Proceedings of the 2007 IEEE Symposium on Artificial Life (CI-ALife 2007)

Repeated Occurrences of the Baldwin Effect Can Guide Evolution on Rugged Fitness Landscapes

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

***Abstract*— The Baldwin effect is known as a possible scenario of interactions between evolution and learning caused by the balances between benefit and cost of learning. It is still controversial how learning can affect evolution on rugged fitness landscapes because previous studies merely focused on a process in which the population reaches a local optimum through a single occurrence of this effect, even though there exist a lot of local optimums on the landscape. Our purpose is to clarify whether and how learning can facilitate the adaptive evolution of population on rugged fitness landscapes in view of the repeated occurrences of the Baldwin effect. For this purpose, we constructed a simple fitness function that represents a multi-modal fitness landscape in which there is a trade-off between the adaptivity of individual and the strength of the epistatic interactions among its phenotypes. 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 evolutionary experiments clearly showed that the Baldwin effect repeatedly occurred through the evolutionary process of the population on this landscape, and facilitated its adaptive evolution as a whole.**

I. Introduction

Evolution and learning are different adaptive mechanisms that occur on different levels (population or individual) in biological populations. There have been various discussions on effects of learning on the course of evolution. Especially, the Baldwin effect [1] is known as a possible scenario caused by interactions between both mechanisms [2]. This effect explains these interactions by paying attention to balances between benefit and cost of learning through the following two steps [3]. In the first step, lifetime learning gives individual agents chances to change their phenotypes. If the learned traits are beneficial 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 [4]. Through these steps, learning can guide the genetic acquisition of learned traits without the Lamarckian mechanism in general.

A fitness landscape is often used so as to visualize and intuitively understand evolutionary dynamics of the population. The height of the landscape is the fitness value of the corresponding genotype (phenotype) on a possible genetic (phenotypic) space, and the adaptive evolution of the population corresponds to a hill-climbing process on the landscape. A well-known effect of learning on the fitness landscape is that learning smoothes the rugged fitness landscape by allowing the population to explore neighboring regions of phenotypic space.

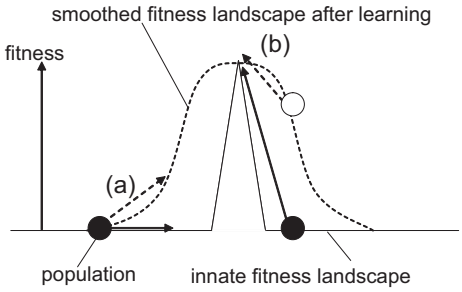

Fig. 1. An example of the smoothed fitness landscape after learning process. The slope of the fitness landscape is (a) increased or (b) decreased after learning process.

Hinton and Nowlan's pioneering work on the Baldwin effect [5] assumed the evolution of the population on a "needle in a haystack" fitness landscape, which is typically illustrated in Fig. 1. By introducing the quantitative evolution of phenotypic plasticity into a simple genetic algorithm, they showed that this effect of learning can guide the evolution of the population toward the spike by increasing the slope of the surface around it (Fig. 1 (a)).

It is also known that learning can bring about a negative effect on the evolution of population. Mayley conducted an evolutionary experiment using Kauffman's NK fitness landscape [6]. He adopted a learning process that searches for an adaptive phenotype in neighboring phenotypes. He pointed out that if the learning process enables the individuals with different innate phenotypes to obtain the same adaptive trait, the ability of the selection to discriminate between them is reduced. Such a negative effect of learning, termed hiding effect, corresponds to the decrease in the slope of the surface around the spike (Fig. 1 (b)). Recently, Paenke *et al.* derived the conditions for the occurrences of these two opposite effects of learning by applying the selection gradient analysis to simple uni-modal fitness functions [7].

