Adaptation toward Changing Environments:
Why Darwinian in Nature?

Takahiro Sasaki *__and Mario Tokoro__†
Department of Computer Science,
Faculty of Science and Technology, Keio University
3-14-1 Hiyoshi, Kohoku-ku, Yokohama 223, Japan
{*sasaki,mario*}*@mt.cs.keio.ac.jp*

Abstract

The processes of adaptation in a multi-agent system consist of two complementary phases: 1) *learning*, occurring within each agent's individual lifetime, and 2) *evolution*, occurring over successive generations of the population. In this paper, we observe the dynamics of such adaptive processes in a simple abstract model, where each neural network is regarded as an individual capable of learning, and genetic algorithms are applied as the evolutionary processes for the population of such agents. By evaluating the characteristics of two different mechanisms of genetic inheritance, i.e *Darwinian* and *Lamarckian*, we show the following results. While the Lamarckian mechanism is far more effective than the Darwinian one under static environments, it is found to be unstable and performs quite poorly under dynamic environments. In contrast, even under dynamic environments, a Darwinian population is not only more stable than a Lamarckian one, but also maintains its adaptability with respect to such dynamic environments.

1 Introduction

Conventional artificial systems are usually defined strictly in a *top-down* manner so as to function precisely and effectively for certain purposes, under specifically closed domains. Therefore, those systems lack the adaptiveness to any uncertain or unexpected situations. On the other hand, for natural systems (ranging from swarms of cells to human societies) in the real world, not only their designs but also their entire behaviours emerge through *bottom-up* processes. However, these natural systems do adapt themselves quite well to the real-world environment that exhibits dynamic and unpredictable characteristics, and also, somehow, cope with a variety of difficulties. Therefore, the underlying mechanisms in nature or human societies may be relevant when we consider, for example, novel information processing mechanisms for artificial intelligence systems or software agents that are to be used in an open environment. A research area called *Artificial Life* [7], which analyses mathematical aspects of the dynamics residing in life by synthesis and simulation, has recently received much attention, and considerable advances have been made in using principles of nature as models for possible methods of adaptive information processing.

Any natural system can be regarded, to some extent, as a *Multi-agent system*, which is a kind of environment populated by multiple (semi-)autonomous subjects referred to as agents. Natural ecosystems or human societies are undoubtedly typical examples of the multi-agent system. In such systems, where each agent possesses a certain degree of autonomy, we should consider the processes of adaptation that have taken place, not only at the population level, but also at an individual level. For example, in the world of natural organisms, the adaptation of the system can be viewed as a process consisting of two complementary phases, each taking place at different spatio-temporal levels: 1) *learning*, occurring within each individual lifetime, and 2) *evolution*, occurring over successive generations of the population. We can also observe these kinds of hierarchical adaptive processes in the human economic world, where an agent may be either an individual or a company. Here arises a naive question: "How should these processes of adaptation at the different levels be connected with each other for a higher advantage?" The main goal of this paper is to provide a possible direction for answers to this question.

1.1 Lamarckism and Darwinism

In the following, we consider a world populated by natural organisms as a typical example of multi-agent systems and focus attention on the processes of adaptation

*Also with "Research for the Future" Project, Faculty of Science and Technology, Keio University, Shin-Kawasaki-Mitsui Building West 3F, 890-12 Kashimada, Saiwai-ku, Kawasaki 221, Japan
†Also with Sony Computer Science Laboratory Inc., 3-14-13 Higashigotanda, Shinagawa-ku, Tokyo 141, Japan

in such system. First of all, the behaviour of an organism is not fixed through its lifetime. It develops a tendency to repeat the actions that produce pleasure or benefit, and to avoid those that cause danger or pain. For basic survival, each organism becomes adapted to the environment through its interactions with the environment, by processes called "learning." On the other hand, organisms are not born in a blank state. The basic structure of the brain, which determines the organism's behaviour, as well as its entire body, is developed according to genetic information that is inherited from its ancestors. Such genetic mechanisms, through which features can be inherited in succeeding generations, may not produce exactly the same features in offspring as in parents, because of genetic mutation and recombination. In general, the genes that succeed in constructing an individual better at survival than others in the population, tend to have more copies of themselves reproduced. The cumulative process consisting of genetic mutation and natural selection, leading to improved adaptation of organisms to their environment, is called "evolution."

