

Language Evolution and the Baldwin Effect

Yusuke WATANABE Reiji SUZUKI Takaya ARITA
Graduate School of Information Science, Nagoya University
Furo-cho, Chikusa-ku, Nagoya 464-8601, Japan

E-mail: watanabe@alife.cs.is.nagoya-u.ac.jp, reiji@is.nagoya-u.ac.jp, arita@nagoya-u.jp

Abstract

Recently, a new constructive approach characterized by the use of computational models for simulating the evolution of language has emerged. This paper investigates the interaction between the two adaptation processes in different time scales, evolution and learning of language, by using a computational model. Simulation results show that the fitness increases rapidly and remains at a high level, while the phenotypic plasticity increases together with the fitness but then decreases and gradually converges to a medium value. This is regarded as the two-step transition of the so-called Baldwin effect. We investigate the evolutionary dynamics governing the effect.

Keywords: language evolution, Baldwin effect, genetic algorithm, recurrent neural network, artificial life.

1 Introduction

Humans are the only species that has evolved sophisticated language. For hundreds of years, many researchers have investigated why and how it could be possible. Recently, a new constructive approach characterized by the use of computational models for simulating the evolution of language has emerged. Language is an emergent system that has been created and maintained through language faculty evolution in a long time scale and cultural change in a short time scale, and thus these models treat either biological evolution or cultural evolution of language. The most obvious purpose of language is to communicate information. If we use natural selection to explain the evolution of language faculty, an individual carrying a “beneficial” grammatical mutation must have a higher fitness. However, how could the mutation be beneficial, if all the other less-evolved individuals in the population could not have understood her [1]. Therefore, it is a very plausible idea that learning combined with evolution played a crucial role in the evolution of language. We focus on the interaction between these two adaptation processes driving the evolution of language in different time scales by using a computational model based on the constructive approach.

The Baldwin effect, which is the focus of this paper, explains the interaction between evolution and learning in general by paying attention to balances between benefit and cost of learning through the two steps [2]. In the first step, life time learning gives individual agents chances to change their phenotypes. If the learned traits are useful for agents and make their fitness increase, they

will spread in the next population. The learning behavior acts as a benefit in this step. In the second step, if the environment is sufficiently stable, the evolutionary path finds innate traits that can replace learned traits (genetic assimilation), because of the cost of learning. Through these steps, learning can accelerate the genetic acquisition of learned traits without the Lamarckian mechanism, which has been clearly demonstrated with a variety of models [3]. When analyzing the interaction between evolution and learning, one of the most important aspects is the cost of learning, because the second step of the Baldwin effect can not occur, if learning is ideal, in other words, there is no cost at all arising from the learning process.

We adopt a speaker-hearer model proposed by Batali [4], in which each agent used a simple recurrent neural network and structured utterance, in other words, partially compositional communication could be obtained by learning from each other. We use the model in a combined framework of cultural learning and genetic evolution. Adopted cultural learning is an extended version of Iterated Learning proposed by Kirby and Hurford [5], which is based on vertical (oblique) communication from adults to children and horizontal communication between adults. Evolution of the weights in the neural network is achieved by a genetic algorithm. In order to examine whether and how the Baldwin effect might occur, we use a mechanism for the evolution of the plasticity (learnability) of each weight in the neural network as we did in [6].

2 Model

A conceptual overview of the model is shown in Figure 1. There are two types of communication: vertical (oblique) communication which is unidirectional transmission from adults to children and horizontal communication which is bidirectional transmission between adults (Figure 2). In the first stage, each child agent learns to interpret the characters produced by a biological parent and randomly-selected cultural parents in each communicative episode. In the second stage, a communicative episode is repeated between a pair of grown-up agents of their generation in which each agent alternates between learning to interpret sequences of the characters produced by other agents and producing sequences of characters. Then, the next generation is produced based on the fitness of agents based on their communicative accuracy in the first and second stages. In the third stage, each




Figure 1: Overview of our model.




Figure 2: Two types of communication.

agent as a parent just produces sequences of characters for their biological and cultural children as their parents do for them in the first stage.

There are two forms of linguistic representation in this model: 1) I-language: Internal language as patterns of connection weights in the neural network, 2) E-language: External language as sequences of uttered characters. Linguistic information in I-language can be inherited from a generation to the next generation via the following two ways: 1) genetic inheritance: initial connection weights of the agents are transmitted to their children through evolutionary operations (Lamarckian inheritance is not adopted), 2) cultural inheritance: E-language is produced from I-language through use and is transmitted to the I-language of the next generation through learning (vertical communication).

Each agent uses a simple recurrent neural network (Elman network) consisting of three layers of neurons (4 character input units each of which corresponds to each of 4 character (a, b, c, or d) and 30 context input units, 30 hidden units and 10 output units). A communicative episode is illustrated in Figure 3. Agents produce sequences of characters to encode structural patterns (vectors) each of which stores 10 real values between 0 and 1. The values in the patterns are partitioned into two groups: the left four of the values are taken as encoding a subject and the right six of the values are taken as encoding a predicate. There are 5 patterns each for the subjects and predicates, and therefore 25 different patterns.

