

On Crossing Fitness Valleys with the Baldwin Effect

Rob Mills and Richard A. Watson

School of ECS, University of Southampton, Southampton, SO17 1BJ, UK

rmm05r@ecs.soton.ac.uk

Abstract

Escaping local optima and crossing fitness valleys to reach higher-fitness regions of a fitness landscape is a ubiquitous concept in much writing on evolutionary difficulty. The Baldwin effect, an interaction between non-heritable lifetime plasticity (e.g. learning) and evolution, has been shown to be able to guide evolutionary change and ‘smooth out’ abrupt fitness changes in fitness landscapes – thus enabling genetic evolution that would otherwise not occur. However, prior work has not provided a detailed study or analysis on the saddle-crossing ability of the Baldwin effect in a simple multi-peaked landscape. Here we provide analytic and simulation studies to investigate the effectiveness and limitations of the Baldwin effect in enabling genotypic evolution to cross fitness valleys. We also discuss how canalisation, an aspect of many prior models of the Baldwin effect, is unnecessary for the Baldwin effect and a hindrance to its valley-crossing ability.

Crossing fitness valleys and the Baldwin effect

One of Darwin’s basic tenets for the operation of evolution by natural selection was that it must be possible for adaptations to arise via monotonic improvements provided by successive slight variations (1859). Local fitness peaks, or local optima, genotypes for which no small genetic change affords a fitness improvement, present a problem for incremental improvement, and the greater the width of a fitness saddle (for our purposes, the distance from the local optimum to the nearest genotype of equal or higher fitness), the greater the difficulty of continued adaptation. Wright described the problem of finding “a trial and error mechanism by which the locus of a population may be carried across a saddle from one peak to another and perhaps higher one” as the “central problem of evolution” (1935, p.264). And in general, there is widespread understanding that the presence of local optima and the likelihood of escaping them is fundamental in understanding the difficulty of both biological evolution and applications of evolutionary algorithms (e.g. Kauffman 1993).

Many mechanisms that effect the likelihood of escaping local optima, or avoiding them, have been proposed and investigated including: neutral networks (Huynen et al. 1996) that increase connectivity of genotypes via neutral evolution; extra-dimensional bypass (Conrad 1990) where the number of features an entity exhibits increases over evolutionary time so a population might be able to move around an impasse in the extra degrees of freedom thus

provided; exaptation (Gould & Vrba 1982) where a collection of features adapted for some purpose is co-opted for some other purpose; sexual selection via mate choice (Todd and Miller 1997) which considers the modifications on fitness landscape made by mate choice, allowing populations to move into new adaptive zones; and genetic operators such as sexual recombination that allow large non-random genetic changes by crossing diverse individuals (Jansen and Wegener 1999; Watson 2004, 2006). In this paper we investigate the influence of a different mechanism on the possibility of escaping local optima and crossing fitness saddles – the Baldwin effect.

The Baldwin effect (Baldwin 1896), modelled in detail in the following sections, is an effect resulting from the interaction of learning (or more generally, any non-heritable lifetime plasticity) with evolution (Hinton and Nowlan 1987). Although the Lamarckian inheritance of acquired characteristics is not involved, the Baldwin effect nonetheless describes a mechanism whereby non-heritable characteristics acquired during an organism’s lifetime can influence the selective pressures on an evolving population, and more specifically, over time thereby cause the population to genetically assimilate the previously non-heritable characteristics.

Many examinations and computational simulations of the Baldwin effect have been undertaken (e.g. see Turney et al. 1996, Belew and Mitchell 1996). A particularly celebrated computational model of the Baldwin effect by Hinton and Nowlan (1987) provided a simple model of the effect on a single-peaked landscape. Although many studies of the Baldwin effect have involved more general fitness landscapes, some of which will certainly involve multiple fitness peaks, studies which specifically address this seem lacking. One exception (Wiles et al. 2001) provides a simulation study of the Baldwin effect on an explicitly multi-peaked landscape – however, the findings of this paper are concerned with the interaction between the Baldwin effect and the operation of genetic crossover on the underlying modular structure that produced these peaks. In this paper we address the influence of the Baldwin effect on crossing fitness valleys using a simple two-peaked example (building on Hinton and Nowlan’s approach), and we provide quantitative analysis of its abilities and limitations in this process.