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. These ideas are called *Lamarckism* and *Darwinism*. The main point of the former is that the motive force of evolution is the effect of "inheritance of acquired characters." Through interactions with the environment or learning, individuals may undergo some adaptive changes, that will then somehow be put in their genes and direct the evolutionary process. On the other hand, the central dogma of the latter 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 through natural selection. It claims that evolution is nothing but these cumulative processes of natural selection. To summarize in other words, while the Lamarckian idea assumes the direct connection between the adaptation at the individual level and that at the population level, the Darwinian idea clearly divides them from each other. As we know, the mainstream of today's evolutionary theory is Darwinian, and Lamarckism is regarded as wrong or as a heresy.

1.2 Adaptation toward Dynamic Environments

Due to the biological background, most of the studies on the issues of learning and evolution have been based on the mechanism of Darwinian genetic inheritance [3, 1, 9]. On the other hand, especially from a pure viewpoint of constructing practical applications, where there is no need to persist in holding the Darwinian model, some studies have attained significant improvements in the system's performance by introducing a Lamarckian mechanism [2, 5]. However, few investigations have attempted a thorough comparison between the Lamarckian mechanism and the Darwinian one, since the advantage of Lamarckian over Darwinian seems rather obvious. As we are going to mention in the following, however, neither of the opinions are satisfactory, not only biologically but also from engineering viewpoint.

With regard to biology, we should be aware that processes that must be regarded as Lamarckian inheritance actually do take place in nature, albeit rarely. For example, a certain kind of water flea develops thorns on its body surface in an environment where many predators exist, and these thorns are transmitted to the offspring through the ovum once this adaptive change occurs [11]. It has been proven that this inheritance is not caused by changes in DNA itself, which are typically involved in what is called evolution, but by changes in the mechanism of genetic switching. However, this Lamarckian process should not be neglected even from the biological viewpoint. With regard to engineering, we would like to point out that most of the previous studies on the application of evolutionary computing took only static environments into consideration, and few observations and discussions have been made on dynamic environments. While it is natural for us to suppose that the Lamarckian mechanism is far more effective than the Darwinian one under static environments, it would not be true under dynamic environments. In such environments, in addition to the requirement of "how well agents can adapt themselves to a certain environment", another requirement arises of "*how well agents can follow the changes in the environment*."

In response to our motivation mentioned above, the following pages focus attention on the adaptive processes of evolutionary agents under dynamic environments. In order to observe its evolutionary dynamics, we construct a simple abstract model, where each *neural network* [10] is regarded as an agent capable of learning, and *genetic algorithms* (GAs) [4] are applied as the evolutionary processes for the population of such agents. Two mechanisms of genetic inheritance are considered, *Darwinian* and *Lamarckian*. We evaluate the adaptability and robustness of the evolutionary agents, and try to clarify the characteristics of mechanisms required for the adaptation toward dynamic environments.

2 Model: A World of Agents

Here, we will describe our experimental framework with a concrete scenario in order to make our discussion more easily understandable.

A hundred agents come into the "world", with 500 units of initial "life energy" for each. In our simulation, each agent is an individual which has a feed-forward neural network that serves as its "brain", meaning that the agent takes action based on the network outputs (Figure

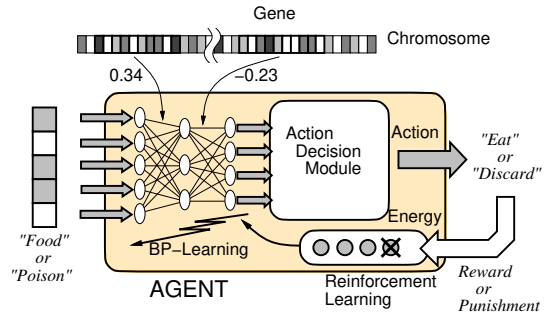

Figure 1: An architecture of an agent

1). The neural network has three layers, each of which contains five or six input neurons, three 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, 12]. Each value of the chromosomes in the initial generation is initialized randomly (between $-0.30 \sim 0.30$), and the better combinations are passed on genetically. It hardly needs to be said that not every neuronal connection of natural organisms can be determined directly from genes. Models closer to the real mechanism of embryogenesis do exist; for example, a grammatical method that applies graph rewriting rules recursively for the development of a network [6]. However, we can still focus on issues of learning and evolution even without such sophisticated models.

