXIMATELY HERE

Figure 6 plots the population's average value for the evolved a_m gene in the two experiments, again with and without imitative learning. In the absence of learning, the costs of immaturity dominate and a_m rapidly evolves to zero. When the young are allowed to learn via imitation, however, a conspicuous delay (from ~20,000 to ~80,000 cycles) in the extinction of the immature phase is observed. Notice that this is an evolutionary delay experienced by the species, not to be confused with the developmental delay controlled by the a_m and experienced by individuals. It is also important to recognize that the inclusion of imitative learning does not keep a_m from becoming zero; it only prolongs this process. The fact is, however, that real, biological environments rarely enjoy the constant environmental conditions of our artificial LEE world. Factors that maintain plasticity, making an organism more responsive to changes in its environment, are therefore of consequence when interpreting these results in natural settings.

FIGURE 6 APPROXIMATELY HERE

In summary, the delay observed represents evidence that there is a trade-off between costs and advantages of delayed maturation. This trade-off is quantitatively estimated by the value -- slightly above 50 cycles -- around which the population's average maturation age oscillates initially (cf. Figure 6). Since in our model the only advantage can be that of offspring learning via parental imitation, we must conclude that cultural transmission

of adapted behaviors is one of the evolutionary factors creating selective pressure in favor of longer immature stages.

5 Analysis

The role learning and other forms of phenotypic plasticity can play in accelerating evolutionary change is increasingly well explored. The “Baldwin effect” (Baldwin 1896) has come to describe a mechanism by which organisms that adapt within their lifetime to the environment (i.e., learn) tend to be successful, leading to higher probability of reproduction and a fixation of their useful characteristics in the population, despite lack of Lamarckian genetic propagation. Hinton and Nowlan (1987) are most often cited for their explanation, according to which individual learning changes the fitness function, without changing the genotype, by smoothing the fitness landscape. Most interpretations of the effect seem to involve the evolution of genetically determined characteristics similar in function to those that arise through phenotypic plasticity (see, e.g., Belew and Mitchell 1995). Another version of the Baldwin effect was given by Waddington (1942), who called it “genetic assimilation” to stress the back-filling role of mutation with respect to phenotypic changes. A range of evidence for these important, albeit subtle connections between learning and evolution has been found by a number of investigators using computer simulations (Hinton and Nowlan 1987, Belew 1990, Ackley and Littman 1992, Gruau and Whitley 1993, Nolfi, Parisi and Elman 1994).

Our experiments are consistent with these accounts. Organisms are able to explore a range of behavioral strategies within their lifetimes through learning. The first evidence of this effect is the observation that learning *phenotypes* can exploit the environment better than their non-learning counterparts, given equal evolutionary times, as shown by the larger population size (cf. Figure 5, between 20,000 and 90,000 cycles; recall that learned behaviors are not transmitted genetically). But by inspecting Figures 5 and 6, we can obtain a stronger demonstration of the Baldwin effect: learning individuals do better

even when learning is turned off. In fact, the foraging capability of the learning population becomes innate after approximately 80,000 cycles (where error bars begin to overlap in Figure 6), while the non-learning population acquires such capability genetically after approximately 90,000 cycles (where error bars begin to overlap in Figure 5). In other words, the Baldwin effect confers a *genotypic* advantage as well. Therefore the evolutionary process is capable of exploiting more information about adaptively favorable characteristics of the environment than is possible without learning.

But if learning by imitation during immaturity provides selective pressure toward delayed maturation, why does a_m eventually converge to zero, as shown in the last phase of its evolutionary course in Figure 6? This fact can be interpreted as evidence of Waddington's genetic assimilation. Modifications that learning causes upon the phenotype during the juvenile period, are eventually re-discovered by the evolutionary process and thereby affect the genotype. Once this happens, learning is no longer useful because offspring are innately capable of the behaviors that in earlier generations they could only acquire by imitating their parents. Therefore learning no longer confers an advantage to immature offspring. The costs of delayed maturation remain the same, and the missing benefit causes the immature period to disappear. The next subsections analyse just how genetic assimilation occurs in our experiments, interpreting the phenomenon in two phenotypic spaces: connection weights and behaviors.