One way to understand the effect of lifetime plasticity on the selective pressures acting on a population is as a ‘smoothing’ of the fitness landscape (Watson et al. 2000). This is demonstrated clearly in Hinton and Nowlan’s model where the smoothing effect produces fitness slopes

around an otherwise abrupt fitness needle standing on a fitness plateau. This model was sufficient for illustrating the Baldwin effect, but although it introduces the idea of smoothing the fitness landscape, it does not (and was not intended to) inform our understanding of the potential for the Baldwin effect to escape local optima and cross fitness valleys.

In this paper we introduce a simple two-peaked fitness landscape, one higher than the other, and examine the likelihood of a population at the low peak escaping and traversing to the high peak, with and without learning. We find that the smoothing concept of the Baldwin effect is useful for understanding the consequences of learning in this scenario, and we find that there are cases where, although reaching the high peak is infeasible for a non-learning population, a learning population can reach the high peak easily. This fits straightforwardly with what we might expect about the operation of the Baldwin effect, but the contributions of this paper exceed this basic result.

Specifically, in addition to showing that the Baldwin effect can enable the crossing of fitness valleys, we provide quantitative analysis of the modified fitness landscape provided by phenotypic plasticity, and examine specific conditions and probabilities describing the strength and limitations of the effect. We show that, whereas in the Hinton and Nowlan model the Baldwin effect merely converts a flat area of the landscape to an inclined area, the Baldwin effect can go further and convert a negative selective gradient into a positive selective gradient for the learning population. This can enable the complete removal of a fitness valley, providing a path of monotonic improvement leading genetic evolution of the learning population across the valley to the higher fitness peak.

We also reiterate the difference between genetic assimilation and canalisation (Mills and Watson 2005), often conflated in prior work, and discuss how, not only is the latter not required for the former, but that canalisation is actually a hindrance to the use of the Baldwin effect for the crossing of fitness valleys. We argue that the incorporation of canalisation in Hinton and Nowlan’s model seems natural only because they use a single-peaked landscape, and in a multi-peaked landscape canalisation would be a hindrance to finding the higher peak. Moreover we see that the particular way in which Hinton and Nowlan represent canalisation in their model is questionable as well as unnecessary.

The next section details the two-peaked landscape and the model of the learning population that we will examine in this paper. This is followed by an analytic study of the effect that lifetime plasticity/learning has in smoothing the fitness landscape and removing fitness valleys in this scenario. Subsequently, simulation studies illustrate the effect that the modified landscape has on an evolving population. Finally we review the conceptual difference between genetic assimilation and canalisation introduced in prior work (Mills and Watson 2005) and discuss how canalisation hinders valley crossing and moreover obfuscates the mechanism of the Baldwin effect in general.

A Model of Learning on a Two-Peaked Landscape

Fitness Landscape

In order to investigate the ability of the Baldwin effect in crossing valleys, we study a simple landscape comprising two peaks on an otherwise flat plateau; that is, all genotypes aside the two peaks have the same fitness, F_0 . We nominate a unique genotype for each peak and, without loss of generality, we choose the two peaks to be from the set of genotypes $G_k=1^k0^{N-k}$ for $0 < k < N$ (where N is the number of bits in a binary genotype), using $k=m$ and $k=n$ for the two selected peaks. The two peaks are straightforwardly represented in a one-dimensional cross-section of the fitness landscape that passes through the two peaks. One should be careful with a one-dimensional representation of an N -dimensional space: this section usefully represents genotypes that are ‘between’ the two peaks (in the sense of being on one of the length $|n-m|$ shortest mutational paths between them), but note that other genotypes exist which are not on this cross-section.

To find the capability of the Baldwin effect we choose the distance between the two peaks to be great enough that if a population were located at the lower peak, random mutation alone would be extremely unlikely (quantified later) to allow the population to cross this valley to the higher peak.




