The Baldwin Effect in the Immune System: Learning by Somatic Hypermutation

Ron Hightower
Dept. of Comp. Sci.
University of New Mexico
Albuquerque, NM 87131 USA
high@cs.unm.edu

Stephanie Forrest
Dept. of Comp. Sci.
University of New Mexico
Albuquerque, NM 87131 USA
forrest@cs.unm.edu

Alan S. Perelson
Theoretical Division
Los Alamos Nat. Lab.
Los Alamos, NM 87545 USA
asp@receptor.lanl.gov

1 Introduction

Almost one hundred years have passed since J. Mark Baldwin proposed [1] that learned or acquired characteristics could become part of the genetic makeup of succeeding generations without Lamarckian inheritance. Individuals that learn or acquire useful characteristics during their lifetimes tend to survive, and, Baldwin claimed, this would cause succeeding generations to have a higher probability of acquiring the same characteristics, even though the characteristics themselves were not genetically propagated. Baldwin called this "a new factor in evolution," and it has come to be known as the *Baldwin effect*.

In recent years the Baldwin effect has been shown to exist in several experiments involving the simulated evolution of learning systems (for example, [2] [3]). These simulations have typically used the genetic algorithm (GA) to act the part of evolution and neural networks to play the role of the learning mechanism. The nervous system, however, is not the only part of an organism that is capable of learning. This paper discusses the Baldwin effect with respect to the *immune system*—a learning mechanism that rivals the complexity and computational power of the nervous system.

In previous work we used a binary model of the immune system to study the effects of evolution on the genetic representation of antibodies [4]. Here we have extended this model to include "clonal selection," which is the learning process used by the immune system. Once this learning mechanism was incorporated into the model it became possible to observe the Baldwin effect in the evolution of our binary immune system. We found that the Baldwin effect is not a universal relationship between evolution and learning, and that the strength of the effect is sensitive to the shape of the fitness landscape.

2 The Immune System and Clonal Selection

The human immune system is responsible for recognizing, and defending against, pathogens, toxins and other foreign molecules, collectively called *antigens*. The immune system produces special molecules, *antibodies*, which bind to antigen and thereby lead to their elimination. Antigen recognition is essentially a form of template matching—when the shape and charge

of the antibody and antigen molecules match, in a complementary fashion, the molecules can bind and the antigen is recognized. The closer the match is between antibody and antigen, the stronger the molecular binding and the better the recognition.

Antibodies are produced by the B cells of the immune system. Each B cell produces a specific type of antibody that can recognize only specific types of antigen. The antibodies produced by a B cell are expressed on the surface of the cell as receptors. When the antibody receptors of a B cell recognize an antigen, that B cell is stimulated to reproduce. In the presence of antigen, the daughter cells of a stimulated B cell will also become stimulated and reproduce. Thus, the presence of antigen will cause the proliferation of those B cells best suited to recognizing that antigen. This process is called clonal selection.

The stimulated proliferation of B cells activates a mechanism called *somatic hypermutation* [5]. Somatic hypermutation affects the genes that encode for the antibody molecule (in particular the portion of the antibody that binds to antigen, called the variable region). The daughter cells of a stimulated B cell, therefore, exhibit wide variation in their ability to recognize antigen.

Due to competition for binding antigen, the better B cells with higher affinity antibodies will be stimulated by the antigen and will grow at the expense of B cells expressing "poorer" or lower affinity antibodies. By repeating this process of mutation and selection a number of times, the immune system "learns" to produce higher affinity antibodies for the antigen stimulating the system[1] .

Learned characteristics in the immune system are not passed directly to offspring. The DNA in a B cell that is modified during the learning process cannot be spliced into the DNA of an egg or sperm cell, and so this DNA cannot be passed on genetically to the next generation. Any contribution that clonal selection and somatic hypermutation have on the evolutionary process must occur via the Baldwin effect.

3 The Binary Immune System

```
 Antigen: 1100100101101001101010000100001111010011001010101001010100110101
Antibody: 1011010101001010010111000101011101011011100011001100100100100111
    XOR:   11111    1   111111 1     1 1  1   1   1 1  11  1 111     1  1
```

Figure 1: Computing the match score between binary molecules.

Our model of the immune system uses bitstrings, with a length of 64 bits, to represent the shape of antibody and antigen molecules. In this bitstring universe, molecular binding takes place when an antibody bitstring and an antigen bitstring "match" each other in a complementary fashion. The *match score* between an antibody bitstring and an antigen

[1]It is interesting to note that somatic hypermutation and clonal selection behave like an internal version of evolution that operates within an individual's body.

