

How Learning Can Guide Evolution

Geoffrey E. Hinton & Steven J. Nowlan

Originally published in 1987 in *Complex Systems*, 1, 495-502.

Reprinted by permission.

Abstract. The assumption that acquired characteristics are not inherited is often taken to imply that the adaptations that an organism learns during its lifetime cannot guide the course of evolution. This inference is incorrect (Baldwin, 1896). Learning alters the shape of the search space in which evolution operates and thereby provides good evolutionary paths towards sets of co-adapted alleles. We demonstrate that this effect allows learning organisms to evolve *much* faster than their nonlearning equivalents, even though the characteristics acquired by the phenotype are not communicated to the genotype.

1. INTRODUCTION

Many organisms learn useful *adaptations* during their lifetime. These adaptations are often the result of an exploratory search which tries out many possibilities in order to discover good solutions. It seems very wasteful not to make use of the exploration performed by the phenotype to facilitate the evolutionary search for good genotypes. The obvious way to achieve this is to transfer information about the acquired characteristics back to the genotype. Most biologists now accept that the Lamarckian hypothesis is not substantiated; some then infer that learning cannot guide the evolutionary search. We use a simple combinatorial argument to show that this inference is incorrect and that learning can be very effective in guiding the search, even when the specific adaptations that are learned are not communicated to the genotype. In difficult evolutionary searches which require many possibilities to be tested in order to discover a complex co-adaptation, we demonstrate that each learning trial can be almost as helpful to the evolutionary search as the production and evaluation of a whole new organism. This greatly increases the efficiency of evolution because a learning trial is much faster and requires much less expenditure of energy than the production of a whole organism.

Learning can provide an easy evolutionary path towards co-adapted alleles in environments that have no good evolutionary path for non-learning organisms. This type of interaction between learning and evolution was first proposed by Baldwin (1896) and Lloyd Morgan (1896) and is sometimes called the Baldwin effect. Waddington (1942) proposed a similar type of interaction between developmental processes and evolution and called it "canalization" or "genetic assimilation." So far as we can tell, there have been no computer simulations or analyses of the combinatorics that demonstrate the magnitude of the effect.

2. AN EXTREME AND SIMPLE EXAMPLE

Baldwinism is best understood by considering an extreme (and unrealistic) case in which the combinatorics are very clear. Imagine an organism that contains a neural net in which there are many potential connections. Suppose that the net only confers added reproductive fitness on the organism if it is connected in exactly the right way. In this worst case, there is no reasonable evolutionary path toward the good net and a pure evolutionary search can only discover which of the potential connections should be present by trying possibilities at random. The good net is like a needle in a haystack.

The evolutionary search space becomes much better if the genotype specifies some of the decisions about where to put connections, but leaves other decisions to learning. This has the effect of constructing

a large zone of increased fitness around the good net. Whenever the genetically specified decisions are correct, the genotype falls within this zone and will have increased fitness because learning will stand a chance of discovering how to make the remaining decisions so as to produce the good net. This makes the evolutionary search much easier. It is like searching for a needle in a haystack when someone tells you when you are getting close. The central point of the argument is that the person who tells you that you are getting close does not need to tell you anything more.

3. A SIMULATION

We have simulated a simple example of this kind of interaction between learning and evolution. The neural net has 20 potential connections, and the genotype has 20 genes[1], each of which has three alternative forms (alleles) called 1, 0, and ?. The 1 allele specifies that a connection should be present, 0 specifies that it should be absent, and ? specifies a connection containing a switch which can be open or closed. It is left to learning to decide how the switches should be set. We assume, for simplicity, a learning mechanism that simply tries a random combination of switch settings on every trial. If the combination of the switch settings and the genetically specified decisions ever produce the one good net we assume that the switch settings are frozen. Otherwise they keep changing.[2]

The evolutionary search is modeled with a version of the genetic algorithm proposed by Holland (1975). Figure 1 shows how learning alters the shape of the search space in which evolution operates. Figure 2 shows what happens to the relative frequencies of the correct, incorrect, and ? alleles during a typical evolutionary search in which each organism runs many learning trials during its lifetime. Notice that the total number of organisms produced is far less than the 2^{20} that would be expected to find the good net by a pure evolutionary search. One interesting feature of Figure 2 is that there is very little selective pressure in favor of genetically specifying the last few potential connections, because a few learning trials is almost always sufficient to learn the correct settings of just a few switches.