Figure 1: Two peaked fitness landscape

Lifetime Plasticity and Learning Model

For a population to exhibit the Baldwin effect it requires some form of phenotypic plasticity. For a non-plastic population we let the N -bit genotype represent the phenotype of the individual directly. A plastic individual, in contrast, produces a number of different phenotypes (each one represented by an N -bit string) in each of its lifetime time-steps. Each phenotype is produced by applying mutation-like variation to the individual’s genotype. That is, each phenotype is produced by copying the genotype then, with some probability, replacing each bit with a new random bit for each bit independently. Note that, in contrast to Hinton and Nowlan’s model and many other studies, our simpler model of lifetime plasticity does not involve any demarcation of which bits are plastic and

which bits are non-plastic, nor any mechanism that varies the number of bits that are plastic over evolutionary time. Accordingly, we refer to our model as a constant plasticity (CP) model. The important reasons for selecting these properties for the CP model are discussed later.

We investigate two variants of how the set of phenotypes produced in an individual's lifetime are mapped to its fitness: one where individuals have a simple form of phenotypic plasticity (without learning) (CP), and a second where individuals have a learning ability (CP-L). In the CP model, an individual's fitness is simply the mean of the fitnesses of all of its phenotypes. In the CP-L model, the fitness of an individual is based on that used by Hinton and Nowlan but extended to suit a multiple-peaked landscape. At each lifetime time-step, we allow individuals to exhibit the most fit phenotype found in all time-steps thus far, then take an average of the exhibited phenotypes. As in Hinton and Nowlan's model, this model of learning represents the ability of an individual to recognise and exploit successful phenotypes when they are discovered. However, ours does not assume that an individual knows it has found the global optimum; thus exploration continues throughout its lifetime.

Note that whilst neither of the constant plasticity model variants exhibit any canalisation, this does not imply that there are no costs to this plasticity (or learning): if no mechanism exists to facilitate canalisation (i.e. reduce plasticity) then the costs of plasticity are irrelevant.

We use a simple example to illustrate the two fitness models (see figure 2). For an individual with $L=200$ phenotypic trials, and a landscape with plateau fitness (F_0) of 1, and peak fitnesses (F_m, F_n) of 10 and 100, and if the lower peak is first found after 60 trials (T_m), and the higher peak is found after 180 trials (T_n), assuming each peak is found only once, we find the overall fitness as follows: for the plastic individual: $F = (198*1 + 10*1 + 100*1) / 200 = 1.54$; for the learning individual: $F = (60*1 + 10*(180-60) + 100*(200-180)) / 200 = 16.3$.

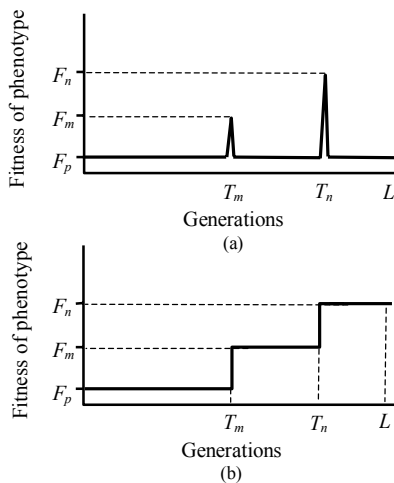


Figure 2: Calculating the fitness for phenotypes of (a) a plastic individual, and (b) a learning individual.

Analytic Results

In this section we provide analytic results of how the plasticity and learning models detailed above 'smooth' the aforementioned two-peak landscape. We find the extent of modifications to this landscape provided by plasticity, by calculating the expected fitness of individuals as a function of their location with respect to the two peaks. This is calculated for any width fitness valley between the two peaks and any heights of those peaks.

CP Expected Fitness Derivation

In this model each of the phenotypic trials is independent and the fitness of an individual is calculated by the mean of the fitness of all of its phenotypes. We calculate the probability of hitting each peak from a genotype, $K = I^k 0^{N-k}$, lying on the section drawn in Figure 1, for a variation rate of μ , position of peaks m and n , fitness of these peaks F_m and F_n , and genotype size N loci.