The world contains two groups of materials both of which have distinctive features (patterns of bits): "edible" materials and "poisonous" materials. When an agent is given any material, the agent inputs the pattern of the material into its neural network and stochastically determines whether to "eat" or "discard" it according to the outputs of the network. If what the agent ate was a food, the agent 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 agent ate a poison, the agent 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 agent discards the material, no learning is conducted. The aim of each agent is to maximize its energy by learning a *rule* that discriminates food and poison through its experiences. We use the *Back Propagation Learning Algorithm*, in combination with a *Reinforcement Learning* framework, to train each agent. Connective weights of the network are modified by applying the expression typically used in various research on neural networks:

$$\boldsymbol{w}(t+1) = \boldsymbol{w}(t) - \eta \left.\frac{\partial E(\boldsymbol{w})}{\partial \boldsymbol{w}}\right|_{\boldsymbol{w}=\boldsymbol{w}(t)} + \alpha \Delta \boldsymbol{w}(t) \quad (1)$$

Here, two constants η and α are the coefficients of learning and inertia, respectively. In this paper, we set these values as $\eta = 0.75$ and $\alpha = 0.8$. The vector of connective weights at learning step t is represented as $\boldsymbol{w}(t)$, and the error function of a network is represented as $E(\boldsymbol{w})$.

When an agent selects an action based on the network outputs, the action is not mapped directly from the pattern of outputs itself. The network outputs are fed once as signals to an "Action Decision Module" (Figure 1), which then finally determines the action of the agent stochastically according to the signals. As shown in the following, a Boltzmann distribution is used by the action decision module to determine the agent's action.

$$p(a_i|s) = \frac{\exp(o_i/T)}{\sum_{j \in possible_actions} exp(o_j/T)} \quad (2)$$

Here, $p(a_i|s)$ represents the probability that an agent takes an action a_i in a situation s, and o_i represents the total value of the network outputs which corresponds to the action a_i. The degree of "adventurousness" of the agent is controlled by the temperature value T. When T is low, the agent tends to determine its own action by faithfully obeying the network outputs, which reflect its own experience of the past, and shows conservative behaviours. On the other hand, as T takes a higher value, the agent becomes more adventurous and its decisions are less affected by the network outputs. We set the value T to 0.3. This kind of stochastic mechanism is necessary to maintain the possibility of seeking more advantageous behaviours, even if an agent has already acquired a certain appropriate behavioural pattern.

Each agent is repeatedly offered a certain number of materials and learning occurs. We regard this number of repeated events as the length of an agent's "lifetime." At the end of each generation, some of the agents are selected as parents by a stochastic criterion proportional to the level of their energy, which is thus regarded as their fitness. Chromosomes of the selected agents undergo the genetic processes of crossing-over and mutation. Thus the selected parents reproduce new offspring, which then undergo lifetime learning in the following generation. The connective weights of the agents' neural networks must have been modified through their lifetime learning, but Darwinian agents do not transmit the results of the modification to the next generation. They just transmit their chromosomes which they inherited from their parents in the process of GAs (Figure 2a). On the other hand, Lamarckian agents re-encode the connection weights that suffered modification into their chromosomes, and transmit them in the process of GAs (Figure 2b). Here, the number of crossing-over points is set randomly from the range $0 \sim 4$, and positions of the points are also determined randomly. A mutation occurs at the rate of 5%, and its value ranges randomly between ± 0.5. Although the values used in this paper

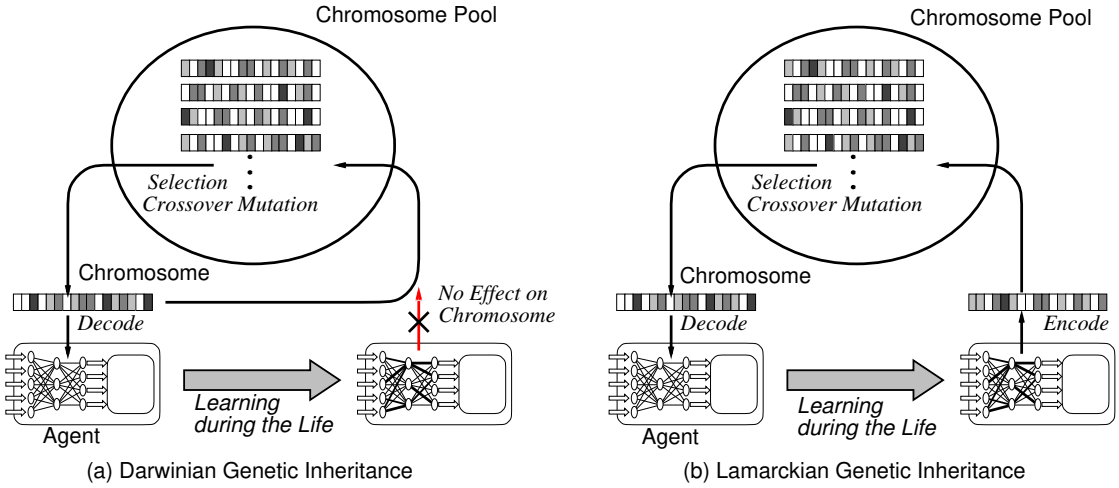

Figure 2: The mechanisms of Darwinian and Lamarckian genetic inheritance

are set heuristically according to some preliminary experiments, we have confirmed that changing these values within a moderate range results in qualitatively similar outcomes.

