

On the influence of lifetime learning on the selection pressure

Ingo Paenke, Tadeusz Kawecki, Bernhard Sendhoff

2006

Preprint:

This is an accepted article published in Proceedings of the ALife X: The Tenth International Conference on the Simulation and Synthesis of Living Systems. The final authenticated version is available online at: [https://doi.org/\[DOI not available\]](https://doi.org/[DOI not available])

On the Influence of Lifetime Learning on Selection Pressure

Ingo Paenke^{1,3}, Tadeusz Kawecki² and Bernhard Sendhoff³

¹Institute AIFB, University of Karlsruhe, D-76128 Karlsruhe, Germany, ingo.paenke@aifb.uni-karlsruhe.de

²Unit for Ecology and Evolution, University of Fribourg, CH-1700 Fribourg, Switzerland, tadeusz.kawecki@unifr.ch

³Honda Research Institute Europe GmbH, D-63073 Offenbach, Germany, bernhard.sendhoff@honda-ri.de

Abstract

Evidence for both, acceleration and deceleration of evolution through learning can be found in the literature. We suggest a selection gradient model that allows to predict whether acceleration or deceleration is predominant. The main idea is that learning alters the genotype-to-fitness landscape, which determines selection pressure. Assuming that fitness can be split up into an innate and a learned component, conditions for the occurrence of the Baldwin and the Hiding effect are derived. We introduce learning curves to analyse what we term lifetime fitness. The influence of the shape of the learning curve on the interaction between evolution and learning is analysed.

Introduction

The interaction between evolution and learning has been studied from different perspectives. Straightforwardly, *evolution influences learning* by resulting in individuals that have the ability to learn. In artificial life, this mechanism is employed e.g. in evolutionary robotics (Harvey et al., 2005) and in virtual environments (Todd and Miller, 1991; Niv et al., 2002). Furthermore, *learning also influences evolution*. Learning may guide a population toward evolutionary paths, which would not have been taken in the absence of learning (Hinton and Nowlan, 1987). Besides this exploratory influence, learning also affects the rate of evolution (Keesing and Stork, 1991) if learning alters the selection pressure. These phenomena are usually subsumed under the terms *Baldwin effect* (Baldwin, 1896; Simpson, 1953) and *Hiding effect* (Mayley, 1997).

In the next section, we briefly review the Baldwin and the Hiding effect and highlight some examples from the literature that have demonstrated either or both of these effects. Thereafter, we utilize a common method from quantitative biology, known as the *selection gradient* analysis (Lande and Arnold, 1983), to mathematically derive conditions under which selection pressure is increased (Baldwin effect) or decreased (Hiding effect). Based on this, we investigate how lifetime learning curves influence the Baldwin and the Hiding effect. We demonstrate that learning, even in a form where the learning curves are identical for all individuals in

a population, can increase or decrease the selection pressure. Furthermore, we discuss the importance of the convexity and concavity of learning curves emphasizing their influence on the selection pressure. The influence of learning curves is important for both biological and artificial systems, whenever the system's behavior is relevant throughout its lifetime, i.e., (already) while it is learning.

Baldwin vs. Hiding Effect

In the literature, the Baldwin effect has been widely discussed, but apparently there exists no clear definition of "the" Baldwin effect. As Depew states:

The Baldwin effect does not reliably refer either to a theory-neutral empirical phenomenon, or to a single hypothesis, or to an identifiable mechanism (Depew, 2003, page 5).

Seemingly, the term Baldwin effect subsumes a number of effects arising from the interdependency of phenotypic plasticity and genetic evolution.

In this paper, we use the term Baldwin effect to describe the *increase in selection pressure as an effect of learning* (unfortunately, a comprehensive review of the Baldwin effect is beyond the scope of the paper). If learning alters the fitness landscape such that only the (genetically) very good individuals of the population reproduce, selection pressure is increased and evolution is accelerated. The other extreme, learning may allow innately weak individuals to catch up with innately strong individuals, and thus selection pressure is reduced. In the latter case, learning causes a reduction of fitness differences between weak and strong individuals and evolution is decelerated (Johnston, 1982). In (Mayley, 1997) this effect is named *Hiding effect*. Figure 1 illustrates both effects.