The probability, p , of a genetic change at a locus is $\mu/2$, (with binary alleles, the probability of a genetic change is half the probability of assigning a new random allele), and the probability of no change, r , is $1-p$. The Hamming distances from peaks m and n to the genotype K are $d_m = |n-k|$ and $d_n = |m-k|$, respectively. The probabilities, Q_m and Q_n , of exactly the correct variation occurring to hit peaks m and n from genotype K are:

$$Q_n = p^{d_n} \cdot r^{(N-d_n)}$$

$$Q_m = p^{d_m} \cdot r^{(N-d_m)}$$

Finally, we calculate the expected fitness, E_k , for a genotype K from the contributions of peaks F_m and F_n and the plateau fitness F_0 :

$$E_k = Q_n F_n + Q_m F_m + (1 - Q_n - Q_m) F_0$$

The number of phenotypic lifetime trials does not affect the mean expected fitness – it only affects the variance in the expected fitness (in practice, the sample provided by lifetime trials will need to be large enough to reflect these probabilities).

This derivation should apply for the expected fitness calculation of any genotype, and for any given peak configuration, not just those described by G_k . The one necessary change is in calculating the Hamming distances d_m and d_n , where a modulus operation would not be sufficient. Instead we can use an xor operation between each peak genotype and the genotype we wish to calculate the expected fitness for. However we choose to consider G_k in order that interpreted of results be straightforward.

CP-L Expected Fitness Derivation

At each lifetime time-step, the individuals in the CP-L model employ the best phenotype they have found thus far, which means that the fitness afforded by each learning

trial is not independent of the trials which have already occurred. This calls for a different approach in calculating the total expected fitness, and for this we extend the method used by Harvey (1993) to analyse Hinton and Nowlan’s model.

Where symbols are reused from the CP analysis they have the same meaning and definition. The probability that peak n has been hit in at least one time-step by time-step t is q_{nt} below, and likewise q_{mt} for peak m .

$$q_{nt} = 1 - (1 - Q_n)^t$$

$$q_{mt} = 1 - (1 - Q_m)^t$$

We use a probability tree to find the expected fitness contribution for each time-step. Since the learning model grants an individual the fitness of the greatest peak found by that time-step, we firstly consider the probability highest peak being hit by this time-step (q_{nt}), in which case the fitness contribution is F_n . If this does not occur ($1 - q_{nt}$) we construct a branch for the alternative events, in fitness order: the probability that the low peak is hit (q_{mt}) by its contribution F_m , and a final branch if the low peak is not found ($1 - q_{mt}$) the fitness contribution is from the plateau (F_0). Finally, the expected fitness is calculated by summing the probability for each peak being hit over all learning trials:

$$E_k = \frac{1}{L} \sum_{t=1}^L (q_{nt} F_n + ((1 - q_{nt})(q_{mt} F_m + (1 - q_{mt}) F_0)))$$

In this section we have extended Hinton and Nowlan’s learning population fitness function, and Harvey’s analysis of it, from a specific case of a one-peaked landscape to a two-peaked landscape. In doing so, this also provides a means to calculate the expected fitness of a learning individual with a given genotype in a landscape with any number of peaks. In principle, this allows the method to apply to any fitness landscape using the limit where a landscape is modelled as a field of neighbouring fitness ‘needles’.

Example study

We now consider the specific landscape introduced above, with $N = 20$, $F_n = 100$, $F_m = 10$, $F_0 = 1$, $\mu = 0.2$, $n = 20$, $m = 15$, $L = 1024$ for each of the models.

Figure 3(a) shows the ‘expected-fitness landscape’ of a non-learning plastic population, which shows a small amount of smoothing around the two peaks; whilst the plasticity affords an advantage in fitness for individuals not on a peak (compared with the plateau fitness), this advantage is neither significant nor far reaching from each peak. Note also that for a plastic individual whose genotype is on a peak, its expected fitness is greatly reduced when compared to a non-plastic individual with the same genotype. Figure 3(b) depicts the expected-fitness landscape of a learning population for a variety of lifetime trial counts. The landscape is smoothed far more significantly even at low trial counts, indicating that learning individuals make much better use of information

found during the lifetime trials than their non-learning counterparts. This confers a learning individual a greatly improved chance of survival a number of bits away from a peak when compared with a non-learning or non-plastic individual with the same genotype.

