Drastic Changes in Roles of Learning in the Course of Evolution

Reiji Suzuki and Takaya Arita
Graduate School of Information Science, Nagoya University
Furo-cho, Chikusa-ku, Nagoya 464-8601, Japan
{reiji, arita}@is.nagoya-u.ac.jp

Abstract

An interaction between evolution and learning called the Baldwin effect is known as the two-step evolutionary scenario caused by the balances between benefit and cost of learning in general. However, little is still known about dynamic evolutions on these balances in complex environments. Our purpose is to give a new insight into the benefit and cost of learning by focusing on the quantitative evolution of phenotypic plasticity under the assumption of epistatic interactions. For this purpose, we have constructed an evolutionary model of quantitative traits by using an extended version of Kauffman's NK fitness landscape. Phenotypic plasticity is introduced into our model, in which whether each phenotype is plastic or not is genetically defined and plastic phenotypes can be adjusted by learning. The simulation results have clearly shown that the drastic changes in roles of learning cause the three-step evolution of the Baldwin effect (Suzuki and Arita, 2003) and also cause the evolution of the genetic robustness against mutations. We also conceptualize four different roles of learning by using a hill-climbing image of a population on a fitness landscape.

Introduction

The Baldwin effect is known as one of the interactions between evolution and learning, which suggests that individual lifetime learning can influence the course of evolution without the Lamarckian mechanism (Baldwin, 1896). This effect explains these interactions by paying attention to balances between benefit and cost of learning through the following two steps (Turney et al., 1996). In the first step, lifetime learning gives individual agents chances to change their phenotypes. If the learned traits are useful for agents and make their fitness increase, they will spread in the next population. In the second step, if the environment is sufficiently stable, the evolutionary path finds innate traits that can replace learned traits, because of the cost of learning. This step is known as genetic assimilation (Waddington, 1942). Through these steps, learning can guide the genetic acquisition of learned traits without the Lamarckian mechanism in general.

Hinton and Nowlan conducted the first computational experiment of the Baldwin Effect (Hinton and Nowlan, 1987). They assumed an extremely simplified version of a network connection model. The essential point of this study is that they introduced the quantitative evolution of phenotypic plasticity into their model, in other words, they allowed a population to adjust how much it depends on these two adaptive mechanisms through evolution. They revealed the existence of the Baldwin effect by showing the increase and subsequent decrease in the phenotypic plasticity. However, the learning mechanism in their model was too simple on the ground that its benefit was approximately proportional to the number of plastic phenotypes and the cost of learning was explicitly introduced. Thus, further investigations were necessary so as to understand this effect in more realistic situations. Since then, this effect has been discussed in various contexts such as the evolution of the strategies for iterated Prisoner's Dilemma (Arita and Suzuki, 2000), the evolution of developmental mechanisms (Downing, 2004) and so on.

The effects of epistasis are of interest in this field because epistatic interactions among loci are ubiquitous in modern genetics and evolutionary biology. For instance, Mayley conducted an evolutionary experiment using the Kauffman's NK fitness landscape (Mayley, 1996). He pointed out that there should be a neighborhood correlation between genotype and phenotype space to guarantee a genetic assimilation to occur. Bull also discussed the evolution on NK landscapes using a different learning mechanism, and then concluded that whether the learning can increase the fitness or not depends on the ruggedness of the landscape, the probability of learning, and the number of learning iterations (Bull, 1999). However, all phenotypes were plastic and the quantitative evolution of phenotypic plasticity was not introduced into their models. In this sense, the two steps of the Baldwin effect were not clearly discussed in these models when compared with Hinton and Nowlan's model.

