Adaptation under Changing Environments with Various Rates of Inheritance of Acquired Characters:
Comparison between Darwinian and Lamarckian Evolution

Takahiro Sasaki[1] 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: sasaki@mt.cs.keio.ac.jp, mario@csl.sony.co.jp

Abstract. In this paper, we study the relationship between learning and evolution in a simple abstract model, where neural networks capable of learning are evolved through genetic algorithms (GAs). The connective weights of individuals' neural networks undergo modification, i.e., certain characters will be acquired, through their lifetime learning. By setting various rates for the heritability of acquired characters, which is a motive force of Lamarckian evolution, we observe adaptational processes of the populations over successive generations. Paying particular attention to behaviours under changing environments, we show the following results. The population with the lower rate of heritability not only shows more stable behaviour against environmental changes, but also maintains greater adaptability with respect to such changing environments. Consequently, the population with zero heritability, i.e., the Darwinian population, attains the highest level of adaptation toward dynamic environments.

1 Introduction

It is obvious that the adaptational processes of natural organisms consist of two complementary phases, each taking place at different spatio-temporal levels: 1) *learning*, occurring within each individual's lifetime, and 2) *evolution*, occurring over successive generations of the population[4, 9, 1]. Here, a simple question arises: "How should these processes of adaptation at the different levels be connected with each other for a greater advantage?" The main goal of this paper is to point to a possible direction for the answers to this question.

In the history of evolutionary theory, there have been two major ideas that give different explanations for the motive force of natural evolution and the phenomenon of genetic inheritance: *Lamarckism* and *Darwinism*. The former regards the effect of "inheritance of acquired characters" as the motive force of evolution. Through interactions with the environment or learning, individuals may undergo some adaptive changes that will then somehow be encoded in their genes and direct the evolutionary process. On the other hand, the central dogma of Darwinism is that the motive force of evolution is "(non-random) natural selection following on random mutation"; mutation itself has no direction, but some individuals with advantageous mutations will have more chance of survival and reproduction through natural selection. It claims that evolution is nothing but these cumulative processes of natural selection. As we know, the mainstream of today's evolutionary theory follows Darwinism, and Lamarckism is regarded as wrong or as a heresy.

In spite of the biological background mentioned above, from the viewpoint of engineering it is not necessary to consider only Darwinian evolution models. Indeed, the possibility of using the heredity of acquired characters would be quite attractive[3, 5]. For this reason, we compared Darwinian and Lamarckian evolution using a simple abstract model in our previous paper[10]. We showed that, under changing environments, the population with Darwinian evolution not only showed more stable behaviour against environmental changes, but surprisingly, could also maintain greater adaptability with respect to such dynamic environments than could the Lamarckian population. While Lamarckian populations could adapt themselves quite quickly to a certain single situation, they had difficulty in leaving the specific state of adaptation once it had taken place owing to their extremely greedy strategy for genetic inheritance. That is why the Lamarckian population in [10] performed poorly under changing environments. Therefore, in this paper, we introduce a new parameter, *heredity rate*, into our Lamarckian model to control the amount of inheritance of acquired characters. We discuss whether there is any appropriate value range of the heredity rate that enables Lamarckian evolution to cope appropriately with changing environments while maintaining quick adaptability toward each single condition.

With other researchers who have recently considered adaptational processes under changing environments[2, 8, 7], we also believe that any evolutionary computation for real-world application must be equipped with adaptability toward dynamic situations. For this reason we have concentrated especially on changing environments.

2 Experimental Model

Here we present our experimental framework and 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). 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[6]. 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, i.e., patterns of bits. For example, in Figure 2, materials in "group A" are food, and those in "group B" are poison. The symbol "$*$" means *don't care* whether the cell is black or white. Thus, food and poison are discriminated by the upper three bits, and the lower three bits are noise.

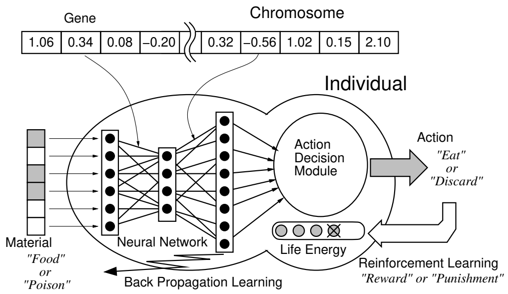

Fig. 1. The architecture of an individual.

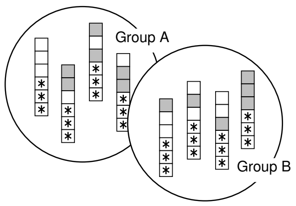

Fig. 2. The materials in the virtual world.

