

## Evolved Sensitive Periods in Learning

Kai Olav Ellefsen

Department of Computer and Information Science,  
Norwegian University of Science and Technology  
email: kaiolae@idi.ntnu.no

### Abstract

In this paper we study age-varying plasticities across different components in an artificial neural network performing a reinforcement learning task. An evolutionary algorithm is given the task of mapping the age of agents to the plasticity levels of different network components. The results show that patterns of plasticity resembling biological sensitive periods appear, and that these periods schedule learning across the components of the network, which leads to a reduction in the total learning effort while retaining the quality of learning. The sequencing of sensitive periods forms a cascade of partially-overlapping learning periods, which has been proposed as a way of organizing sensory development of abilities that depend on several interrelated brain functions.

### Introduction

Periods in the life of an individual where environmental stimuli are of particular importance for the development of a certain ability are called *critical periods* or *sensitive periods*. They were originally called critical periods to emphasize how lack of the correct stimuli would lead to the sensory system not developing as in normal individuals. In other words, the sensory input is *critical* for neural development. Hubel and Wiesel's classic paper (Hubel and Wiesel (1970)) illustrates this: One eye of a kitten was sutured in various periods throughout life, and it was found that visual deprivation of one eye early in life prevented it from following the regular path of development, leaving the cat blind in that eye, even when it was opened later in life. The period where sutures had this effect was found to have a very specific beginning and end (about four weeks and three months of age, respectively), and this finding is typical of how a critical period was interpreted: A time period with a strict beginning and end, where sensory stimuli have large effects on neural development, and where sensory deprivation leads to abnormal development.

Later research (see for instance Lewis and Maurer (2005)) has shown that sensory systems have *several* critical periods that affect different parts of the sensory system at different stages of development. It has also been shown that these periods can be flexible, and their timing may be controlled by

*experience* rather than age. For example, dark-reared kittens can have the ending of their critical period for vision delayed due to the lack of visual stimuli (Trotter et al. (1981)). These findings have led many researchers to adopt the term *sensitive period*, to emphasize that the period shows a great deal of variation a) across *individuals* that have different experiences during development, b) across *sensory systems* in the same individual and c) across *different parts of a single sensory system*.

An example of the last type of variation was studied by Harwerth et al. (1986). The authors studied the effect of monocular deprivation on different visual functions in rhesus monkeys to investigate the timing of sensitive periods of different functions within a single sensory system. They found there to be *several partially-overlapping* sensitive periods within the visual system, and basic functions (such as spectral sensitivity) were found to have shorter sensitive periods than more complex functions (such as binocular vision). Knudsen (2004) suggests that this property is likely to be found also in other parts of cognitive development, such as in the development of language and social skills. It is logical that low-level behaviors should finish their sensitive period before high-level behaviors, because the low-level outputs will be noisy until these systems have matured, and it will not be possible for the high-level systems to learn from these noisy signals.

Werker and Tees (2006) reviewed findings about sensitive periods in speech processing. The authors argued that speech processing, like vision, depends on a number of inter-related, *hierarchically ordered* brain functions. Based on a review of studies in language development, the authors suggested a possible way for different levels of language learning to be organized through development, where low-level functions once again tend to stabilize before higher-level functions can initiate their sensitive period. Figure 1 shows how they envisioned a *cascade* of sensitive periods in speech processing.

We will mainly use the term *sensitive period* for the remainder of the paper. The term is meant to refer to a period of *heightened plasticity* and sensitivity to environmen-

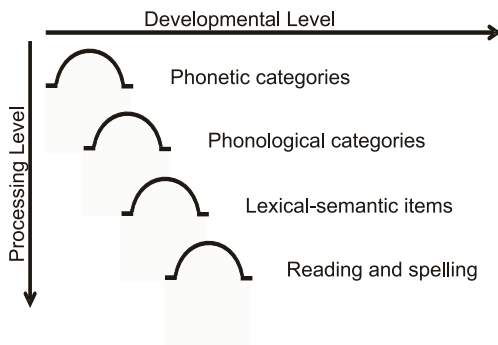


Figure 1: Sensitive periods in speech processing, as suggested by Werker and Tees (2006).

tal stimuli for a given function. If we want to emphasize that the period has a strict ending and that there is no plasticity outside the period, we will use the term *critical period*.