Epistatic interactions among loci, which are ubiquitous in modern genetics and evolutionary biology [8], can bring about strong effects on the Baldwin effect. Mayley pointed out that there should be a neighborhood correlation between genotype and phenotype space to guarantee a genetic assimilation to occur [9]. Watson *et al.* also showed the tendency of genetic assimilation to occur increases as the complexity of the learning task decreases and the environmental stability increases

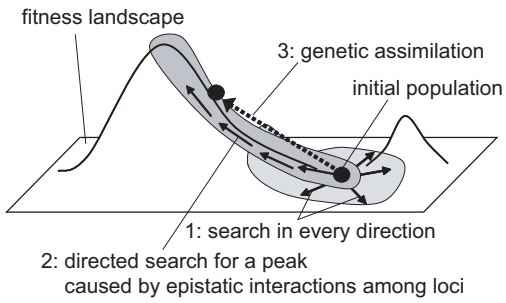

Fig. 2. Three-step evolution through the Baldwin effect observed in [12]. The initial population existed on the black filled circle on the right hand. The gray region around it represents the potential area where the current population can reach through learning. Their experiments suggest that the role 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) If epistasis limits the maximum size of the potential area, the simultaneous learning with smaller number of phenotypes enables the population to get to more adaptive phenotypic combinations by transforming the shape of the potential area. 3) The learning guides the genetic combination to approach the maximum because of the cost of learning.

[10]. Recently, Suzuki and Arita constructed an evolutionary model of quantitative and plastic traits with epistasis by using an extended version of Kauffman's NK fitness landscape [11], [12]. They found that epistatic interactions among loci during learning process caused three-step evolution through the Baldwin effect, which includes the evolution of a more directed search mechanism caused by the reduction of residual phenotypic plasticity. The dynamic changes in roles of learning were conceptualized by using a hill-climbing image of a population and change in the searchable area through learning process on a fitness landscape as shown in Fig. 2.

Previous discussions above mainly focused on the process in which the population reaches a peak of the fitness through a single occurrence of the Baldwin effect. However, such a peak is not always the global optimum when the fitness landscape is rugged because there are a lot of local optimums on the landscape. In this case, we can suppose a further scenario in which the population reached a local peak through the Baldwin effect can subsequently move toward more adaptive peak through another occurrence of the Baldwin effect. Mills and Watson reported that learning enables the population to cross valleys on the fitness landscape if the benefit of learning is sufficient [13]. Thus, we can expect that the population acquires more and more adaptive phenotypes through repeated occurrences of the Baldwin effect. That is, the increase and subsequent decrease in the phenotypic plasticity through the Baldwin effect is expected to create a scaffold for a further adaptive evolution through another Baldwin effect.

Our purpose is to clarify whether and how learning can facilitate the adaptive evolution of population on rugged fitness landscapes in view of the repeated occurrences of the Baldwin effect. Especially, we focus on how these processes can facilitate the adaptive evolution of the population depending on the state of the population on the rugged fitness landscape such as the degree of epistasis. For this purpose, we constructed

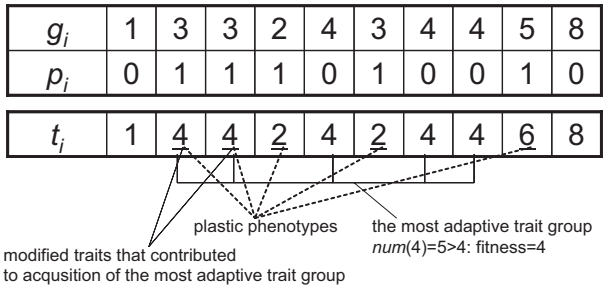

Fig. 3. An example of genetic representation and traits (M=10).

a simple fitness landscape model that represents a multi-modal fitness function in which there is a trade-off between the adaptivity of individual and the strength of the epistatic interactions among its phenotypes. 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 evolutionary experiments clearly showed that the Baldwin effect repeatedly occurred through the hill-climbing process of the population on the multi-modal landscape, and facilitated the evolution of the population as a whole.