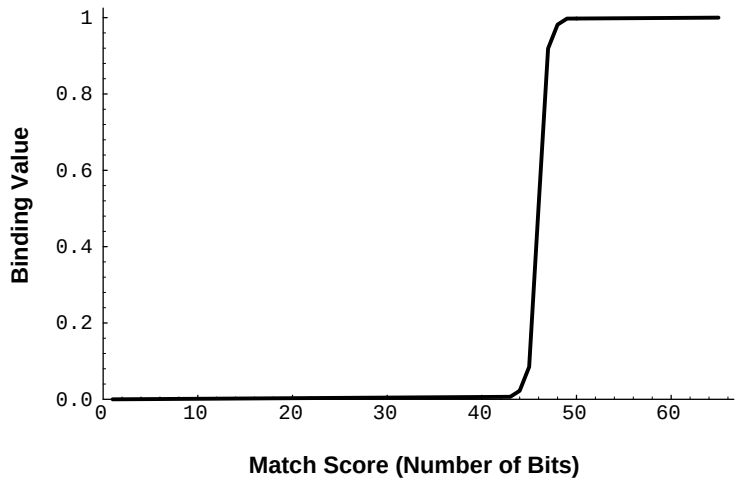

Figure 2: Relation of binding value to match score.

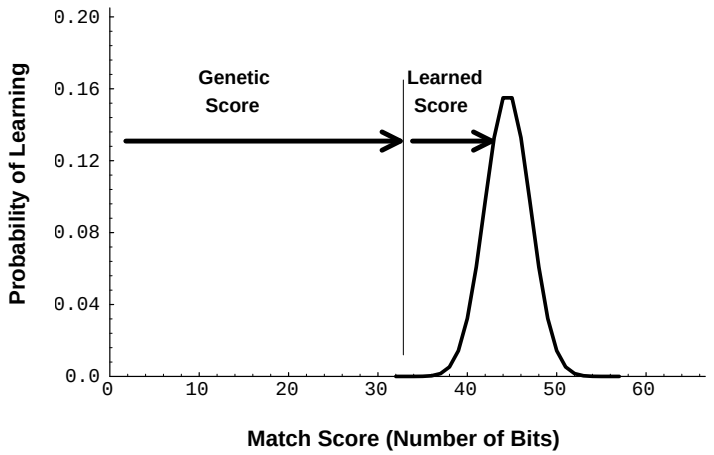

Figure 3: Probability of learning, with G = 25 guesses and genetic component of 32.

bitstring is the number of complementary bits. As shown in Figure 1, the match score can be computed by applying the exclusive-or operator (XOR). The expected match score between two randomly chosen bitstrings is 32 (half the length of the 64-bit strings).

The *binding value*, derived from the match score, represents how well two molecules bind. In reality, two molecules must match each other over a sufficiently large surface area before a stable bond can form. In our bitstring universe this is emulated by requiring the match score to exceed a certain threshold before binding takes place. The curve in Figure 2 shows the relationship of binding value to match score. For match scores below the threshold (in these experiments the threshold was equal to 45 bits) the binding value is essentially zero. For match scores above the threshold, the binding value is one. Note that the binding value function makes a soft transition on either side of the threshold, rather than an abrupt shift from zero to one.

We chose a non-linear binding value function for two reasons. First, it is a more accurate description of the actual molecular binding process. Second, it appears to be a necessary

condition for the Baldwin effect to occur (as will be discussed later).

In our original model of the immune system [4] no learning (somatic hypermutation) took place during fitness evaluation, and the match score was determined entirely by the genes. Here we extend that model and assume that the match score of an antibody has both a genetic component and a learned component, as shown in Figure 2. The genetic component of the match score is computed by comparing an antibody bitstring to the antigen before any learning occurs. The antibody is then allowed to make G guesses in an attempt to improve its score. (This learning-as-guessing process follows that used by Hinton and Nowlan [2].) Each guess has a 50% probability of being correct, so the probability distribution for the learned component is a binomial distribution, with a mean of $Genetic\ Component + (Allowed\ Guesses)/2$. Comparing Figure 2 and Figure 3, we see that if the genetic component is insufficient to produce a strong binding value, a small amount of learning may push the match score above the necessary threshold.

4 The Experiment

The experiment described in this paper tests for a modified statement of the Baldwin effect. Instead of testing whether "learning guides evolution," the experiment tests whether "learning accelerates evolution." The approach is to vary the learning rate—the number, G, of allowed guesses—and to measure the corresponding amount of evolutionary progress after a fixed amount of time has passed. This is more similar to the experiments of Keesing and Stork [3] than the Hinton and Nowlan experiment [2].

Before looking at the results, we detail the genetic algorithm portion of the experiment. Grefenstette's GENESIS [6] was used as the genetic algorithm, with default settings for mutation and crossover probabilities. The population size was set at fifty. The GA experiments all ran for exactly one thousand generations. Individuals are initially set to all zero bits for the first generation[2].