In the beginning of each communicative episode, a subject and a predicate are randomly selected. In order to choose which character to send at each point in a sequence, the speaker agent determines which of the four characters would bring its own output pattern closest to the structural pattern being conveyed. She stops sending if all the speaker’s output units are correct for the structural pattern. If the sequence of characters which the

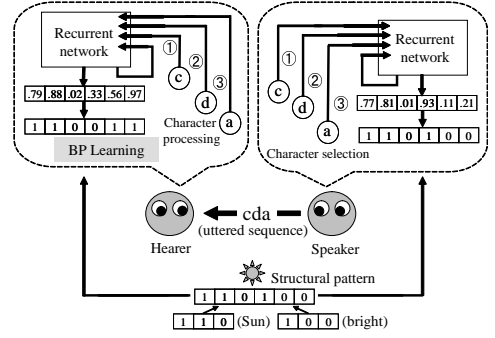


Figure 3: A communicative episode.

agent produces does not reach a limit length of ten characters, the agent succeeds in producing the sequence of characters. The hearer agent then processes the sequence of characters sent by the speaker, and produces an output pattern. The back-propagation algorithm is conducted to modify the weights of the network using the difference between the speaker’s and the hearer’s structural patterns. The network is trained until it converges.

Biological evolution is achieved by a genetic algorithm as follows. Each agent has a pair of chromosomes containing the same number of genes initially assigned to random values. Each gene in the chromosome GW encodes the initial connection weight in the neural network, and each gene in the chromosome GP represents whether the corresponding connection weight in the neural network is plastic (“1”) or not (“0”). If a gene of GP is 0, the corresponding connection weight is invariable in the lifetime. GW consists of a real value within the range $[-1.0; 1.0]$. Agents obtain a reward when they correctly interpret a sequence of characters or when they successfully produce a sequence of characters in a communicative episode regardless of the hearer’s success. Total rewards when the second stage is completed are used as their fitness values. A new population is generated by the tournament selection, and then a mutation is applied with a prespecified probability. A mutation in GW changes the current value into a randomly generated value within the range $[-1.0; 1.0]$ and a mutation in GP flips the current binary value.

3 Experiments

We conducted an experiment for 140 generations. The following parameters were used: N (number of agents) = 100, N_p (number of parents) = 5, r (reward) = 1, m (mutation probability) = 0.01, s (tournament size) = 2, L_v (number of learning iterations for vertical communication) = 990000, L_h (number of learning iterations for horizontal communication) = 1485000. The initial population was generated on condition that initial values in GW were taken at random within the range $[-1.0; 1.0]$ and the proportion of “1” in GP for each agent was uniformly distributed within the range $[0.05; 0.95]$ at intervals of 0.05.

Figure 4 shows the transitions of the fitness that is the




Figure 4: Fitness and plasticity of population.

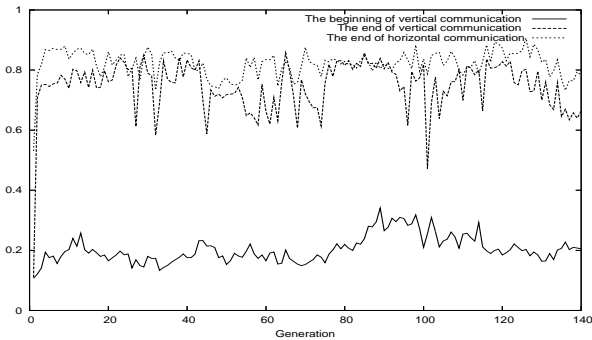


Figure 5: Linguistic coherence among the same generation.

average reward of a agent per communicative episode and the “plasticity of population” which is the ratio of “1” in all GPs of the population. We see that the fitness increased rapidly during the first several generations and kept high afterward, which means the agents have developed an accurate communication system through evolution and learning. Plasticity increased together with the fitness, but then decreased and gradually converged to some medium value (genetic assimilation). This is a typical two-step evolution caused by the Baldwin effect, a key concept clarifying the interaction between evolution and learning.

Figure 5 shows the transitions of the coherence at the beginning or end of each stage. Coherence is the average




Figure 6: Linguistic coherence among two successive generations.




Figure 7: Correlation diagram of fitness and plasticity of population.

ratio of agents uttering the majority sequence for every possible structural pattern. The coherence both after the first and second stages increased rapidly and remained at a high level around 0.85. Also, the coherence before the first stage (innate coherence) moved from 0.1 to 0.2 in the first step. The difference in coherence between before the second stage and the third stage is supposed to mean the diversity amplified by the selected cultural parents in the first stage. Figure 6 shows that the coherence among two successive generations tends to increase while it shows a chaotic oscillation.

Table 1: A part of the characters shared most in the 140th generation. In most of the sequences, we observed that subjects correspond to suffixes (bold font) and predicates correspond to prefixes (underline).

subjects and predicates	1000	1011	0101
011100	<u>bb</u> d	<u>bb</u> c	<u>b</u> d
101001	<u>aa</u>	<u>ac</u>	<u>ad</u>
100011	<u>ca</u>	<u>cc</u>	<u>cd</u>

Table 1 shows a part of the sequences used by a majority of the population in the 140th generation. The agents in the 140th last generation tended to share a little shorter sequences than previous generations. Also, syntactic regularities in the order of token sequence tended to be observed more clearly.