Our purpose is to give a new insight into the dynamic evolution of the benefit and cost of learning in complex environments. Especially, we focus on the effects of epistatic interactions on the quantitative evolution of phenotypic plasticity. As a first approach, we have investigated the quantitative evolution of phenotypic plasticity in a neural network as a

Table 1: Three-step evolution of the Baldwin effect. The arrows represents the transitions of indices in each step as follows: "steady (→)", "increasing (↗)", "slightly increasing (⇀)" and "decreasing (↘)".

step	1st	2nd	3rd
lifetime fitness	↗	↗	⇀
innate fitness	→	→	↗
phenotypic plasticity	↗	↘	→
phenotypic variation	→	↗	↘
the standard interpretation	1st	1st and 2nd	2nd

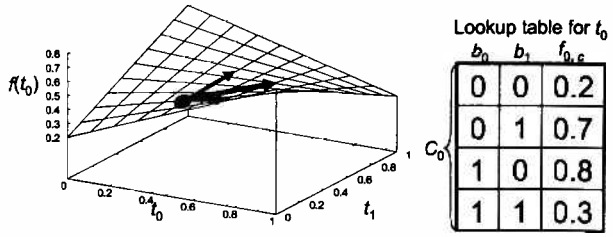

Figure 1: The example of the interpolated fitness of the trait t_0 for N=2 and K=1.

realistic situation (Suzuki and Arita, 2003). The transitions of the phenotypic plasticity and the phenotypic variation revealed that the evolutionary scenario consists of three steps unlike the standard interpretation of the Baldwin effect.

The next approach, discussed in this paper, is to clarify the dynamic changes in roles of learning through the course of evolution by paying attention to the effects of epistasis and to genetic robustness against mutations. For this purpose, we have constructed an evolutionary model based on Kauffman's NK fitness landscape (Kauffman, 1993) in which we can explicitly adjust the degree of epistasis. We discuss the evolution of quantitative traits by extending the fitness evaluation of the NK model. We introduced the phenotypic plasticity into our model, in which whether each phenotype is plastic or not is genetically defined and the plastic phenotype can be adjusted by a simple learning process. By conducting experiments with various degree of epistasis, we show that the drastic changes in roles of learning cause the three-step evolution of the Baldwin effect and then the evolution of the genetic robustness against mutations.

Three-step Evolution of the Baldwin Effect

We investigated the evolution of connection weights in a neural network as a situation where there are epistatic interactions among loci (Suzuki and Arita, 2003). It was observed that the evolutionary scenario consists of three steps by focusing on the transitions of four indices as shown in Table 1. The *lifetime fitness* represents the actual fitness after learning in the population and the *innate fitness* is the potential fitness before learning based on initial phenotypes. The *phenotypic plasticity* represents the proportion of plastic phenotypes in the population. The *phenotypic variation* is the absolute difference in phenotypic values between before and after learning among plastic phenotypes.

The first step, that is the increase in both lifetime fitness and phenotypic plasticity, was simply caused by the benefit of learning. It is noteworthy that the second step has both properties of the first and second step in the standard interpretation of the Baldwin effect. The decrease in the phenotypic plasticity corresponds to the second step in the standard interpretation of the Baldwin effect in the sense the increased fitness by learning becomes dependent on fewer plastic phenotypes. At the same time, the increase in phenotypic variation means that the population becomes strongly dependent on the remaining plastic phenotypes. Thus, we can also say that the population was still in the first step in this point of view. This phenomenon is supposed to be due to the implicit cost of learning caused by the epistatic interactions among plastic phenotypes through the learning processes. The third step corresponds to the second step in the standard interpretation because the genetic assimilation occurred in the remaining plastic phenotypes.

Model

NK Landscape with Real Valued Traits

We have constructed an evolutionary model based on Kauffman's NK fitness landscape (Kauffman, 1993), so as to discuss the evolution of phenotypic plasticity in quantitative traits with / without epistatic interactions among loci. There are P individuals in a population and each individual has N traits of which initial phenotypes are determined by genes in a N-length chromosome GI. Each gene represents the quantitative trait t_i (i=0, ···, N-1) which consists of a real value within the range [0.0, 1.0]. We adopt NK fitness landscapes for evaluation of fitness because we can explicitly adjust the degree of epistasis by using the parameter K. It represents the number of other traits that affect the fitness contribution of each trait. However, the standard NK fitness landscape only assumes the binary traits ("0" or "1"). Then we extended the definition of the fitness evaluation so as to deal with the fitness contributions of quantitative traits.