## Background

### Factors behind sensitive periods

Much work has been done in studying which factors *drive* sensitive periods. Armstrong et al. (2006) reviewed some of this work, pointing out two categories of changes that govern their initiation and termination: genetically-mediated changes and experientially-mediated changes. In the former category, sensitive periods will begin and end at a predetermined age as a consequence of innate traits. The latter category concerns factors affecting sensitive periods differently depending on the individual’s sensory-input history. Hensch (2005) reviews findings regarding sensitive-period plasticity in the primary visual cortex. He shows evidence from genetically altered mice that the sensitive period in this part of the visual pathway is related to the maturation of *inhibitory circuits*.

Johnson (2005) presents three views on what causes plasticity to decrease as a sensitive period ends: (a) Maturation changes, (b) self-termination of learning, and (c) stabilization of the constraints of plasticity. The first explanation describes chemical processes in the brain that terminate learning independently of experience. The second explanation describes how the learning itself may lead to a decrease in plasticity. One example of a factor leading to self-termination of learning is limited computational resources: when a lot has been learned, there is simply not the same capacity for adding new knowledge. The third explanation regards how sensitive periods may end due to a stabilization in external factors, such as bodily growth, rather than an actual decrease in plasticity.

### Sensitive periods in food preference formation

The experiments in this paper use food-preference formation as the domain for studying sensitive periods. The reason is

that this domain gives us a natural way of splitting the learning task into several subtasks at different levels in a hierarchy, each depending on subtasks below. This is essential if we hope to observe sensitive periods forming in a cascading manner.

Sensitive periods in food preference learning have been studied among animals, for instance snapping turtles (Burghardt and Hess (1966)), lynx spiders (Punzo (2002)) and cuttlefish (Darmaillacq et al. (2006)). Also among human children, a sensitive period of food preference learning has been suggested (Cashdan (1994)).

### Evolving sensitive periods

Bullinaria (2003) studied sensitive periods of learning, as part of a simulation of the human oculomotor system. By the use of an evolutionary algorithm, *age-dependent neural plasticity* was generated. The type of age-dependent plasticity arising from these experiments had parallels with biological sensitive periods. The evolved sensitive periods had the effect of letting individuals be plastic as their sensory systems underwent development, and less plastic after their development was done. Bullinaria describes two simplifying assumptions made in this work that will be important to address in the future. First, plasticities are fully determined by the genotype, meaning *experience* has no effect on the mapping from age to learning rates. Second, the evolved age-dependent plasticity is the same for all parts of the network. In this paper, we remove the second simplification, to study the effect of different evolved learning rates on different parts of an agent’s behavior. In a later paper (Bullinaria (2009)) Bullinaria found that a longer period of parental protection gave a lengthening of learning periods in children.

Kirby and Hurford (1997) studied *incremental learning* in language acquisition. By using an evolutionary algorithm to decide the timing of increments to learning capacity, they aimed to study *when* and *why* sensitive periods would form as an effect of incremental learning abilities. They enabled their evolutionary algorithm to shape incremental learning in two ways. It could determine learning resources on the basis of the *age* of an individual or the *experience level* of the individual. Kirby and Hurford found that evolving learning resources based only on the age of an individual gave extreme sensitive periods, similar to what has traditionally been called *critical periods*: A slight delay in the expected stimuli made individuals miss the evolved window for learning, unable to learn language at all. Evolving resources based only on the experience level of an individual gave the opposite effect: No sensitive periods were formed at all – experience could be postponed indefinitely and learning would proceed as normal. It was finally found that letting evolution *combine* the two forms of learning control would give sensitive periods similar to the ones seen in language acquisition in humans.

Hurford (1991) set up an evolutionary algorithm to model




Figure 2: A model of neuromodulated plasticity. The activity of modulatory neurons affects the plasticity of connections between regular neurons.