The flow of our experiments is summarized as follows.

1. A population of the first generation is generated ($g = 1$).
2. Each agent of the g-th generation conducts learning by taking actions during a certain period defined as its lifetime.
3. Fitness of each agent is calculated from the energy which it possesses.
4. Based on their fitness, some agents are selected stochastically as the parents and reproduce offspring of $(g+1)$-*th* generation through genetic processes such as crossing-over and mutation.
5. Increment g by unity ($g = g + 1$), and return to 2.

3 Experimental Evaluations: Darwinian vs. Lamarckian under Dynamic Environments

We have confirmed from the preliminary experiments that the Lamarckian mechanism is far more effective than the Darwinian one under static environments. These results are intuitively understandable, since Lamarckian agents can continue the learning process that their parents have suspended halfway in the previous generation, while Darwinian agents must make a fresh start in each generation.

However, a real-world agent such as a natural organism should cope adaptively with dynamic and complex changes in the environment. In such environments, the most advantageous rule for an agent to learn may change accordingly.

3.1 Experiment 1 – An environment where only partial information is available

First, we consider a situation with a dynamic characteristic in a low degree, where the discrimination rule between food and poison itself does not change, but not all the information necessary to learn the complete rule is available at one time. Moreover, an available piece of information changes with time. This corresponds to a situation where some kinds of unknown materials suddenly appear and other kinds of materials disappear from the world.

Now, let us consider a world where food and poison are characterized by arrays of six bits, as shown in Figure 3. White and black cells represent "0" and "1", respectively. The symbol "*" means *don't care* whether it is "0" or "1". That is to say, food and poison are discriminated by the upper three bits, and the lower three are noise bits. Note that each agent does not "know" of the existence of noise, nor which bits are noise. The agent tries to maximize the chance of its own survival by acquiring the discrimination rule that corresponds to a parity problem of three bits. However, not all the information necessary for the agents to learn the complete rule is supplied at one time. The world considered here contains only four types of materials at one time, where two are food and the other two are poison (Figure 3), and where the constituent materials of the world change with time. We consider a situation where constituent materials of the world change at intervals of 20 generations.

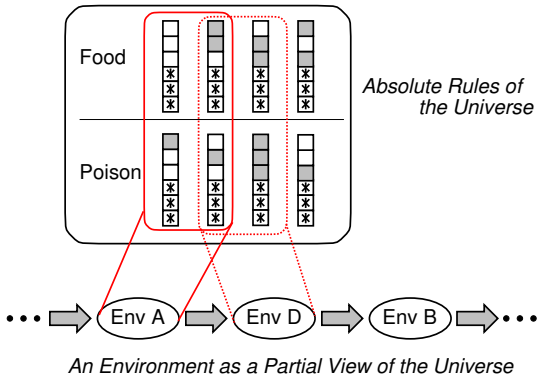

	Food				Poison			
Env A	○	○	×	×	○	○	×	×
Env B	×	×	○	○	×	×	○	○
Env C	○	×	○	×	○	×	○	×
Env D	×	○	×	○	×	○	×	○
Env E	○	×	×	○	×	○	○	×
Env F	×	○	○	×	○	×	×	○

Figure 3: Experiment 1 – An environment where only partial information is available

Both Figures 4a and 4b show the change in the average fitness of the populations. While Figure 4a shows the range from initial generation to 1000*th* generation with a magnified scale, Figure 4b shows the results for the longer span. As is evident from these figures, the fitness of Lamarckian agents oscillates violently as constituent materials change, while the fitness of Darwinian agents oscillates less and is more stable. These results indicate that the mechanism of Darwinian inheritance is superior to that of Lamarckian with regard to robustness against changes in the environment.

