

The Influence of Learning on Evolution: A Mathematical Framework

Ingo Paenke^{*,**,†}
Institute AIFB
Honda Research Institute Europe
GmbH

Tadeusz J. Kawecki[‡]
University of Fribourg

Bernhard Sendhoff[†]
Honda Research Institute Europe
GmbH

Abstract The *Baldwin effect* can be observed if phenotypic learning influences the evolutionary fitness of individuals, which can in turn accelerate or decelerate evolutionary change. Evidence for both learning-induced acceleration and deceleration can be found in the literature. Although the results for both outcomes were supported by specific mathematical or simulation models, no general predictions have been achieved so far. Here we propose a general framework to predict whether evolution benefits from learning or not. It is formulated in terms of the gain function, which quantifies the proportional change of fitness due to learning depending on the genotype value. With an inductive proof we show that a positive gain-function derivative implies that learning accelerates evolution, and a negative one implies deceleration under the condition that the population is distributed on a monotonic part of the fitness landscape. We show that the gain-function framework explains the results of several specific simulation models. We also use the gain-function framework to shed some light on the results of a recent biological experiment with fruit flies.

Keywords

Evolution, learning, Baldwin effect, selection pressure, gain function

1 Introduction

Evolution and learning are two important adaptation processes for natural systems that operate on different spaces and with different time scales. Evolution is a gradual process that occurs on the genotype level from one generation to the next. Learning is a fast process that occurs on the phenotype level within the lifetime of an individual. The two processes interact in many ways. The most direct interaction, the genetic fixation of learned phenotypic characteristics, is not possible in nature. However, it has been used successfully in artificial evolutionary systems (see, e.g., [7, 15, 19]), while at the same time it has been demonstrated that this so-called Lamarckian inheritance has an adaptive disadvantage in quickly changing environments [30, 33].

The more indirect interaction can be observed in both artificial and natural systems. The Baldwin effect, which was first suggested by Baldwin [3] and which received its name from Simpson [34],

* Contact author.

** Institute AIFB, University of Karlsruhe, D-76128 Karlsruhe, Germany. E-mail: ingo.paenke@aifb.uni-karlsruhe.de

† Honda Research Institute Europe GmbH, Carl-Legien-Strasse 30, D-63073 Offenbach/Main, Germany. E-mail: bernhard.sendhoff@honda-ri.de

‡ Section of Ecology and Evolution of Fribourg, Chemin du Musee 10, CH-1700 Fribourg, Switzerland. E-mail: tadeusz.kawecki@unifr.ch

describes the influences that learning has on the evolutionary process because it changes the evolutionary fitness of individuals without the need for translating acquired characteristics back into the genome.

Whether changes of fitness due to learning accelerate or decelerate evolution cannot be predicted in general. Evidence for both learning induced acceleration [4, 6, 8, 10, 11, 13, 14, 17, 22–25] and deceleration [1, 8, 11, 12, 16–18, 20, 22, 24, 31] of evolutionary change can be found in the literature. Explanations for both effects have been based on the analysis of *in silico* experiments [4, 6, 8, 11, 13, 14, 17, 18, 20, 22, 23, 25, 31, 36] and of *in vivo* experiments [24]. Furthermore, mathematical models [1, 10, 16] have been proposed, and theoretical analyses [12, 23] carried out. Several properties have been identified that affect the interaction of learning and evolution, such as *epistasis* [22, 36], (implicit or explicit) *cost of learning* [8, 22, 36], the amount of learning [11, 17, 6], and the shape of the learning curves [28].

However, there have only been a few attempts to develop a general model to predict the influence of learning on evolution. In [5] Borenstein et al. show that several forms of learning smooth a multimodal fitness landscape in such a way that on average evolution proceeds quicker to a global optimum. Their model predicts that the expected time to reach a higher fitness value starting from a local maximum depends on the difference between the local fitness maximum and the following local fitness minimum on the evolutionary path to the global optimum. This quantity, which is named *drawdown* in [5] and *negative extent* in [26], is reduced through learning. In order to arrive at this clear and general conclusion, Borenstein et al. employed an abstract model of evolution, in which the genotype space is discrete and one-dimensional. The population is represented by one value of the genotype space (whether this value represents the average genotype value of the population or something similar is not specified), and the population's mutation-selection movement is modeled by a one-dimensional nonsymmetric random walk [35].

In this article, we will outline a general framework to study the indirect interaction between evolution and learning that will allow us to predict whether evolutionary change will benefit from the interaction or whether it will be penalized.

The framework, which is based on the definition of a gain function $g(x)$, was first introduced by the authors in [28, 29]. In order to derive the gain function in [29], we represented a population as a probability distribution and had to limit the analysis to symmetric distribution functions. Furthermore, we had to include approximations (second-order Taylor expansion) in the proof. While a continuous representation of the population and the required assumptions are common in quantitative genetics, we will show in this article that they are not needed if we introduce a different representation, which is less common in theoretical biology, but which is standard in simulated evolution and artificial life. We model the population as a set of discrete entities (individuals). No assumption on how these entities are distributed is required. The new gain function quantifies the proportional change of fitness due to learning, dependent on the genotype value x . An inductive proof that does not require any approximations shows that the influence of learning on evolution depends on the derivative of g w.r.t. x .

The gain-function framework applies to directional selection, that is, scenarios in which the population climbs up a fitness hill. It cannot be transferred to a multimodal fitness landscape; instead it makes exact short-term predictions on monotonic landscapes and may be used for approximate predictions in nearly monotonic landscapes. The drawdown model in [5] deals with multimodal fitness landscapes; however, it cannot predict the influence of learning in a unimodal fitness landscape.

In the next section, we outline the idea of the gain-function framework in detail, and then provide a mathematical analysis and proof that generally shows under which conditions learning accelerates or decelerates evolution. We apply the gain-function analysis in Section 3 to models from ALife and evolutionary biology and to a new model incorporating evolution and learning, thereby demonstrating the generality of the approach. The comparison of different models (with different results with regard to the influence of learning on evolution) within one mathematical framework highlights under which conditions learning accelerates and decelerates evolution. In Section 3.4, we also apply the gain function to data from a biological experiment that investigates the evolution of resource

preference in fruit flies. Based on this, a simulation model of the fruit fly experiment is developed in which the digital fruit flies indeed evolve a similar resource preference. The results of this article are discussed in Section 4.

2 Theoretical Analysis

The rate of evolution increases with the relative differences in fitness among different individuals. Learning can affect evolution by influencing the fitness of phenotypes that have a certain genetic predisposition; for example, learning may amplify fitness differences between strong and weak genetic predispositions. In other words, the genetically strong individuals benefit (or *gain*) more from learning than their genetically weak rivals (of course, the opposite case may occur as well). This is an example of a positive gain-function derivative, which we will introduce in this section.

In biology, the transformation from genotype to phenotype is usually enormously complex. Development and learning are parallel processes during the whole lifetime of individuals. There is no transition when one ceases and the other one starts. Nevertheless, in order to allow a simple mathematical analysis, in our model we want to distinguish between the two processes in a sequential fashion: first development (ontogenesis) and then learning (epigenesis). Genotypic information is used during development to produce an innate phenotype, which is modified through learning, resulting in the learned phenotype; see Figure 1. Since we concentrate on the influence of learning, we keep the first transition phase as simple as possible in our model (for a recent article on the developmental phase see, e.g., [9]). In particular, an individual is characterized by a real-valued genotype variable x and a real-valued phenotype variable z . As a mapping from the genotype to the innate phenotype, we assume the identity function. An individual changes its innate phenotype via a learning function l . Thus, as a result of learning, an individual's genotype value x is mapped to its phenotype z via a learning function $z = l(x)$. In the absence of learning the phenotype equals the genotype: $z = x$. The fitness of an individual is assigned using a fitness function $f(z)$, defined on the phenotype space. Thus, the fitness in the presence of learning is given by $f(l(x))$, and in the absence of learning by $f(x)$. For simplicity we will often write $f_l(x)$ for $f(l(x))$. We assume that fitness functions $f(x)$ and $f_l(x)$ are positive and monotonic (the sign of $f'(x)$ is constant) within the range of population variability.