*language acquisition*. This is a task suggested to have a sensitive period, based on findings from, among other, language recovery in children and adults suffering from aphasia. The evolved learning efforts showed plasticity of the language system peaking in the first period of the individual's life, and gradually falling off to zero. Hurford hypothesizes that the ending of the sensitive period is not happening because it is beneficial to turn off language learning. Rather, as individuals master the language fully at a young age, the pressure to boost language learning is simply gone, so there is no pressure to drive the sensitive period into adulthood. This hypothesis is strengthened by the observations that individuals subjected to a high chance of “language amnesia” at some stage in life tend to evolve lifelong sensitive periods.

### Neuromodulated learning

In the experiments presented here, we study plasticity regulation *across several modules* of a neural network. To exchange reinforcement signals between these modules, we use *neuromodulated learning*.

Neuromodulated neural networks are networks that include another type of signal in addition to the traditional activity-propagating signals. Fellous and Linster (1998) present a review of work on these kinds of networks, revealing that modulatory signals have been used to affect network function in diverse ways, most of them biologically plausible, or at least biologically inspired.

Of particular interest in this context, is the use of neuromodulation to allow efficient reinforcement learning, as that is the role modulation plays in the experiments presented in this paper. In the model of neuromodulated reinforcement learning we employ (Figure 2), modulatory neurons affect the *plasticity* of connections between regular neurons. A single modulatory neuron affects all connections in a *link*, a concept which will be defined in the experimental setup-

section. The model we used is similar to that used in (Soltoggio et al. (2008)), which allowed an agent to learn from sparse events, by letting the modulatory signals have a multiplicative effect on neural plasticity.

### Research Questions

The experiments discussed herein study the emergence of sensitive periods in simple agents in an ALife environment. The environment was set up with a “hierarchical” learning task, where learning of high-level behaviors is dependent on the low-level behaviors already being stable. An evolutionary algorithm (EA) was used to search for optimal sensitive periods for different behaviors. The questions we want to answer are:

- Will the agents show sensitive periods that sequence learning in order of the complexity of behaviors (from low-level to high-level)?
- Will the evolved sensitive periods be able to reduce the agents' learning efforts, while still allowing them to learn the correct behavior?

### Experimental Setup

#### SEVANN

To evolve neural networks with age-varying plasticity, the system SEVANN (Script-Based Evolution of Artificial Neural Networks) (Downing (2010)) was used. SEVANN is a system that lets the user form an *underspecified script* defining an ANN, and then searches for good values for the unspecified parts by use of an evolutionary algorithm. For the experiments reported here, the architecture and initial weights of the ANN were fully specified in the scripts, to allow SEVANN to focus on evolving plasticity values through life, and to allow analyses of results to only depend on the plasticity.

In SEVANN, networks consist of *layers* and *links*: Layers are a number of *neurons* that share common attributes, as well as common inputs and outputs. Links are a number of *arcs* (single connections between neurons) that are grouped together, because they have common attributes, and connect the same layers. When working with age-varying plasticities, we encode learning rates on the link level, allowing evolution to differentiate learning rates between links, but not between single arcs within a link. This granularity was chosen because different links learn different behaviors, and thereby we may allow the sensitive periods to *sequence* behavior learning.

#### Food-Gathering Task

To study how sensitive periods may evolve to sequence the learning of several behaviors, a task with dependencies between different behaviors is needed. The task must have

different levels of behavior, where one level affects performance of the level above - this way, a *cascade* of sensitive periods, as suggested by Werker and Tees (2006), may emerge as a solution to the problem.

Agents are placed on a toroidal grid with food and poison randomly scattered around. To make the hierarchy of behaviors more complex, food is placed in *nuts*, which need to be cracked open to access the food inside. So, the agent has to make three kinds of decisions: where to navigate, which nuts to open, and which food elements to eat. The decisions depend upon each other in a bottom-up fashion: To learn which nuts to open, the agent must first have understood which foods are healthy, then link these to the nuts containing them. And to learn which nuts to navigate to, the agent must first have learned which nuts it wants to open.

Eating food, opening food nuts and visiting food locations increases the agent’s fitness, whereas eating poison, opening poison nuts and visiting poison locations decreases it by the same amount. The layout of the food-gathering grid was randomly initialized in each fitness evaluation, with a given probability of each cell containing a poison nut, a food nut or nothing.

