Adaptability of Darwinian and Lamarckian Populations toward an Unknown New World

Yuka Yamamoto[1], Takahiro Sasaki[2] and Mario Tokoro[2]

[1] Department of Computer Science, Faculty of Science and Technology, Keio University 3-14-1 Hiyoshi, Kohoku-ku, Yokohama 223-0061, Japan
[2] Sony Computer Science Laboratory Inc., 3-14-13 Higashigotanda, Shinagawa-ku, Tokyo 141-0022, Japan
E-mail: yuka@mt.cs.keio.ac.jp, {sasaki,mario}@csl.sony.co.jp

Abstract. In this paper, we describe adaptive processes of populations with two distinct mechanisms of evolution, *Darwinian* and *Lamarckian*. We use a simple abstract model where neural networks capable of learning are evolved through GAs. Each individual in the populations tries to maximize its life energy by learning certain rules that distinguish between two groups of materials: food and poison. The best-performing individuals are selected to reproduce offspring according to their mechanism of genetic inheritance, which is either Darwinian or Lamarckian, and the offspring conduct lifetime learning in the succeeding generation. In particular, we examine the adaptability of both populations toward a new unknown world, which is given after some evolutionary steps have taken place under the original world. As the main result, we show that only Darwinian populations can adapt to the new world.

1 Introduction

Through interactions with the environment, or learning, natural organisms may undergo some adaptive changes. However, these *acquired characters* will not be re-encoded in their genes and therefore will not be directly transmitted to their offspring. While *Lamarckism*, which has been one of the major ideas concerned with natural evolution, regards the effect of "inheritance of acquired characters" as the motive force of evolution, *Darwinism* claims that evolution is nothing but the cumulative processes of natural selection with random mutation, and denies the possibility of inheritance of acquired characters[1]. As we know, the mainstream of today's evolutionary theory follows Darwinism, and Lamarckism is regarded incorrect or as a heresy.

Due to this biological background, most of the discussions about learning and evolution from an artificial life research view are based on Darwinian models of evolution[6, 9, 1]. On the other hand, from an engineering view, it is not necessary

[1] Darwin himself accepted the possibility of inheritance of acquired characters, while Weismann, whose fundamental ideas have had a great influence on Neo-Darwinism, completely denied the possibility.

to consider only Darwinian models. Indeed, the possibility of using the heredity of acquired characters would be quite attractive and some studies have shown that significant improvements in the performance of problem-solving systems can be attained by introducing Lamarckian mechanisms[5, 2, 7]. In these works, however, the effectiveness and superiority of Lamarckian evolution have been demonstrated only for static environments (or more generally for fixed tasks).

In our previous papers[10, 11], we compared Darwinian and Lamarckian evolution under non-stationary environments, and showed that the Darwinian population not only showed more stable behaviour in the face of environmental changes but could also maintain greater adaptability regarding such dynamic environments than could the Lamarckian population. While Lamarckian population could adapt itself quickly to a certain single situation, it had difficulty in leaving the specific state of adaptation once it had taken place, owing to its extremely greedy strategy for genetic inheritance. In this paper, which offers an extended discussion using the same model as in [10], we observe the adaptability of both populations toward a new unknown world that is given after some evolutionary steps have taken place under the original world. The populations are first cultivated under a certain world (original world), either static or non-stationary, for a rather long period of generations. After that the populations are given a completely new world, and their adaptability toward this new world is evaluated in terms of learning ability. Note that in this paper, a non-stationary environment, where environmental conditions change repeatedly, is itself regarded as one type of world. We discuss the relation between the adaptability to a new world and the level of innate errors of individuals' neural networks.

2 Experimental Model

Here, we describe our experimental settings. A hundred individuals come into a virtual "world", with 500 units of initial "life energy" for each. Each individual has a feed-forward neural network that serves as its "brain", meaning that the individual takes action based on the network outputs (Figure 1). The neural network has three layers, each of which contains six input neurons, four hidden neurons, and eight output neurons. Each neuron is fully connected to all neurons of the next layer. We take an array of real numbers as a "chromosome" from which the neural network is developed. The chromosome directly encodes all the connective weights of the network[8]. Values of the chromosomes in the initial generation are set randomly, from the range $-0.30 \sim 0.30$.