We now consider a finite population of n individuals, where the genotype values are labeled $x_i, i = 1, \dots, n$. The rate of evolution is measured as the distance that the population's mean genotype $\bar{x} = \frac{1}{n} \sum_{i=1}^n x_i$ moves toward the optimum in one generation. We will later use simulations to

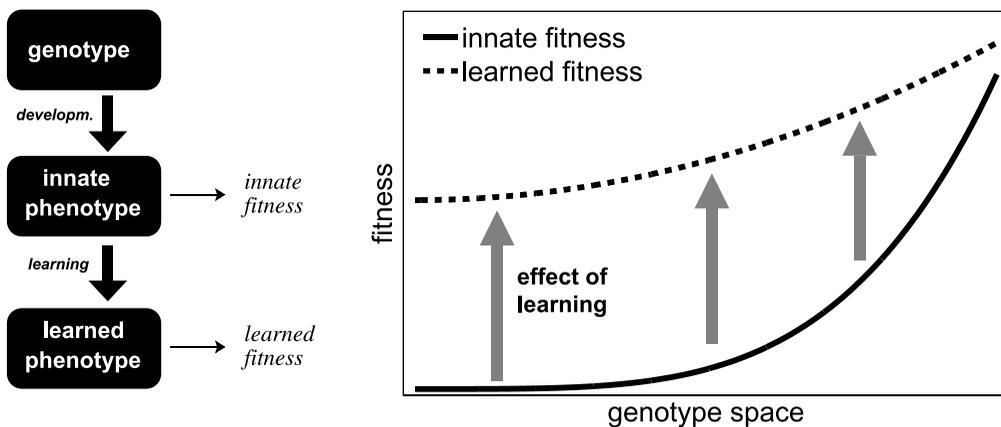


Figure 1. The basic model to analyze the influence of learning on evolution. By changing the phenotype (left), learning also changes the mapping from genotype to fitness (right).

demonstrate that this is indeed an appropriate indicator for the evolutionary velocity over the course of many generations. An individual's reproduction probability is assumed to be proportional to its fitness value. If we look at the biological concept of fitness, which corresponds to the number of offspring produced by an individual, this is the most reasonable selection model. Note, that in the field of evolutionary computation, this selection method is known as *fitness-proportional selection*. With this assumption the expected mean genotype after selection \bar{x}^* can be calculated as follows:

$$\bar{x}^* = \frac{\sum_{i=1}^n x_i f(x_i)}{\sum_{i=1}^n f(x_i)}. \quad (1)$$

Assuming an unbiased, symmetric mutation, this is equal to the mean genotype of the next generation, and the expected change of the mean genotype in one generation is given by

$$\Delta\bar{x} = \frac{\sum_{i=1}^n x_i f(x_i)}{\sum_{i=1}^n f(x_i)} - \frac{1}{n} \sum_{i=1}^n x_i. \quad (2)$$

The mean genotype change in the case of learning, $\Delta\bar{x}_l$, is derived analogously by replacing f with f_l in Equation 2. Thus, learning accelerates (decelerates) evolution if

$$\text{sign}(\Delta\bar{x}_l - \Delta\bar{x}) = \text{sign}\left(\frac{\sum_{i=1}^n x_i f_l(x_i)}{\sum_{i=1}^n f_l(x_i)} - \frac{\sum_{i=1}^n x_i f(x_i)}{\sum_{i=1}^n f(x_i)}\right) \quad (3)$$

is positive (negative). We now define the *gain function* as the quotient between the genotype-to-fitness function with learning and the genotype-to-fitness function without learning, that is,

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

Under the assumption that $g(x)$ is monotonic over the range of population variation, we show with an inductive proof (see Appendix) that

$$g'(x) \begin{cases} > 0 \Leftrightarrow \Delta\bar{x}_l - \Delta\bar{x} > 0, \\ < 0 \Leftrightarrow \Delta\bar{x}_l - \Delta\bar{x} < 0, \\ = 0 \Leftrightarrow \Delta\bar{x}_l - \Delta\bar{x} = 0. \end{cases} \quad (5)$$

Equation 5 shows that whether learning accelerates or decelerates evolution is determined by the sign of the derivative of the gain function. A positive derivative implies acceleration, a negative derivative implies deceleration, and a constant gain function implies that learning has no effect on evolution. If we find that learning has accelerated (decelerated) evolution, we know that the gain-function derivative is positive (negative), under the above given assumptions.

3 Application of the Gain-Function Framework

In this section, we apply the gain-function framework to models from ALife and evolutionary biology and to our own model coupling evolution and learning. The aim is to highlight the wide applicability of the gain-function approach and to understand under which conditions learning-induced acceleration or deceleration of evolution appears.

3.1 Hinton and Nowlan's Model: A Positive Gain-Function Derivative

The first computational model that demonstrated that (Baldwinian) learning can accelerate evolution was published by Hinton and Nowlan [14] (we will refer to it as the H&N model) 20 years ago. They demonstrated “how learning can guide evolution” toward a global optimum. We briefly summarize the H&N model: In the first scenario, the absence of learning, a genotype is given by 20 genes with alternative values (alleles) $\{0, 1\}$. A phenotype of the same structure (a 20-bit string) is produced using the identity mapping, that is, $0 \mapsto 0$ and $1 \mapsto 1$. There exists exactly one *good* phenotype, which, without loss of generality, can be set to 11111111111111111111 (all ones) with the high fitness value of 20, and all others have the same low fitness of 1. In this needle-in-haystack fitness landscape, there exists no smooth path to the all-ones genotype (equal to the phenotype), and simulated evolution fails to identify it. In the second scenario, the presence of learning, the allele set is extended to three $\{0, 1, ?\}$, and again $0 \mapsto 0$ and $1 \mapsto 1$ in the genotype-phenotype mapping. However, the phenotypic characteristic of the $?$ gene is assigned after a learning period. Learning is a sequence of random guesses, where the individual stops when it finds the all-ones phenotype, or after 1000 trials. Of course, the optimal phenotype cannot be learned by individuals that carry one or more 0 alleles. Hinton and Nowlan's simulations show that in contrast to the non-learning population, the learning population finds the global optimum.

The gain-function framework assumes that learning and non-learning individuals have a genotype of equal structure, which is not directly given in H&N, where learning individuals have the $?$ as an additional allele. We circumvent this with a formal reformulation of the H&N model: In both scenarios, presence and absence of learning, all genes have four alleles: $\{0, 1, ?_0, ?_1\}$. Alleles 0 and 1 encode the phenotype directly ($0 \mapsto 0$ and $1 \mapsto 1$). In contrast, alleles $?_0$ and $?_1$ map either to 0 or 1 after a learning period, but learning starts at 0 in the case of $?_0$ and at 1 in the case of $?_1$. The difference between learning and non-learning individuals in this example is that learning individuals are allowed to perform 1000 random guesses, whereas for non-learning individuals the genotype translates directly to the phenotype, that is, alleles $\{0, ?_0\}$ encode a phenotypic 0, while alleles $\{1, ?_1\}$ encode a phenotypic 1, and no further improvement is possible.

This modification does not substantially change the H&N model and allows us to apply the gain-function approach. We distinguish three scenarios. First, if there exists one or more 0 alleles in the genotype, the optimal phenotype will not be found in either case, with or without learning. This means the gain function is constant and equal to one. Second, if the genotype is composed of alleles 1 and $?_1$, the optimal phenotype will be generated in both cases with or without learning, which also implies a constant gain function equal to one. In both scenarios, learning has no influence on evolution, as the constant gain function shows (cf. Equation 5). In the third scenario, the genotype is composed of alleles 1, $?_1$, and at least one allele $?_0$. In this situation, there is a difference between learning and non-learning. We estimate the gain function $g(x) = f(l(x))/f(x)$ (see Equation 4) for the third scenario as follows. The denominator is a constant, 1.0 in the setting of H&N. The numerator is the mean fitness achieved after learning and can be derived by summing the possibilities that the first correct guess is made exactly on the k th trial [4]:

$$\bar{f}(q) = \sum_{k=1}^{1000} f_{\text{H\&N}}(1000 - k) p_{\text{H\&N}}(k, q), \quad (6)$$

where q is the number of $?_0$ alleles ($?_0$'s and $?_1$'s) in the genotype (the remaining genes carry allele 1). According to [4], the fitness landscape and the probability function are given by $f_{\text{H\&N}}(n) = 1 + 19n/1000$, where n is the number of remaining trials when the reference vector has been found, and $p_{\text{H\&N}}(k, q) = (1 - 2^{-q})^{k-1} \cdot 2^{-q}$. Figure 2 shows the gain function $g(q)$ and estimated gain-function derivative $g'(q)$, where q is the number of $?_0$ alleles in the genotype. As the figure illustrates, the gain function has a positive derivative that predicts the qualitative outcome of the H&N model, namely, that learning will accelerate evolution.

