

On the Baldwin Effect

Larry Bull

Faculty of Computer Studies
and Mathematics
University of the West of
England
Bristol BS16 1QY, U.K.
larry@ics.uwe.ac.uk

Abstract In this article the effects of altering the rate and amount of learning on the Baldwin effect are examined. Using a version of the abstract tunable NK model, it is shown that the adaptation process is sensitive to the rate of learning, particularly as the correlation of the underlying fitness landscape varies. Typically a high learning rate proves most beneficial as landscape correlation decreases. It is also shown that the amount of learning can have a significant effect on the adaptation process, where increased amounts of learning prove beneficial under higher learning rates on uncorrelated landscapes.

Keywords

epistasis, evolution, landscape correlation, learning, NK model

1 Introduction

In 1896 both Baldwin and Lloyd-Morgan defined “a new factor” in evolution which states that an evolutionary process can be guided by a learning process [2], [18]. That is, local search/learning allows individuals to simulate effectively the alteration of some of their genes to fitter alleles than may exist in the gene pool and hence indicate their closeness to fitter regions of the gene space. Over time beneficial alleles tried out by individuals come to exist in the gene pool; learned alleles become assimilated [25] in subsequent generations.

In this article a version of Kauffman’s NK model [17] of fitness landscapes is used to examine the progress of a species given the ability to do local random learning at various rates, and of differing amounts, in comparison to an equivalent species without the ability to do learning. It is shown that the complexity of the underlying fitness landscape, the amount of learning, and the rate of learning are all significant to the Baldwin effect. The ability to learn becomes increasingly beneficial as the landscape complexity increases, where the range of beneficial learning rates is typically small and the amount of learning beneficial at such rates generally increases with increasing complexity.

The paper is arranged as follows: the next section briefly reviews some of the previous work on the use of the Baldwin effect, Section 3 describes the version of the NK model used in this paper and Section 4 presents the results from its use.

2 The Baldwin Effect

Hinton and Nowlan [16] were the first to investigate the Baldwin effect in a complex environment. Using a “needle-in-a-haystack” problem, they showed that giving genetically specified neural networks the ability to alter inter-neuron connections randomly during their lifetime enabled the evolutionary system to find the isolated optimum, something the system without learning was unable to do. That is, the ability to learn “smoothed” the fitness landscape into a unimodal hill/peak. They also found that over time more and more correct connections became genetically specified and hence less

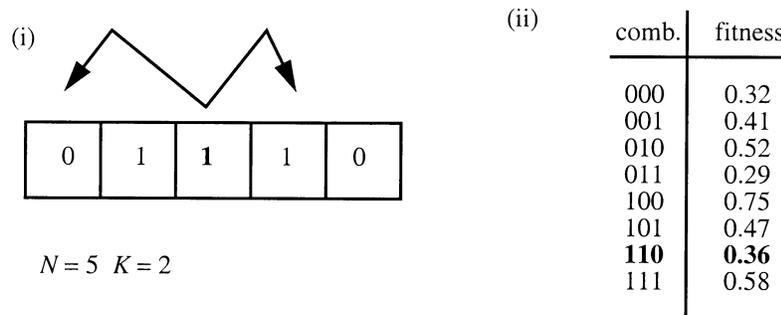


Figure 1. Shows a NK model example. (i) shows how the fitness contribution of each gene depends on K random genes. Therefore there are $2^{(K+1)}$ possible allele combinations, each of which is assigned a random fitness as shown in (ii). Each gene of the genome has such a table created for it. Total fitness of a given genome is the normalized sum of these values.

and less random learning was necessary; the evolutionary process was guided toward the optimum by the learning process. Belew et al. [6] (after [3]) added the Baldwin effect via backpropagation to his work on the evolution of neural networks for various classes of problem, finding that the search process was greatly improved (see also [22]). Since then a body of work has built up on the examination and/or use of learning in conjunction with evolutionary computing techniques, for example [1], [4], [7], [9], [10], [11], [12], [13], [26], and so forth. More recently it has also been suggested that the Baldwin effect can prove beneficial even when there is little correlation between what is learned and the underlying fitness landscape [21] (see [20] for an opposing case). That is, under some circumstances the perturbations due to learning a different task can compensate for unavoidable sub-optimal genetic mutations [14]. In this article direct correlation between the phenotype (learning) space and genotype (evolution) space is assumed [19].