On each occasion when given any material, an individual inputs the pattern of the material into its neural network and stochastically determines whether to "eat" or "discard" it according to the network outputs. These actions 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[10].

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, which in this case corresponds to a parity problem of three bits. We use the *Back Propagation Learning (BP Learning)* algorithm, in combination with a *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.

Each individual is repeatedly offered a certain number of materials, 400 in the current experiments, and learning occurs. We regard this number of repeated events as the length of an individual's "lifetime." 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, i.e. their fitness. Parents re-encode their network connective weights, that suffered modification through their lifetime learning, into their chromosomes according to a given *heredity rate* $\boldsymbol{\tau}$ (Figure 3). In Figure 3, $\boldsymbol{w_0}$ and $\boldsymbol{w_L}$ represent the vectors of connective weights at the time of birth and at the time of reproduction, i.e., after the lifetime learning, respectively. Chromosomes $\boldsymbol{c_e}$ and $\boldsymbol{c_i}$ represent the one from which an individual

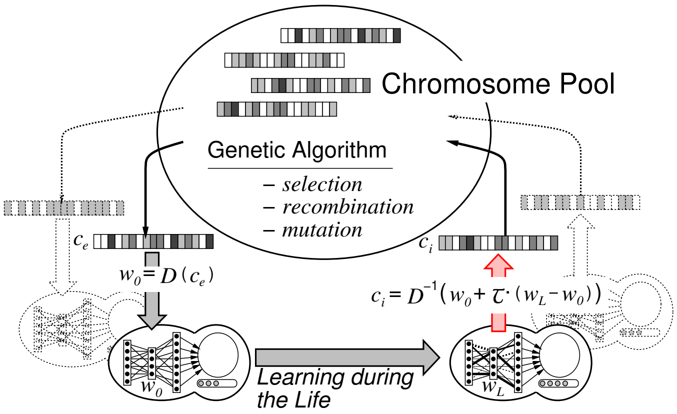

Fig. 3. The mechanisms of genetic inheritance

develops and the one which the individual produces after learning, respectively. $\boldsymbol{D}$ is a mapping from genotype to phenotype, and $\boldsymbol{D}^{-1}$ is its inverse. Individuals with $\tau = 0$ do not re-encode any acquired characters into their chromosomes, but just hand the chromosomes that they inherited from their parents to the process of GAs. On the other hand, with $\tau = 1$ all the acquired characters are re-encoded into the chromosomes. Therefore, we refer to the populations with $\tau = 0$ and those with $0 < \tau < 1$ as *Darwinian* and *Lamarckian*, respectively, and especially refer to those with $\tau = 1$ as *full-Lamarckian*. 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 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.

3 Experimental Evaluations

Now consider a world where food and poison are characterized by arrays of six bits, as shown in Figure 2. At any given time, one group of the two is set up as food and the other as poison. We referred to an environment where "group A" is food as "Env A," and the other where "group B" is food as "Env B." To consider changing environments, we make the world switch between "Env A" and "Env B," so that food and poison swap their roles repeatedly at each particular interval, which is here set at 20 generations. Although a situation such as this may seem to be rather arbitrary or unrealistic, it can actually happen that characters advantageous to survival are totally overturned. A well-known

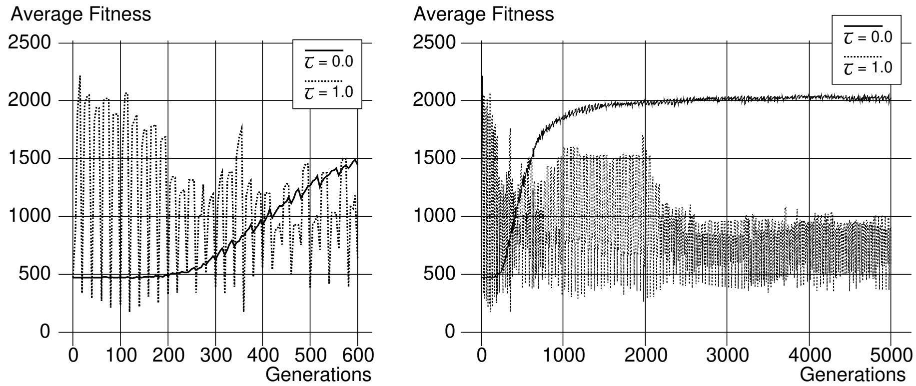

(a) 0 – 600 generations (b) 0 – 5000 generations

Fig. 4. The average fitness: Darwinian ($\tau = 0.0$) vs. full-Lamarckian ($\tau = 1.0$)