In the artificial life and biological literature, examples can be found for both effects. In Hinton and Nowlan's computer experiments learning leads to an acceleration of evolution (Hinton and Nowlan, 1987) and an analytical treatment supports the simulation results (Fontanari and Meir, 1990). Dopazo et al. find a *halting effect* of learning in an




Figure 1: Illustration. **Baldwin effect (left):** Learning increases fitness differences; **Hiding effect (right):** Learning decreases fitness differences.

extension of the Hinton and Nowlan model (Dopazo et al., 2001). Keesing and Stork identify both an acceleration and a deceleration effect in an evolving neural network simulation (Keesing and Stork, 1991). They show that learning with a larger amount of *plasticity* leads to faster evolution than learning with less plasticity. Furthermore, both forms of learning (large and small amount of plasticity) can accelerate evolution, but they also demonstrate that *too much* learning can slow down the rate of evolution. Similarly, French and Messinger show empirically that the strength of the Baldwin effect depends on the level of phenotypic plasticity (French and Messinger, 1994). In (Papaj, 1994) a simple simulation model and a neural network model are used to show how learning inhibits and how learning facilitates evolutionary change. In 2004, Mery and Kawecki have shown that both acceleration and deceleration effects can even be found in *real* biological experiment with fruit flies (Mery and Kawecki, 2004).

A Selection Gradient Analysis

In the following, we use the *selection gradient* (Lande and Arnold, 1983) as a method to analytically treat the Baldwin and Hiding effect. The selection gradient measures the strength of selection, and can be approximated by

$$\beta = \frac{1}{f} \frac{\partial f}{\partial x}, \quad (1)$$

where $f(x)$ is the *fitness function*, i.e., the mapping from genotype x to fitness. The selection gradient describes the relative fitness change caused by a marginal change in the genotype and is therefore a measure for the *selection pressure*. In evolutionary computation, the corresponding selection method is the *fitness proportional selection*.

In the following analysis x is treated as a one-dimensional real number and can be interpreted as one particular dimension of the complex high-dimensional genotype space as found in biological and artificial systems. The effect of learning on the rate of evolution w.r.t. the entire gene set, can be obtained straightforwardly by combining the effects of all individual dimensions.

The basic idea underlying our analysis is to compare how the fitness landscape is altered through learning, and how this influences the selection gradient. We compare a learning population and a non-learning one, and derive conditions for acceleration (Baldwin effect) and deceleration (Hiding effect) of the rate of evolution. In the absence of learning, we call the resulting landscape *innate landscape* f and in case of learning simply *learning landscape* f_l . We assume both landscapes f and f_l to be positive and monotonically increasing. The influence of learning on evolution is given by

$$\beta_l(x) - \beta(x) \begin{cases} > 0 \Rightarrow \text{acceleration (Baldwin effect)} \\ < 0 \Rightarrow \text{deceleration (Hiding effect)} \\ = 0 \Rightarrow \text{no effect,} \end{cases} \quad (2)$$

where β_l denotes the selection gradient in case of learning and β the one in the absence of learning. We reformulate

$$\begin{aligned} \beta_l(x) - \beta(x) &= \frac{f'_l(x)}{f_l(x)} - \frac{f'(x)}{f(x)} \\ &= \frac{f'_l(x)f(x) - f'(x)f_l(x)}{f_l(x)f(x)} \\ &= \frac{f(x)}{f_l(x)} \frac{f'_l(x)f(x) - f'(x)f_l(x)}{f^2(x)} \\ &= \frac{f(x)}{f_l(x)} \left(\frac{f_l(x)}{f(x)} \right)' = \frac{g'(x)}{g(x)}, \end{aligned} \quad (3)$$

with

$$g(x) = \frac{f_l(x)}{f(x)}. \quad (4)$$

We call g the *gain function* of learning. With sgn denoting the sign function, Equation (3) implies

$$\text{sgn}(\beta_l(x) - \beta(x)) = \text{sgn}(g'(x)), \quad (5)$$

under the assumption of positive $f(x)$, $f_l(x)$. If $g'(x) > 0$ evolution is accelerated, and if $g'(x) < 0$ evolution is decelerated. In Figure 2, a scenario with a negative gain func-




Figure 2: Learning alters the fitness landscape: An example of the Hiding effect.

tion gradient is shown. Here the landscape is smoothed through learning, i.e., the Hiding effect dominates. In other scenarios the opposite effect may appear, i.e., learning alters the fitness landscape such that the Baldwin effect dominates.