The Baldwin effect has also been connected with the immune system [15], sexual selection [23], and to the emergence of both multicellularity and eusociality [8]. The reader is referred to [5] and [24] for recent overviews of the Baldwin effect.

In this article the performance of the Baldwin effect under varying rates and amounts of learning is examined using a version of the NK model.

3 The NK Model

Kauffman [17] has presented a tunable family of fitness landscapes specified by two parameters: N , the length of the genome; and K , the number of genes that has an effect on the fitness contribution of each (binary) gene. Thus increasing K with respect to N increases the epistatic linkage, increasing the ruggedness/complexity of the fitness landscape. Kauffman shows that the increase in epistasis increases the number of optima, increases the steepness of their sides, and decreases their correlation. The model assumes all intragenome interactions are so complex that it is only appropriate to assign random values to their effects on fitness. Therefore for each of the possible K interactions a table of $2^{(K+1)}$ fitnesses is created for each gene with all entries in the range 0.0 to 1.0, such that there is one fitness for each combination of traits (Figure 1). The fitness contribution of each gene is found from its table. These fitnesses are then summed and normalized by N to give the selective fitness of the total genome. The reader is referred to [17] for full details of the properties of the NK family of fitness landscapes.

4 Learning

4.1 The Model

Kauffman used a mutation-based hill-climbing algorithm, where the single point in the fitness space is said to represent a converged species, to examine the properties and evolutionary dynamics of the NK model. Here a species evolves by making a random change to one gene per generation, and the population is said to move to the new genetic configuration if the fitness is infinitesimally greater than its current fitness.

In this paper the evolutionary progress of two initially identical species is compared on fitness landscapes with varying complexity (K). The learning species is able to do local random learning at a rate $1/L$. If the probability $1/L$ is satisfied on a given generation, after the random alteration of a gene allele to form the new possible genetic configuration for the species, a second random gene is altered. If the fitness of this “learned” configuration is greater than that of the original, the species is said to move to the *first* mutant configuration but assigned the fitness of the *second* configuration. This form of learning, based on Hinton and Nowlan’s model [16], is envisaged as being the most simple, and effectively without restriction (unlike Hinton and Nowlan’s where some genes/network connections were not alterable by the learning process in a given individual). The purely evolutionary species does two rounds of mutation-based search on these generations, that is mutate-mutate-evaluate-select. If the learning probability is not satisfied, both species evolve as in Kauffman’s model, that is mutate-evaluate-select.

Genomes of length 16 bits ($N = 16$) are used here and all results reported are the average of 10 runs (random start points) on each of 10 NK functions, that is 100 runs, after 10,000 generations.

4.2 Results

The results obtained indicate that the Baldwin effect can be beneficial to the search process depending on the rate of learning ($1/L$), for all $K > 0$. In the simplest case of $K = 0$, the underlying fitness landscape contains a large, single peak (unimodal) and hence the ability to learn offers no advantages over purely genetics-based search (not shown—see also [19]). For very low K , the fitness landscape is said to be correlated, that is the highest optima are nearest one another—an effect that becomes insignificant around $K > 3$, regardless of N [17, p. 60]. When $K = 1$, it has been found that learning is neither beneficial nor detrimental to the adaptation process at higher rates (lower L), in comparison to the equivalent purely evolutionary species, but confers a slight benefit at lower rates, around $L > 30$ (Figure 2). When $K = 2$, the Baldwin effect gives a small benefit for high learning rates ($L < 5$), which reduces to that found for $K = 1$ for all $L > 30$. When $K = 3$, the Baldwin effect gives a significant benefit (> 0.01) for high learning rates ($L < 5$), which decreases to that found under $K = 2$ for all other L . This general behavior is seen thereafter, with the size of the relative benefit from the Baldwin effect increasing with increasing K . That is, the Baldwin effect gives most advantage for high rates of learning, reducing to a small advantage for all other rates, for all $K > 2$ (Figure 2); as the ruggedness of the landscape increases and the correlation decreases, a high learning rate proves most advantageous. However, when $K > 13$, learning at the highest rate ($L = 1$) becomes detrimental (not shown) and drops in relative benefit are seen around $K > 11$ (see Figure 2 where $L = 1$ for $K = 13$ is less than for $K = 3$).