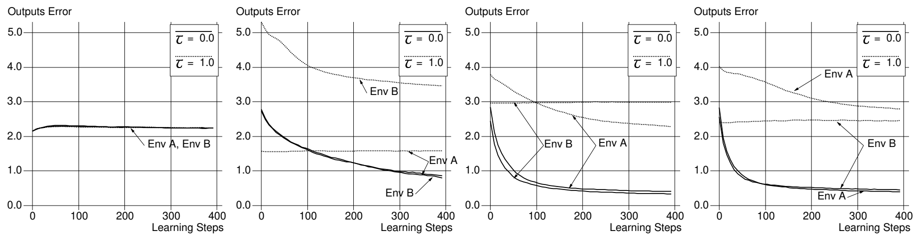

(a) Initial generation (b) 500th generation (c) 2000th generation (d) 5000th generation

Fig. 5. Learning curves: Darwinian ($\tau = 0.0$) vs. full-Lamarckian ($\tau = 1.0$)

example is the *industrial melanism* of certain moths in the Industrial Revolution era in England[11]. Although we obtain qualitatively similar results even if we use a more difficult and random situation, we show here the results for this type of simple case, where the environment is repeatedly overturned, for clarity. All the experimental results, shown hereafter, are the average of 20 runs.

3.1 Darwinian ($\tau = 0.0$) versus full-Lamarckian ($\tau = 1.0$)

Let us first compare the two extreme cases: Darwinian ($\tau = 0.0$) and full-Lamarckian evolution ($\tau = 1.0$). Figures 4(a) and (b) show the changes in the average fitness of the populations. As is evident from the figures, the fitness of the full-Lamarckian population oscillates violently as the environment is overturned, while that of the Darwinian population hardly oscillates and is more stable. The point that we should especially emphasize is that the fitness of the Darwinian population rises over successive generations. This suggests that a population that can cope with both the rules of "Env A" and "Env B" is formed through Darwinian evolution. To confirm this practically, we let four groups of

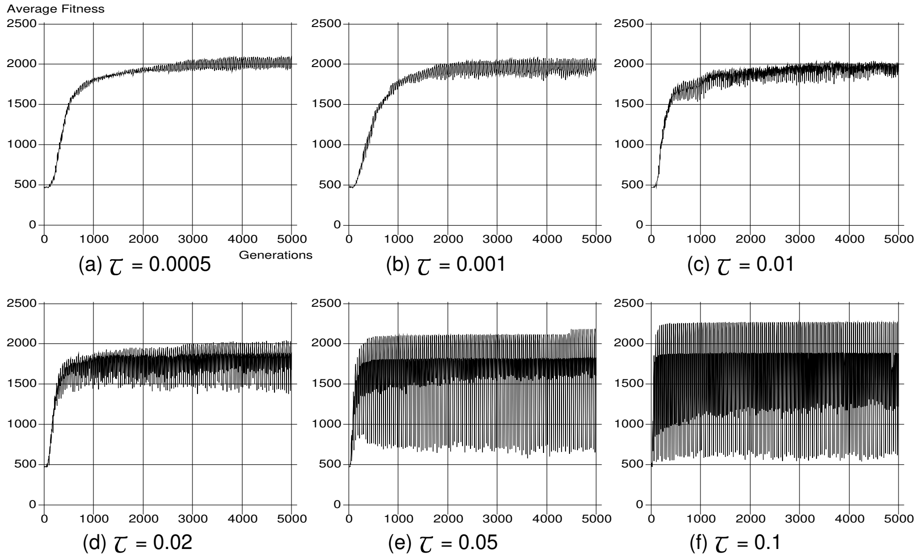

Fig. 6. The average fitness with various heredity rates

populations (an initial generation, $500th$, $2000th$, and $5000th$ generation) conduct learning under each of the two environments. Figures 5(a) – 5(d) show the learning curves of both populations under each environment. The figures show the changes in average output errors for the discrimination ability learned during their lifetime. The mean squared error is used to measure the difference between the actual outputs and the ideal outputs. We can confirm that the Darwinian mechanism forms a population of individuals that learn both rules more appropriately as the generations proceed. In contrast, full-Lamarckian evolution produces individuals that cannot appropriately learn either rule. The two learning curves for Lamarckian individuals in the later generation differ from each other, which means that they cope with one rule better than the other. However, even if the preferred rule is given, the Lamarckian population cannot learn it better than the Darwinian one.