In the literature, several articles have commented on Hinton and Nowlan's results. A gain-function interpretation is given as follows. In the absence of learning, differences between genetic predispositions are invisible, but learning amplifies or actually unveils these differences. As described in the beginning of Section 2, this—the learning-induced amplification of genetic predispositions—is exactly the conclusion that follows from a positive gain-function derivative. We conclude that in extreme fitness landscapes with large plateaus, learning potentially accelerates evolution.

3.2 Papaj's Model: A Negative Gain-Function Derivative

With a simulation model, Papaj [31] studied the interaction of evolution and learning in insects that need to adapt to new environmental conditions in which only one host species (a plant) is available. In contrast to Hinton and Nowlan [14], he concluded that learning decelerates evolution. In his model, an insect's behavior (the phenotype) is represented by a real-valued response number ξ , which is under the control of a genotypic value x ($x \in [0, 1]$); a prespecified learning parameter L ($L \in [0, 0.1]$ in [31]); and the number of learning trials made so far, t ($t = 0, \dots, T$, where T is the total number of learning trials in an insect life):

$$\xi(x, L, t) = x + (1 - x)(1 - e^{-Lt}) = 1 + (x - 1)e^{-Lt}. \tag{7}$$

This function is shown in Figure 3(a) for $L = 0.06$, for five different genotypic values x . Papaj presumably chose this type of learning curve because it guarantees that insect behavior at birth is solely specified by the genotype [i.e., $\xi(x, L, 0) = x$] and because in the T consecutive learning trials ξ converges asymptotically toward the optimal phenotype $\xi = 1$, which is a typical animal learning curve according to [31]. All individuals make strong progress in learning; those with higher genotypic values approach the learning target quicker, but the genetically weak ones seem to catch up during

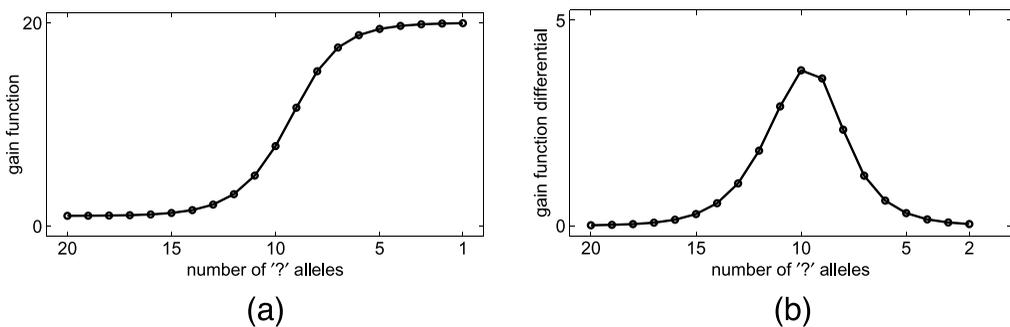


Figure 2. Hinton and Nowlan's model [14]. Fitness gain achieved through learning in the scenario where the genotype is composed of 1 's, $?_1$'s, and at least one $?_0$. The x axis is in reverse order to illustrate the direction of evolution. (a) Gain function, (b) differential $f'(q - 1) - f'(q)$ as an estimate of the gain derivative. The reader is also referred to [4], which takes a similar approach.

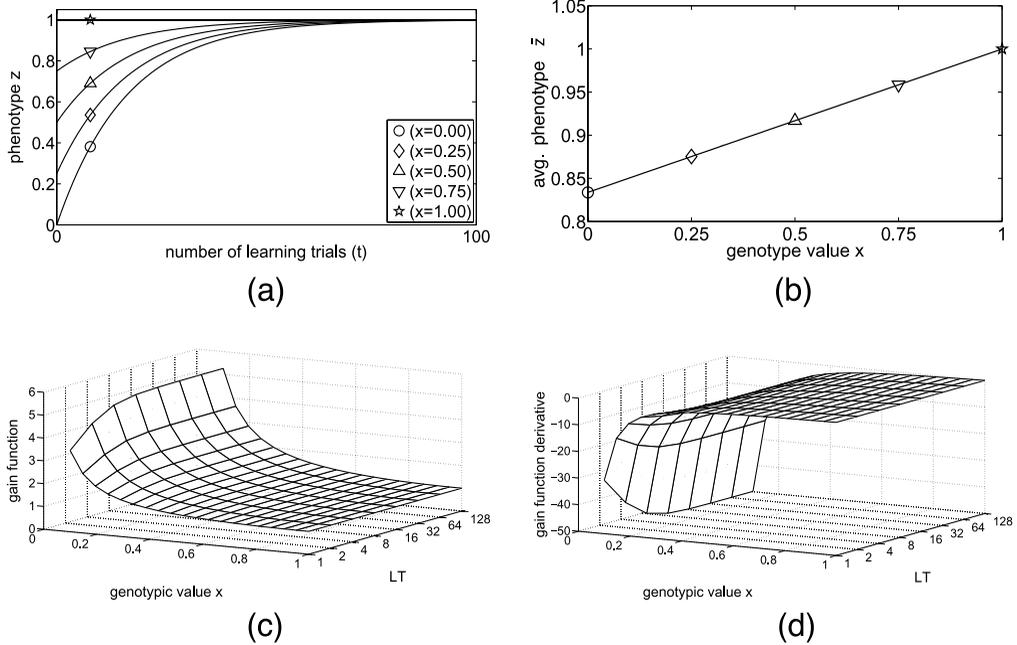


Figure 3. Papaj’s model of evolution and learning in insects [31]. (a) shows learning curves for a learning parameter $L = 0.06$ and different genotypic values (equal to innate phenotypes) $x \in \{0.0; 0.25; 0.5; 0.75; 1.0\}$; cf. Equation 7. With all genetic predispositions (different x values), individuals make strong progress in learning; those with higher genotypic values approach the learning target $z = 1$ quicker, but the genetically weak ones seem to catch up during learning. In (b), the average phenotype over $T = 100$ learning trials with learning parameter $L = 0.06$ is shown, as calculated using Equation 8. (c) shows the gain function $g(x)$ plotted against the genotypic value x and the product of lifetime and learning parameter, LT (logarithmic scale), and (d) shows its derivative with respect to x . For all possible parameter combinations LT , the gain function is negatively sloped toward the optimum at $x = 1$, which corresponds to a negative gain-function derivative as proven in Equation 11. Parameter combinations for very small values of LT and x are omitted to avoid numerical difficulties, since the gain function is not defined for $x = 0$ and $LT = 0$.

learning. In order to account for this, the lifetime fitness of learning individuals is determined by the average phenotype \bar{z} , which we approximate as

$$\bar{z}(x, L, t) = \begin{cases} x & \text{if } T = 0, \\ \frac{1}{T} \int_{t=0}^T \bar{z}(x, t, L) dt = 1 + \frac{1-x}{LT} (e^{-LT} - 1) & \text{if } T > 0. \end{cases} \tag{8}$$

The resulting average phenotype (for $T = 100$ and $L = 0.06$) is shown in Figure 3(b). Papaj assumed a concave fitness landscape on $[0; 1]$,

$$f(\bar{z}) = 1 - (1 - \bar{z})^2 \tag{9}$$

(a turned parabola), with maximum at $\bar{z} = 1$. Using Equation 8, we obtain the gain function

$$g(x) = \frac{f(x)}{f(\bar{z})} = \frac{f(\bar{z}(x, L, T > 0))}{f(\bar{z}(x, L, T = 0))} = \frac{1 - \left((x - 1) \frac{e^{-LT} - 1}{LT} \right)^2}{1 - (1 - x)^2}. \tag{10}$$

Differentiation with respect to x yields the gain-function derivative, and after some straightforward calculations we get

$$g'(x) = \frac{2(1-C)}{(x^2-2x)^2}(x-1) \quad \text{with } C = \left(\frac{e^{-LT}-1}{LT}\right)^2. \quad (11)$$

Since $L > 0$ and $T \geq 0$, the product $LT \geq 0$ can be interpreted as one variable. Since $C \in]0; 1[$ for $LT > 0$, we see that $g'(x) < 0$ for all $x \in]0, 1[$. The gain function (Equation 10) is visualized in Figure 3(c), and its derivative (Equation 11) in Figure 3(d).

For all combinations of LT , the gain function is negatively sloped toward the optimum (located at $x = 1$), that is, the gain-function derivative is negative (as proven in Equation 11).

In the model chosen by Papaj, learning allows individuals with a genetic predisposition toward a weak innate phenotype to catch up with innately strong individuals. Considering the whole lifetime, an advantage of the innately strong individuals remains. However, learning strongly reduces the selective pressure toward good genes. This phenomenon, which some years after Papaj's simulation was named the *hiding effect* [22], appears if innately weak individuals gain proportionally more from learning than innately strong individuals, thus *hiding* genetic differences. Exactly this is revealed by the negative derivative of the gain function.