An agent resulting from a successful run of the EA typically opens all food nuts and eats all foods. It eats one or a few poison items before learning these should be avoided. Thereafter it still *opens* poison nuts until it learns they are associated with poisons. Next, it passes over poison nuts without opening them, before finally learning to steer away from poison nuts. The order of learning here reflects the hierarchical ordering of the task, and sensitive periods are expected to form in the same order. Each individual was tested *five times* during each fitness evaluation, on grids with different positions of food and poison. This was done to make agents form *general* sensitive periods, instead of sensitive periods adapted to one particular environment.

### Food-Gathering Network

The network structure was chosen to enable learning to propagate from “low-level” (food preferences) to “high-level” (movement preferences) behaviors. The network structure is depicted in Figure 3. It utilizes *neuromodulatory neurons* to be able to transfer what it learned at lower levels to higher levels.

The network mirrors the hierarchical arrangement of the task: Higher levels of behavior are learned by use of reinforcing signals from lower levels. For instance, to learn which nuts to crack (in the link between **InputNut** and **DecisionCrack**), the output from evaluating food is used as a reinforcing signal, with a delay of one timestep: if the food evaluated in the current timestep is good, the decision to crack the nut in the previous timestep was obviously good, and should be repeated in the future. Otherwise, the decision should be avoided, and the arc causing it should weaken.

The bottom layer in each behavior is used to *scale* the out-




Figure 3: The Neural Network that controls food-gathering agents. Rounded rectangles in the figure correspond to **layers** of neurons. Arrows are **links** between the layers. Dotted lines are **plastic links**. Links targeting other links are **modulatory links**, regulating their plasticity. The link ending in a square signifies a negative modulatory activity, whereas the other modulators give positive reinforcement. The plastic links are all initialized with positive values, to facilitate initial exploration.

put activity of that behavior, so that reinforcing modulation on each level is of the same magnitude. Without this scaling, it would be more difficult to compare learning rates across layers, because layers with a very strong modulatory input could learn with a very low learning rate.

### Learning

Arcs in the network are updated by the following learning rule:

$$\Delta w_{ij} = \eta * mod * |x_i x_j| \tag{1}$$

where  $\eta$  is the learning rate, *mod* is the strength of incoming neuromodulation and  $x_i x_j$  is the product of pre-synaptic and post-synaptic activity, in other words a regular Hebbian update term.

As the equation shows, it is the *absolute value* of the hebbian update that is used in the calculation of the new weight value, since we want the *modulatory signal* to decide the direction of the weight change: negative modulation means whatever action was taken was a bad idea, so the weight of the link causing the action should be decreased. Positive modulation should have the opposite effect. In the absence of modulation (in other words, if  $mod = 0$ ), weights are not updated.

The age-dependent learning rate is encoded as a sequence of real values ( $\bar{\eta}$ ) in the genome. One such sequence is evolved for each plastic link in the network. Each value in the sequence describes the *change* in learning rate for the current age. The first value in the sequence encodes the *initial* learning rate for the link. A separate parameter decides the *delta age*,  $\delta_{age}$ , of the link. This parameter tells the algorithm how often it should update its current learning rate. For instance,  $\delta_{age} = 5$  would mean the rate is updated every

Parameter	Value
Generations	100
Adults	30
Children	50
Crossover probability	0.1
Mutation probability	0.005
Genes per individual	75
Bits per gene	8
Elite fraction	0.2
Culling fraction	0.25

Table 1: Parameters of the Evolutionary Algorithm

fifth timestep, giving a total of 20 updates for a 100-timestep run. A reasonable setting of this parameter allows us to reduce the complexity in finding a good mapping from age to plasticity. For the experiments reported here,  $\delta_{age} = 4$ .

For the environment and network presented here, it was found that a range of numbers from -0.25 to 0.25 was a good selection for the available values for  $\bar{\eta}$ . This gave the evolutionary algorithm the opportunity to tune the learning rate quite finely, as well as having the possibility to move fairly rapidly towards high or low rates when needed.