In the following, we assume that the fitness $f_l(x)$ of a learning individual x is composed of its innate fitness $f(x)$ and a learning component $l(x)$:

$$f_l(x) = f(x) + l(x). \quad (6)$$

Using the selection gradient approach, we derive conditions for the occurrence of the Baldwin respectively the Hiding effect (recall that we assume $f(x) > 0, f'(x) > 0$):

$$\begin{aligned} g'(x) &= \left(\frac{f(x) + l(x)}{f(x)} \right)' \\ &= \left(1 + \frac{l(x)}{f(x)} \right)' \\ &= \frac{l'(x)f(x) - l(x)f'(x)}{f^2(x)} \\ &= \frac{l(x)}{f(x)} \left(\frac{l'(x)}{l(x)} - \frac{f'(x)}{f(x)} \right) \\ &= \frac{l(x)}{f(x)} (\log(l(x)) - \log(f(x)))'. \end{aligned} \quad (7)$$

If $l(x) > 0$ and $f(x) > 0$, then

$$\text{sgn}(g'(x)) = \text{sgn}((\log(l(x)))' - (\log(f(x)))'). \quad (8)$$

We see that the *combination* of $l(x)$ and $f(x)$ determines whether evolution is accelerated or decelerated. In case of positive functions $l(x), f(x)$, the result of Equation (8) can also be interpreted in terms of the selection gradient terminology:

Since $(\log(l(x)))' = \frac{l'(x)}{l(x)}$, the term $(\log(l(x)))'$ can be interpreted as the selection gradient of $l(x)$. Thus, if the selection gradient of $l(x)$ is larger than the selection gradient of $f(x)$, evolution is accelerated, and in the opposite case, evolution is decelerated. If $f(x) = l(x)$, learning has no influence on evolution. In the following, we define categories of functions $l(x)$ and derive general conclusions for functions of each category:

Positive, decreasing $l(x)$: A positive, decreasing $l(x)$ implies that weak individuals benefit more from learning than strong individuals. Since

$$\begin{aligned} l(x) > 0 \wedge l'(x) < 0 &\Rightarrow (\log(l(x)))' < 0 \\ f(x) > 0 \wedge f'(x) > 0 &\Rightarrow (\log(f(x)))' > 0, \end{aligned} \quad (9)$$

and using Equation (8), we obtain $g'(x) < 0$. Therefore, for all scenarios with positive, decreasing function $l(x)$, the Hiding effect dominates and evolution is decelerated.

Constant $l(x)$: Next, we consider the case when learning causes a constant fitness change, i.e., $l(x) = a$. Using Equation (7), we obtain

$$\text{sgn}(g'(x)) = \text{sgn}\left(-\frac{a}{f(x)}(\log(f(x)))'\right) = \text{sgn}(-a). \quad (10)$$

Therefore, in case of constant fitness *increase* (positive a), evolution is decelerated (Hiding effect) while for a constant fitness *decrease* (negative a) fitness is accelerated (Baldwin effect). This might at first sight be counterintuitive, but recall that relative fitness changes count. A constant fitness increase implies a larger *relative* fitness gain for a weak individual (with small innate fitness) than for a strong individual (with large fitness), and vice versa.

Positive, increasing $l(x)$: Finally, we consider the case of positive, increasing $l(x)$. For such functions, strong individuals always benefit more from learning than weak individuals (in terms of absolute fitness gain). Unfortunately, no simpler formulation than Equation (8) can be derived for this case, without specifying either $l(x)$ or $f(x)$. However, an example illustrates that functions of this category can both accelerate or decelerate evolution: If $l(x) = x^b$ and $f(x) = x^c$, then $\text{sgn}(b - c)$ determines whether evolution is accelerated ($b - c > 0$) or decelerated ($b - c < 0$).

The Influence of Learning Curves - An Extension of the Fitness Landscape Model

In the literature, the concept of *fitness landscape* has “many different meanings and associated mathematical representations” (Stadler and Stephens, 2003, page 390). To model evolutionary selection, an individual’s fitness has to represent the number of its offsprings. However, a proper definition has to take the stochastic nature of the evolutionary process into account, thus a more favorable notion would be the *average number of offsprings*. This definition of fitness was also assumed in the selection gradient analysis of the previous section.