5.1 Correlation of learned and evolved tasks

The behavior that is eventually evolved turns out to be identical to the one learned by parental imitation, as we have verified by inspection. This may seem trivial, given the direct correlation between the capability of approaching atoms and fitness. Recall that in the original construction, an offspring's neural net learns a mapping from the shared input to the parent's output. We can imagine that the effect of learning on evolution might be weaker if the offspring were trained on a different task, less correlated with fitness.

To explore these questions, we consider a variant of our model in which, in addition to imitating their parents, immature offspring are trained to predict their parents' sensory input following the current move. This task, while different from the approaching task, might still be helpful to survival because offspring can take advantage of their parents' experience (Nolfi *et al.* 1994). Genetically, the single 'maturation gene' a_m is replaced by two new genes. The duration of the prediction learning phase is regulated by a 'prediction gene' a_p while the imitation learning phase is regulated by an 'imitation gene' a_i . During the prediction phase the offspring performs prediction learning using the parent's input, while imitation is as described in Section 3. Thus imitation and prediction are effectively decoupled, and prediction learning becomes an independent mechanism to model the cultural transmission of a parent's experience to its immature offspring. Between birth and age $\min(a_i, a_p)$, both imitation and prediction learning occur. Between $\min(a_i, a_p)$ and $\max(a_i, a_p)$, only one of the two tasks is learned, namely, the one corresponding to the greater gene value. Both genes evolve analogously to a_m . The model stipulates that an organism is able to reproduce only after age $\max(a_i, a_p)$. This way both the prediction and imitation processes have associated costs as necessary to model the trade-off. Furthermore, assuming that the selective advantage of prediction learning may be smaller than the one of imitation learning, we have reduced to zero the parental cost of carrying offspring on one's shoulders. The other costs associated with immaturity, of course, remain (cf. Section 3).

FIGURE 7 APPROXIMATELY HERE

The result of the experiment is shown in Figure 7. The a_p gene goes to zero very quickly, much like a_m without imitation learning (cf. Figure 6). On the other hand, the a_i gene evolves like a_m with imitation learning enabled (cf. Figure 6). This result suggests that prediction learning is not advantageous during immaturity. Note that the prediction task using ambient sensors is very difficult, due to their non-invertible mapping from world

states to neural net inputs; for a more detailed discussion of this problem see (Menczer and Belew 1994). Therefore we do not have evidence that prediction learning is detrimental, however we are unable to demonstrate its usefulness in this context.

5.2 Genetic assimilation in weight and behavior spaces

We now return to consider our standard model, in which juveniles are trained by their parent's actions and this imitation is controlled by a single a_m gene. The most straightforward analysis of genetic assimilation consists of verifying that phenotypic modifications made by learning have become unnecessary due to behaviorally equivalent changes in the genome. In our simulations, learning is being modeled by changes made to the neural networks weights as the result of error back-propagation. One simple measure of the magnitude of this change is given by the Euclidean distance in the space of the network connection weights:⁵

$$\Delta \underline{w} = \sqrt{\sum_{i,j} (w_{ij}^{after} - w_{ij}^{before})^2}. \quad (2)$$

Here the superscripts indicate the neural network's initial, genetically determined weights before learning and the final phenotypic weights after learning. Genetic assimilation would imply that this distance decreases to zero as learned changes are transferred to the genome. Since our experiments typically allow the amount of learning to vary with the a_m gene, care must be taken to evaluate the magnitude of these changes after a constant amount of learning.

We therefore consider all the organisms in a typical simulation run. The initial genotypic weights of their neural nets are recorded when the organisms are immature (at birth), and then subjected off-line to a fixed period of imitation learning (100 cycles of back-propagation) based on the parents' behaviors. This procedure guarantees that if the

⁵ In fact, this measure depending only on initial and final positions is too simple. The distance travelled by a learning network also depends critically on features of the learning surface, type of search performed by the learning algorithm, etc. (Belew *et al.* 1991)

Euclidean distance decreases, this is due to genetic assimilation rather than an artefact of variable immaturity periods. With this caveat, the learning process takes place under identical environmental conditions as during the actual simulation.

FIGURE 8 APPROXIMATELY HERE

Figure 8 shows a scatter plot of the average values of neural network weight change from Equation (2) versus the corresponding a_m gene value. The two are very well correlated ($\mu \sim 0.936$): earlier maturation corresponds consistently with smaller phenotypic change in weight space. This confirms our interpretation that genetic assimilation makes imitation learning unnecessary and thus the immature stage disappears.