Negative learning rates are not allowed. If the learning rate goes into negative values, it is treated as a learning rate of zero. Still, the negative value of the learning rate is remembered when calculating the learning rate for the next age. This makes it possible for evolution to drive the learning rate far into the negative domain, meaning it can effectively “shut off” learning by adjusting the learning rate. This way, the encoding allows evolution to generate sensitive periods that are quite resistant to disruptive mutations.

### Evolutionary Parameters

Table 1 gives the parameters of the evolutionary algorithm for the experiments reported here. Crossover probability gives the probability of crossover *per individual*, and mutation probability gives the probability of mutation *per bit* in the individuals’ genotype.

The same parameters were used for all runs, except for the case where *static* learning rates were evolved. For these runs, we needed only *three* genes to specify the learning strategy followed by an individual (one gene per learning link). Those three genes were each encoded by 20 bits, to allow evolution to fine-tune the static learning rate. The remaining parameters were not altered.

## Results

### Learning costs and sensitive periods

Figure 4 shows average evolved age-dependent plasticities for the three different behaviors in the task when there is no cost of learning. As expected, no sensitive periods are




Figure 4: Evolved plasticities with no cost of learning. The figure displays the plasticity in the three links of the network as a function of the agent’s age. – Averages over 50 runs. Error bars show one standard deviation.

formed, as there is no pressure to form them imposed by the evolutionary algorithm.

A cost of learning on the fitness of individuals is documented in biology (see for instance Mery and Kawecki (2003) for an example from the common fruit fly). We simulate the costs associated with plasticity to see how this affects the evolution of age-dependent plasticity. The cost of learning is implemented as a term subtracted from an individual’s total fitness. It is proportional to the sum of areas under the agent’s three plasticity graphs.

Figure 5 shows the evolved plasticities when adding this cost to fitness evaluations. As the figure shows, we see cascading sensitive periods moving from lower to higher levels of behavior. Averaging over 50 individuals smooths the sensitive periods, meaning the cascades show more overlap than what is normally present in individuals. We will see later in this section that the evolved plasticities within individuals do show the expected ordering of learning.

Another reason for the large degree of overlap is that the individual behaviors are very simple to learn: they can be learned by a single observation of the relevant association. Had each behavior been more complex, taking more time to learn, it is expected that sensitive periods would be further separated in time.

**Sensory deprivation** A common way of studying sensitive and critical periods, is to subject individuals to sensory deprivation up until a certain age, and study their development after this age. If they are unable to follow a regular path of development, we have an indication that the plasticity of the considered system underwent a critical period before that age. In our deprivation study, we subjected the winner individual of each evolutionary run to sensory deprivation on all three levels of behavior, and measured its performance on the considered behavior. For instance, to test

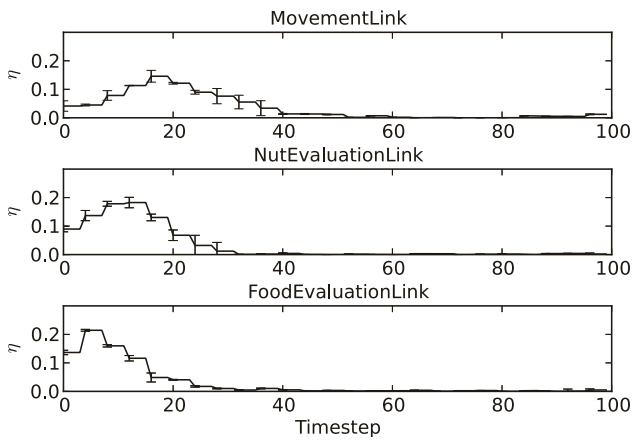


Figure 5: Evolved plasticities with a cost of learning – Averages over 50 runs. Error bars show one standard deviation