So far, we have considered the mapping from genotype to fitness as a black box and have not yet discussed how lifetime fitness (the average number of offsprings) is actually attained. For the selection gradient analysis, there has been no need for this discussion, because so far we have only been interested in the result of learning w.r.t. lifetime fitness. In (Mayley, 1996) this type of fitness assessment is named *posthumous fitness assessment*.

However, in the following, we explicitly introduce *learning curves* as an important influence factor of lifetime fitness. Learning curves describe the progress of an individual’s effort to improve its fitness during lifetime. This applies to many biological and artificial systems. We term this type of fitness assessment *lifetime fitness assessment*. In this case, an individual is evaluated throughout its lifetime including the learning period. Thus, not only the *re-*

sult of learning is relevant but also the *learning curve*. For example, a population of “early learners” might influence evolution and in particular selection pressure in a different way than “late learners”, although the “result” of learning is identical. In the following, we investigate the influence of learning curves on the Baldwin and Hiding effect for lifetime assessment, compared to the case of posthumous fitness assessment.

Extension of the Fitness Landscape Model The basic idea is to extend the traditional fitness landscape model (Figure 2) by a lifetime learning dimension, and use this model to compare environments with lifetime learning (learning curves), posthumous assessment and without learning.

We define

$$(x, t) \mapsto f_t(x, t), \quad (11)$$

where $f_t(x, t)$ describes some kind of “partial” lifetime fitness of an individual with genotype x at time t . The overall fitness is the accumulation of an individual’s partial fitness values over its lifetime $t \in [0; T]$. An individual’s lifetime fitness F_l is then given by

$$F_l(x) = \int_0^T f_t(x, t) dt. \quad (12)$$

At first sight the notion of partial fitness $f_t(x, t)$ might be counterintuitive, because an individual either reproduces at time t or it does not. However, we want to interpret Equation (11) as the average offspring generation of an *average* individual with certain phenotypic characteristics at time t . The genotype directly encodes the innate phenotype of an individual with corresponding partial fitness at time $t = 0$. Without loss of generality, we set the maximal lifetime, $T = 1$. An example of such a landscape is given in Figure 3. In the absence of learning the partial fitness is constant (in t -direction). In this case, the lifetime fitness $F(x)$ is given by the dark-gray area, which we call *non-learning area*. In case of learning, lifetime fitness $F_l(x)$ of an individual x is given by the *learning area*, which is the sum of the non-learning area and the light-gray area. In Figure 3 all learning curves




Figure 3: Lifetime fitness assessment model.

are linear. Posthumous assessment can also be visualized in the landscape extension as shown in Figure 3. Since learning curves are not taken into account in this case, the maximum partial fitness is achieved immediately after birth. In the figure, the light-gray (learning curve) triangle becomes a rectangle. The fitness in case of lifetime assessment can be related to posthumous assessment as follows

$$F_l(x) = \int_0^1 (f(x) + l_t(x, t)) dt = f(x) + \int_0^1 l_t(x, t) dt, \quad (13)$$

$$F(x) = f(x), \quad (14)$$

where $f(x)$ equals the dark-gray area of Figure 3, and $\int_0^1 l_t(x, t) dt$ equals the light-gray area. $l_t(x, t)$ describes the shape of the learning curve for a given x with

$$l_t(x, 0) = 0 \wedge l_t(x, 1) = f_l(x) - f(x). \quad (15)$$

Identical Learning Curves for All Genotypes Firstly, we consider the case, where for all genotype values the “shape” of the corresponding learning curves are identical. The shape is defined by an arbitrary function $\hat{l}_t(t)$ with $\hat{l}_t(0) = 0$ and $\hat{l}_t(1) \neq 0$, so that

$$l_t(x, t) = \alpha(x) \hat{l}_t(t), \quad \alpha(x) = \frac{f_l(x) - f(x)}{\hat{l}_t(1)} \quad (16)$$

is a valid learning function w.r.t. the conditions defined in Equation (15). Since $\hat{l}_t(t)$ is independent of x , the learning function $l_t(x, t)$ has the same shape for all individuals. After reformulation of

$$\begin{aligned} F_l(x) &= f(x) + \alpha(x) \int_0^1 \hat{l}_t(t) dt \\ &= f(x) + (f_l(x) - f(x)) \int_0^1 \frac{\hat{l}_t(t)}{\hat{l}_t(1)} dt, \end{aligned} \quad (17)$$