An individual in the population was represented as a genome of 512 bits. Using a compression technique found in the real immune system, these 512 bits encode for a total of 4096 antibodies, each of length 64-bits. The details of this genetic representation can be found in [4].

The fitness of an individual, as used by the GA, is found by testing all 4096 antibodies against a small number, K, of randomly chosen antigens (here $K = 8$). The K antigen bitstrings are randomly selected from a larger set of bitstrings called the *antigen universe*. For the experiments described here the antigen universe contained 32 randomly generated antigen bitstrings.

Individual fitness is calculated as follows. For each of the eight antigens, the individual being evaluated selects the single antibody that best matches the antigen. The match score of this

[2]Initializing the population to all zeros increases the difficult of the learning task. There is also some biological justification: the genes encoding for antibodies probably arose through a process of gene duplication, so the DNA would have had a high degree of self-similarity early in its evolutionary history.

antibody with the antigen is the genetic component of the score. The chosen antibody then undergoes the learning process and makes G guesses to determine the learned component of the match score. The genetic component and the learned component are added (see Figure 3), and used to find the corresponding binding value (using the function shown in Figure 2). Thus a binding value is found for each of the $K = 8$ antigens. The fitness of the individual is the average binding value, averaged over the eight randomly chosen antigen. This fitness is used by the GA to determined which individuals will survive into the next generation.

No Lamarckian inheritance is performed in these experiments. The learning that takes place does not modify the genetic character of the individual being evaluated. Therefore the offspring of an individual inherit no direct information about the results of the learning process.

5 The Results

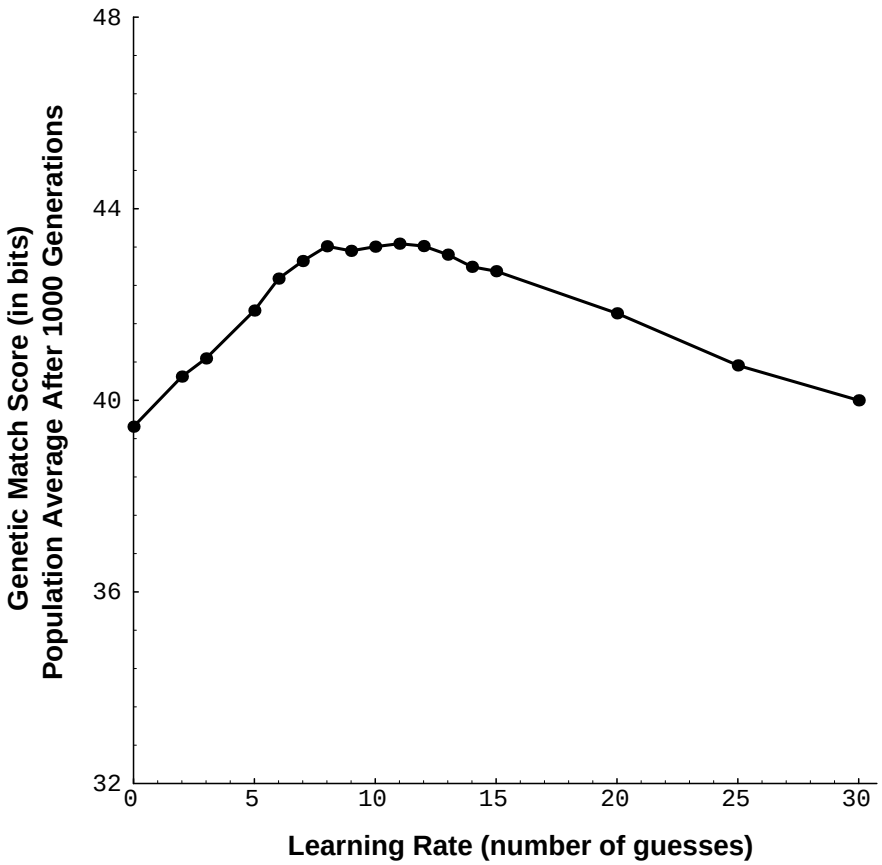

Figure 4: The Baldwin effect and the mastery effect. The population average match score at generation = 1000. 30 experiments per data point. Population size = 50.

Figure 4 shows the results of these experiments for a number of different learning rates, G, between 0 and 30. Each point is an average over thirty GA runs. The vertical axis shows the

genetic component of the match score, averaged for the entire population (and the thirty GA runs). This average genetic match score is essentially the "fitness at birth," before learning takes place, and shows the genetic goodness of the population. The data curve shows that the highest average match scores (around 43 bits) were found by the genetic algorithm when the learning rate was between 8 and 12 guesses per lifetime.