Here, we investigate the evolutionary dynamics which governs the Baldwin effect. The agents in the first generation varies greatly both in the amount of the plastic phenotypes and the connection weights of the network. Agents with more plastic connections could communicate with others successfully in this situation and therefore could occupy the population within several generations. Figure 7 shows the correlation between plasticity and fitness in the 1st generation. This also supports the scenario that the plasticity drives the evolution in this step.

In the second step, plasticity gradually decreased to about 0.55 around the 140th generation while keeping the fitness high. This shows a dramatic change between the both steps in the necessity of the plasticity caused by the change in genetic organization. We conducted several additional experiments in order to clarify the factors




Figure 8: The distribution of the proportion of the plastic connection weights in each locus.

driving the evolution in the second step. As a result, the following factors drastically decreased the selection pressure for the evolution of plasticity compared with the evolution of the connection weights, and thus it gradually decreased to 0.55 by a random drift in the second step.

The first factor is the decrease in the necessity of learning caused by a linguistic shift towards easier language. The variation in the initial connection weights decreased rapidly in the first stage, which made the learning in the second stage easier because the language that should be learned and shared among agents became nearer to the innate language of each agent. The fact that the variation in the plasticity in the population decreased rapidly in the first stage also decreases the selection pressure for plasticity. We see that the coherence was about 0.1 in the first generation while it was 0.2 in the generations of the second step, which supports this explanation.

There is another language-specific factor. Figure 8 shows the distribution of the proportion of the plastic connection weights in each locus. It is shown that the architecture of the network (the location of the plastic connection weights) evolved to be identical. Also, there is a possibility that regularization and compactification in the uttered sequences played a role to become easier for agents to acquire. The fact that the coherence among two successive generations tended to keep high afterward is supposed to show this possibility. These factors are specific to language evolution and seem parallel with the idea by Deacon that language evolves to be adaptive to human cognitive capacity [7].

The second factor is implicit cost associated with learning. In our experiments, there is no explicit learning cost in learning such as fitness tax which is proportional to the learning period. Also, the length of each communicative episode seems enough for learning to converge, which is supported by the result of the experiment (not shown) in which the length of the episode was doubled. However, the interactions among weights or genetic epistasis based on abundant plasticity could cause the bad effect in the back-propagation learning, leading to the decrease in the reward. In this sense, phenotypic plasticity evolves under such selection pressure [6]. This type of cost becomes larger, as the necessity of learning decreases through the phase of evolution. Note that epistasis could have an opposite function to repress genetic assimilation

by making the relationship between genotype and phenotype less correlated [8].

Another aspect of implicit cost, which is purely specific to linguistic evolution, is related with the variation in parents. In the first stage, each child learns sequences uttered from the biological and randomly selected cultural parents. Overlearning to the parents with new local dialect owing to mutation could cause a decrease in rewards both in vertical communication and succeeding horizontal communication.

4 Conclusion

This paper investigates the interaction between evolution and learning of language by using a computational model which we believe to be a minimal model to capture the essence of it. We have found that the factors specific to language evolution or linguistic behavior might have a crucial role in shaping its evolutionary pathways. Specifically, it has been shown that the second step in the Baldwin effect (genetic assimilation) could be driven by the random drift caused indirectly by the adaptive shift in language or overlearning to a variety of parents. It should be noted that genetic assimilation in this evolutionary scenario does not necessarily need unchanged linguistic environment.

References

- [1] Geschwind, N.: Some Comments on the Neurology of Language in *Biological Studies of Mental Processes* (ed. D. Caplan), MIT P, 1980.
- [2] Turney, P., Whitley, D. and Anderson, R. W.: Evolution, Learning, and Instinct: 100 Years of the Baldwin Effect, *Evolutionary Computation*, 4(3), 4-8, 1996.
- [3] Hinton, G. E. and Nowlan, S. J.: How Learning Can Guide Evolution, *Complex Systems*, 1, 495-502, 1987.
- [4] Batali, J.: Computational Simulations of the Emergence of Grammar in *Approach to the Evolution of Language*, 405-426, Cambridge UP, 1998.
- [5] Kirby, S. and Hurford, J. R.: The Emergence of Linguistic Structure: An Overview of the Iterated Learning Model in *Simulating the Evolution of Language* (ed. A. Cangelosi and D. Parisi), Springer Verlag, 2002.
- [6] Suzuki, R. and Arita, T.: The Baldwin Effect Revisited: Three Steps Characterized by the Quantitative Evolution of Phenotypic Plasticity, *Proc. of 7th Eur. Conf. on Artificial Life*, 395-404, 2003.
- [7] Deacon, T. W.: *The Symbolic Species: The Co-evolution of Language and the Brain*, Norton, 1997.
- [8] Yamauchi, H.: The Difficulty of the Baldwinian Account of Linguistic Innateness, *Proc. of the 6th Eur. Conf. on Advances in Artificial Life*, 391-400, 2001.