II. Model

A. A multi-modal fitness function

There are N individuals in a population and each individual has M traits as shown in Fig. 3. Each gene g_i (i=0, $\cdots$, M-1) in a M-length chromosome GI represents the initial value of the corresponding trait t_i (i=0, $\cdots$, M-1) which consists of an integer value within the range $[1, M]$.

So as to evaluate the fitness of individual, we adopted a simple fitness function defined as follows:

$$fitness = \arg\max(f(n)), \tag{1}$$

$$f(n) = \begin{cases} n & \text{if } num(n) \geq n, \\ 0 & \text{otherwise,} \end{cases} \tag{2}$$

where $num(n)$ is the number of traits of which phenotypic value is n. This function represents a situation as follows: All traits of the individual are divided into several groups each in which the phenotypic values are identical. The trait group of n becomes adaptive and yields the fitness value n if its group size ($num(n)$) is greater than or equals to n. The fitness of individual is taken as the highest fitness value among adaptive trait groups. Figure 3 shows an example of the fitness of an individual. The trait group of 2 and 4 satisfied the condition $num(n) \geq n$, and the actual fitness of the individual becomes 4 in this case.

Eq. (2) shows that the higher the fitness of a trait group is, the larger its minimum size that is needed for its adaptivity to express becomes. The increase in the minimum size means that such a group becomes difficult to be acquired because it needs epistatic interactions with larger number of phenotypes. Thus, there is a trade-off between the adaptivity of the trait groups and their degree of epistatic interactions. In addition, as the

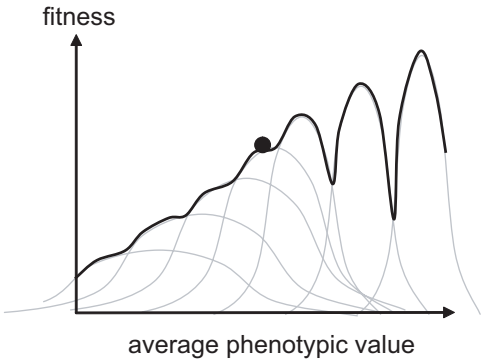

Fig. 4. A rough image of the fitness landscape. The horizontal axis corresponds to the average phenotypic value among all phenotypes. Each peak (in gray) corresponds to the fitness which can be acquired when each trait group becomes adaptive. The black line is the actual fitness of the population. As the fitness of the population increases, it tends to need to cross the valley to reach the next optimum.

fitness of individual becomes larger, it becomes difficult for it to acquire larger fitness without decreasing the current fitness temporarily. As the size of the current adaptive trait groups n increases, it becomes difficult to satisfy both condition $num(n) \geq n$ and $num(n+1) \geq n+1$ at the same time, because of the limitation of the total number of phenotypes M. To be exact, it is impossible to satisfy both conditions above when $n > N/2$. In this sense, there exist valleys on this fitness landscape as roughly illustrated in Fig. 4.

B. Learning

Each individual has another M-length chromosome GP which decides whether the corresponding phenotype of GI is plastic ("1") or not ("0") as shown in Fig. 3. Each trait whose corresponding bit in GP equals to "1" can be changed through lifetime learning procedures as follows: First, each individual calculates the innate fitness which is determined by initial phenotypic values g_i. Then, each individual iterates the learning processes for L time steps. For each step l, the individual evaluates its own traits t_i all of which phenotypic values are determined by the equation as follows:

$$t_i = \begin{cases} g_i + rand() & \text{if } p_i = 1, \\ g_i & \text{otherwise,} \end{cases} \quad (3)$$

where $rand()$ is the function that returns a randomly selected value from {-1, 0, 1}. Note that, if a generated phenotypic value exceeds its domain, another randomly selected value is added to the initial value. This equation shows that the values of plastic traits can slightly deviate from their genetically specified values for each step.