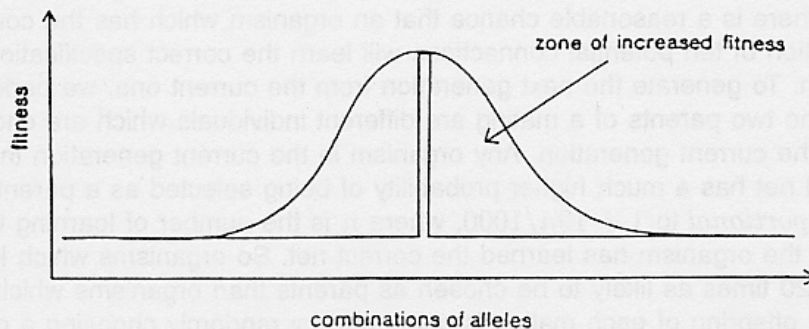


FIGURE 1 The shape of the search space in which evolution operates. The horizontal axis represents combinations of alleles and so it is not really one-dimensional. Without learning, the search space has a single spike of high fitness. One cannot do better than random search in such a space. With learning, there is a zone of increased fitness around the spike. This zone corresponds to genotypes which allow the correct combination of potential connections to be learned.

The same problem was never solved by an evolutionary search without learning. This was not a surprising result; the problem was selected to be extremely difficult for an evolutionary search, which relies on the exploitation of small co-adapted sets of alleles to provide a better than random search of the space. The spike of fitness in our example (Figure 1) means that the only co-adaptation that confers improved fitness requires simultaneous co-adaptation of all 20 genes. Even if this co-adaptation is discovered, it is not easily passed to descendants. If an adapted individual mates with any individual

other than one nearly identical to itself, the co-adaptation will probably be destroyed. The crux of the problem is that only the one good genotype is distinguished, and fitness is the only criterion for mate selection. To preserve the co-adaptation from generation to generation it is necessary for each good genotype, on average, to give rise to at least one good descendant in the next generation. If the dispersal of complex co-adaptations due to mating causes each good genotype to have less than one expected good descendant in the next generation, the co-adaptation will not spread, even if it is discovered many times. In our example, the expected number of good immediate descendants of a good genotype is below 1 without learning and above 1 with learning.