we obtain the gain function $G(x)$ for lifetime assessment

$$G(x) = \frac{F_l(x)}{F(x)} = 1 + (g(x) - 1) \int_0^1 \frac{\hat{l}_t(t)}{\hat{l}_t(1)} dt, \quad (18)$$

where $g(x)$ is the gain function for posthumous assessment (cf. Equation 4) and the gradient $G'(x)$ as

$$G'(x) = g'(x) \int_0^1 \frac{\hat{l}_t(t)}{\hat{l}_t(1)} dt. \quad (19)$$

This shows, that w.r.t. the sign of the gain function gradient, the lifetime assessment approach leads to the same result as the posthumous assessment approach:

$$\text{sgn}(G'(x)) = \text{sgn}(g'(x)), \quad (20)$$

if the integral over the normalized learning curve of $\hat{l}_t(t)$ is positive (which can indeed be assumed since otherwise

learning could be detrimental in terms of lifetime fitness for the individual). However, the shape of the learning curve might have a strong influence on the magnitude of the acceleration/deceleration effect.

The following short simulation study provides empirical evidence to confirm the results of the theoretical analysis (cf. Equation 19). We simulate evolution with a standard evolutionary algorithm with asexual populations of 100 individuals. The genotype x is a 1-dimensional real number confined to $[0; 1]$. Mutation is realized by adding a $N(\mu = 0, \sigma = 10^{-4})$ normally distributed random number to the genotypic value, and is clipped at the boundaries of the genotype interval. Stochastic Universal Sampling (Baker, 1987) is employed as a selection method. Initially, populations are randomly distributed on $[0; 0.1]$.

In case of posthumous assessment, an environment is defined completely by (a) a fitness function $f(x)$ that in the absence of learning maps the genotype value to fitness, and (b) a learning function $f_l(x)$ that maps a genotype value to fitness in case of learning. For lifetime assessment, additionally, (c) the learning curve $\hat{l}_i(t)$ has to be defined. For these three categories, we define sets of functions and combine them to describe the environments.

In Table 1, six lifetime assessment environments are defined, which all have the same innate fitness function $f_1(x)$. Figure 4 shows the resulting landscapes and learning functions. For comparison, the upper gray line shows the fitness assessment in case of posthumous assessment, i.e., the maximum fitness is achieved at birth ($t = 0$). The “no learning” curve at the bottom illustrates the fitness assessment in the absence of learning. For each environment, we analysed the selection gradient using equations (4,19). However, due to space limitations, we omit these calculations here. The results are shown in the last column of Table 1(b). For all combinations of innate fitness f_1 with learned fitness f_{l1} (Environments E1-E3), evolution is decelerated through learning. The magnitude of this deceleration effect is strongest in case of E1 and weaker for E2 and E3. For all combinations of innate fitness f_1 with learned fitness f_{l2} (E4-E6), evolution is accelerated through learning, and again this effect is strongest for E4 and weaker for E5 and E6.

In Figure 5, we compare the evolution with the three learning curves to the cases of posthumous assessment and “no learning” using simulations with the evolutionary algorithm introduced above.

Since the differences between the evolutionary progresses can be small, we also show the evolution of relative (percentage) genotype differences normalized to the “no learning” case. The simulation results confirm our theoretical considerations summarized in Table 1(b): Evolution can both be accelerated as well as decelerated. The effects are strongest in case of posthumous assessment, and weakened through learning curves. The more convex the learning curve is, the

Table 1: Settings of the first simulation study

(a) Functions	(b) Environments	Expected effect
$f_1(x) = x$	E1: (f_{l1}, \hat{l}_{l1})	strong decel.
$f_{l1}(x) = f_1(x) + x^{\frac{1}{2}}$	E2: (f_{l1}, \hat{l}_{l2})	medium decel.
$f_{l2}(x) = f_1(x) + x^2$	E3: (f_{l1}, \hat{l}_{l3})	weak decel.
$\hat{l}_{l1}(t) = t^{\frac{1}{4}}$	E4: (f_{l2}, \hat{l}_{l1})	strong accel.
$\hat{l}_{l2}(t) = t$	E5: (f_{l2}, \hat{l}_{l2})	medium accel.
$\hat{l}_{l3}(t) = t^4$	E6: (f_{l2}, \hat{l}_{l3})	weak accel.