The fitness of the individual at time step l is the highest value of the fitness among previously evaluated l fitness values and the innate fitness. It means that the individual evaluates a set of generated phenotypes, and then adopts the most adaptive traits so far. The lifetime fitness of the individuals is the average among L fitness values evaluated during its learning period and its innate fitness.

C. Evolution

After all individuals have finished their learning processes described above, the offsprings in the next generation is selected by the "roulette wheel selection" (in which the probability that an individual will be chosen as an offspring is proportional to its fitness) from the current population. Then, every gene of all offsprings is mutated with a probability p_{mi} for GI and p_{mp} for GP respectively. A mutation in GI adds a randomly selected value from {-1, 1} to the current value. If a generated value exceeds its domain, another mutation is operated on the original value again. A mutation in GP flips the current binary value.

III. Experiments

We conducted evolutionary experiments using the following parameters: N=400, M=12, p_{mi}=0.002 and p_{mp}=0.005. The initial population was generated on condition that initial values in GI were 1 and the genetic values in GP were randomly decided. We adopted this initial condition so as to observe the hill-climbing process of the population from around the bottom of the fitness landscape.

A. Experiments without learning

First, we conducted experiments without learning process (L=0). Figure 5 shows a typical example of the evolution of innate fitness over 15000 generations. The horizontal axis represents the generation, and the line shows the average innate fitness among all individuals. In this case, the innate fitness is the same as the lifetime fitness.

Starting from the initial population of which the innate fitness was a very small value 1.0, we observed several rapid increases in the innate fitness. Each increase occurred when more adaptive individuals, whose adaptive trait group consisted of $n+1$, appeared by mutations and rapidly occupied the current population in which the individuals' adaptive trait group consisted of n ($n \geq 1$). We also see that the intervals between these rapid increases became longer as the innate fitness increased. It reflects that the more adaptive a trait group is the more difficult it becomes to be acquired by mutations only.

The innate fitness reached the moderate value 6.0 until around 6200th generation, and it converged and never exceeded 6.0 in the case without learning. It is due to the fact that the population could not obtain a more adaptive trait group without discarding the current adaptive trait group as its fitness increases. All the other trials with the same condition also showed the similar evolutionary dynamics except for small differences in the timing of the rapid increase in the innate fitness.

B. Experiments with learning

1) Repeated occurrences of the Baldwin effect: Next, we conducted experiments with learning process. Figure 6 and 7 show a typical example of the trial in case of L=100. In addition to the *innate fitness* and *lifetime fitness*, the *phenotypic plasticity* represents the average proportion of "1" in GP.

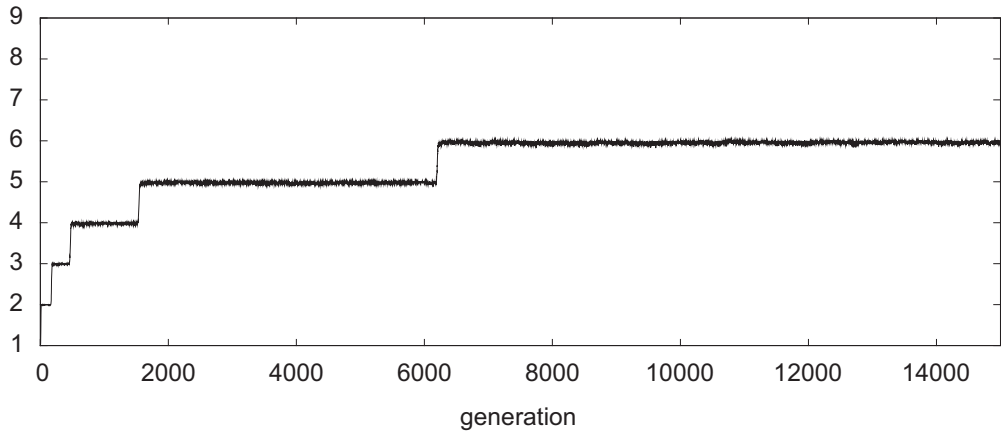

Fig. 5. Evolutionary dynamics of fitness when L=0.