Another way of analysing genetic assimilation is by directly inspecting phenotypic behaviors during evolution. Movement is designed to be the canonical behavior in LEE organisms. A necessary component of optimal foraging behavior in the environments used in the experiments of this paper is to move forward when the contact sensor signals the presence of an atom in the cell immediately facing the organism. An indicative measure of *poor* behavior is therefore $P(\sim F)$, the fraction of times a mature organism presented with an atom in the next cell does *not* move forward to consume it. When the population converges to *low* values of $P(\sim F)$ at birth, we have an indication that the *good* behavior has been encoded into the genotypic weights, making it unnecessary to learn during life.

FIGURE 9 APPROXIMATELY HERE

Figure 9 shows a scatter plot of the average values of $P(\sim F)$, measured off-line and without learning, vs. the corresponding a_m gene value. The two are again well correlated, though less well than in the previous experiment ($\mu \sim 0.765$). Through genetic assimilation, organisms are born with progressively better behavior. The fraction of times organisms move forward in the presence of an atom in the facing cell increases in the

course of evolution, and the average maturation age correspondingly decreases due to the costs of immaturity.

6 Conclusion

The simulation experiments reported in this paper have shown that behavioral improvements of the immature phenotype may be among the factors determining the trade-off between costs and benefits of delayed maturation. We have shown that learning during immaturity provides selective pressure toward delayed maturation. We have also demonstrated that this extended window of juvenile plasticity can be used, by the Baldwin effect, to good evolutionary advantage over non-learning populations: non-Lamarckian learning accelerates the evolutionary discovery of adaptive behaviors as well as their genetic assimilation.

We believe that our experiments bring an important new dimension to the discussion of maturation as an evolved life history, i.e., the important role *cultural* influences can play. Following Boyd and Richerson (1985), we mean “culture” to refer to: “the transmission from one generation to the next, via teaching and imitation, of knowledge, values, and other factors that influence behavior” [p. 2]. In the present case the “knowledge” being transmitted is the adult experience, whereby impressionable young learners are trained via imitation of their parents' behaviors. Our model also provides an interesting contrast to more tangible forms of transmission, for example those mediated by external representational media (Hutchins and Hazelhurst 1991). Evidence for cultural transmission is most clear when it leaves a trace of such external artefacts, but our experiments demonstrate that culture can also take the form of behavioral guidance across generations. The more plastic the juvenile, the more possible it is that environmental exigencies can perturb its “normal” developmental trajectory, with increasingly direct consequence on population fitness.

Prior quantitative approaches to the study of this important life-history trait are often subject to the limitations of field experiments, and therefore these have mainly focused on structural benefits of delayed maturation, such as body size, weight, etc. Animal psychology studies of cultural transmission, on the other hand, have difficulty in casting the phenotypic nature of learned behaviors into the more general evolutionary framework. The simulation approach employed here provides a bridge between the two sides.

Modeling the cultural transmission of learned behaviors from parents to offspring by forced imitation of motor responses to sensory stimuli is not intended to point to imitation as the only or the best mechanism by which an organisms can improve their reproductive fitness. One direction for further research is to develop a model of imitation more amenable to testing on the field, or more plausible as a model of the learning that occurs in many animal species. For example, juveniles may receive sensory input that is statistically correlated --- rather than identical --- to that of their parents, if they keep within some distance while foraging in patchy environments. In our model, we were able to make a correspondence between costs and benefits of learning by forcing parents to carry immature offspring on their shoulders. Proximity between parents and offspring could originate, in more realistic models, from foraging in groups --- for more successful predation, better defense, or sharing of large resource patches. Another direction is to consider alternative forms of learning, such as instruction and collaboration (Tomasello *et al.* 1993), to gain a more general understanding of the mechanisms by which cultural transmission can influence the evolution of life-history traits.

A related approach is to study the adaptive value of maturation using a task that must be learned by an organism during its life and which cannot evolve directly (Todd and Miller 1993). In such a situation the Baldwin effect must be limited. This is also the object of current research by the authors, with preliminary results that confirm our expectation from the results illustrated here: maturation age evolves to a value greater than zero, and

remains at such value indefinitely (Cecconi and Parisi 1994). Another adaptive advantage of phenotypic plasticity may be found when the population evolves in a non-stationary environment, as considered by many researchers. Menczer (1994) illustrates the idea in the same LEE context used in this paper. If environmental changes occur at a rate faster than the time necessary for evolution to encode adaptive changes in genotypes, then we observe selective pressure for maintaining a period of phenotypic adaptation, thus allowing evolution to track the environment.