The point we would like to especially emphasize here, is that each time the environment changes, the fitness of Darwinian agents does not make a fresh start but gradually increases as the generation proceeds. That is to say, even though only a partial piece of information is available at a time, Darwinian agents seem to be gradually acquiring the complete rule by integrating those pieces of partial information through successive generations. To confirm this practically, we carried out a further experiment. In the first place, agents of the initial generation, 2000*th* generation, 4000*th* generation, and 6000*th* generation were preserved while conducting the experiment whose results are shown in Figure 4. Next, each of the four groups of agents was trained in an environment with the complete set of the rule. That is to say, all the eight patterns shown in Figure 3 were presented to agents, and their learning abilities were evaluated. The learning curves of each generation are shown in Figure 5. Each of the figures shows the average output error curves 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. As shown in Figure 5a, neither the initial generation of Lamarckian agents nor that of Darwinian agents can learn an appropriate discrimination rule. As the generation reaches 2000 (Figure 5b), Lamarckian agents come to output innately somewhat better values, yet the errors are not reduced much through their lifetime learning. Conversely, Darwinian agents tend to output innately somewhat worse values than that of Lamarckian agents, yet they reduce their errors during their lifetime learning. As the generations proceed further (4000*th* generation in Figure 5c and 6000*th* generation in 5d), we can see that a population of agents which appropriately learn the complete rule can be formed through the Darwinian mechanism. On the other hand, the Lamarckian mechanism still produces populations of agents that cannot appropriately learn the complete rule.

The explanation for the Lamarckian agents' unstable behaviour is as follows. For example, with regard to "Env A" shown in Figure 3, agents do not need to learn the perfect rule of the three-bits parity problem, but that of the two-bits parity problem (the XOR problem) is sufficient for discriminating between food and poison, since the third bit of any material takes the same value which in this case is "0". However, let us consider a situation where the world suddenly changes its condition from "Env A" to "Env B". The knowledge of the two-bits parity problem which agents have acquired in "Env A" now becomes not only useless, but also even harmful for their survival, since the acquired knowledge will have the opposite meaning in "Env B". Through the Lamarckian mechanism, agents directly transmit to their offspring the modification of network connections caused by their lifetime learning, thus adapting themselves too deeply to a specific situation. It is difficult for a population to escape from this deep adaptation. In contrast, through the Darwinian mechanism, while each agent can quickly adapt itself to specific situations in the short term with learning at an individual level, agents do not commit to the specific situations, but gradually approach universality in the long term through evolution at the population level.

3.2 Experiment 2 – An environment where the rule changes

In an environment with a dynamic characteristic in a higher degree, the discrimination rule itself may change. Here, we consider a situation where the rules are reversed, so that food and poison swap their roles repeatedly at each particular interval (Figure 6). Although a