The *plasticity contribution* also represents the proportion of plastic phenotypes which actually contributed to the increase in the fitness of individual. Specifically, for each individual, we counted the number of modified traits in the most adaptive trait group which was first acquired through learning processes as shown in Fig. 3, then divided it by the number of plastic traits of the individual. We adopted this index in order to measure how phenotypic plasticity actually contributed to acquisition of the adaptive traits through learning. These indices are averages among all individuals in each generation.

These graphs clearly show that the Baldwin effect occured repeatedly through generations. We can roughly grasp the evolutionary dynamics by focusing on the transition of the differences between the lifetime fitness and the innate fitness. As shown in Fig. 6, we see that this difference repeated the increasing and subsequently decreasing processes, and each process corresponds to an occurrence of the Baldwin effect. The increase in this difference corresponds to the first step of the Baldwin effect in that the population was getting much dependant on the learning process to acquire higher fitness. Its decrease also corresponds to the second step of the Baldwin effect because the innate fitness increased and finally caught up with the lifetime fitness, and thus, the genetic assimilation occurred.

Here, take the evolution of the population from around 4300th to 6000th generation for example. Around the 4300th generation, the lifetime fitness is almost 6.0 and the average innate fitness almost the same or slightly smaller than the lifetime fitness. It shows that most individuals innately possessed an adaptive trait group of 6 by satisfying the condition $num(6) \geq 6$. The phenotypic plasticity fluctuated around 0.5 that is the expected value when there is no selection pressure on the phenotypic plasticity. Also, the plasticity contribution is quite small (about 0.05). These mean that the population was not dependant on the learning process.

From the 4700th generation, the lifetime fitness began to gradually increase to 7.0. The innate fitness gradually decreased until around the 5500th generation. At the same time, the phenotypic plasticity increased to around 0.7 and the plasticity contribution also became about 0.5. This process shows the occurrence of the first step of the Baldwin effect because more plastic individuals could obtain higher fitness and could occupy the population due to the acquisition of more adaptive trait group by satisfying the condition $num(7) \geq 7$ through learning process.

Subsequently, the innate fitness began to increase and approached to the lifetime fitness until about the 6000th generation. We can say that the second step of the Baldwin effect occurred during this period. In this case, the number of innate traits of 7 increased and finally satisfied the condition $num(7) \geq 7$ innately, which was satisfied through learning process in the first step. Such genetic assimilation occurred because that the lifetime fitness is defined as the average fitness over all learning processes in this model.

After that, the evolutionary dynamics became stable for a thousand generations, and then another first step of the Baldwin effect started around the 7300th generation in this case. This is because that the innately acquired adaptive trait group of 7 enabled the individuals to acquire more adaptive trait group of 8 through learning processes because the value of plastic traits could be increased or decreased from the innate value only by 1. In this sense, we can say that the result of the Baldwin effect became the scaffold for the next Baldwin effect to occur.

As a whole, the innate and lifetime fitness finally reached almost 8.0, which is larger than the previous case without learning (6.0). It means that the repeated occurrences of the Baldwin effect enabled the population to reach higher peak on this rugged fitness landscape.

2) The effects of epistasis on the Baldwin effect: The evolutionary scenarios occurred through each Baldwin effect were basically similar to the case explained in the previous section. Furthermore, we also observed several differences between them, which were caused by the gradual increase in the epistasis of adaptive trait group that was going to be acquired by learning process. The first thing we notice is that