Two things should be pointed out about the scaling of this plot. First, the match score between a randomly chosen antibody and a randomly chosen antigen has an expected value of 32 bits (half the length of the 64-bit strings), and this determined the minimum coordinate for the graph. Second, the binding value for all match scores above 48 is 1.0, so there is no selection pressure for the population to exceed a score of 48. Hence 48 was used as the maximum coordinate of this plot. The plot is scaled between the minimum and maximum obtainable values.

The upward slope on the left side of the curve is due to the Baldwin effect. It shows that increased learning leads to an increase in evolutionary progress during the course of one thousand generations. The explanation for this acceleration of evolution is that learning rewards those individuals that are nearer to the threshold in the binding value function. Without learning, individuals near the threshold have essentially the same poor fitness of those further away. This is similar to the situation described by Hinton and Nowlan [2] in their experiment where learning put "shoulders" around the single solution, providing fitness information to the GA that would not have been available otherwise.

The ramp on the right side of the curve shows, what we call the *mastery effect*. The antigen recognition task has only a finite difficulty, so it is possible for an individual to completely master the task. Given a sufficient amount of learning, even the worst members of the population can learn the correct solution in 1000 generations, which effectively hides their genetic disabilities from the GA. This reduction in information slows the progress of evolution, and the slowing appears to occur linearly with an increase of learning rate. Presumably, an infinite amount of learning would absolutely halt evolution, given a task with finite difficulty. Keesing and Stork [3] demonstrated a similar effect with a neural network model.

A small modification to this experiment removed the Baldwin effect completely. Figure 2 shows a binding value function that is non-linear with respect to the match score. When the binding value function is made to be a linear function of match score, then changes in the learning rate G do not improve evolutionary progress. With a linear binding value function there is no Baldwin effect, and only the mastery effect remains. This makes good sense with respect to Hinton and Nowlan's "shoulders" explanation. Without the non-linearity, there is nothing for learning to put shoulders around, and therefore no evolutionary advantage to learning.

6 Discussion

We have taken an existing model of the immune system and shown the presence of the Baldwin effect. The explanation of the effect is that learning allows the population to perform local search of the fitness landscape during evolution. As noted by Hinton and Nowlan, this

is like putting shoulders around the solution. In these experiments the "shoulders" are put around the non-linearity in the binding value function. Learning allows evolution to discover which individuals are nearest to the threshold of success.

The Baldwin effect is not a universal relation between evolution and learning, and appears to be sensitive to the shape of the fitness landscape. As we discovered, the Baldwin effect disappears when the non-linearity in the binding value function is removed. Without the non-linearity, there is nothing for learning to put shoulders around, so local search (learning) offers no evolutionary advantage. This observation was never made explicit in Hinton and Nowlan's explanation of the Baldwin effect.

Our experiment combines aspects of the work by Hinton and Nowlan and the work by Keesing and Stork. As with the Hinton and Nowlan, our genetic representation is binary and learning is implemented as a process of bit guessing. The experiment itself, however, follows the "learning accelerates evolution" model of Keesing and Stork, and the resulting curve has a shape similar to their experimental results.

The immune system experiment differs, however, in the representational correctness of our model. Learning and evolution operate on the same genetic representation, and this is true for both our model immune system and the real immune system. Contrast this with the nervous system models. Hinton and Nowlan used genes to represent neural connections. Keesing and Stork used genes to represent weights on neural connections. In real nervous systems, there is no direct genetic representation for either weights or connections. Weights and connections derive from a complicated process of growth and development. The point is that evolution and learning operate on different levels of representation for real nervous systems, but in the simulated models, evolution and learning operate on the same thing. This may be an important distinction when we begin looking for the specific conditions that promote or inhibit the Baldwin effect.

Acknowledgments

We thank the Center for Nonlinear Studies, Los Alamos National Laboratory and the Santa Fe Institute for ongoing support of this project. Forrest also acknowledges the support of the National Science Foundation (grant IRI-9157644). Perelson acknowledges the support of the National Institutes of Health (grant AI28433)

References

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

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

[3] R. Keesing and D.G. Stork. Evolution and learning in neural networks: the number and distribution of learning trials affect the rate of evolution. *Proceedings of Neural*

Information Processing Systems 3. Morgan Kaufmann Publishers, Inc., San Mateo, CA, 1991.

[4] R. Hightower, S. Forrest, and A. S. Perelson. Evolution of Secondary Organization in Immune System Gene Libraries. *Proceedings of the Second European Conference on Artificial Life.* 1993.

[5] T.B. Kepler and A.S. Perelson. Somatic hypermutation in B cells: An optimal control treatment. *J. Theoret. Biol.*, **164**:37–64, 1993.

[6] J. J. Grefenstette. *GENESIS: A System for Using Genetic Search Procedures.* In *Proceedings of a Conference on Intelligent Systems and Machines.* 161–165, 1984.