Figure 4: Environments of the first simulation study. **Left:** Innate landscape f_1 and learning landscapes f_{l1}, f_{l2} for posthumous assessment. **Right:** Learning curves $\hat{l}_{l1}, \hat{l}_{l2}, \hat{l}_{l3}$.

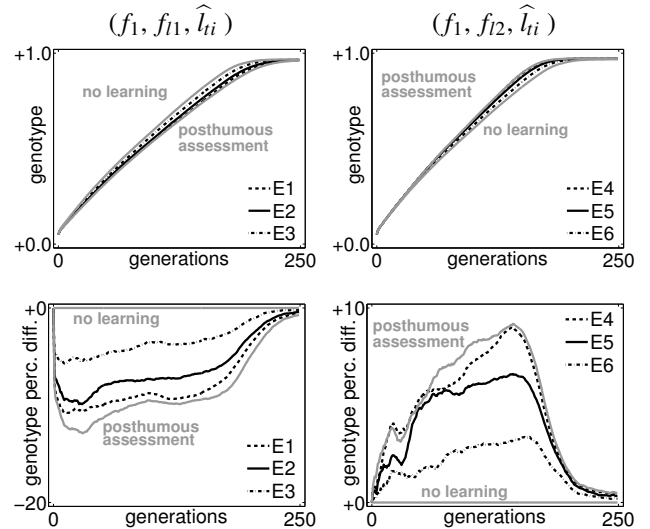


Figure 5: Results of the first simulation study (average over 100 independent runs). **Top-left:** Genotype evolution for the combinations (f_1, f_{l1}) and the learning curves $\hat{l}_{l1}, \hat{l}_{l2}, \hat{l}_{l3}$, representing environments E1-E3. **Bottom-left:** Evolution of the corresponding relative genotype differences (in percent). **Top-right:** Same as top-left but using f_{l2} instead of f_{l1} , i.e., representing environments E4-E6. **Bottom-right:** Evolution of the corresponding relative genotype differences.

weaker is the acceleration (Baldwin effect) respectively deceleration (Hiding effect).

Different Learning Curves for Different Genotypes We now consider the case where learning curves are different for different genotype values. In the notation of Equation (16), this means that the shape of the learning curve \hat{l}_t now depends on both t and x , and we get $\hat{l}_t(x, t)$. The following example illustrates how learning curves can influence the rate of evolution for this case. We choose $f_l(x) = \alpha f(x)$ so that without consideration of learning curves (i.e., posthumous fitness assessment), there is no influence of learning on evolution. As in the first simulation, we choose learning curves of the form $\hat{l}_t(t) = t^d$. However, now the learning curves are influenced by genotype x . We define

$$\hat{l}_t(x, t) = t^{e(x)}, \quad (21)$$

with $e(x) = k^{2x-1}$, $k > 0$ and $x \in [0; 1]$.

We compare two environments (cf. Figure 6), setting 1 with $(\alpha = 3, k = 0.25)$ and setting 2 with $(\alpha = 3, k = 4.0)$. In setting 1, learning curves are convex for small genotypic values x and concave for large x , and vice versa in setting 2. The results are shown in Figure 6. In setting 1, learning accelerates evolution and in setting 2, learning decelerates evolution. A detailed analysis of the selection gradient can not be included here, however, qualitatively the results can be understood as follows. In general, an individual with concave learning curve gains more from learning than an individual with convex learning curve. Therefore, if genotype values of innately strong individuals are associated with *concave*

learning curves, and genotype values of innately weak individuals with *convex* learning curves, the strong ones gain proportionally more from learning, thus the Baldwin effect dominates (cf. Figure 6, setting 1). In the opposite case, the Hiding effect dominates (cf. Figure 6, setting 2).