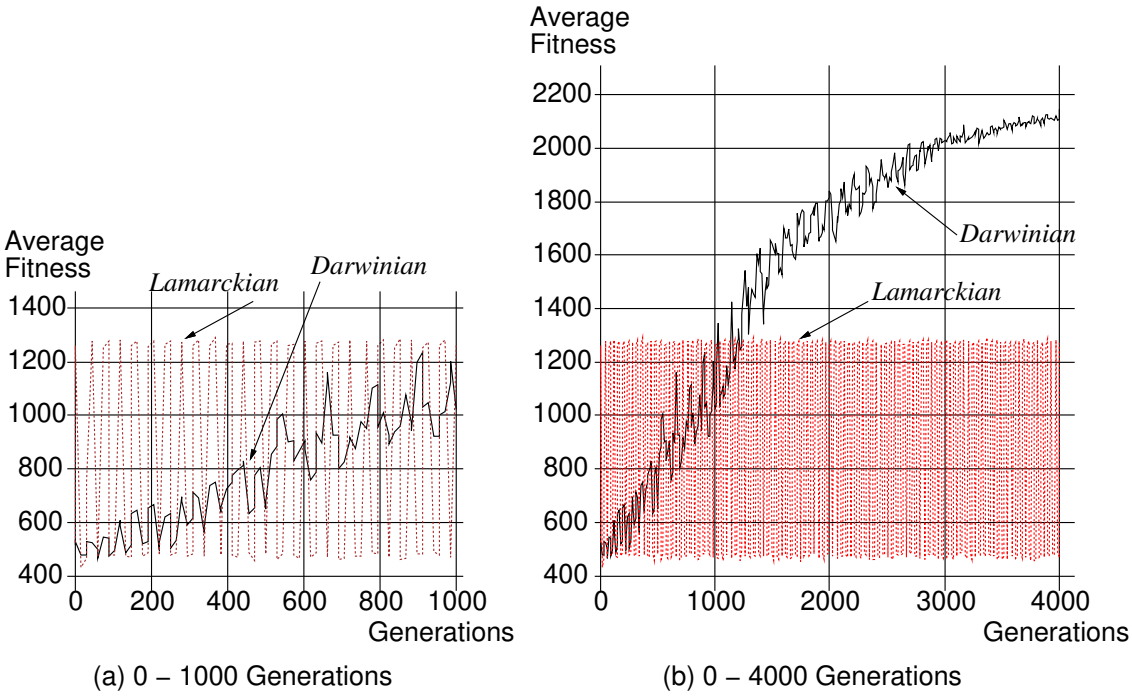

Figure 4: Experiment 1 – The average fitness

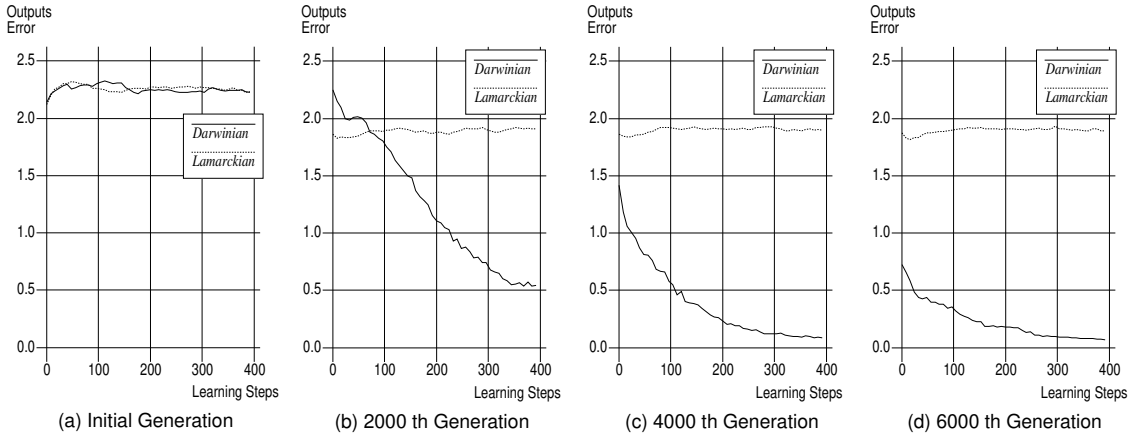

Figure 5: Experiment 1 – The changes in learning curves through generations

situation where conditions advantageous to survival are overturned, as considered here, may seem to be rather arbitrary, it can actually happen. A well-known example is the *industrial melanism* of certain moths, which occurred in the Industrial Revolution era in England [11]. Although details of the example are omitted in this paper due to limitations of space, it indicates that rules which affect survival chances are not eternal but fluid, and may sometimes suffer drastic changes.

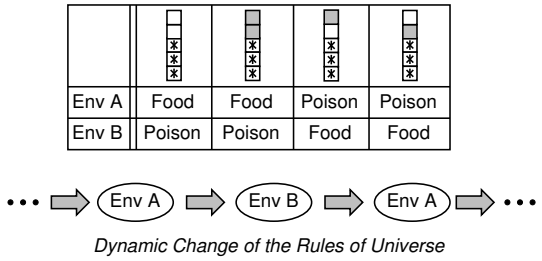

Figure 6: Experiment 2 – An environment where the rule changes

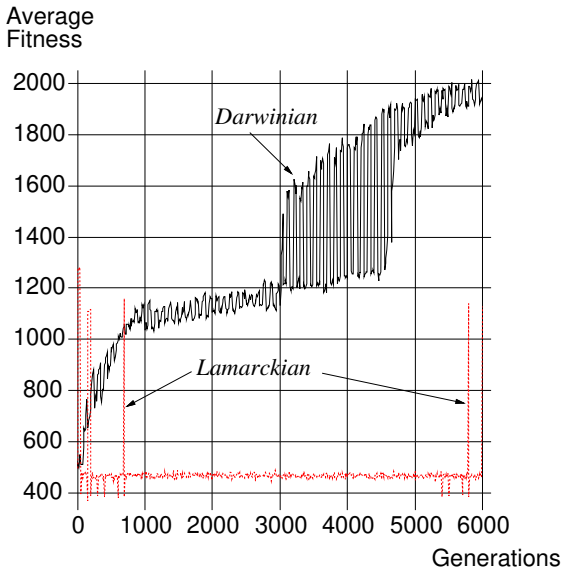

Figure 7: Experiment 2 – The average fitness

As shown in Figure 6, let us consider here a situation where each material is represented as a pattern of five bits. Neglecting the three noise bits, agents can discriminate between food and poison based on the rule of the XOR problem, yet the semantics of the patterns change with time. We consider a situation where the discrimination rule is overturned at intervals of 50 generations.

Figure 7 shows the change in the average fitness of the populations. As one can intuitively imagine, Lamarckian agents cannot adapt themselves to the environment. On the other hand, the point that we should especially emphasize is that the fitness of Darwinian agents rises over successive generations, although oscillation is observed. This seems counterintuitive, since here the discrimination rule itself suffers changes.