The world contains two groups of materials, "food" and "poison", both of which have distinctive features as shown in Figure 3, for example. Each material is represented by an array of six bits as shown in Figure 3. Food and poison are distinguished by the upper three bits, and the lower three are set randomly for every new input as noise. Thus, each individual tries to maximize the chance of its own survival by learning the distinction rule, which in this case corresponds to a parity problem of three bits. On each occasion when given any material, an individual inputs the pattern of the material into its neural net-

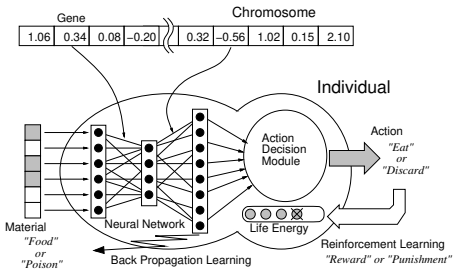

Fig. 1. The architecture of an individual.

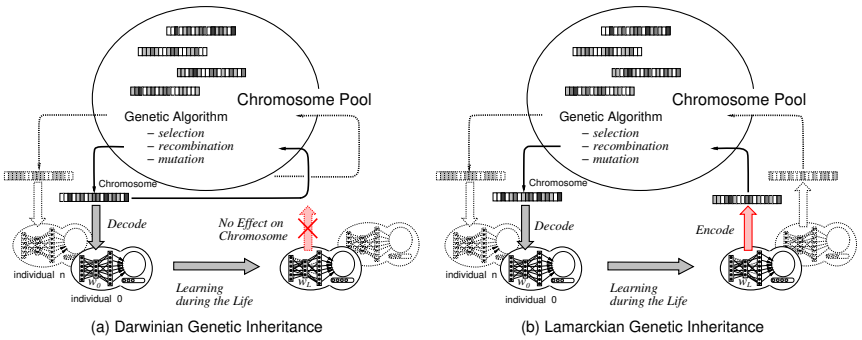

Fig. 2. The genetic inheritance for Darwinian and Lamarckian evolution.

work and stochastically determines whether to "eat" or "discard" it according to the network outputs. These actions, however, are not mapped directly from the outputs themselves. The network outputs are fed once as signals to an "Action Decision Module" (Figure 1), which then finally determines the action of the individual stochastically, according to a Boltzmann distribution. This type of stochastic mechanism is necessary to maintain the possibility of seeking more advantageous behaviours, even if an individual has already acquired a certain adequate behavioural pattern.

If what the individual ate was food, it receives 10 units of energy and tries to train itself to produce the "eat" action with a higher probability for that pattern. Conversely, if the individual ate poison, it loses a comparable amount of energy and tries to train itself to produce the "discard" action with a higher probability for that pattern. When the individual discards the material, no learning is conducted. The aim of each individual is to maximize its life energy by learning a *rule* that discriminates food and poison. We use the *Back Propagation Learning (BP Learning)* algorithm, in combination with a simple *Reinforcement Learning* framework, to train each individual. The coefficients of learning and inertia of BP Learning are set at $\eta = 0.75$ and $\alpha = 0.8$, respectively.

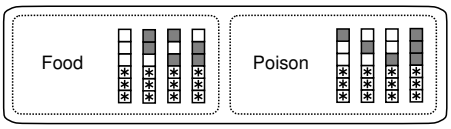

Fig. 3. A static environment. Whether the material is food or poison is determined by the parity property of upper three bits. The material is food if the number of black cells in the upper three bits is even, otherwise poison. The symbol "$*$" means *don't care* whether the cell is black or white; these cells have no effect on the distinction rule.