Each trait t_i has epistatic interactions among other K traits $t_{i+jmodN}$ (j=1, ···, K). For each t_i, we prepare a lookup table which defines its fitness corresponding to all possible (2^{K+1}) combinations of interacting traits when these phenotypes consist of only binary values ("0" or "1"). The value of each fitness in the lookup table is randomly set within the range [0.0, 1.0]. These tables are similar to those of the standard NK landscape.

The fitness for quantitative trait is defined as the linearly interpolated value among the fitness for binary combinations of interacting phenotypes using the following equation:

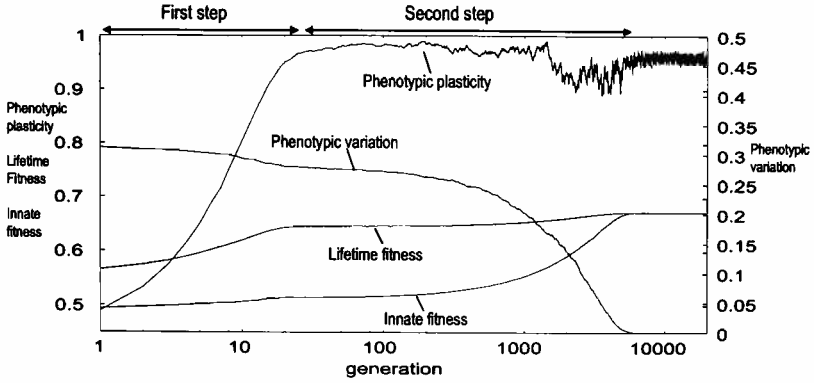

Figure 2: Evolutionary dynamics of fitness and phenotypic plasticity for K=0.

$$f(t_i) = \sum_{c \in C_i} [f_{i,c} \cdot \prod_{j=i}^{i+K} \{(1.0 - b_{jmodN}) \cdot (1.0 - t_{jmodN}) + b_{jmodN} \cdot t_{jmodN}\}], \quad (1)$$

where $f(t_i)$ is the fitness of the trait t_i, C_i denotes the all possible 2^{K+1} combinations of binary traits, $f_{i,c}$ is the fitness of t_i when the combination of traits is c. b_j represents the j th binary phenotype in c. Figure 1 shows an example of the interpolation of the fitness of quantitative traits for $N = 2$ and $K = 1$. The table on the right side represents the lookup table which determines the fitness of t_0 corresponding to four binary combinations of t_0 and t_1. The left figure shows the interpolated fitness of t_0 generated by the right table and the equation (1). The individual fitness is regarded as the average fitness over all traits. Note that if we assume only binary phenotypic values, this model is equivalent to the standard NK fitness landscape.

Learning

Each agent has another N-length chromosome *GP* which decides whether the corresponding phenotype of *GI* is plastic ("1") or not ("0"). Each trait whose corresponding bit in *GP* equals to "1" is adjusted by repeating the following procedure L times. First, for each plastic trait t_i, we calculate the difference in t_i between time t and t+1 (Δt_i) using the following equation:

$$\Delta t_i = \begin{cases} -\beta(F_0 - F_c) & \text{if } max(F_0, Fc, F_1) = F_0, \\ \beta(F_1 - F_c) & \text{if } max(F_0, Fc, F_1) = F_1, \\ 0 & \text{otherwise,} \end{cases} \quad (2)$$

where F_c represents the individual fitness of the current combinations of traits and F_0 is the individual fitness when t_i is set to 0, F_1 is the individual fitness when t_i is set to 1. Next, we actually adjust all values of the plastic traits by adding Δt_i at the same time. This process means that the individual gradually adjust its own plastic phenotypes toward fitter extreme phenotypic value ("0.0" or "1.0") in proportion to the increase in the fitness. The gray arrows in Figure 1 show examples of Δt_0 and Δt_1. The black arrow corresponds to the resultant direction and distance of learning process.