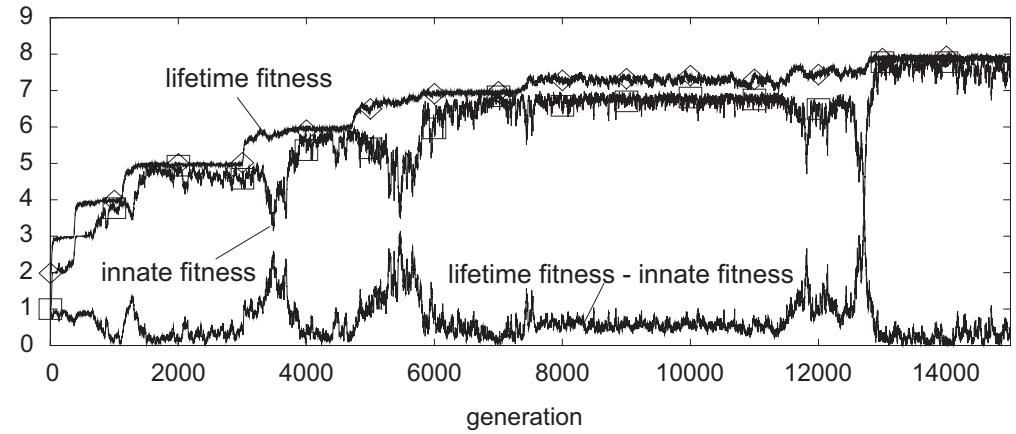

Fig. 6. Evolutionary dynamics of fitness when L=100.

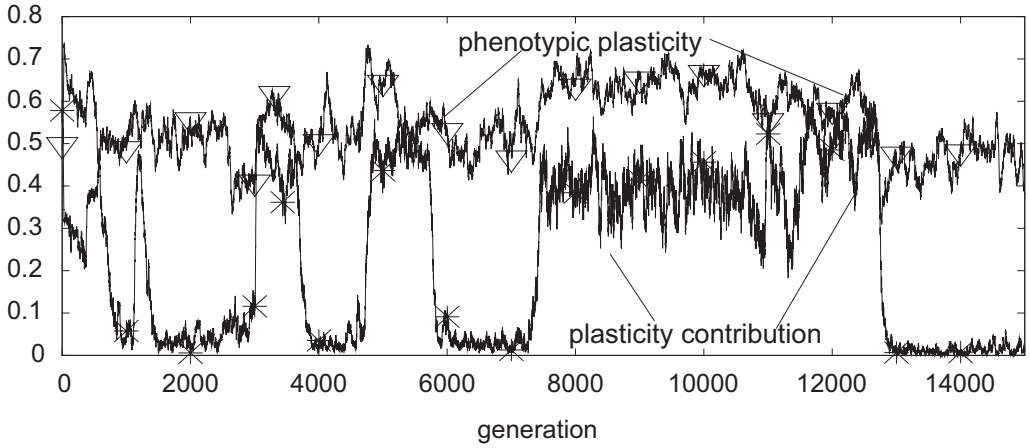

Fig. 7. Evolutionary dynamics of plasticity when L=100.

the Baldwin effect continuously occurred without interval until the 2000th generation. The first step of the Baldwin effect occurred before or immediate after the end of the second step in the previous Baldwin effect as observed around the 400th and 1000th generation. This is because that if the innate fitness is relatively small, the epistasis of the adaptive trait group which is needed to be acquired by learning process is also small. Thus, the individuals can quickly find more adaptive trait group if the genetic assimilation almost finished. Oppositely, as the innate fitness increases, it takes longer generations for the next first step to occur, and there were stable periods between them as shown around the 2000th, 4300th and 6500th generation.

The second thing is that the innate fitness temporarily decreased through the Baldwin effect when the fitness of the population is relatively high. We see such decreases around the 3500th, 5500th and 12700th generation. These clearly show that learning enabled the population to cross the valley on the landscape of the innate fitness. As explained before, it becomes difficult for the individual to acquire more adaptive trait group without discarding the current adaptive trait group as its fitness increases. However, if changes in the innate traits contribute to the acquisition of more adaptive trait group during learning process, such changes can invade the population even if they decrease the innate fitness. Actually, the change in the values of the innate phenotypes from 6 to 7 occurred around the 5500th generation.