Each individual is offered a certain number of materials one at a time and learning occurs. We regard this number of repeated events as the length of an individual's "lifetime", which is set at 400 in this paper. At the end of each generation, some of the individuals are selected as parents by a stochastic criterion proportional to the level of their energy, which is thus regarded as their fitness. Selected parents then reproduce new offspring according to their genetic strategies, which is either Darwinian or Lamarckian. The connective weights of the individuals' neural networks must have been modified through their lifetime learning. The result of the modification through the learning is what we here call "acquired characters". Darwinian individuals do not transmit the acquired characters to the next generation, instead they simply pass their chromosomes, which they inherited from their parents, to the process of GAs (Figure 2(a)). On the other hand, Lamarckian individuals re-encode the acquired characters into their chromosomes, and pass them to the process of GAs (Figure 2(b)). Chromosomes of the selected individuals undergo the genetic processes of recombination and mutation. Here, the number of crossing-over points is set randomly from the range $0 \sim 4$. Each mutation at each locus of chromosome occurs at the rate of 5%, with a variation range between ± 0.5. Thus, the selected parents reproduce new offspring, which then undergo lifetime learning in the following generation.

Although the parameters in this paper are set heuristically according to some preliminary experiments, we have confirmed that changing these values within a moderate range results in qualitatively similar outcomes. All the results shown in the following are averages of ten trials.

3 Experimental Evaluations

3.1 Experiment 1: Static environment

Firstly, we consider a static world where the distinction rule between food and poison does not change, and observe the adaptive processes of the Darwinian and the Lamarckian populations.

Figure 4 shows the changes in the average fitness of the populations through generations. As shown by the figure, we see that the Lamarckian population can adapt themselves toward this static world more quickly and effectively than can

the Darwinian population. The result meets our expectations quite well, since the Lamarckian population can continue the learning process that their parents suspended part-way in the previous generation, while the Darwinian population are forced, to some extent, to make a fresh start in each generation.

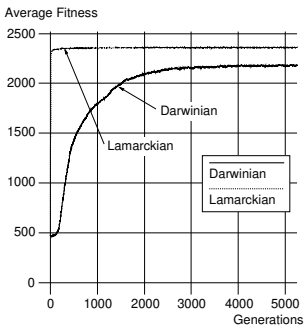

Fig. 4. Average fitness (Static environment)

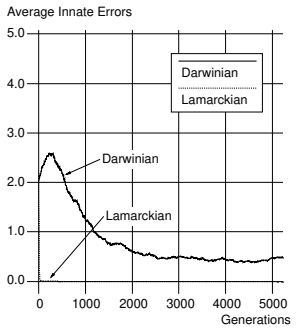

Fig. 5. Average innate errors (Static environment)

Figure 5 shows the changes in the output errors of innate neural network through generations. Network output error is measured by the mean squared error, over every pattern, between the actual outputs and the ideal outputs. Innate output errors are the errors that individuals show before any learning. Therefore, curves of innate errors declining through generations means that individuals came to behave more appropriately from their birth in the later generations. As shown in Figure 5, the Lamarckian individuals decrease their innate errors far more quickly than do the Darwinian individuals.

We must note here (in Figure 5) that the innate errors do decrease even for the Darwinian population, where acquired characters cannot be transmitted from parents to children through genetic processes. This case can be explained from the point of view of the "*Baldwin Effect*"[3]. As Hinton and Nowlan showed in their historical work[6], even if evolution is Darwinian, characters first acquired through learning can then be *assimilated* in the genotype as the generations proceed. One more thing we must note here is that the innate errors for the Darwinian population do not decrease monotonously from the initial generation, but there is a short period where their innate errors increase and thus the curve has a bump. We will discuss this point in Section 4.

3.2 Experiment 2: Non-stationary Environment

Next, we consider a non-stationary world where each material is the same as Experiment 1. Now, however, the distinction rules suffer irregular changes as shown in Figure 6. The following shows the results for a situation where the distinction rule is changed randomly at intervals of 20 generations.