Evolution

After all individuals have finished their learning processes, the population in the next generation is generated by a simple genetic operation as follows: First, the worst individual's chromosomes (*GI* and *GP*) are replaced by copies of the best individual's. Then, every gene for all individuals is mutated with a probability p_m. A mutation in *GW* adds a randomly generated value within the range $[-d, d]$ to the current value and a mutation in *GP* flips the current binary value. If a mutated phenotypic value in *GI* exceeds the domain of the phenotypic space, another mutation is operated on the original value again. We adopted these procedures so as to observe the gradual transitions of four indices explained previously.

Experiments

Experiments without Epistasis (K=0)

We have conducted evolutionary experiments using the following parameters: P=20, N=15, K=0 or 4, L=5, β=10.0, p_m=0.003 and d=0.03. The initial population was generated on condition that initial values in *GI* were taken at random within the range $[0, 1]$ and the proportion of "1" in *GP* for each individual was uniformly distributed also within the range $[0, 1]$.

First, we have conducted the experiments without epistatic interactions among loci. Figure 2 shows the course of evolution over 20000 generations with K=0. The results shown are averages over 10 trials. The horizontal axis represents the generation in logarithmic scale. The lines represent the four indices that we have explained above. Specifically, *lifetime fitness* denotes the average actual fitness

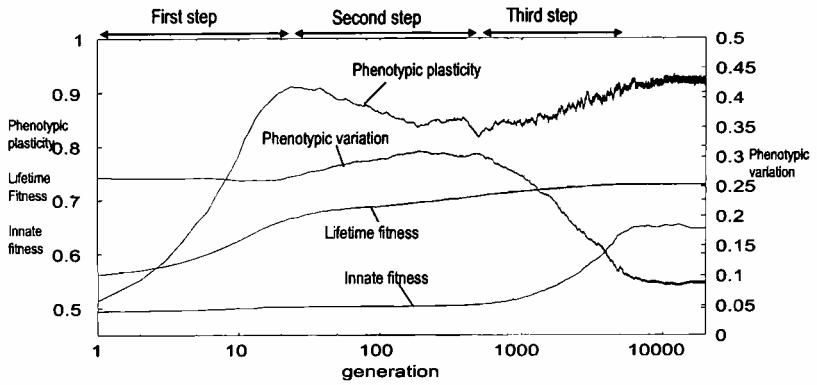

Figure 3: Evolutionary dynamics of fitness and phenotypic plasticity for K=4.

among all individuals calculated after the learning process, and the *innate fitness* is the average potential fitness calculated before the learning process using initial phenotypic values. The *phenotypic plasticity* represents the average proportion of "1" in all *GP* s and the *phenotypic variation* is the average absolute difference between the initial value and the resultant value adjusted by the learning process among all plastic phenotypes.

As shown in the transitions of these indices, the evolutionary process basically consists of the standard two-step evolution of the Baldwin effect. From the initial population, we observe an increase in both lifetime fitness and phenotypic plasticity while the innate fitness remained steady. The phenotypic plasticity rapidly rose and exceeded 0.97 at around the 28th generation. This means that more plastic individuals could obtain higher fitness and could occupy the population due to the benefit of learning.

Next, the innate fitness slowly increased and the phenotypic variation gradually decreased until around the 6000th generation. We can regard that the genetic assimilation occurred on the learned traits because the initial phenotypic values were getting closer to resultant phenotypic values after learning. The main reason for this phenomenon is due to the limitation in the number of iterated learning processes (L). Besides, as the innate fitness got closer to the lifetime fitness, the phenotypic plasticity was slightly decreased until around the 2300th generation. It is due to the genetic drift because there is no explicit cost of learning when $K = 0$.