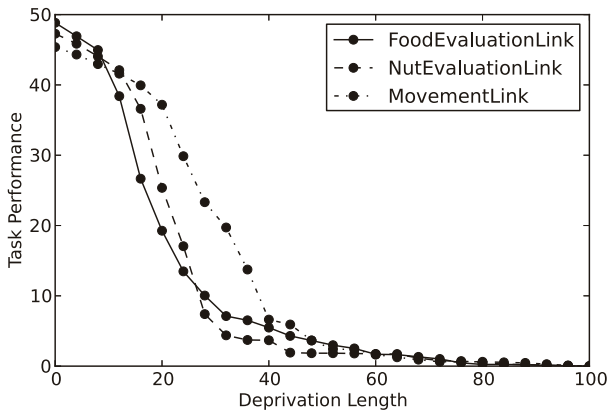


Figure 6: The result of sensory deprivation. – Averages over 50 runs. Dots show measured values, lines interpolate between them.

whether there was a sensitive period for learning about nuts, we waited for a given time before presenting the individual to nut stimuli, and measured how well it was able to learn the nut association task from this age.

Task performance was measured as the amount of correct associations made during the rest of the individuals’ life (after finishing sensory deprivation), reduced by the amount of incorrect associations made. An inverse relationship between deprivation length and task performance was observed for individuals with a constant learning rate throughout life, because for longer deprivation lengths, individuals had less time to accumulate performance points.

For individuals allowed to evolve age-dependent plasticity, the sensory deprivation resulted in the task performance shown in Figure 6. All behaviors approach zero for sensory deprivation above a certain age, indicating that learning the associated behavior was not possible after this age. Further, we see the cascading of learning observed in Figure 5 af-

fect the timing of the age when sensory deprivation disrupts further learning.

The fact that performance seems to fall off *gradually* with increasing durations of sensory deprivation is due to differences in the evolved timing of sensitive periods between individuals. Studying single individuals, we observed that the cut off is much more dramatic. In other words, the term *critical period* describes this learning better than *sensitive period*. The reason is of course that *age* is the factor controlling plasticity, so delaying stimuli beyond a certain age will prevent all learning.

### Learning Order

This section presents statistics about the *ordering* of sensitive periods. We evaluated 50 individuals, noting how many of these formed the *expected order* of sensitive periods. The order of sensitive periods was measured by calculating the *center of mass* for plasticity in each behavior, and ordering them in order of ascending center of mass. An early center of mass for a behavior means a lot of learning is going on early in the life of the agent, corresponding to an early sensitive period.

The order of the centers of mass was then compared to the ordering that would be generated by random processes. For three different behaviors, 6 different orderings are possible, and the expected number of times each would be seen is  $1/6 * 50 = 8.33$ . For a statistically significant indication that the expected ordering is preferred, we need to see the expected order of sensitive periods at least 14 times, which is associated with a p-value lower than 0.04. For comparisons between *pairs* of behaviors, we need to see the expected order of sensitive periods at least 32 times, again indicating a p-value lower than 0.04.

Table 2 shows statistics from runs of the experiment with and without a learning cost. The table shows that there is *no significant evidence* that plasticity schedules learning in the expected order, when there is *no cost of plasticity*. However, when individuals are evolved *with* a plasticity cost, they are forced to being more cost-effective, thus scheduling their learning.

**Shuffling** After observing that sensitive periods tended to form in an order from “lower” to “higher” levels of behavior, we wanted to test just how important this ordering was. To do this, we shuffled sensitive periods in each of the winner individuals from our 50 runs of evolution. This was done by extracting evolved learning strategies from one layer, and inserting it into another within the same individual. The resulting individuals were each tested on 50 new mazes, and their average fitness was stored as an indication of how well they were able to learn preferences in these mazes. Notice that shuffling does not affect the learning costs of an individual, so a difference in fitness indicates a difference in learning performance.

Learning Cost	p(F-N-M)	p(F-N)	p(N-M)	p(F-M)
Yes	< <b>0.001</b>	<b>0.008</b>	< <b>0.001</b>	< <b>0.001</b>
No	0.317	0.839	0.556	0.336

Table 2: The ordering of sensitive periods. Letters F, N and M are used to indicate the different learned associations: **F**ood, **N**uts and **M**ovement associations. The following columns show the probability of observing the evolved ordering of sensitive periods by chance, first for all three behaviors, then their pairwise orderings. Results in bold signify statistically significant evidence that evolution prefers to arrange sensitive periods in the expected order. All results were obtained by doing 50 runs of the EA.