Figures 7(a) and 7(b) show the changes in the average fitness of the Darwinian and the Lamarckian populations, respectively. As one can intuitively

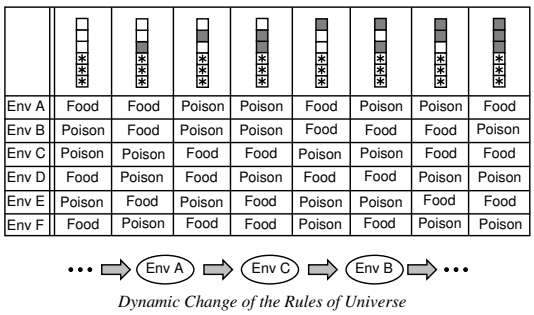

Fig. 6. Non-stationary environment

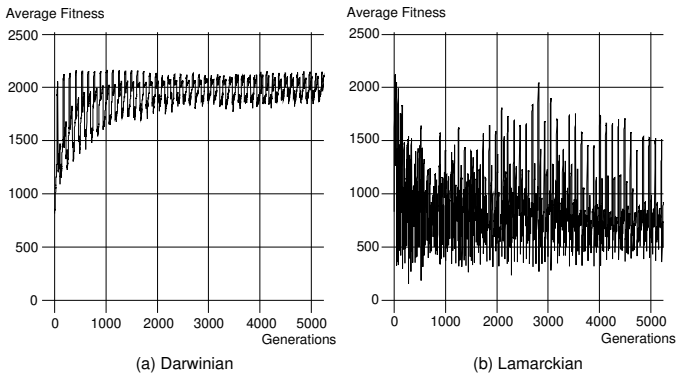

Fig. 7. Average fitness (Non-stationary environment)

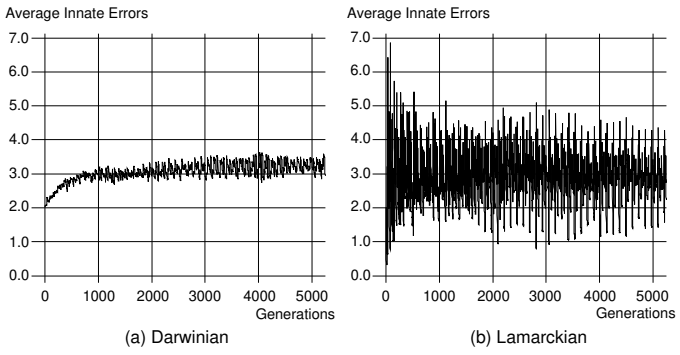

Fig. 8. Average innate errors (Non-stationary environment)

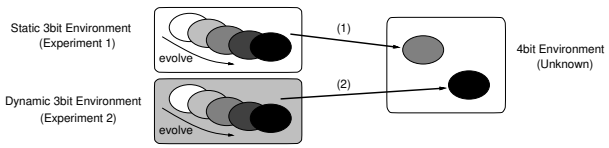

Fig. 9. Setup for Experiment 3. Picking up populations that have been cultivated for some generations under either (1)the static world or (2)the non-stationary world, then put them into a new world.

imagine, the fitness of the Lamarckian population oscillates violently whenever environmental conditions change and thus they can not adapt themselves to this non-stationary world. On the other hand, the point we should especially emphasize is that the fitness of the Darwinian population rises as the generations proceed, although oscillation is observed. We have confirmed in further experiments that the Darwinian genetic mechanism forms, through evolution, a population of individuals that can cope appropriately with any of the rules given in Figure 6(See [10] for more details).

Figures 8(a) and 8(b) show the changes in the average innate errors of the Darwinian and the Lamarckian populations, respectively. As we can see, innate errors of the Darwinian population show smaller oscillations than those of the Lamarckian population. It is interesting here to note that innate errors of the Darwinian population in the earlier generation firstly rise a little and are subsequently maintained at a rather higher level, unlike the results of the previous experiment where the errors decrease via *genetic assimilation* in the later generation(Figure 5). That is to say, individuals become worse from the point of view of their innate behaviour. However, their learning ability becomes better and thus their fitness through their whole life increases as the generations proceed.