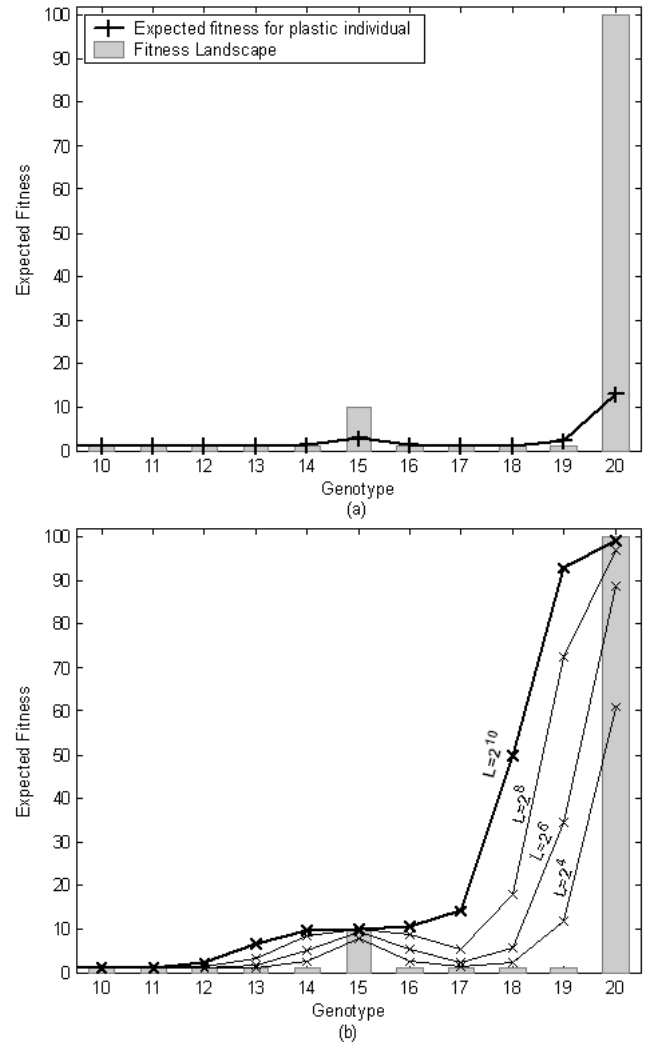


Figure 3: Expected fitness for (a) plastic and (b) learning populations

When the fitness landscape is modified to the extent that a positive gradient toward peak n exists at all genotypes between peak m and peak n , a population will be able to cross the valley in very few generations. In the specific case shown in Figure 3(b); the learning population with 2^{10} learning trials, this condition is satisfied and consequently we expect the valley to be crossed very quickly. In contrast, for a non-plastic population, there is no fitness gradient leading the population to the high peak, of course, and accordingly such a population will not discover that a higher-fitness peak exists unless by chance.

The expected time for an individual situated at the low

peak to hit the high peak by random mutation is $1/(p^{d_n} r^{(N-d_n)})$.

This time increases exponentially with both the width of the valley and, more importantly the size of the search space, N . Thus, a population will have to wait an exponentially large number of trials before its first hit on the high peak, in order to cross the valley. Accordingly, we see that the Baldwin effect has a dramatic influence on the expected time of a population to cross a fitness valley – changing it from infeasible for a non-learning population, to easy for a learning population.

However, it should be noted that the number of lifetime trials required to modify the expected-fitness landscape sufficiently for valley-crossing, is dictated by the same probabilities as those that required for a non-learning population to cross the valley via genetic variation alone. That is, the Baldwin effect provides valley-crossing ability only if we assume a number of lifetime trials that are exponential in N . Thus from an engineering perspective, the Baldwin effect is not an efficient means for crossing valleys in terms of the number of fitness evaluations.

Note that in some cases the fitness of a non-plastic population is greater than for a plastic or learning population when the population exists on the lower peak (see figures 3(a) and 3(b)). However in other cases (depending on height of the high peak, distance from the high peak, number of lifetime trials and other costs of learning) the learning population out-competes a non-plastic population.