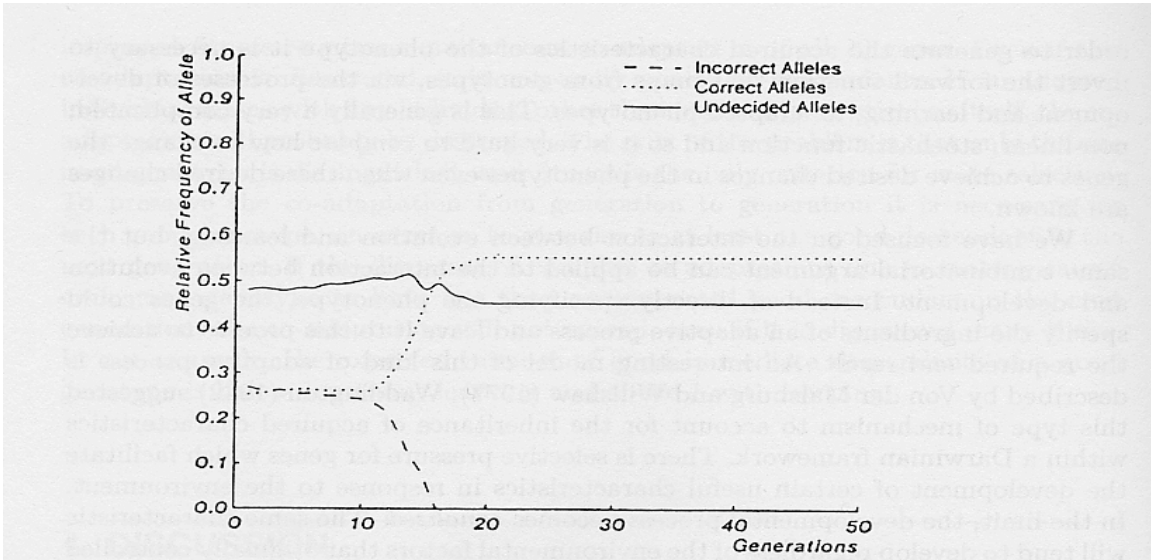


FIGURE 2 The evolution of the relative frequencies of the three possible types of allele. There are 1000 organisms in each generation, and each organism performs 1000 learning trials during its lifetime. The initial 1000 genotypes are generated by selecting each allele at random with a probability of 0.5 for the ? allele and 0.25 for each of the remaining two alleles. A typical genotype, therefore, has about ten decisions genetically specified and about ten left to learning. Since we run about 2^{10} learning trials for each organism, there is a reasonable chance that an organism which has the correct genetic specification of ten potential connections will learn the correct specification of the remaining ten. To generate the next generation from the current one, we perform 1000 matings. The two parents of a mating are different individuals which are chosen at random from the current generation. Any organism in the current generation that learned the good net has a much higher probability of being selected as a parent. The probability is *proportional* to $1 + 19n/1000$, where n is the number of learning trials that remain after the organism has learned the correct net. So organisms which learn immediately are 20 times as likely to be chosen as parents than organisms which never learn. The single offspring of each mating is generated by randomly choosing a cross-over point and copying all alleles from the first parent up to the cross-over point, and from the second parent beyond the cross-over point.

4. DISCUSSION

The most common argument in favor of learning is that some aspects of the environment are unpredictable, so it is positively advantageous to leave some decisions to learning rather than specifying them genetically (e.g. Harley, 1981). This argument is clearly correct and is *one* good reason for having a learning mechanism, but it is different from the Baldwin effect which applies to complex co-adaptations to *predictable* aspects of the environment.

To keep the argument simple, we started by assuming that learning was simply a random search through possible switch settings. When there is a single good combination and all other combinations are equally bad a random search is a reasonable strategy, but for most learning tasks there is more structure than this and the learning process should make use of the structure to home in on good switch configurations. More sophisticated learning procedures could be used in these cases (e.g. Rumelhart, Hinton, and Williams, 1986). Indeed, using a hillclimbing procedure as an inner loop to guide a genetic search can be very effective (Brady, 1985). As Holland (1975) has shown, genetic search is particularly good at obtaining evidence about what confers fitness from widely separated points in the search space. Hillclimbing, on the other hand, is good at local, myopic optimization. When the two techniques are combined, they often perform much better than either technique alone (Ackley, 1987). Thus, using a more sophisticated learning procedure only strengthens the argument for the importance of the Baldwin effect.

For simplicity, we assumed that the learning operates on exactly the same variables as the genetic search. This is not necessary for the argument. Each gene could influence the probabilities of large numbers of potential connections and the learning would still improve the evolutionary path for the genetic search. In this more general case, any Lamarckian attempt to inherit acquired characteristics would run into a severe computational difficulty: To know how to change the genotype in order to generate the acquired characteristics of the phenotype it is necessary to invert the forward function that maps from genotypes, via the processes of development and learning, to adapted phenotypes. This is generally a very complicated, non-linear, stochastic function and so it is very hard to *compute* how to change the genes to achieve desired changes in the phenotypes even when these desired changes are known.