3.2 Controlling Heredity Rate (τ) of Acquired Characters

We now control the amount of heredity by setting τ at various values and observe the evolutionary processes of the populations. Figures 6(a) – 6(f) show the changes in the average fitness of the populations with the heredity rate set at 0.0005, 0.001, 0.01, 0.02, 0.05, 0.1, respectively. As shown by the figures, populations with a higher heredity rate become more unstable than those with a lower rate. The fitness oscillations of the populations with heredity rates smaller than 0.02 are within a rather tolerable range, while the oscillations grow intolerable with larger heredity rates. Figures 7(a) – 7(f) show the learning curves for

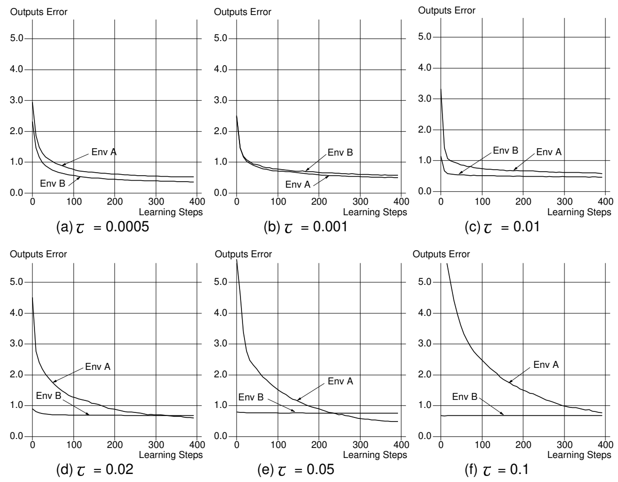

Fig. 7. Learning curves with various heredity rates

each environment of the $5000th$ populations with the heredity rate set at 0.0005, 0.001, 0.01, 0.02, 0.05, 0.1, respectively. As we can see from the figures, the two learning curves become more different from each other as the heredity rate gets higher, which indicates that the evolutionary processes with lower heredity rates produce individuals that can learn both rules more appropriately.

4 Discussion

Under a dynamic environment, the ability to cope with various situations becomes more important than the ability to cope appropriately only with a specific one. The direct Lamarckian effect that greedily transmits the "ability to perform something" is not only useless but also harmful under a changing environment. On the other hand, the indirect *Baldwin effect* that transmits the "ability to learn something" plays a crucial role. Thus, the Darwinian population became most adapted toward the changing environments in our experiment.

In the real biological world, while a phenotype is developed dynamically through quite complex chemical processes according to the information of a genotype, it is very difficult to do the reverse, that is, to determine and compose the corresponding formation of the genotype for a certain phenotype. This is why Lamarckian inheritance is generally said to be infeasible. However, from our experimental results we may suggest another explanation for the essential reason why creatures selected the Darwinian strategy of genetic inheritance in the earlier stages of their evolution. The real world is an environment with strong dynamic characteristics; therefore Darwinian inheritance itself has been an advantageous strategy for adaptation to the real world.

5 Conclusions

We evaluated on a simple abstract model how learning with inheritance of acquired characters affects the evolution of the population, especially under changing environments. By controlling the amount of inheritance, we have shown that populations with lower heritability not only showed more stable behaviour against environmental changes, but also maintained greater adaptability with respect to such changing environments.

Although it must be considered whether our experimental model and results are sufficiently general, we have clarified possible fundamental characteristics that are required for adaptation toward changing environments. Therefore, we believe that the results obtained here may give helpful suggestions in, for example, designing artificial intelligence systems or software agents that will be brought into play under dynamic environments.

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. Helen G. Cobb and John J. Grefenstette. Genetic Algorithms for Tracking Changing Environments. In *Proceedings of 5th International Conference on Genetic Algorithms and their applications (ICGA-93)*, pages 523–530, 1993.
3. 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.
4. G. E. Hinton and S. J. Nowlan. How Learning Can Guide Evolution. *Complex Systems*, 1:495–502, 1987.
5. 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.
6. 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.
7. Naoki Mori, Seiji Imanishi, Hajime Kita, and Yoshikazu Nishikawa. Adaptation to Changing Environments by Means of the Memory Based Thermodynamical Genetic Algorithm. In *Proceedings of 7th International Conference on Genetic Algorithms and their applications (ICGA-97)*, pages 299–306, 1997.
8. Stefano Nolfi and Domenico Parisi. Learning to adapt to changing environments in evolving neural networks. Technical Report 95-15, Institute of Psychology, National Research Council, Rome, 1995.
9. Domenico Parisi, Stefano Nolfi, and Federico Cecconi. Learning, Behaviour and Evolution. In *Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*, pages 207–216. MIT Press, 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. John Maynard Smith. *Evolutionary Genetics*. Oxford University Press, 1989.