The last thing we found is the occurrence of the three-step evolution through the Baldwin effect when the fitness of the population was relatively high. We observe the long term occurrence of the first step of the Baldwin effect from the 7200th to 12500th generation. During this period, we see the slight decrease toward around 0.6 in the phenotypic plasticity that once exceeded 0.7 at the beginning the first step. At the same time, we also see that the plasticity contribution increased from 0.4 to 0.6. The transition of these two indices indicates that learning process became more directed by decreasing residual phenotypic plasticity. Thus, this process corresponds to the second step in the three-step evolution through the Baldwin effect [12].

C. Quantitative analyses on genetic acquisition of the adaptive phenotypes

Finally, we conducted the quantitative analyses on the effects of learning process on the genetic acquisition of the adaptive phenotypes, so as to understand how repeated occurrences of the Baldwin effect facilitated the adaptive evolution in detail. Fig. 8 shows the increasing speed of the average innate fitness in the cases with or without learning. The horizontal axis represents the average innate fitness, and the vertical axis shows the generation at which the innate fitness exceeded the corresponding fitness value. For each case with / without learning, the results are averages over 15 trials[1] conducted for 30000 generations respectively.

From this figure, when the innate fitness was smaller than 3.0, the population with learning took longer generations to acquire innate adaptive phenotypes compared with the population without learning. For example, the population with learning needed about 350 generations to increase its innate fitness to around 3.0, while the population without learning needed about 170 generations only. During this period, the adaptive trait group can be acquired easily with evolution only as shown in Fig. 5. Thus, the introduction of learning rather decreases the increasing speed of the innate fitness. This corresponds to the Mayley's hiding effect of learning [6] explained before.

Oppositely, as the innate fitness became larger, the population without learning took longer generations to acquire adaptive phenotypes until the innate fitness reached around 6.0. For example, the population without learning needed about 11100 generations to increase its innate fitness to around 6.0, while the population with learning needed about 3500 generations only. This phenomenon corresponds to the Hinton and Nowlan's guiding effect [5] in that learning facilitated the genetic acquisition of adaptive traits. Furthermore, the population with learning successfully increased its innate fitness to almost 8.0 until 16000th generation, while the population without learning never reached because it finally got stuck on the local optimum. This phenomenon corresponds to Mills and Watson's crossing valley process [13].

In addition, we conducted additional experiments with various settings of the number of learning iterations L. The results showed that the effects of learning described above became strong as L increased. For example, when L=200, the fitness of the population slowly increased due to the hiding effect in comparison with the cases with L=100. Instead, it successfully reached more adaptive value 9.0 in several trials due to the guiding effect of learning. Oppsitely, when L=10, the fitness of the population rapidly increased but it converged to 7.0 in most trials.

These results clearly show that whether the Baldwin effect facilitates evolution or not strongly depends on the environmental condition in which the population exists. In this model, the epistasis of adaptive phenotypes is an important factor which decides whether learning can guide or hide evolution.

[1]Actually, we conducted 20 trials with and without learning respectively. We found that a few trials in which the average fitness did not reach 7.95 in the former case and 5.95 in the latter within 30000 generations. Then, we excluded these exceptional trials and randomly picked up 15 trials for each case.

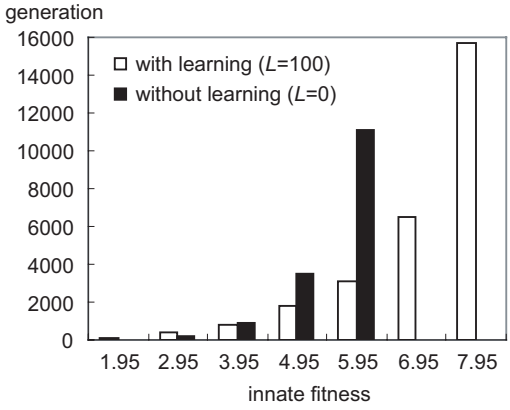

Fig. 8. The increasing speed of the innate fitness. Note that the innate fitness never reached 6.95 in the cases without learning.