Learning by imitation is a relatively simple paradigm within which we have been able to model the costs of delayed maturity on both mature organisms (parents) and immature ones (offspring). The adaptive advantage of transferring a parent's experience onto its offspring by “carrying them on its shoulders” has been shown to be large enough to permit a significant delay in maturation age, at least until the adaptive behaviors are transferred to the genotypes of the population via genetic assimilation. Our data, therefore, support behavioral and cultural models of maturation age as a life-history trait. Since imitation seems to be a very basic form of cultural transmission in social animals, this result not surprisingly points to one account for the long immaturity stages in these species. Clearly, many other factors determine the trade-off for optimal age at maturity, but critical data about cultural learning are available for only a very few species, e.g., apes and humans (see Tomasello *et al.* 1993). We hope our research stimulates field studies aimed at providing experimental evidence for (or against!) the behavioral benefits of delayed maturation illustrated here.

Acknowledgements

We thank Domenico Parisi, Jeff Elman, and everyone in the Cognitive Computer Science Research Group of the CSE Department at UCSD for useful discussions. We are also grateful to Lashon Booker and two anonymous reviewers for helpful comments on an earlier draft of the paper. The first author is also grateful to the Institute of Psychology of the National Research Council (Rome) and the CSE Department (UCSD) for support.

Figure captions

Figure 1. The basic LEE algorithm. The constant α is the reproductive threshold.

Figure 2. The sensory system of an organism. The ranges of the different sensors are shown. An immature offspring is also represented "on the shoulders" of its parent.

Figure 3. Imitative learning architecture. Labels on input and output units identify sensors ("c" for contact, "a" for ambient) and motors ("m").

Figure 4. Population size for different fixed ages at maturity. Data are collected from two single simulation runs. Extinction occurs in the case of non-zero maturation age. The line labeled "random walk" in this and the following figure is from Equation (1). At the beginning of the simulations, there is a large amount of environmental resources that results in an exponential peak in population size (not shown in the plot for clarity). Then density dependence sets in because the rate of replenishment of resources is held constant, and the population is limited by the carrying capacity of the environment.

Figure 5. Population size with evolving maturation age gene. The case with learning from imitation results in faster fitness improvement. Errors correspond to one standard deviation (a 68% confidence interval) over repeated simulation runs.

Figure 6. Evolution of maturation age gene. When imitation occurs, maturation is delayed noticeably; the disappearance of immaturity is also delayed evolutionarily. Errors correspond to one standard deviation (a 68% confidence interval) over repeated simulation runs.

Figure 7. Evolution of prediction and imitation genes. Errors correspond to one standard deviation (a 68% confidence interval) over repeated simulation runs.

Figure 8. Scatter plot showing correlation between a_m and Δw (measured over a fixed number of back-propagation cycles at different evolutionary times between 50,000 and 120,000 cycles). For clarity, both statistics are averaged over the population and over 10 different simulation runs. Linear regression fit of the data is also shown.

Figure 9. Scatter plot showing correlation between a_m and $P(\sim F)$ (measured at different evolutionary times between 0 and 150,000 cycles). For clarity, both statistics are averaged over the population and over 10 different simulation runs. With respect to the linear regression fit of the data, there are more outliers here compared to Figure 8; the main group (bottom, center) illustrates how maturation age decreases *following* the discovery of useful behavior.