As noted above, as K increases the number of optima in the fitness landscape increases, and the correlation of such peaks decreases. The ability to do local search (learning) becomes increasingly useful since it allows low fitness regions between op-

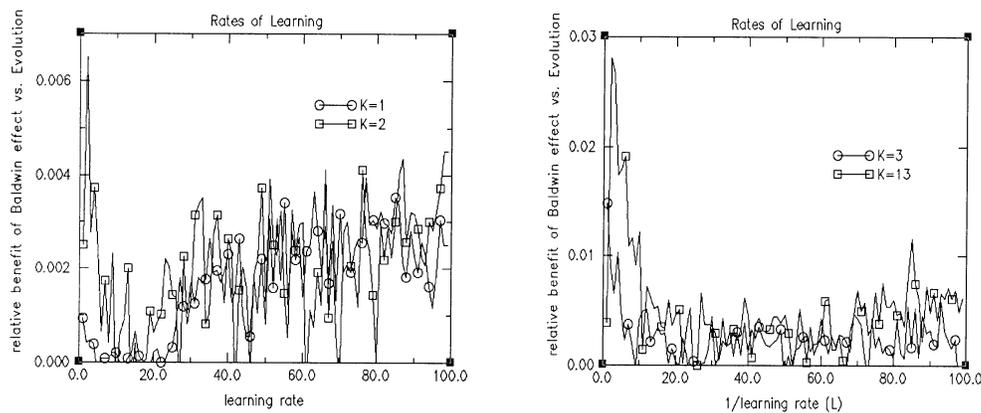


Figure 2. Shows examples of how as the complexity of the function (K) increases, the more beneficial the Baldwin effect becomes in relation to the equivalent purely evolutionary species, particularly at higher rates of learning (low L). N.B. Y-axis shows fitness of the Baldwinian species, at a given learning rate, minus the fitness of the equivalent species adapting via evolution only, after 10,000 generations.

tima to be “smoothed” over and hence crossed by the evolving species. The fact that this is most beneficial at higher learning rates can be attributed to the average time taken to reach an optimum. As the correlation of the fitness landscape decreases, the time taken to reach an optimum decreases (from $N/2$ at $K = 0$ to around $\ln N$ for $K = N - 1$ [17, p. 49]), and so for higher rates the “guiding process” has a greater effect on the evolutionary progress of a species; at lower rates, learning does not have a chance to influence evolution before the species becomes stuck on an isolated (suboptimal) peak. That very frequent learning proves detrimental on the most uncorrelated landscapes implies that there must be a balance between learning and evolution (selection) at the extreme; when all local peaks become evenly spaced and roughly the same height, evolution can be disrupted by almost constant guiding *at the amount of learning used here*.

From this, the effects of altering the amount of learning on a given evaluation have also been examined. For example, if the number of mutations after the first genetic mutation is increased from one (as above) to three, the relative benefits of the Baldwin effect are found to be altered considerably. On the most highly correlated rugged landscapes, that is $0 < K < 4$, the increased amount of learning proves detrimental to the search process, in comparison to the smaller amounts, at all rates apart from $L = 1$ (Figure 3). Thereafter, increasing K increases the number of rates at which the larger amount of learning proves beneficial, such that the larger amount is beneficial at most rates around $K > 11$.