This result shows that a population of agents which can cope with both rules of "Env A" and "Env B" is formed through the Darwinian mechanism. To confirm this practically, we let four populations (an initial generation, $2000th$, $4000th$, and $6000th$ generation), which were preserved in the experiment, conduct learning under each of the two rules, and observed their learning curves. Figure 8 shows the results. Neither the initial population of Lamarckian agents nor of Darwinian agents can conduct appropriate learning (Figure 8a). However, by the $2000th$ generation, the difference between the two populations becomes apparent (Figure 8b). As shown in the figure, we can confirm that the Darwinian mechanism forms a population of agents that learns both rules to some extent. As the generations proceed further and reach the $6000th$, Darwinian agents come to learn each of the two rules more appropriately (Figure 8d). In contrast, Lamarckian agents cannot appropriately learn either rule. The two learning curves for Lamarckian agents in the $6000th$ generation differ from each other, since agents cope with one rule better than the other. However, even if the preferred rule is given, Lamarckian agents cannot learn it better than Darwinian agents.

A possible explanation for the surprising behaviour of the Darwinian population is that they may have gradually grasped something common to both rules through the generations. In the environment considered here (Figure 6), materials can at least be grouped into two sets according to the pattern of first and second bits, although the rule of which group corresponds to food and which corresponds to poison changes with time. Agents may acquire the abstract "grouping rule" genetically at a population level, and then learn the details of the discrimination rule that differs from generation to generation at an individual level.

4 Summary

The mechanism of Lamarckian genetic inheritance is far more effective than that of Darwinian agents under static environments, since it merges both processes of learning and evolution in a direct manner, and thus enables agents to adapt quickly toward a given situation. However, since Lamarckian agents adapt themselves to the situation too greedily, they have difficulty in leaving the specific state of adaptation once it has taken place. Therefore, under dynamic environments where rules may suffer changes, the Lamarckian mechanism performs poorly and turns out to behave quite unstably. On the other hand, the Darwinian mechanism maintains stability. In

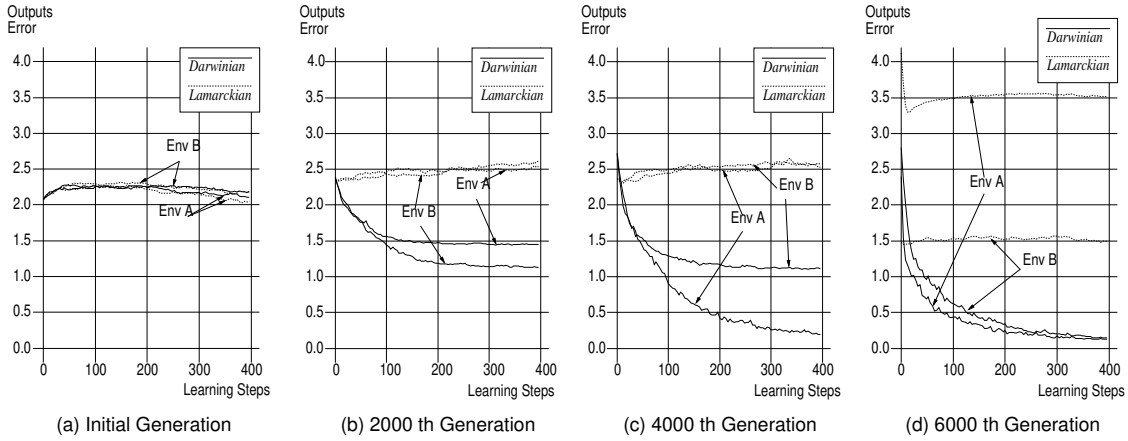

Figure 8: Experiment 2 – The change of learning curves through generations

such environments, the Darwinian mechanism, where the processes of learning and evolution are kept clearly divided, realizes more stable and better behaviour compared to the Lamarckian one. Darwinian agents cope with the detailed changes of rules at an individual level of learning, while to some extent keeping the generality. Moreover, Darwinian agents shows not only stability but also gradual improvements in their fitness over successive generations, even under dynamic environments. That is to say, Darwinian agents have adaptability toward dynamic environments.

Related to the above, there is one thing that we would like to point out concerning the Darwinian mechanism, from the learning ability curves shown in Figure 5 and 8. In Experiment 1, where the rule itself did not change, output errors of innate neural networks decreased in later generations (Figure 5). In short, individuals came to behave appropriately from their birth. On the other hand, in Experiment 2, where the rule itself changed with time, output errors of innate neural networks increased instead (Figures 8). This indicates that, under environments with dynamic characteristics to a higher degree, the ability to behave appropriately from the beginning is not so important; but, to have possibilities for learning to cope with a variety of situations becomes much more important. That is to say, rather than the "ability to *perform* something", the "ability to *learn* something" plays a more important role under the dynamic environments. Since the Lamarckian mechanism transmits the former ability too greedily, it does not work well under the dynamic environments.