References

- Ackley DH and Littman ML (1992) Interactions between learning and evolution. In Langton CG, Taylor C, Farmer JD and Rasmussen S, eds.: *Artificial Life II*, Reading, MA: Addison-Wesley, 487-507
- Baldwin JM (1896) A new factor in evolution. *American Naturalist* 30, 441-451
- Belew RK (1990) Evolution, learning, and culture: Computational metaphors for adaptive algorithms. *Complex Systems* 4, 11-49
- Belew RK, McInerney J and Schraudolph NN (1991) Evolving networks: Using the genetic algorithm with connectionist learning. In Langton CG, Taylor C, Farmer JD and Rasmussen S, eds.: *Artificial Life II*, Reading, MA: Addison-Wesley, 511-547
- Belew RK and Mitchell M, eds. (1995) *Adapting individuals in evolving populations: Models and algorithms*. Santa Fe Institute Studies in the Sciences of Complexity. Reading, MA: Addison Wesley
- Boyd R and Richerson PJ (1985) *Culture and the evolutionary process*. Chicago, IL: Chicago University Press
- Cavalli Sforza LL and Feldman MW (1981) *Cultural transmission and evolution: a quantitative approach*. Princeton, NJ: Princeton University Press
- Cecconi F and Parisi D (1994) Learning during reproductive immaturity in evolving populations of neural networks. Technical Report PCIA-94-12 Institute of Psychology of the National Research Council, Rome, Italy
- De Jong KA and Sarma J (1993) Generation Gaps Revisited. In L.D. Whitley, ed.: *Foundations of Genetic Algorithms 2*, San Mateo, CA: Morgan Kaufmann
- Denaro D and Parisi D (1994) Imitation and cultural transmission in populations of networks. Technical Report PCIA-94-17 Institute of Psychology of the National Research Council, Rome, Italy
- Gruau F and Whitley D (1993) Adding learning to the cellular development of neural networks: Evolution and the baldwin effect. *Evolutionary Computation* 1(3), 213-233

- Hinton GE and Nowlan SJ (1987) How learning can guide evolution. *Complex Systems* 1, 495-502
- Hutchins E and Hazelhurst B (1991) Learning in the cultural process. In Langton CG, Taylor C, Farmer JD and Rasmussen S, eds.: *Artificial Life II*, Reading, MA: Addison-Wesley, 689-706
- Johnston TD (1982) Selective costs and benefits in the evolution of learning. *Advances in the study of behavior* 12, 65-106
- Kozlowski J and Wiegert RG (1987) Optimal age and size at maturity in annuals and perennials with determinate growth. *Evolutionary Ecology* 1, 231-244
- Menczer F (1994) Changing Latent Energy Environments: A case for the evolution of plasticity. Technical Report CS94-336, University of California, San Diego
- Menczer F and Belew RK (1994) Evolving sensors in environments of controlled complexity. In Brooks R and Maes P, eds.: *Artificial Life IV*. Cambridge, MA: MIT Press, 210-221
- Menczer F and Belew RK (1995) Latent Energy Environments. In Belew RK and Mitchell M, eds.: *Adapting individuals in evolving populations: Models and algorithms*. Santa Fe Institute Studies in the Sciences of Complexity. Reading, MA: Addison Wesley
- Nolfi S, Parisi D and Elman JL (1994) Learning and evolution in neural networks. *Adaptive Behavior* 3(1), 5-28
- Rumelhart DE, Hinton GE and Williams RJ (1986) Learning internal representations by error propagation. In Rumelhart DE and McClelland JL, eds.: *Parallel Distributed Processing*, Vol. 1, Ch. 8 Cambridge, MA: Bradford Books (MIT Press)
- Stearns SC (1992) *The evolution of life histories*. Oxford University Press
- Todd PM and Miller GF (1993) Parental guidance suggested: How parental imprinting evolves through sexual selection as an adaptive learning mechanism. *Adaptive Behavior* 2(1), 5-47

Tomasello M, Kruger AC and Ratner HH (1993) Cultural learning. *Behavioral and Brain Science* 16, 495-511

Waddington CH (1942) Canalization of development and the inheritance of acquired characters. *Nature* 150, 563-565

```

for each cycle {
  for each alive organism {
    sense world;
    compute actions;
    move;
     $Energy \leftarrow Energy + \sum_{\{actions\}} Benefits - \sum_{\{actions\}} Costs ;$ 
    if (  $Energy > \alpha$  ) {
      reproduce;      /* copy genotype      */
      mutate;         /* change new genotype */
      develop;        /* get new phenotype   */
    }
    else if (  $Energy < 0$  ) die;
  }
  replenish world;
}

```

Figure 1. Basic LEE algorithm. The constant α is the reproductive threshold.