With a larger amount of learning the effective local search range increases. On more correlated fitness landscapes the use of such “long jump” learning is less useful because all peaks are clustered together. That is, the results above suggest that closely-packed high fitness optima can be missed by larger amounts of learning. The larger amount of learning proves useful at a very high rate presumably because only through near constant use can high optima be found. As the correlation decreases, the peaks begin to spread out and hence the ability to search further away becomes increasingly beneficial at all rates. That is, even once a species moves to an isolated (sub-optimal) peak, there is a chance that such learning will find another (higher) peak and the species will move in its general direction—unlike equivalent species with smaller amounts (above) of or no [17, p. 74] learning.

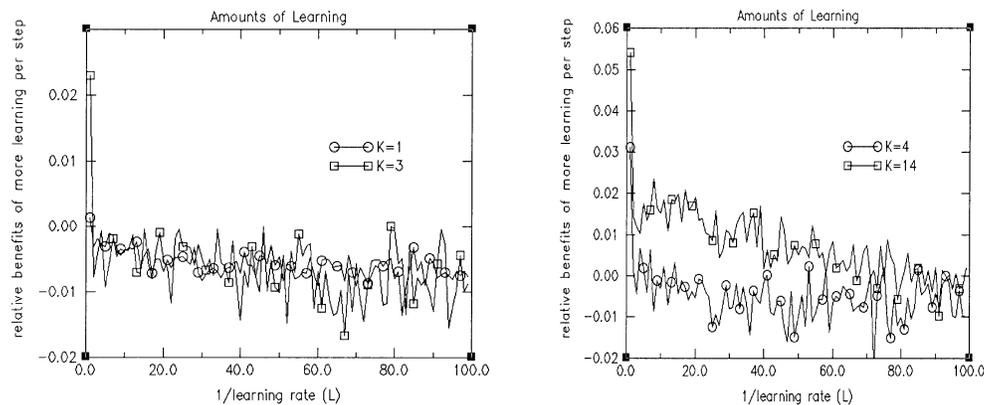


Figure 3. Shows examples of how as the complexity of the function (K) increases, the more beneficial the Baldwin effect becomes with a larger amount of learning used. N.B. Y-axis shows the fitness of the species with the higher amount of learning minus the fitness of the species with the lower amount, after 10,000 generations.

5 Conclusions

In this article the effects of altering the rate and amount of learning on the Baldwin effect have been examined. Using a version of the abstract tunable NK model it has been shown that the adaptation process is sensitive to the rate of learning, particularly as the correlation of the underlying fitness landscape varies, and that typically a high learning rate proves most beneficial. It has also been shown that the amount of learning, at a given rate, can have a significant effect on the adaptation process, with larger amounts proving most beneficial on less correlated landscapes.

Acknowledgments

Thanks to Peter Todd for many useful suggestions on an earlier draft of this manuscript.

References

1. Ackley, D., & Littman, M. (1991). Interaction between learning and evolution. In C. G. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds.), *Artificial Life II* (pp. 487–509). Redwood City, CA: Addison-Wesley.
2. Baldwin, J. M. (1896). A new factor in evolution. *American Naturalist*, *30*, 441–451.
3. Belew, R. K. (1989). When both individuals and populations search. In J. D. Schaffer (Ed.), *Proceedings of the Third International Conference on Genetic Algorithms* (pp. 34–41). San Mateo, CA: Morgan Kaufmann.
4. Belew, R. K. (1990). Evolution, learning and culture: Computational metaphors for adaptive algorithms. *Complex Systems*, *4*, 11–49.
5. Belew, R. K., & Mitchell, M. (1996). *Adaptive Individuals in Evolving Populations*. Redwood City, CA: Addison-Wesley.
6. Belew, R. K., McInerney, I., & Schraudolph, N. (1991). Evolving networks: Using the genetic algorithm with connectionist learning. In C. G. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds.), *Artificial Life II* (pp. 511–548). Redwood City, CA: Addison-Wesley.
7. Boers, E., Borst, M., & Sprinkhuizen-Kuyper, I. (1994). Evolving artificial neural networks using the Baldwin effect. In D. W. Pearson, N. C. Steele, & R. F. Albrecht (Eds.), *Artificial Neural Nets and Genetic Algorithms II* (pp. 333–336). Berlin: Springer-Verlag.