5 Discussion

Although the model used in this paper has taken some ideas from the mechanisms of real life, it is an extremely abstract one that simplifies a number of biologically important factors. Therefore, our results may not have a direct impact, for example, on biology. Nevertheless, we may be able to find a number of important suggestions even from the results on our simple model used here.

For example, it is an evident fact that most organisms evolve in the Darwinian manner, although it is reported that processes that can be regarded as Lamarckian inheritance also take place in nature, as we have mentioned in section 1. While a *phenotype* is developed through quite complex processes according to the information of a *genotype* which is encoded on the *DNA*, it is very difficult to determine and compose "in reverse" the corresponding formation of the genotype for a certain phenotype. It is often said that these facts have made Lamarckian inheritance impossible (or strictly speaking, quite rare). However, from our experimental results, we may suggest another explanation for the essential reason why creatures would select the Darwinian strategy of genetic inheritance in the earlier stages of their evolution. Needless to say, the real world is an environment with strong dynamic characteristics; therefore the Darwinian inheritance itself has been an advantageous strategy for the adaptation to the real world.

We can go further with our discussion in connection with this point of view. The immortality of artificial intelligent systems is often considered to be one of their greatest merits, yet our experimental results urge us to reconsider this naive assumption. The Lamarckian mechanism can be regarded in a sense as a mechanism that enables never-ending learning, since the agents can continue the learning process that their parents suspended halfway[1]. The experimental results indicate that, under dynamic environments, the immortality of

[1]Although the Lamarckian mechanism considered here suffers natural selection, chromosomal crossing-over, and gene mutation, we can consider it as a mechanism for never-ending learning, in a rough sense.

artificial systems will turn out to be a flaw. Rather than living forever, the alternation of generations with appropriate intervals will play an important role, and the mechanism of genetic inheritance should be of a Darwinian style, where successive generations conduct learning independently of their parents.

6 Conclusions

Through some simulations using neural networks and genetic algorithms, we evaluated how learning at the individual level with two different inheritance mechanisms affects the evolutionary processes at the population level. The experimental results are summarized as follows:

1. Under a dynamic environment, agents with the Darwinian mechanism are more robust and show more stable behaviour than Lamarckian agents.

2. Moreover, agents with the Darwinian mechanism not only possess stability but also maintain adaptability even toward the dynamic environment itself.

We have clarified the fundamental characteristics which are required for the adaptation toward dynamic environments by using a model of an artificial organisms' world. Although we have used a number of biological terms in this paper, the essential processes concerned are the *collection, exploitation, modification and transmission* of information. Therefore, 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.

Acknowledgments

The authors wish to thank everyone in the Sony Computer Science Laboratory for the fruitful discussions that helped shape our work. Special thanks are due to Dr. Hiroaki Kitano, Dr. Eiichi Osawa, Dr. Jun Tani, and Dr. Toru Ohira for their helpful suggestions with respect to the direction of this work.

References

[1] David H. Ackley and Michael L. Littman. Interactions between Learning and Evolution. In Christphor G. Langton, Charles Taylor, J. Doyne Farmer, and Steen Rasmussen, editors, *Artificial Life II, SFI Studies in the Sciences of Complexity, vol.X*, pages 487–509. Addison-Wesley, 1992.

[2] 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.

[3] G. E. Hinton and S. J. Nowlan. How Learning Can Guide Evolution. *Complex Systems*, 1:495–502, 1987.

[4] J. H. Holland. *Adaptation in Natural and Artificial Systems.* The University of Michigan Press, 1975.

[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] Hiroaki Kitano. Designing Neural Networks using Genetic Algorithms with Graph Generation System. *Complex System*, 4(4):461–476, 1990.

[7] Christopher G. Langton, editor. *Artificial Life: An Overview.* MIT Press, 1995.

[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. MIT Press, 1991.

[10] David E. Rumelhart, James L. McClelland, and the PDP Research Group. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition. (Volume 1: Foundations, Volume 2: Psychological and Biological Models).* MIT Press, 1986.

[11] John Maynard Smith. *Evolutionary Genetics.* Oxford University Press, 1989.

[12] Darrell Whitley and Thomas Hanson. Optimizing Neural Networks Using Faster, More Accurate Genetic Search. In *Proceedings of 3rd International Conference on Genetic Algorithms and their applications (ICGA-89)*, pages 391–396, 1989.