IV. Conlusion

There have been many discussions on the effects of learning on evolution such as the Baldwin effect, and various interpretations about roles of learning were proposed recently. However, previous studies mainly focused on a single occurrence of the Baldwin effect, which corresponds to the hill-climbing process of the population toward a local minimum if the fitness landscape is rugged.

We constructed a simple fitness landscape model that represents a multi-modal fitness function in which there is a trade-off between the adaptivity of phenotypes and the strength of the epistatic interactions among them. The evolutionary experiments with / without learning process clearly showed that the Baldwin effect can occur many times through the course of evolution on such a rugged fitness landscape. It should be emphasized that each occurrence of the Baldwin effect has a role of creating a scaffold for another Baldwin effect to occur. However, whether it actually guides the genetic evolution or not depends on the degree of epistasis. If it is difficult for the population to acquire the adaptive phenotypes simply with evolution due to the strong epistasis among adaptive phenotypes, learning can guide evolution by creating the selection pressure toward more adaptive phenotypes. This also enables the population to cross valleys on the innate fitness landscape. Otherwise, learning can hide evolution by decreasing the selection pressure on the innately adaptive phenotypes, and prevents another Baldwin effect to occur. Also, we found the three-step evolution through the Baldwin effect was observed when the degree of the epistasis was large.

Future work includes analyses with other types of rugged fitness landscapes.

REFERENCES

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

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

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

[4] C. H. Waddington, "Genetic assimilation of an acquired character", *Evolution*, 7:118–126 (1953).

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

[6] G. Mayley, "Guiding or hiding: explorations into the effects of learning on the rate of evolution", *Proceedings of the Fourth European Conference on Artificial Life* (pp. 135–144), Berlin: Springer (1997).

[7] I. Paenke, T. Kaercki and B. Sendhoff, "On the influence of lifetime learning on selection pressure", In L. M. Rocha, L. S. Yaeger, M. A. Bedau, D. Floreano, R. L. Goldstone and A. Vespignani (Eds.), *Proceedings of the Tenth International Conference on the Simulation and Synthesis of Living Systems* (pp. 500–506), Cambridge, MA: MIT Press (2006).

[8] J. B. Wolf, E. D. Brodie III and M. J. Wade (Eds.), *Epistasis and Evolutionary Process*, New York: Oxford University Press (2000).

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

[10] J. R. Watson, N. Geard and J. Wiles, "Stability and task complexity: a neural network model of evolution and learning", In R. K. Standish, M. A. Bedau and H. A. Abbass (Eds.), *Proceedings of the Eighth International Conference on the Simulation and Synthesis of Living Systems* (pp. 153–156), Cambridge, MA: MIT Press (2002).

[11] R. Suzuki and T. Arita, "Drastic changes in roles of learning in the course of evolution", In J. Pollack, M. A. Bedau, P. Husbands, T. Ikegami and R. A. Watson (Eds.), *Proceedings of the Ninth International Conference on the Simulation and Synthesis of Living Systems* (pp. 369–374), Cambridge, MA: MIT Press (2004).

[12] R. Suzuki and T. Arita, "The dynamic changes in roles of learning through the Baldwin effect", *Artificial Life, 13(1)*, 31-43 (2007).

[13] R. Mills and R. A. Watson, "On crossing fitness vallys with the Baldwin effect", In L. M. Rocha, L. S. Yaeger, M. A. Bedau, D. Floreano, R. L. Goldstone and A. Vespignani (Eds.), *Proceedings of the Tenth International Conference on the Simulation and Synthesis of Living Systems* (pp. 493–499), Cambridge, MA: MIT Press (2006).