We have focused on the interaction between evolution and learning, but the same combinatorial argument can be applied to the interaction between evolution and development. Instead of directly specifying the phenotype, the genes could specify the ingredients of an adaptive process and leave it to this process to achieve the required end result. An interesting model of this kind of adaptive process is described by Von der Malsburg and Willshaw (1977). Waddington (1942) suggested this type of mechanism to account for the inheritance of acquired characteristics within a Darwinian framework. There is selective pressure for genes which facilitate the development of certain useful characteristics in response to the environment. In the limit, the developmental process becomes *canalized*: The same characteristic will tend to develop regardless of the environmental factors that originally controlled it. Environmental control of the process is supplanted by internal genetic control. Thus, we have a mechanism which as evolution progresses allows some aspects of the phenotype that were initially specified indirectly via an adaptive process to become more directly specified.

Our simulation supports the arguments of Baldwin and Waddington, and demonstrates that adaptive processes within the organism can be very effective in guiding evolution. The main limitation of the Baldwin effect is that it is only effective in spaces that would be hard to search without an adaptive process to restructure the space. The example we used in which there is a single spike of added fitness is clearly an extreme case, and it is difficult to assess the shape that real evolutionary search spaces would have if there were no adaptive processes to restructure them. It may be possible to throw some light on this issue by using computer simulations to explore the shape of the evolutionary search space for simple neural networks that do not learn, but such simulations always contain so many simplifying assumptions that it is hard to assess their biological relevance. We therefore conclude with a disjunction: For biologists who believe that evolutionary search spaces contain nice hills (even without the restructuring caused by adaptive processes) the Baldwin effect is of little interest,^[3] but for biologists who are suspicious of the assertion that the natural search spaces are so nicely structured, the Baldwin effect is an important mechanism that allows adaptive processes within the organism to greatly improve the space in which it evolves.

ACKNOWLEDGMENTS

This research was supported by grant IST-8520359 from the National Science Foundation and by contract N00014-86-K-00167 from the Office of Naval Research. We thank David Ackley, Francis Crick, Graeme Mitchison, John Maynard-Smith, David Willshaw, and Rosalind Zalin for helpful discussions.

REFERENCES

Ackley, D. H. "Stochastic Iterated Genetic Hillclimbing." Doctoral dissertation. Carnegie-Mellon University, Pittsburgh, PA, 1987.

Baldwin, J. M. "A New Factor in Evolution." *American Naturalist* 30 (1896): 441-451.

Brady, R. M. "Optimization Strategies Gleaned From Biological Evolution. *Nature*, **317** (1985): 804-806.

Harley, C. B. "Learning the Evolutionary Stable Strategy." *J Theoretical Biology* **8** (1981): 611-633.

Holland, J. H. *Adaptation in Natural and Artificial Systems*. Ann Arbor, MI: University of Michigan Press, 1975.

Lloyd Morgan, C. "On Modification and Variation." *Science* **4** (1896): 733-740.

Rumelhart, D. E., G. E. Hinton, and R. J. Williams. "Learning Representations by Back-Propagating Errors" *Nature* **323** (1986): 533-536.

Von der Malsburg, C., and D. J. Willshaw. "How to Label Nerve Cells so that They Can Interconnect in an Ordered Fashion. *Proc. Natl. Acad. Sci. U.S.A.* **74** (1977): 5176-5178.

Waddington, C. H. "Canalization of Development and the Inheritance of Acquired Characters. *Nature* **150** (1942): 563-565.

Footnotes

[1] We assume, for simplicity, that each potential connection is controlled by its own gene. Naturally, we do not believe that the relationship between genes and connections is so direct.

[2] This implicitly assumes that the organism can "recognize" when it has achieved the good net. This recognition ability (or an ability to tell when the switch settings have been improved) is required to make learning effective and so it must precede the Baldwin effect. Thus, it is possible that some properties of an organism which are currently genetically specified were once behavioral goals of the organism's ancestors.

[3] One good reason for believing the search space must be nicely structured is that evolution works. But this does not show that the search space would be nicely structured in the absence of adaptive processes.