3.3 Accelerated and Decelerated Evolutionary Phases on the Sigmoid Fitness: A Nonmonotonic Gain Function

In both preceding examples, the gain function turned out to be monotonic. Evolution was either accelerated (Section 3.1) or decelerated (Section 3.2) at any time of the evolutionary process. We now look at a scenario where it depends on the learning function whether the gain function is monotonic or not. The fitness landscape (mapping phenotype value z to fitness) is the sigmoid function $f(z) = (1 + \exp(-z))^{-1}$ [see Figure 4(a)], which is monotonic, convex for negative genotype values, and concave for positive genotype values. In the absence of learning, the phenotype value z equals the genotype value x , [$z(x) = x$]; in the presence of learning, $z(x) = l(x)$, where l is a learning function. Thus, fitness is given by $f(x)$ in the absence of learning, and $f(l(x))$ in the presence of learning.

In addition to the gain-function analysis, we ran some experiments to verify the analytical results. These experiments were set up in the following way¹: We simulated an asexual population of 100 individuals, each characterized by a one-dimensional (real-valued) genotypic value x , and initialized uniformly in the vicinity of -3 (in the interval $[-3.1, -2.9]$). In the absence of learning, the phenotype z equals x ; with learning the phenotype $z = l(x)$. To simulate selection, we used an algorithm known as *stochastic universal sampling* [2]. This algorithm implements sampling (with replacement) of n offspring from n parents, where the probability of an individual being sampled is proportional to its fitness $f(z)$ (i.e., $f(x)$ without learning, $f(l(x))$ with learning). Biologically, this algorithm is equivalent to assuming that each parent produces a very large number of offspring, and the survival probability of the offspring is proportional to the fitness. To simulate mutation, a random number from a normal distribution with parameters $\mu = 0$ and $\sigma = 10^{-3}$ was added to the genotypic value x of each offspring.

In the first scenario, learning is defined as $l_1(x) = x + 0.25$. Learning moves an individual a constant distance toward the optimum in phenotype space, regardless of the genotypic value (which equals the innate phenotype, as in Papaj's model in Section 3.2). We refer to this type of learning as *constant learning*. In combination with the sigmoid fitness function, this results in a gain function that decreases during the course of evolution (from low to high genotype values); see the solid line in Figure 4(b). Therefore, we expect that learning decelerates evolution on both the convex and the concave part of the sigmoid fitness function. To verify this, we ran some experiments for this fitness function.

¹ The C++ source code for these experiments is available at http://www.aifb.uni-karlsruhe.de/EffAlgo/ipa/gainfunc_sigmoid_experiment.zip

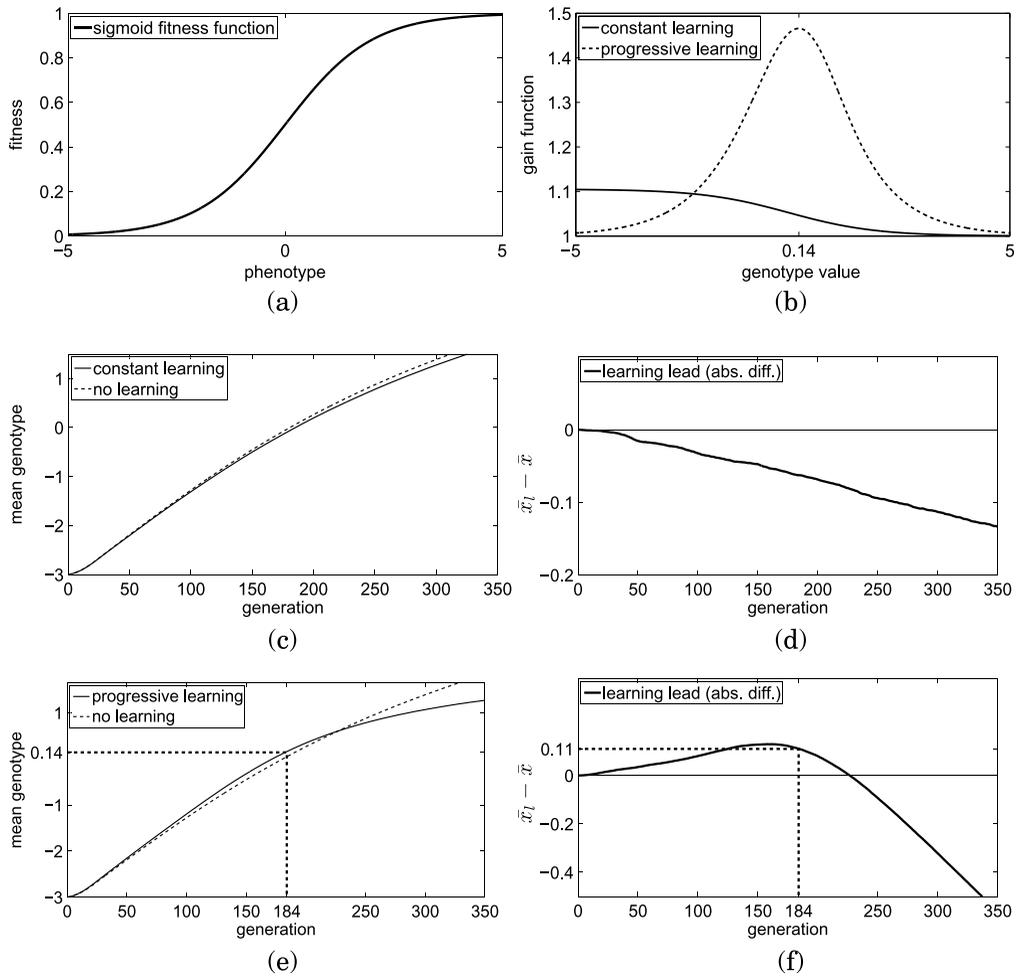


Figure 4. Evolution and learning on the sigmoid fitness function. (a) The sigmoid fitness function. (b) Gain functions for constant learning and progressive learning. (c–f) Averaged results of 1000 independent simulation runs with the sigmoid fitness function. In particular, (c) shows the mean genotype evolution with constant learning and no learning, and (d) the absolute difference of the curves in (c) (i.e., the mean genotypes in the presence and in the absence of learning), $\bar{x}_t - \bar{x}$; we call this *learning lead*. (e) is the same as (c), but with progressive learning, and (f) is the same as (d), but with progressive learning.

The simulation results for the constant learning case are shown in Figure 4c,d. Constant learning indeed decelerates evolution, throughout the evolution—on the convex as well as on the concave part of the fitness landscape. Again the hiding effect [22] can be observed. This might be counter-intuitive, because a constant phenotype shift yields a larger fitness increase in every convex fitness landscape. Thus, one would expect that evolution is accelerated on the convex landscape segment. However, what matters is the strength of the relative fitness increase, and this is what is taken into account in the gain-function framework.

The question that arises from this observation is, whether there exist fitness landscapes for which this basic form of learning (constant learning) accelerates evolution. We use the gain-function framework to find a general answer to this. Formally, constant learning is defined by

$$l_{\delta}(x) = x + \delta, \tag{12}$$

where δ is a positive constant. Assuming a monotonic and continuously differentiable fitness landscape f , the sign of the gain function derivative satisfies

$$\begin{aligned} \text{sign}(g'_\delta) &= \text{sign}\left(\frac{[f(l_\delta(x))]' }{f(l_\delta(x))} - \frac{f'(x)}{f(x)}\right) \\ &= \text{sign}([\ln f(l_\delta(x))] - [\ln f(x)]') \\ &= \text{sign}([\ln f(x + \delta)]' - [\ln f(x)]') \\ &= \text{sign}([\ln f(x)]''). \end{aligned} \tag{13}$$

The last equality follows from the relationship $\text{sign}(F'(x)) = \text{sign}((x_2 - x_1)[F(x_2) - F(x_1)])$, which holds for any monotonic function (here $F(x) = [\ln f(x)]'$) and arbitrary x_1, x_2 with $x_1 \neq x_2$. The influence of constant learning on evolution solely depends on the second derivative of logarithmic fitness: Positive (negative) $[\ln f(x)]''$ implies learning-induced acceleration (deceleration) for this type of learning. Indeed, the second derivative of the logarithmic sigmoid function $-e^{-x}(1 + e^{-x})^{-2}$ is negative for all x .