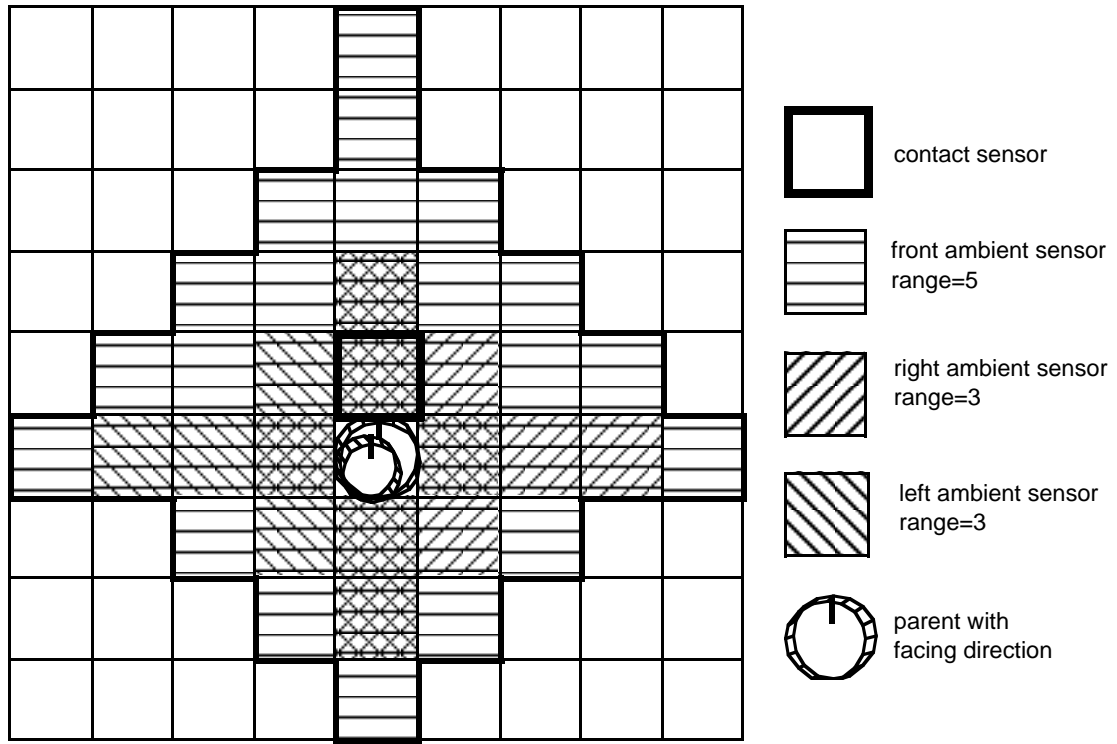


Figure 2. The sensory system of an organism. The ranges of the different sensors are shown. An immature offspring is also represented "on the shoulders" of its parent.