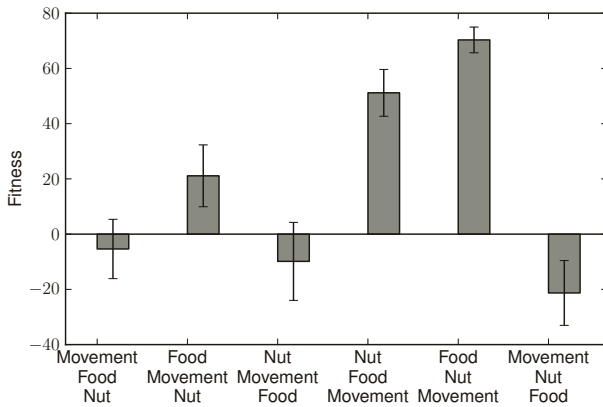


Figure 7: How learning is affected by shuffling sensitive periods. – Averages over 50 runs. Error bars show a 95% confidence interval of the mean.

Figure 7 shows the result of the shuffling of sensitive periods. The labels below axes indicate what new ordering of learning the shuffling corresponds to. (Food - Nut - Movement) was the original ordering, so this label indicates no shuffling. We see that increased shuffling gives a decrease in learning ability (seen as a decrease in fitness). Shuffling only two of the evolved sensitive periods, gives less of a fitness decrease than shuffling all three. The lowest fitness is observed when the sensitive periods are completely reversed, indicating that a proper sequencing is essential for learning in the winner individuals.

### The utility of age-dependent plasticity

So far, we have seen how age-dependent plasticity allows the formation of sensitive periods. In this section, we will analyze exactly what is the *utility* of these periods. We will do that by comparing the fitness of evolved agents with static plasticity levels and with plasticity levels that vary throughout life.

Figure 8 shows the fitness of individuals *with* and *without* a plasticity cost. During evolution, the plasticity cost was active in both cases, meaning all individuals evolved to form as cost-effective learning strategies as possible. As the right figure shows, the agents allowed to utilize age-dependent

plasticity become more *cost effective* – their fitness is higher because they pay a lower cost. However, as can be seen in the left figure, this more efficient learning strategy does not degrade their learning ability. When not applying a cost of plasticity, the fitness is only based on what the agents can learn, and this value is not significantly different for the two types of learners.

### Conclusion

In this paper we have shown how sensitive periods in learning can emerge for a reinforcement learning task by allowing an evolutionary algorithm to tune the mapping from ages of individuals to plasticity values. The task to be solved required several levels of interdependent behavior to be learned, and a different age-plasticity mapping was evolved for each level. The evolved sensitive periods showed the ability to *sequence* learning in a bottom-up fashion, allowing the network to learn the simpler behaviors first, before learning the higher-level behaviors that depend upon them.

A condition for observing the sensitive periods was that plasticity had an associated *cost*. This cost made evolution form solutions that learn in a cost-effective way. It is the balance between having the ability to learn while paying as low a cost as possible, that drove evolution to finding *cascading sensitive periods*.

In summary, we have seen how evolved sensitive periods can *sequence* the learning of sub-behaviors, and that the factor that drives this sequencing is achieving a good balance between the costs and benefits of learning. Currently, we are working on extending this model to account for experience as a regulator of plasticity, and to include the possibility of evolved hard-coded preferences.

Another interesting direction for future work is to investigate critical periods in more complicated tasks, perhaps tasks that cannot be solved with a constant learning rate. In other words, tasks that would normally be approached with a different technique such as incremental learning. We believe sensitive periods could be useful in such tasks, because they offer a way of scheduling learning activities between different behaviors.

Finally, the complexity of individual behaviors in the current study is quite low. It would be interesting to see how sensitive periods are affected by scaling up the complexity

Downloaded from http://direct.mit.edu/jaal/proceedings-pdf/ecal2013/25/409/1901241/978-0-262-31709-2-ch059.pdf by guest on 30 March 2023




Figure 8: Box plots of fitness values calculated with and without plasticity cost. Nonoverlapping box notches indicate significantly different means with a 0.05 significance level. – Averages over 50 runs.

of behaviors, and also how they are affected by tasks where the progress of learning each individual behavior is less predictable.