In the second scenario, learning is defined as $l_2(x) = x + e^x$. The larger the genotype value (equal to the innate phenotype), the more learning shifts the phenotype toward the optimum. We call this form of learning *progressive learning*. The corresponding gain function is shown as a dashed line in Figure 4b. The gain function increases before the turning point $x = 0.14$ and decreases after it. Thus, we expect that in the early phase of evolution, learning accelerates evolution, and in the later phase it decelerates evolution. These predicted dynamics are qualitatively confirmed by the simulation results for the progressive learning case (Figure 4e, f). The mean genotype of the learning population reaches the genotype that corresponds to the gain-function maximum ($x = 0.14$) in generation 184 (Figure 4e). The maximum difference between the learning and the non-learning population has been reached already 25 generations earlier (at generation 159; see Figure 4f). However, during these 25 generations, the learning population has largely maintained its distance from the non-learning population.

The gain-function analysis only allows an approximate prediction of the population dynamics over time. An exact prediction based on the gain function assumes that both the learning and the non-learning population have the same distribution in genotype space. However, during the early phase of evolution, the learning population moves quicker toward higher genotype values; thus, the learning individuals populate a different region in genotype space than the non-learning ones. Despite a positive gain-function derivative the selection pressure might be stronger in the region of the non-learning population than in the region of the learning population.

Nevertheless, the evolutionary dynamics are quite well described by the gain function, as the example has demonstrated. We conclude that the gain-function approach can approximately predict the evolutionary dynamics even in the case where acceleration is followed by deceleration.

3.4 Biological Data: An Inverse Gain-Function Application

In the models that we have investigated so far, knowledge about the fitness landscape and the learning algorithm was given, and we used this knowledge in the gain-function framework to predict the evolutionary dynamics. However, the logical equivalence in Equation 5 tells that an *inverse* approach is also possible. Given some evolutionary data (in the absence and the presence of learning), we can deduce the sign of the gain function. In other words, we learn something about the effect of learning on fitness.

In the following, we do this in a rather qualitative way with data from the first biological experiment that demonstrated the Baldwin effect [24]. In this experiment Mery and Kawecki studied the effect of learning on resource preference in fruit flies (*Drosophila melanogaster*). For details of the experiment, we refer to [24]. Here we only give a brief qualitative description: The flies had the choice

between two substrates (pineapple and orange) to lay their eggs on, but the experimenters took only the eggs laid on pineapple to breed the next generation of flies, which are (once grown up) given the same choice for their eggs.

Measuring the proportion of eggs laid on pineapple, one could see that a stronger preference for pineapple evolved, from 42% in the first generation to 48% in generation 23. To test the Baldwin effect, another experiment was done, where also eggs laid on pineapple were selected to breed the next generation, but flies could previously learn that pineapple is the good substrate. To allow for learning, several hours before the experimenter took away the eggs for breeding, the disfavored orange was supplemented with a bitter-tasting chemical for some time (and replaced with a fresh orange after that). If flies learned to avoid orange, they would lay fewer eggs on it later, that is, show a preference for pineapple. After 23 generations of learning flies, the innate preference (measured in the absence of the bitter chemical) evolved to 55%, significantly more than the 48% that evolved in the absence of learning. Thus, in this experiment learning accelerated evolution. According to Equation 5, the gain function has a positive derivative.

Mery and Kawecki did the same experiment with orange as the favored substrate, that is, eggs for breeding were taken from orange, and pineapple was supplemented with the bitter-tasting chemical for learning. In 23 generations the innate preference for orange evolved from an initial 58% to 66% in the presence of learning, but to even more, 72%, in the absence of learning. Thus, in this setting, learning decelerated evolution. According to Equation 5, the gain function has a negative derivative. The first row of Table 1 summarizes the experimental results. As in [24], we refer to the cases when pineapple was the favored resource as *learning pineapple* in the presence of learning and *innate pineapple* in the absence of learning, and correspondingly *learning orange* and *innate orange* when orange was the favored resource.

We want to shed some light on these—seemingly contradictory—results. If the relationship between innate resource preference and success of the resource preference learning is independent of what the high-quality resource currently is, the experimental results can be interpreted as follows: When evolution starts from a weak innate preference for the favored fruit (42%, as in the first experiment with pineapple as the high-quality resource), this leads to learning induced acceleration. However, if evolution starts from a strong innate preference for the favored fruit (58%, as in the second experiment with orange as the high-quality resource) this leads to learning-induced deceleration of evolution. Therefore, if evolution started further away from the evolutionary goal, then learning accelerated evolution, implying an increasing gain function, and if it started closer to the evolutionary goal, learning decelerated evolution, implying a decreasing gain function. Thus, in principle, we can expect a gain function that is increasing for a weak innate preference for the target fruit and decreasing for a strong innate preference for the target fruit. This implies a maximum gain-function value at an intermediate innate preference for the target fruit and lower gain-function values for weak and strong innate preferences.

Table 1. Experimental results for the natural evolution [24] and the artificial evolution. For both cases the average innate preference for orange after 23 generations is shown. The numbers in parenthesis are normalized w.r.t. the initial preference.

Selection for	Evolution	Preference		
		Initial	Evolved	
			Innate	Learning
Orange	Natural	.58 (100%)	.72 (124%)	> .66 (114%)
	Artificial	.58 (100%)	.61 (105%)	> .59 (102%)
Pineapple	Natural	.42 (100%)	.48 (114%)	< .55 (130%)
	Artificial	.42 (100%)	.46 (109%)	< .48 (114%)

Recalling that the gain function $g(x) = f_l(x)/f(x)$ reflects the relative fitness gain due to learning, we deduce that learning seems to be not very effective when the starting point of learning is far away from or very close to the learning goal (low gain-function values), and is probably most effective for a starting point with an intermediate distance to the learning goal.

Besides these conclusions from the experimental results, there are other arguments for such a relationship.

For an individual that already shows strong innate preference for a high-quality resource, its learning success might be low because perfection is usually difficult (and requires large resources), or simply because the preference cannot be increased beyond 100%.

In contrast, there is scope for a large effect of learning in individuals that show a weak preference for the high-quality resource, that is, a strong preference for the low-quality resource. However, there are two reasons why such individuals with strong innate preference for the low-quality resource might be slow in changing their preference toward the high-quality resource. Firstly, because of their strong initial preference for the one resource, individuals will only rarely sample the other one, and thus rarely have a chance to find that the other resource is in fact better. Secondly, even if they occasionally sample the other resource, their strong innate preference for the first one may be difficult to overwrite. This argument is supported by experiments with (phytophagous) insects (e.g., [32]) and also with humans [27].

To test our conclusions, we simulate the biological experiment using an artificial evolutionary system of resource preference.² In the simulation model, the innate preference for orange is genetically encoded as $x \in [0; 1]$ and represents the probability of choosing orange in a Bernoulli trial. If the individual fails to choose the high-quality resource, it does not produce offspring. However, if the high-quality resource is chosen, the digital fly receives a fitness score of 1, which results in a high probability of producing offspring for the next generation (assuming a linearly proportional selection scheme).

Thus, if pineapple is the high-quality resource, the expected fitness in the absence of learning, f^P , is given by $f^P(x) = 1 - x$ (innate pineapple). Since learning is on average beneficial, the fitness in the presence of learning, $f_l^P(x)$, must be larger, that is, $f_l^P(x) \geq f^P(x)$ (learning pineapple). Correspondingly, if orange is the high-quality resource, we obtain $f^O(x) = x$ (innate orange) and $f_l^O(x)$ (learning orange), where $f_l^O(x) \geq f^O(x)$.

In the model, populations are initialized with $x \in [0.55, 0.61]$, and with an average orange preference of $\bar{x} = 0.58$. This is the same mean preference as observed in the initial generation of the biological experiment [24]. For the simulation, we choose a population size of 150, which is similar to the biological experiment. Mutation is simulated by adding a random number from a normal distribution with mean 0 and standard deviation 5×10^{-5} , that is, we assume a small effect of mutation on resource preference. What remains to be defined is the expected fitnesses in the case of learning, $f_l^P(x)$ and $f_l^O(x)$. Recalling Equation 4, they can be derived if we know the corresponding gain function g .