However, in contrast with the evolutionary scenario in Table 1, the phenotypic plasticity increased again and kept high even after the genetic assimilation had completely finished, despite the fact that the learning did not increased the fitness of the population at all. Thus, another role of learning must occur after the Baldwin effect. This will be discussed later in detail.

Experiments with Epistasis (K=4)

Figure 3 shows the course of evolution for K=4. From the initial population, we observe approximately the same transitions as those for K=0 during the first step, but the peak value of the phenotypic plasticity, 0.91 (at around the 24th generation), was relatively smaller than that for K=0.

However, a clearly different scenario caused by epistatic interactions among loci was observed further on. While the lifetime fitness still slowly increased, the phenotypic plasticity gradually decreased to about 0.82 and then the phenotypic variation increased until around the 500th generation. This phenomenon corresponds to the second step in the three-step evolution of the Baldwin effect, in which the benefit and cost of learning worked together as previously described. The cost of learning is considered to bring about the decrease in the phenotypic plasticity. A contribution of each phenotypic value to the individual's fitness strongly depends on the other phenotypic values when there are epistatic interactions. Similarly, the learning in a plastic phenotype also affects the learning processes of the other plastic phenotypes. However, when we calculate Δt_i for each plastic trait t_i respectively, we do not consider any changes in the other plastic traits. Thus, the learning in too many plastic phenotypes does not always yield an effective increase in the whole fitness. This is the implicit cost of learning caused by the epistatic interactions among loci. In fact, when we maximized the plasticity of individuals in this step, their lifetime fitness tended to be smaller than that of the original individuals. At the same time, the benefit of learning is reflected in the steady transition of the innate fitness and increase in the phenotypic variation, because these transitions mean that the lifetime fitness increased by learning was getting more strongly dependent on the remaining plastic phenotypes.

Finally, the innate fitness eventually began to increase, however in contrast, the phenotypic variation decreased. Thus, the genetic assimilation occurred in the remaining plastic phenotypes because these initial phenotypic val-

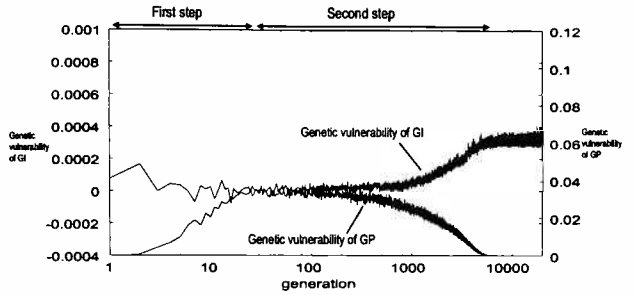

Figure 4: The evolution of the genetic vulnerability for K=0.

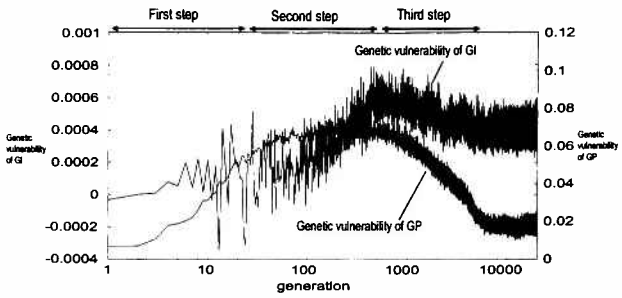

Figure 5: The evolution of the genetic vulnerability for K=4.

ues were getting closer to resultant phenotypic values after learning. In comparison with K=0, the innate fitness converged to around 0.64 and the genetic assimilation did not occurred completely. This step approximately corresponds to the third step in our three-step evolution of the Baldwin effect, except that the phenotypic plasticity gradually increased again to high values as observed for K=0.