3.3 Experiment 3: Adaptability toward New World

Now, we observe the adaptability of the Darwinian and the Lamarckian populations toward a new, unknown world, which is given after some evolutionary steps have taken place under the original world. Firstly, we pick up populations that have been cultivated for some generations under the original world; either the static world (Experiment 1) or the non-stationary world (Experiment 2). Then we put them into a completely new world, and evaluate their adaptability toward this new world in terms of their learning ability (Figure 9). In the new world, materials are distinguished by the upper four bits, and the lower two bits are noise. Individuals cannot solve the parity problem of four bits with the three-bits rule they have acquired under the previous world, so they have to learn the new rule for their survival.

Figures 10(a) and 10(b) show the learning curves under the new world for both populations, which were originally cultivated under the static world (Figure 9(1)). Figure 10 shows the changes in the network output errors through their

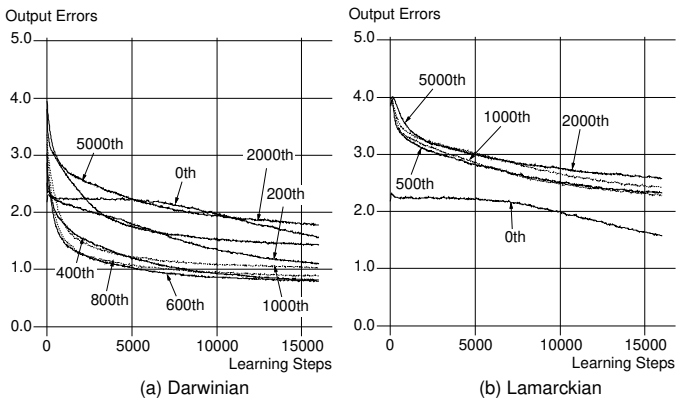

Fig. 10. Learning curves: Static 3 bits world → Unknown 4 bits world

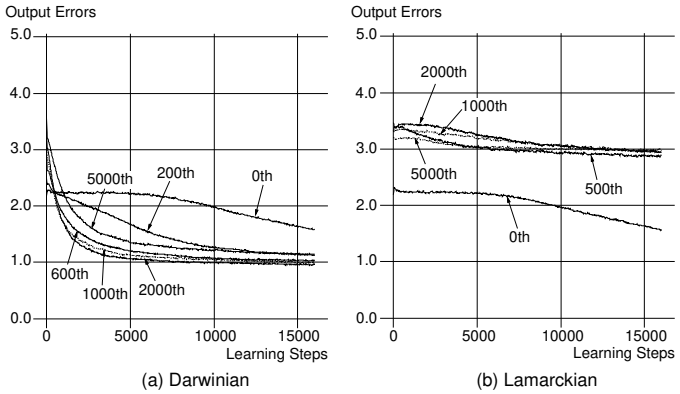

Fig. 11. Learning curves: Non-stationary 3 bits world → Unknown 4 bits world

learning steps. If individuals have adaptability toward the new world and can learn the new rule, their learning curve declines as the learning steps proceed. As expected, Lamarckian individuals cannot learn the new rule at all, since they have deeply adapted for the specific environment under which they were cultivated. Darwinian individuals that have been cultivated for long generations under the original world cannot learn the new rule either. However, those that have been cultivated for rather short periods, say up to 1000 generations, increase their adaptability even toward the new world (Figure 10(a)). The population that has been cultivated for about 600 generations shows the best adaptability.

Figures 11(a) and 11(b) show the learning curves under the new world for both populations, which were originally cultivated under the non-stationary world (Figure 9(2)). The Lamarckian individuals of any generations cannot learn the new rule at all. On the other hand, Darwinian populations that have been

cultivated for some generations, say about 600 generations, can learn the new rule to some extent. Note here, that in this case, unlike the case of populations were originally cultivated under static world, the Darwinian individuals maintain the adaptability toward the new world even after they have been cultivated in the original world for a higher number of generations.