A gain function that is increasing for weak, maximal for intermediate, and decreasing for strong innate preference for the high-quality resource is given by a linear transformation of the Gaussian function $\phi(x, \sigma)$:

$$g(x, \alpha, \sigma) = a_1(\alpha, \sigma) + a_2(\alpha, \sigma) \phi(x, \sigma), \quad (14)$$

where $a_1(\alpha, \sigma) = 1 - \frac{\alpha\phi(0, \sigma)}{\phi(0.5, \sigma) - \phi(0, \sigma)}$ and $a_2(\alpha, \sigma) = \frac{\alpha}{\phi(0.5, \sigma)}$, such that g is 1 at the genotype boundaries and maximal at the center of the genotype space ($x = 0.5$). The parameter a reflects the maximum relative fitness gain (at $x = 0.5$) that can be achieved through learning. In the biological experiments of Mery and Kawecki [24], the fitness gain due to learning was assessed by comparing the innate preference and the preference after learning (given by the proportion of eggs on the fruits)

² The C++ source code for these experiments is available at http://www.aifb.uni-karlsruhe.de/EffAlg/ipa/gainfunc_digital_fruitflies.zip

at generation 23. Depending on whether and what the ancestor populations have learned, and what the target resource in the assay was, the fitness gain varied widely in the biological experiment. Among the different settings, the maximum fitness gain due to learning was an increase from 45% to 57% of eggs laid on the high-quality resource, that is, a fitness gain of $(57 - 45)/45 = 0.27$. For the gain function of the simulation, Equation 14, we choose a similar value $\alpha = 0.25$. The only remaining parameter, σ , was tuned to get a maximally steep gain function in the preference region where evolution starts (satisfying the requirement that $f_i(x)$ still be monotonic), resulting in $\sigma = 0.075$. Figure 5a shows how learning influences the fly's probability of choosing orange and the resulting gain function. Figure 5(b) shows the evolution of the mean innate preference for orange. The innate preference for orange evolves faster in the absence of learning (innate orange) than with learning (learning orange). However, the innate preference for pineapple evolves faster in the presence of learning (learning pineapple) than in the absence of learning (innate pineapple). The short error bars (of length two standard errors) indicate the statistical significance of the difference in evolved preferences. This qualitatively confirms the results of the biological experiment of [24]. In Table 1, the experimental results of the artificial evolution are directly compared with the results of the biological evolution. The numbers in parentheses are normalized w.r.t. the initial preference. First of all, we see that the effects of acceleration and deceleration are *qualitatively* identical. In both cases, with and without learning, and for both orange and pineapple selection, evolution proceeds quicker in the natural evolution experiment. However, with regard to the normalized values, the relative difference between evolution with and without learning is very similar in the natural and artificial evolution.

The aim of this experiment was not to quantitatively replicate the results of the biological experiment. Too many assumptions need to be made in order to simulate evolution of real fruit flies realistically. For example, we simply chose a Gaussian function as the gain function with a maximum at $x = 0.5$. The biological data suggested that the maximum of the gain function lies between 0.42 and 0.58. We did not attempt to tune the simulation model, but simply chose the midpoint, 0.5. If evolution starts at $x = 0.42$ (selection for pineapple), this means that the genotype interval in which evolution is accelerated is rather small. Certainly a larger optimal x value allows one to produce stronger learning-induced acceleration. Furthermore the biological gain function may not be symmetric. Thus acceleration (selection for pineapple) may have a different magnitude than deceleration (selection for orange). We have no direct knowledge about the mutation strength and the mutation

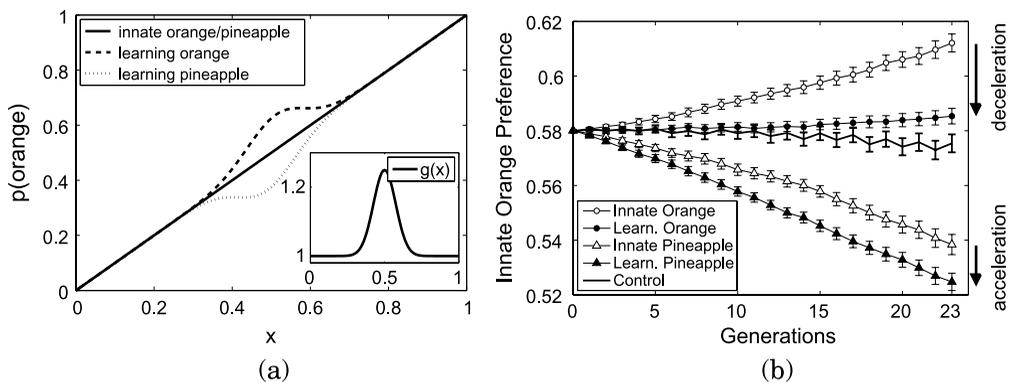


Figure 5. Simulation of the fruit fly experiment. Part (a) shows how learning influences the fly's probability of choosing orange for different values of the innate preference for orange, x (the probability of choosing pineapple is $1 - p(\text{orange})$). The inset in (a) shows the gain function, which is identical for learning orange and learning pineapple. The horizontal axis shows the genetic predisposition of the target fruit. Part (b) shows the evolution of the mean innate preference for orange (averaged over all individuals and 50 independent evolutionary runs, with bars for ± 1 standard error). Note that the preference for pineapple is one minus the preference for orange. If orange is the high-quality resource, learning decelerates evolution, but if pineapple is the high-quality resource, learning accelerates evolution. As in the biological experiment, a set of control runs have been carried out in which the high-quality food changes every generation between orange and pineapple.

symmetry in the biological experiment, but assume the same strength of symmetric mutation over the entire genotype space in the artificial evolution. This may not correspond to reality. For example, in the absence of learning in the biological experiment, selection for orange produced a preference shift from 0.58 to 0.72, while selection for pineapple produced a shift from 0.42 to only 0.48 (in 23 generations).

Despite this, the gain-function argument may not be the only explanation. Mery and Kawecki [24] discuss several other reasons in detail.

We see that the gain-function approach can be applied inversely in order to get a better understanding of the effects of learning on fitness. Of particular interest might be the insect learning pattern that produced a gain function of the type shown in Figure 5a, which might also apply to many artificial learning systems.

4 Discussion

In the literature, evidence for both learning-induced acceleration and deceleration of evolution can be found. In this article, we have presented a general framework—the gain function—to explain and predict under which conditions learning accelerates or decelerates evolutionary change. The gain function is formulated in terms of the effect of learning on the mapping from genotype space to fitness space. Learning is predicted to accelerate evolution if the proportional gain of fitness due to learning is greater for genotypes that would already be fitter without learning. In contrast, if the genetically less fit individuals gain proportionally more from learning, the relative differences in fitness between genotypes become reduced, and selection becomes less effective. Figure 6 illustrates this principle.

This general and quite intuitive result can be used to make predictions for specific models. Since we are not constrained by a particular type of learning, these predictions may be made not only for artificial, but also for biological systems, as we have shown in Section 3.4. It can even be applied if the specific learning algorithm is not known, as can be the case in complex artificial systems, and obviously in natural systems. All that is needed is an estimation of the gain function.

From the examples that we analyzed, we observed that learning is likely to accelerate evolution in extreme fitness landscapes (as in Section 3.1), and decelerate it if individuals with weak genetic predisposition learn very effectively (as in Section 3.2). The gain function predicts the evolutionary dynamics accurately, even if it is not monotonic (Section 3.3). If learning causes a phenotypic change of the same magnitude for all individuals (constant learning), it depends on the second derivative of the logarithm of the fitness function whether evolution is accelerated or decelerated. However, if

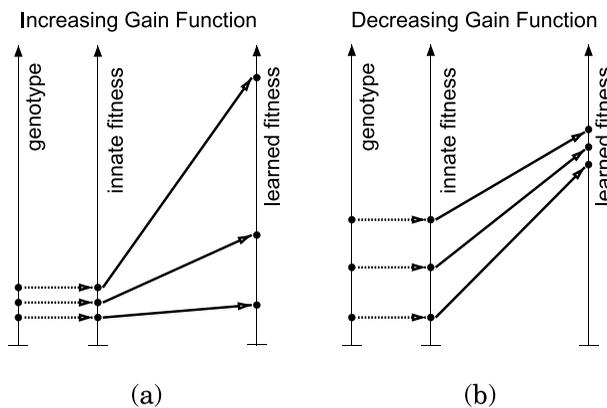


Figure 6. The principle of the gain function. (a) An increasing gain function indicates that relative fitness differences between genetically weak and strong individuals are enlarged through learning. (b) A decreasing gain function indicates that relative fitness differences between genetically weak and strong individuals are reduced through learning.

learning is progressive (genetically fit individuals benefit more), then acceleration is more likely, as Section 3.3 has shown.

The analysis of the gain-function model presented in this article is more appropriate for computational evolutionary models than the continuous one outlined in [29]. It is—in contrast to [29]—based on a population of discrete entities, and no assumption on the population distribution is required. This allows us to directly investigate simulation models. In principle, computational simulation models can also be investigated with the biological gain-function framework [29] if a learning parameter is available. However, we can expect that the necessary assumptions (e.g., symmetrically distributed populations) are violated in most simulation models. The analysis of to what extent this would cause erroneous predictions or, in more general terms, under which circumstances discrete models are more powerful than continuous models (or vice versa) is beyond the scope of this article.³ The analysis presented in this article complements our earlier approach.