Evolution of Genetic Robustness

As discussed in the previous section, we found that the three-step evolution of the Baldwin effect emerges when there are epistatic interactions among loci. However, it is still open to question why the phenotypic plasticity increased again through the last step in both cases of K=0 and K=4. Here, we focus on another different role of learning, that is, the genetic robustness against mutations (DeVisser, 2003).

Instead of measuring the genetic robustness, we measured the genetic vulnerability in view of its adaptive property based on the following procedures: First, for every individual in each generation, we generated a copy of the individual. Then, we conducted the mutational operations on its randomly selected genes in *GI* (or *GP*) for 5 times. We defined the genetic vulnerability of *GI* (or *GP*) as the average difference between the average lifetime fitness in the original population and in the mutated individuals. Thus, the genetic vulnerability becomes smaller when the genetic robustness gets larger.

Figure 4 and 5 shows the evolution of the genetic vulnerability of *GI* and *GP* in the same experiments as Figure 2 and Figure 3 respectively. As a whole, genetic vulnerabilities of *GI* and *GP* for K=4 are larger than those for K=0. We see that the genetic vulnerability of *GP* increased and subsequently decreased in both cases. It should be noticed that its peak exists between the second and third step when K=4. This implies that the population became strongly dependant on learning despite the decrease in the phenotypic plasticity in the second step.

When K=0, the genetic vulnerability of *GI* increased and converged to around 0.0003 along with the increase in the innate fitness. This is because that as the initial phenotypic values get closer to the optimal values, mutations tend to become detrimental for maintaining adaptive traits. However, owing to the recovering adaptive phenotypes by learning with the large phenotypic plasticity, the increase in the genetic vulnerability kept relatively small in comparison with a non-plastic population. Actually, when we eliminated the plasticity of all traits of all individuals in the last generation, the genetic vulnerability of *GI* became much larger than that of the original population.

When K=4, we observe a peak of the genetic vulnerability of *GI* between the second and third step of the Baldwin effect. The increase in the second step implies that the initial values of phenotypes became more important factors for the learning processes in the other plastic phenotypes due to the epistatic effects than the previous step. Its gradual decrease through the third step was accompanied by the increase in the phenotypic plasticity which was restrained by the implicit cost of learning in the previous step. Thus, the increase in phenotypic plasticity was caused by the selective pressure for the evolution of the genetic robustness against mutations on *GI*. The reason why the implicit cost of learning vanished through the third step is due to the fact that epistatic effects among plastic phenotypes got smaller as the initial phenotypes approached to the learned phenotypes through the genetic assimilation in this step.

Conclusion

In the literature of Darwinian evolution, effects of non-genetic factors on genetic evolution (such as Waddington's genetic assimilation or the Baldwin effect) had not been treated as important mechanisms of possible evolutionary change for a long time (Weber and Depew, 2003), while these effects have been investigated based on theoretical or constructive approaches in the field of artificial life or complex systems for more than a decade. However, recent progresses in the molecular and developmental biology have experimentally demonstrated that these mechanisms actually exist and play important roles for genetic evolution in many aspects. They have also revealed that the complex genetic regulations are fundamental basis for developmen-

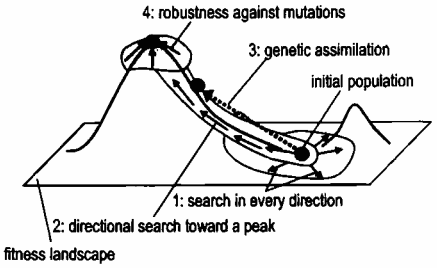

Figure 6: The roles of learning on fitness landscape.

tal process or plasticity of traits. Thus, it is now essential to investigate into evolutionary models based on theoretical or constructive approaches with epistatic effects in conjunction with experimental biology in order to understand these mechanisms in real environments.