4 Discussion

From the results for the Darwinian population, it seems that evolution under the static world falls into two stages. Firstly, in the earlier stage of evolution, individuals that have better learning ability for the given rule are selected, and thus selection pressure develops the "ability to learn". The increase of the innate errors in the earlier stage of evolution, shown in Figure 5, means that the "ability to learn" is more important than the "ability to perform (innately)" at this moment, and hence selected for evolution. We state here that the "ability to learn" may not be specific for the given rule but rather a general one, since the population increases adaptability even toward the unknown new world (Figure 10(a)). Secondly, once the whole population is occupied by individuals that have high learning ability, selection pressure comes to prefer individuals that show innately better behaviours and thus develops the "ability to perform (innately)", also this is shown from decrease of the innate output errors(Figure 5). Therefore, the population comes to adapt itself deeply to the given world, and comes to lose adaptability toward the new world as the generations proceed (Figure 10(a)). Although the generation of the most adaptable population toward the new world shown in Figure 10(a) does not completely coincide with the generations around the bump of the innate output errors shown in Figure 5, there may be some relations between the non-monotonicity in both cases.

On the other hand, under the non-stationary world where environmental conditions suffer changes, individuals cannot devote themselves to a certain specific condition, but they should maintain the "ability to learn". As the "ability to learn" is a general ability, they also maintain adaptability toward the unknown new world, and the second phase, which occurs under the static world, i.e., selection for the "ability to perform (innately)" does not occur in this case (Figure 11(a)). The innate output errors are maintained at a certain higher level as shown in Figure 8. The above fact also suggests that there may be some relations between adaptability toward the new world and the level of innate output errors.

5 Conclusion

We evaluated how learning with inheritance of acquired characters affects the evolution of the population in a simple abstract model where neural networks capable of learning are evolved through GAs. The results obtained were as follows: Firstly, while Lamarckian individuals could adapt themselves toward a static world quite quickly and effectively, they had difficulty in leaving the specific

state of adaptation once it had been reached and therefore could not cope with another new world at all. Secondly, Darwinian individuals that had been cultivated under a static world, surprisingly increased their adaptability even toward another new world if the new world was given in a relatively early generation, but their adaptability declined as the change occurred in later generations. Lastly, Darwinian individuals that had been cultivated under a non-stationary world also increased their adaptability toward another new world, and could maintain the adaptability even after many generations in the earlier world.

References

1. David H. Ackley and Michael L. Littman. Interactions between Learning and Evolution. In *Artificial Life II, SFI Studies in the Sciences of Complexity, vol.X*, pages 487–509. Addison-Wesley, 1992.
2. David H. Ackley and Michael L. Littman. A Case for Lamarckian Evolution. In *Artificial Life III, SFI Studies in the Sciences of Complexity, vol.XVII*, pages 3–10. Addison-Wesley, 1994.
3. J. Mark Baldwin. A new factor in evolution. *American Naturalist*, 30:441–451, 1896.
4. Richard K. Belew and Melanie Mitchell. *Adaptive Individuals in Evolving Populations: Models and Algorithms*. Addison-Wesley, 1996.
5. John J. Grefenstette. Lamarckian Learning in Multi-agent Environments. In *Proceedings of 4th International Conference on Genetic Algorithms and their applications (ICGA-91)*, pages 303–310, 1991.
6. G. E. Hinton and S. J. Nowlan. How Learning Can Guide Evolution. *Complex Systems*, 1:495–502, 1987.
7. Akira Imada and Keijiro Araki. Lamarckian evolution of associative memory. In *Proceedings of 1996 IEEE The Third International Conference on Evolutionary Computation (ICEC-96)*, pages 676–680, 1996.
8. David J. Montana and Lawrence Davis. Training Feedforward Neural Networks Using Genetic Algorithms. In *Proceedings of the 11th International Conference on Artificial Intelligence (IJCAI-89)*, pages 762–767, 1989.
9. Domenico Parisi, Stefano Nolfi, and Federico Cecconi. Learning, Behavior and Evolution. In *Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*, pages 207–216, 1991.
10. Takahiro Sasaki and Mario Tokoro. Adaptation toward changing environments: Why darwinian in nature? In *4th European Conference on Artificial Life (ECAL-97)*, pages 145–153, 1997.
11. Takahiro Sasaki and Mario Tokoro. Adaptation under changing environments with various rates of inheritance of acquired characters: Comparison between darwinianm and lamarckian evolution. In *2nd Asia-Pacific Conference on Simulated Evolution And Learning (SEAL-98)*, 1998.