8. Bull, L. (1999). On the evolution of multicellularity and eusociality. *Artificial Life*, 5, 1–15.
9. Bull, L., & Fogarty, T. C. (1994). An evolution strategy and genetic algorithm hybrid: An initial implementation and first results. In T. C. Fogarty (Ed.), *Evolutionary Computing* (pp. 95–102). London: Springer-Verlag.
10. Cecconi, F., Menczer, F., & Belew, R. K. (1996). Maturation and the evolution of imitative learning in artificial organisms. *Adaptive Behaviour*, 4, 29–50.
11. French, R. M., & Messinger, A. (1994). Genes, phenes and the Baldwin effect: Learning and evolution in a simulated population. In R. A. Brooks & P. Maes (Eds.), *Artificial Life IV* (pp. 277–282). Cambridge, MA: MIT Press.
12. Gruau, F., & Whitley, D. (1994). Adding learning to the cellular development of neural networks: Evolution and the Baldwin effect. *Evolutionary Computation*, 1, 213–233.
13. Hart, W. E., Kammeyer, T. E., & Belew, R. K. (1995). The role of development in genetic algorithms. In L. D. Whitley & M. D. Vose (Eds.), *Foundations of Genetic Algorithms 3* (pp. 315–332). San Mateo, CA: Morgan Kaufmann.
14. Harvey, I. (1997). Another new factor in evolution? *Evolutionary Computation*, 4, 313–329.
15. Hightower, R., Forrest, S., & Perelson, A. S. (1996). The Baldwin effect in the immune system: Learning by somatic mutation. In R. K. Belew & M. Mitchell (Eds.), *Adaptive Individuals in Evolving Populations* (pp. 159–167). Redwood City, CA: Addison-Wesley.
16. Hinton, G. E., & Nowlan, S. J. (1987). How learning can guide evolution. *Complex Systems*, 1, 495–502.
17. Kauffman, S. A. (1993). *The Origins of Order: Self-Organisation and Selection in Evolution*. New York, NY: Oxford University Press.
18. Lloyd-Morgan, C. (1896). On modification and variation. *Science*, 4, 733–740.
19. Mayley, G. (1996). Landscapes, learning costs, and genetic assimilation. *Evolutionary Computation*, 4, 213–234.
20. Menczer, F., & Belew, R. (1994). Evolving sensors in environments of controlled complexity. In R. Brooks & P. Maes (Eds.), *Artificial Life IV* (pp. 210–221). Cambridge, MA: MIT Press.
21. Nolfi, S., Elman, J., & Parisi, D. (1994). Learning and evolution in neural networks. *Adaptive Behaviour*, 3, 5–28.
22. Stork, D. G., & Keesing, R. (1991). Evolution and learning in neural networks: The number and distribution of learning trials affect the rate of evolution. In R. Lippmann, J. E. Moody, & D. S. Touretzky (Eds.), *Proceedings of Neural Information Processing Systems - NIPS 3* (pp. 804–810). San Mateo, CA: Morgan Kaufmann.
23. Todd, P. (1996). Sexual selection and the evolution of learning. In R. K. Belew & M. Mitchell (Eds.), *Adaptive Individuals in Evolving Populations* (pp. 159–167). Redwood City, CA: Addison-Wesley.
24. Turney, P., Whitley, D., & Anderson, R. W. (1996). Evolution, learning and instincts: 100 years of the Baldwin effect. Special issue of *Evolutionary Computation*, 4, iv–viii.
25. Waddington, C. (1942). Canalization of development and the inheritance of acquired characters. *Nature*, 150, 563–565.
26. Whitley, L. D., Scott-Gordon, V., & Mathias, K. (1994). Lamarckian evolution, the Baldwin effect and functional optimisation. In Y. Davidor, H-P. Schwefel, & R. Manner (Eds.), *Parallel Problem Solving From Nature - PPSN III* (pp. 6–15). Berlin: Springer-Verlag.