As the new first approach, we have discussed the quantitative evolution of phenotypic plasticity based on an extended version of the NK fitness landscape. By conducting the evolutionary experiments with various degree of epistasis, we have found that a three-step evolution of the Baldwin effect emerges when the degree of epistasis is relatively large. It also turned out that the phenotypic plasticity brings about the genetic robustness against mutations.

In conclusion, what needs to be emphasized is that the drastic changes in roles of learning emerged through the course of evolution in order, and then each role was the main selective pressure that guided the complex evolution of phenotypic plasticity. Here, we conceptualize this phenomenon by using a hill-climbing image of a population on a fitness landscape. Figure 6 shows an example of a fitness landscape which consists of all possible phenotypic configurations. Let us assume that the initial population existed on the black filled circle on the right hand. The gray region around it represents the potential area where the current population can reach through learning.

Our experiments suggest that the role of learning changes as follows: 1) The learning in many phenotypes allows the population to search adaptive phenotypes in every direction on the phenotypic space owing to the benefit of learning. 2) The learning in less phenotypes enables the population to get to more adaptive phenotypic configurations by transforming the shape of the potential area due to the implicit cost of learning. This corresponds to the decrease in the phenotypic plasticity and increase in the phenotypic variation. 3) If the potential area reaches a maximum phenotypic configuration, the learning guides the genetic configuration to approach the maximum because of the cost of learning resulted from the limit of the learning ability. This phenomenon corresponds to genetic assimilation. 4) When the genetic configuration completely reaches the maximum, the learning in every direction prevents mutations from dropping down the popula-

tion from the optimum. This state continues until the population looses its stability due to some kind of internal or external factors.

In this scenario, the Baldwin effect corresponds to 1)-2), and 2) is clearly observed when there exists implicit cost of learning such as epistatic interactions. We believe that these detailed investigations into the dynamic evolution of the benefit and cost of learning can help further understanding of the phenotypic plasticity in real biological systems.

References

Arita, T. and Suzuki, R. (2000). Interactions between learning and evolution: The outstanding strategy generated by the baldwin effect. In Bedau, M., McCaskill, J., Packard, N., and Rasmussen, S., editors, *Artificial Life VII*, pages 196–205. Cambridge, MA: MIT Press.

Baldwin, J. M. (1896). A new factor in evolution. *American Naturalist*, 30:441–451.

Bull, L. (1999). On the baldwin effect. *Artificial Life*, 5(3):465–480.

DeVisser, J. A. G. M. e. a. (2003). Perspective: Evolution and detection of genetic robustness. *Evolution*, 57(9):1959–1972.

Downing, K. L. (2004). Development and the baldwin effect. *Artificial Life*, 10(1):39–63.

Hinton, G. E. and Nowlan, S. J. (1987). How learning can guide evolution. *Complex Systems*, 1:495–502.

Kauffman, S. (1993). *The Origins of Order: Self-Organization and Selection in Evolution*. New York: Oxford University Press.

Mayley, G. (1996). Landscapes, learning costs and genetic assimilation. *Evolutionary Computation*, 4(3):213–234.

Suzuki, R. and Arita, T. (2003). The baldwin effect revisited: Three steps characterized by the quantitative evolution of phenotypic plasticity. In Banzhaf, W., Christaller, T., Dittrich, P., Kim, J. T., and Ziegler, J., editors, *Seventh European Conference on Artificial Life*, pages 395–404. Springer.

Turney, P., Whitley, D., and Anderson, R. W. (1996). Evolution, learning, and instinct: 100 years of the baldwin effect. *Evolutionary Computation*, 4(3):4–8.

Waddington, C. H. (1942). Canalization of development and the inheritance of acquired characters. *Nature*, 150:563–565.

Weber, B. H. and Depew, D. J., editors (2003). *Evolution and Learning - The Baldwin Effect Reconsidered* -. Cambridge, MA: MIT Press.