In biology, the gain function only applies to *directional* selection, that is, selection that moves the population toward higher fitness (as opposed to *disruptive* or *stabilizing* selection). The gain-function analysis is expectation-based and does not account for the variability of the population movement. Thus, the gain function does not allow one to make exact predictions on the influence of learning on the time needed to cross a fitness valley toward a region with higher fitness. Such a prediction cannot be made expectation-based, since fitness valley crossing requires an unlikely event. A stochastic analysis is more appropriate to predict the time needed to cross a fitness valley. A first approach to such an analysis can be found in [5], which has, however, some drawbacks. Firstly, in [5] the population movement is modeled by a one-dimensional nonsymmetric random walk (cf. Section 1). Secondly, the derived drawdown as an indicator for the time needed to reach the fitness maximum does not allow for directional selection. In ALife an experimental study for a particular fitness landscape (a bimodal version of Hinton and Nowlan's fitness landscape [14]) has been published in [25].

In a multimodal fitness landscape, an alternative interpretation of the gain function is the following: If the gain function is decreasing toward a local optimum, learning reduces selection pressure toward this local optimum. Hence, a population movement away from the optimum (possibly toward the global optimum) becomes more likely.

Furthermore, the gain-function analysis may provide a valuable interpretation if there are monotonic global trends with only low local optima, which can be interpreted as noise.

We have shown that the gain function makes *exact short-term predictions* of the mean genotype movement. If a population that initially populates a fitness landscape region with positive gain-function derivative and then moves on to a region with negative gain function derivative (at some point, the gain function is not monotonic within the range of the population), the gain function framework does not allow exact predictions. It does, however, allow *approximate long-term predictions* of the mean genotype movement, as we have seen in Section 3.3.

Despite a long history of the concept, the interaction between learning and evolution remains only partially understood. The analysis outlined in this article offers a general framework to study the effects of learning on evolution, and an explanation of the results of previously published models. It also provides a theoretical underpinning of biological data.

References

1. AnceL, L. (2000). Undermining the Baldwin expediting effect: How phenotypic plasticity influences the rate of evolution. *Theoretical Population Biology*, 58, 307–319.
2. Baker, J. (1987). Reducing bias and inefficiency in the selection algorithm. In J. Grefenstette (Ed.), *Proceedings of the Second International Conference on Genetic Algorithms and their Applications* (pp. 14–21). Mahwah, NJ: Lawrence Erlbaum Associates.

³ The interesting relation between continuous and particulate models is not restricted to theoretical biology and computational biology, but can also be found in other disciplines, such as physics, where the Navier-Stokes equation is a continuous model of flow interactions and where it has been shown that for complex systems particle-based models can be more appropriate [21].

3. Baldwin, J. (1896). A new factor in evolution. *American Naturalist*, *30*, 441–451.
4. Belew, R. (1989). When both individuals and populations search: Adding simple learning to the genetic algorithm. In J. Schaffer (Ed.), *Proceedings of the Third International Conference on Genetic Algorithms* (pp. 34–41). San Mateo, CA: Morgan Kaufmann.
5. Borenstein, E., Meilijson, I., & Ruppin, E. (2006). The effect of phenotypic plasticity on evolution in multi-peaked fitness landscapes. *Journal of Evolutionary Biology*, *19*, 1555–1570.
6. Bull, L. (1999). On the Baldwin effect. *Artificial Life*, *5*, 241–246.
7. Cortez, P., Rocha, M., & Neves, J. (2002). A Lamarckian approach for neural network training. *Neural Processing Letters*, *15*, 105–116.
8. Dopazo, H., Gordon, M., Perazzo, R., & Risau-Gusman, S. (2001). A model for the interaction of learning and evolution. *Bulletin of Mathematical Biology*, *63*, 117–134.
9. Downing, K. (2004). Development and the Baldwin effect. *Artificial Life*, *10*, 39–63.
10. Fontanari, J., & Meir, R. (1990). The effect of learning on the evolution of asexual populations. *Complex Systems*, *4*, 401–414.
11. French, R., & Messinger, A. (1994). Genes, phenes and the Baldwin effect. In R. Brooks & P. Maes (Eds.), *Artificial Life IV* (pp. 277–282). Cambridge, MA: MIT Press.
12. Gordon, D. (1992). Phenotypic plasticity. In E. Lloyd & E. Kell (Eds.), *Keywords in evolutionary biology* (pp. 255–262). Cambridge, MA: Harvard University Press.
13. Harvey, I. (1997). Is there another new factor in evolution? *Evolutionary Computation*, *4*, 311–327.
14. Hinton, G., & Nowlan, S. (1987). How learning can guide evolution. *Complex Systems*, *1*, 495–502.
15. Hüsken, M., Gayko, J., & Sendhoff, B. (2000). Optimization for problem classes—Neural networks that learn to learn. In X. Yao & D. Fogel (Eds.), *IEEE Symposium on Combinations of Evolutionary Computation and Neural Networks* (pp. 98–109). Piscataway, NJ: IEEE Press.
16. Johnston, T. (1982). Selective costs and benefits in the evolution of learning. *Advances in the Study of Behavior*, *12*, 65–106.
17. Keesing, R., & Stork, D. (1990). Evolution and learning in neural networks: The number and distribution of learning trials affect the rate of evolution. In R. Lippmann, J. Moody, & D. Touretzky (Eds.), *Proceedings of the 1990 Conference on Advances in Neural Information Processing Systems* (pp. 804–810). San Mateo, CA: Morgan Kaufmann.
18. Ku, K., Mak, M., & Siu, W. (1999). Adding learning to cellular genetic algorithms for training recurrent neural networks. *IEEE Transactions on Neural Networks*, *10*, 239–252.
19. Ku, K., Mak, M., & Siu, W. (2000). A study of the Lamarckian evolution of recurrent neural networks. *IEEE Transactions on Evolutionary Computation*, *4*, 31–42.
20. Ku, K., Mak, M., & Siu, W. (2003). Approaches to combining local and evolutionary search for training neural networks: A review and some new results. In A. Ghosh & S. Tsutsui (Eds.), *Advances in evolutionary computing: Theory and applications* (pp. 615–642). Berlin: Springer.
21. Malevanets, A., & Kapral, R. (1999). Mesoscopic model for solvent dynamics. *Journal of Chemical Physics*, *110*, 8605–8613.
22. Mayley, G. (1997). Guiding or hiding: Explorations into the effects of learning on the rate of evolution. In P. Husbands & I. Harvey (Eds.), *Proceedings of the Fourth European Conference on Artificial Life 97* (pp. 135–144). Cambridge, MA: MIT Press.
23. Maynard-Smith, J. (1987). When learning guides evolution. *Nature*, *329*, 761–762.
24. Mery, F., & Kawecki, T. (2004). The effect of learning on experimental evolution of resource preference in *Drosophila melanogaster*. *Evolution*, *58*, 757–767.
25. Mills, R., & Watson, R. (2006). On crossing fitness valleys with the Baldwin effect. In L. Rocha, L. Yaeger, M. Bedau, D. Floreano, R. Goldstone, & A. Vespignani (Eds.), *Proceedings of Artificial Life X* (pp. 493–499).
26. Noskovicz, S., & Goldhirsch, I. (1990). First passage time distribution in random walk. *Phys. Rev. A*, *42*, 2047–2064.

27. Ohman, A., & Dimberg, U. (1978). Facial expressions as conditioned stimuli for electrodermal responses: A case of “preparedness”? *Journal of Personality and Social Psychology*, 36, 1251–1258.
28. Paenke, I., Kawecki, T., & Sendhoff, B. (2006). On the influence of lifetime learning on selection pressure. In L. Rocha, L. Yaeger, M. Bedau, D. Floreano, R. Goldstone, & A. Vespignani (Eds.), *Proceedings of Artificial Life X* (pp. 500–506). Cambridge, MA: MIT Press.
29. Paenke, I., Sendhoff, B., & Kawecki, T. (2007). Influence of plasticity and learning on evolution under directional selection. *American Naturalist*, 170(2), E54–E58.
30. Paenke, I., Sendhoff, B., Rowe, J., & Fernando, C. (2007). On the adaptive disadvantage of Lamarckianism in rapidly changing environments. In F. Almeida e Costa (Ed.), *Advances in Artificial Life, 9th European Conference on Artificial Life* (pp. 355–364). Berlin: Springer.
31. Papaj, D. (1994). Optimizing learning and its effect on evolutionary change in behavior. In L. Real (Ed.), *Behavioral Mechanisms in Evolutionary Ecology* (pp. 133–154). Chicago: University of Chicago Press.
32. Potter, D., & Held, D. (1999). Absence of food-aversion learning by a polyphagous scarab, *Popillia japonica*, following intoxication by geranium, pelargonium × hortorum. *Entomologia Experimentalis et Applicata*, 91, 83–88.
33. Sasaki, T., & Tokoro, M. (2000). Comparison between Lamarckian and Darwinian evolution on a model using neural networks and genetic algorithms. *Knowledge and Information Systems*, 2, 201–222.
34. Simpson, G. (1953). The Baldwin effect. *Evolution*, 7, 110–117.
35. Spitzer, F. (2001). *Principles of random walk* (2nd ed.). Berlin: Springer.
36. Suzuki, R., & Arita, T. (2007). The dynamic changes in roles of learning through the Baldwin effect. *Artificial Life*, 13, 31–43.