### References

- Armstrong, V. L., Brunet, P. M., He, C., Nishimura, M., Poole, H. L., and Spector, F. J. (2006). What Is so Critical?: A Commentary on the Reexamination of Critical Periods. *Developmental Psychobiology*, 48(May):337–44.
- Bullinaria, J. (2009). Lifetime learning as a factor in life history evolution. *Artificial Life*, 15(4):389–409.
- Bullinaria, J. A. (2003). From biological models to the evolution of robot control systems. *Philosophical transactions. Series A, Mathematical, physical, and engineering sciences*, 361(1811):2145–64.
- Burghardt, G. M. and Hess, E. H. (1966). Food imprinting in the snapping turtle, *Chelydra serpentina*. *Science*, 151(7):108–109.
- Cashdan, E. (1994). A sensitive period for learning about food. *Human Nature*, 5(3):279–291.
- Darmaillacq, A.-S., Chichery, R., and Dickel, L. (2006). Food imprinting, new evidence from the cuttlefish *Sepia officinalis*. *Biology letters*, 2(3):345–7.
- Downing, K. L. (2010). A Script-Based Approach to Evolving Neural Networks. In *Proceedings of the second Norwegian Artificial Intelligence Symposium*, pages 29–36.
- Fellous, J.-M. and Linster, C. (1998). Computational models of neuromodulation. *Neural computation*, 10(4):771–805.
- Harwerth, R. S., Smith, E. L., Duncan, G. C., Crawford, M. L. J., and von Noorden, G. K. (1986). Multiple Sensitive Periods in the Development of the Primate Visual System. *Science*, 232(4747):235–238.
- Hensch, T. K. (2005). Critical period plasticity in local cortical circuits. *Nature Reviews Neuroscience*, 6(11):877–888.
- Hubel, D. H. and Wiesel, T. N. (1970). The period of susceptibility to the physiological effects of unilateral eye closure in kittens. *Journal of Physiology*, 206(2):419–436.
- Hurford, J. R. (1991). The evolution of the critical period for language acquisition. *Cognition*, 40(3):159–201.
- Johnson, M. H. (2005). Sensitive periods in functional brain development: problems and prospects. *Developmental psychobiology*, 46(3):287–92.
- Kirby, S. and Hurford, J. R. (1997). The evolution of incremental learning: language, development and critical periods. Technical report, Language Evolution and Computation Research Unit, University of Edinburgh.
- Knudsen, E. I. (2004). Sensitive Periods in the Development of the Brain and Behavior. *Journal of Cognitive Neuroscience*, 16(8):1412–25.
- Lewis, T. L. and Maurer, D. (2005). Multiple sensitive periods in human visual development: evidence from visually deprived children. *Developmental psychobiology*, 46(3):163–183.
- Mery, F. and Kawecki, T. J. (2003). A fitness cost of learning ability in *Drosophila melanogaster*. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 270(1532):2465 – 2469.
- Punzo, F. (2002). Food imprinting and subsequent prey preference in the lynx spider, *Oxyopes salticus* (Araneae: Oxyopidae). *Behavioural Processes*, 58(3):177–181.
- Soltoggio, A., Bullinaria, J. A., Mattiussi, C., Dürr, P., and Floreano, D. (2008). Evolutionary Advantages of Neuromodulated Plasticity in Dynamic, Reward-based Scenarios. In *Artificial Life XI*, volume 11, pages 569–576.
- Trotter, Y., Gary-Bobo, E., and Buisseret, P. (1981). Recovery of orientation selectivity in kitten primary visual cortex is slowed down by bilateral section of ophthalmic trigeminal afferents. *Developmental Brain Research*, 1(3):454–450.
- Werker, J. F. and Tees, R. C. (2006). Speech perception as a window for understanding plasticity and commitment in language systems of the brain. *Developmental Psychobiology*, 46(3):233–51.