Conclusion

In the literature, evidence can be found that learning can both speed up and slow down the rate of evolution. In this paper, we employed a commonly used approach – analysis of the selection gradient – to get a better understanding of the conditions that lead to a positive or negative influence of learning on evolutionary progress. We evaluate how learning alters the genotype-to-fitness landscape, and how this influences selection pressure. Under the assumption that fitness can be split up into an innate and a learned part, conditions for the occurrence of the Baldwin and the Hiding effect have been derived. In the second part of the paper, we introduced the concept of the learning curve and analysed its influence on the rate of evolution. It turns out that in the case when the learning curves are identical for all genotypes, the *result* of lifetime learning determines whether the Baldwin or the Hiding effect appears. However, the *shape* of the learning curve influences the *magnitude* of these effects. In case of different learning curves for different genotypes, not only the magnitude is influenced, but also the kind of effect that can be observed. Whether evolution is accelerated or decelerated depends on how the learning curve varies with the genotype value.

The selection gradient approach allows to quantify the Baldwin and the Hiding effect and has the potential to explain (at least partially) results from the related literature. At the same time, we have to keep in mind that the effect of learning on the selection pressure is only one type of influence that learning can have on evolution.

Acknowledgments

I. Paenke wants to thank H. Schmeck, J. Branke and Y. Jin for their support.

References

- Baker, J. (1987). Reducing bias and inefficiency in the selection algorithm. In *Proceedings of the Second International Conference on Genetic Algorithms and their Application*, pages 14–21.
- Baldwin, J. (1896). A new factor in evolution. *American Naturalist*, 30:441–451.
- Depew, D. (2003). Baldwin and his many effects. In Weber, B. and Depew, D., editors, *Evolution and Learning*, pages 3–31.
- Dopazo, H., Gordon, M., Perazzo, R., and Risau-Gusman, S. (2001). A model for the interaction of learning and




Figure 6: **Top-row:** Partial fitness landscapes for the second simulation study. Genotype x encodes the innate phenotype, which is modified over lifetime t . **Bottom-row:** Simulated average genotype evolution for settings 1 and 2 compared to the case of no learning (average over 100 independent runs). **Bottom-left:** Plain representation of genotype evolution. **Bottom-right:** Percentage genotype difference between setting 1 respectively setting 2 and average genotype evolution in case of no learning.

- evolution. *Bulletin of Mathematical Biology*, 63:117–134.
- Fontanari, J. and Meir, F. (1990). The effect of learning on the evolution of asexual populations. *Complex Systems*, 4:401–414.
- French, R. and Messinger, A. (1994). Genes, phenes and the Baldwin effect. In *Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*, pages 277–282.
- Harvey, I., Di Paolo, E., Wood, R., Quinn, M., and Tuci, E. (2005). Evolutionary robotics: A new scientific tool for studying cognition. *Artificial Life*, 11(1-2):79–98.
- Hinton, G. and Nowlan, S. (1987). How learning can guide evolution. *Complex Systems*, 1:495–502.
- Johnston, T. (1982). Selective costs and benefits in the evolution of learning. *Advances in the study of behavior*, 12.
- Keesing, R. and Stork, D. (1991). Evolution and learning in neural networks: The number and distribution of learning trials affect the rate of evolution. In *Proceedings of Neural Information Processing Systems (NIPS-3)*, pages 804–810.
- Lande, R. and Arnold, S. (1983). The measurement of selection on correlated characters. *Evolution*, 37:1210–1226.
- Mayley, G. (1996). The evolutionary cost of learning. In *From Animals to Animats: From Animals to Animats: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, pages 458–467.
- Mayley, G. (1997). Guiding or hiding: Explorations into the effects of learning on the rate of evolution. In *Fourth European Conference on Artificial Life (ECAL'97)*, pages 135–144.
- Mery, F. and Kawecki, T. (2004). The effect of learning on experimental evolution of resource preference in *Drosophila melanogaster*. *Evolution*, 58(4):757–767.
- Niv, Y., Joel, D., Meilijson, I., and Ruppin, E. (2002). Evolution of reinforcement learning in uncertain environments: A simple explanation for complex foraging behaviors. *Adaptive Behavior*, 10(1):5–24.
- Papaj, D. (1994). Optimizing learning and its effect on evolutionary change in behavior. In *Behavioral Mechanisms in Evolutionary Ecology*, pages 133–154.
- Simpson, G. (1953). The Baldwin effect. *Evolution*, 7:110–117.
- Stadler, P. and Stephens, C. (2003). Landscapes and effective fitness. *Comments on Theoretical Biology*, 8:389–431.
- Todd, P. and Miller, G. (1991). Exploring adaptive agency II: Simulating the evolution of associative learning. In *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior*, pages 306–315.