Appendix: Inductive Proof for the Gain-Function Approach

In this appendix we prove that Equation 5 is true, provided that there is genetic variation (variation in x) in the population, $f'(x) > 0$, and the learning function $l(x)$ is such that the sign of $g'(x)$ is constant within the range of variation ($x_{\min} \leq x \leq x_{\max}$) present in the population. For convenience, we substitute $f_j(x)$ for $f(l(x))$. In the following, we outline the proof for the first case of Equation 5 ($g'(x) > 0$) and omit the other cases, because the proofs are analogous and the transfer from the first case is straightforward.

Recalling Equation 1, we define the statement $S(n)$ as

$$S(n) := \frac{\sum_{i=1}^n x_i f(x_i)}{\sum_{i=1}^n f(x_i)} - \frac{\sum_{i=1}^n x_i f'(x_i)}{\sum_{i=1}^n f'(x_i)} = \bar{x}_f^* - \bar{x}^* > 0. \quad (15)$$

Recalling the gain-function definition $g(x) = f(l(x))/f(x)$, we obtain

$$\forall x, x_i, x_j \in [x_{\min}; x_{\max}], \quad x_i < x_j : \quad g'(x) > 0 \iff \frac{f(x_i)}{f(x_i)} < \frac{f(x_j)}{f(x_j)}. \quad (16)$$

Without loss of generality we further assume the x_i to be arranged in ascending order:

$$\forall (i, j) : \quad i < j \implies x_i \leq x_j. \quad (17)$$

Initialization: For $n = 2$, $S(n)$ can be written and reformulated as

$$\begin{aligned}
 S(2) &\Leftrightarrow \frac{x_1 f(x_1) + x_2 f(x_2)}{f(x_1) + f(x_2)} > \frac{x_1 f(x_1) + x_2 f(x_2)}{f(x_1) + f(x_2)} \\
 &\Leftrightarrow \frac{x_1(f(x_1) + f(x_2)) + (x_2 - x_1)f(x_2)}{f(x_1) + f(x_2)} > \frac{x_1(f(x_1) + f(x_2)) + (x_2 - x_1)f(x_2)}{f(x_1) + f(x_2)} \\
 &\Leftrightarrow x_1 + \frac{(x_2 - x_1)f(x_2)}{f(x_1) + f(x_2)} > x_1 + \frac{(x_2 - x_1)f(x_2)}{f(x_1) + f(x_2)} \\
 &\Leftrightarrow \frac{f(x_2)}{f(x_1) + f(x_2)} > \frac{f(x_2)}{f(x_1) + f(x_2)} \tag{18} \\
 &\Leftrightarrow \frac{f(x_1)}{f(x_2)} + 1 < \frac{f(x_1)}{f(x_2)} + 1 \\
 &\Leftrightarrow \frac{f(x_1)}{f(x_1)} < \frac{f(x_2)}{f(x_2)} \\
 &\Leftrightarrow g(x_1) < g(x_2),
 \end{aligned}$$

which is true according to Equation 16.

Inductive step: Assuming $S(n)$ is true, we show that $S(n + 1)$ is true:

$$\begin{aligned}
 S(n + 1) &\Leftrightarrow \frac{\sum_{i=1}^{n+1} x_i f(x_i)}{\sum_{i=1}^{n+1} f(x_i)} - \frac{\sum_{i=1}^{n+1} x_i f(x_i)}{\sum_{i=1}^{n+1} f(x_i)} > 0 \\
 &\Leftrightarrow \left(\sum_{i=1}^{n+1} x_i f(x_i) \right) \left(\sum_{i=1}^{n+1} f(x_i) \right) > \left(\sum_{i=1}^{n+1} x_i f(x_i) \right) \left(\sum_{i=1}^{n+1} f(x_i) \right) \tag{19} \\
 &\Leftrightarrow L_1 + L_2 + L_3 + L_4 > R_1 + R_2 + R_3 + R_4,
 \end{aligned}$$

where

$$\begin{aligned}
 L_1 &= \sum_{i=1}^n x_i f(x_i) \sum_{i=1}^n f(x_i), & R_1 &= \sum_{i=1}^n x_i f(x_i) \sum_{i=1}^n f(x_i), \\
 L_2 &= f(x_{n+1}) \sum_{i=1}^n f(x_i) x_i, & R_2 &= f(x_{n+1}) \sum_{i=1}^n f(x_i) x_i, \\
 L_3 &= x_{n+1} f(x_{n+1}) \sum_{i=1}^n f(x_i), & R_3 &= x_{n+1} f(x_{n+1}) \sum_{i=1}^n f(x_i), \\
 L_4 &= x_{n+1} f(x_{n+1}) f(x_{n+1}), & R_4 &= x_{n+1} f(x_{n+1}) f(x_{n+1}).
 \end{aligned}$$

With $L_1 > R_1$ (according to the inductive assumption $S(n)$) and $L_4 = R_4$, we can obtain

$$S(n) \wedge (L_2 + L_3 \geq R_2 + R_3) \Rightarrow S(n + 1). \tag{20}$$

Thus, it is sufficient to show

$$\begin{aligned} L_2 + L_3 \geq R_2 + R_3 &\Leftrightarrow f(x_{n+1}) \sum_{i=1}^n f(x_i)x_i + x_{n+1}f(x_{n+1}) \sum_{i=1}^n f(x_i) \\ &\geq f(x_{n+1}) \sum_{i=1}^n f(x_i)x_i + x_{n+1}f(x_{n+1}) \sum_{i=1}^n f(x_i) \\ &\Leftrightarrow f_i(x_{n+1}) \left(\sum_{i=1}^n x_{n+1}f(x_i) - \sum_{i=1}^n x_i f(x_i) \right) \\ &\geq f(x_{n+1}) \left(\sum_{i=1}^n x_{n+1}f_i(x_i) - \sum_{i=1}^n x_i f_i(x_i) \right) \\ &\Leftrightarrow f_i(x_{n+1}) \sum_{i=1}^n (x_{n+1} - x_i)f(x_i) - f(x_{n+1}) \sum_{i=1}^n (x_{n+1} - x_i)f(x_i) \geq 0 \\ &\Leftrightarrow \sum_{i=1}^n (x_{n+1} - x_i) \frac{f(x_i)}{f(x_{n+1})} - \sum_{i=1}^n (x_{n+1} - x_i) \frac{f_i(x_i)}{f_i(x_{n+1})} \geq 0 \\ &\Leftrightarrow \sum_{i=1}^n (x_{n+1} - x_i) \left(\frac{f(x_i)}{f(x_{n+1})} - \frac{f_i(x_i)}{f_i(x_{n+1})} \right) \geq 0 \\ &\Leftrightarrow \sum_{i=1}^n A_i B_i \geq 0, \end{aligned} \tag{21}$$

with

$$A_i = x_{n+1} - x_i, \quad B_i = \frac{f(x_i)}{f(x_{n+1})} - \frac{f_i(x_i)}{f_i(x_{n+1})}.$$

According to Equation 17,

$$\forall i, \quad A_i \geq 0. \tag{22}$$

We reformulate

$$B_i \geq 0 \Leftrightarrow \frac{f(x_i)}{f(x_{n+1})} \geq \frac{f_i(x_i)}{f_i(x_{n+1})} \Leftrightarrow \frac{f(x_{n+1})}{f(x_{n+1})} \geq \frac{f_i(x_i)}{f(x_i)} \Leftrightarrow g(x_{n+1}) \geq g(x_i), \tag{23}$$

which is true for all i according to Equations 16 and 17. Thus, with Equations 22 and 23, Equation 21 is also true, which in turn proves the first case of Equation 5.