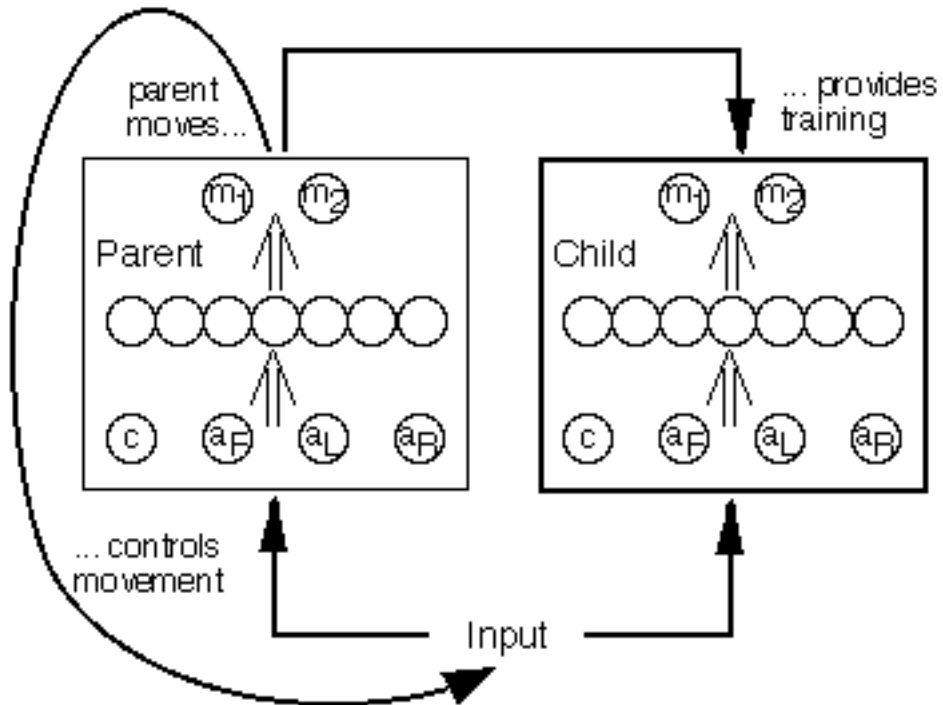


Figure 3. Imitative learning architecture. Labels on input and output units identify sensors (“c” for contact, “a” for ambient) and motors (“m”).




Figure 4. Population size for different fixed ages at maturity. Data are collected from two single simulation runs. Extinction occurs in the case of non-zero maturation age. The line labeled “random walk” in this and the following figure is from Equation (1). At the beginning of the simulations, there is a large amount of environmental resources that results in an exponential peak in population size (not shown in the plot for clarity). Then density dependence sets in because the rate of replenishment of resources is held constant, and the population is limited by the carrying capacity of the environment.




Figure 5. Population size with evolving maturation age gene. The case with learning from imitation results in faster fitness improvement. Errors correspond to one standard deviation (a 68% confidence interval) over repeated simulation runs.

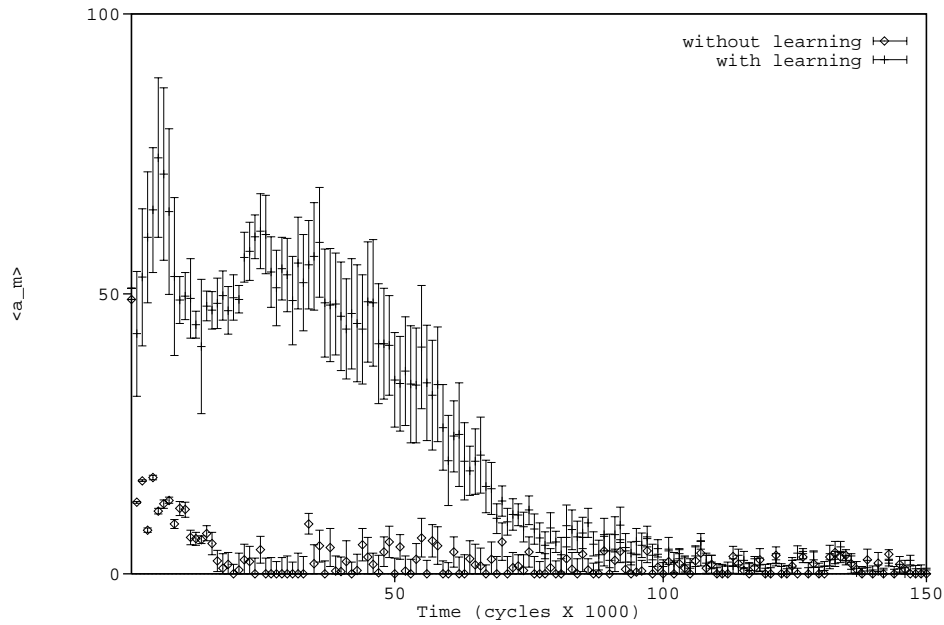


Figure 6. Evolution of maturation age gene. When imitation occurs, maturation is delayed noticeably; the disappearance of immaturity is also delayed evolutionarily. Errors correspond to one standard deviation (a 68% confidence interval) over repeated simulation runs.




Figure 7. Evolution of prediction and imitation genes. Errors correspond to one standard deviation (a 68% confidence interval) over repeated simulation runs.




Figure 8. Scatter plot showing correlation between a_m and Δw (measured over a fixed number of back-propagation cycles at different evolutionary times between 50,000 and 120,000 cycles). For clarity, both statistics are averaged over the population and over 10 different simulation runs. Linear regression fit of the data is also shown.




Figure 9. Scatter plot showing correlation between a_m and $P(\sim F)$ (measured at different evolutionary times between 0 and 150,000 cycles). For clarity, both statistics are averaged over the population and over 10 different simulation runs. With respect to the linear regression fit of the data, there are more outliers here compared to Figure 8; the main group (bottom, center) illustrates how maturation age decreases *following* the discovery of useful behavior.