Although in prior work we have shown that the CP model is sufficient for exhibiting the Baldwin effect in a single-peaked landscape (Mills and Watson 2005), here we see that the ability of the CP model is very limited compared to that of the CP-L model and not likely to facilitate valley crossing except in very restricted circumstances. For this reason we disregard the CP model for further investigation in this study.

Simulated Experiments

In the previous section we have shown the effective smoothing on the fitness landscapes provided by plasticity and learning, given by an analytic result for the expected fitness of a genotype under these models. We have argued that the crossing times for the learning population will be very fast when the selective gradient to the high peak is monotonically improving, as it is with 2^{10} trials for example, and that the crossing time without the Baldwin effect will be very long (exponential in N). In this section we illustrate these crossing-times using a basic form of genetic algorithm (as is common in other work on the Baldwin effect, e.g. Hinton and Nowlan 1987).

Our evolving population is modelled using a constant population size of 200 N -bit genotypes, initialised on the

low-fitness peak. Each new generation is formed by fitness-proportional reproduction. Mutation is the only genetic variation operator used (with a 0.1 probability of assigning a new random allele to each bit independently). Note that this genetic variation therefore has the same variation neighbourhood (Mayley 2000) as the mutation-like phenotypic variation we have been studying.

We compare the performance of a population with learning phenotypes against a population without any form of phenotypic plasticity (as previously shown a plastic population without learning is expected to have a similar performance to the non-plastic case so is not investigated further).

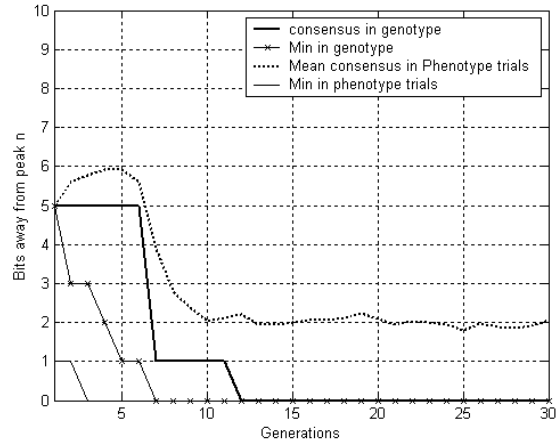


Figure 4: Simulation results of learning populations

We can see in Figure 4 that the phenotypes of the learning population found the higher peak within 3 generations, which is followed by the genotypes of some individuals (see the ‘min in genotype’ line). Within 12 generations the consensus genotype has completely moved to the high peak (the consensus genotype takes the modal allele across the population for each loci). The consensus in phenotypic trials remains approximately 2 bits away from the higher peak even after the consensus genotype is on the peak, indicating lifetime search continues. In 100 additional runs the population moves over to the new peak in a mean of 13.3 generations.

Simulations of the non-plastic population over 3000 generations do not find the high peak until a mean time of at least 2382 generations. We can see the difference in performance is in the order of 100 times in favour of the learning population. To make the most demanding comparison of the learning population with the non-plastic population we might ask which genetic mutation rate would give the non-plastic population the best chance of mutating a genotype at the low peak to hit the genotype at the high peak. The optimal new-random-allele mutation rate to get from the low peak to the high peak that is 5 bits away in our example landscape is 0.5; this rate gives a mean of 5 genetic changes (since new random alleles have a 50% chance of changing the value of their locus) in 20 bits. With this mutation rate the probability of hitting the

high peak from the low peak is 0.0000131 and the expected number of trials before this occurs is therefore 76,000 or 380 generations with a population size of 200 – clearly still far greater than the 13.3 generations that the learning population requires to reach the high peak¹ (also note that we have not optimised the learning model performance carefully).

Moreover, it should be noted that a mutation rate of 0.5 is too high for a population of 200 to maintain its position on the high peak if it were found (or maintain its position on the low peak it starts on for that matter) since the genetic drift is too strong. In principle, a small amount of elitism, retaining the fittest individuals in the population without variation, would alleviate this problem – but this is obviously not an option for a natural population. In contrast, in the learning population the genetic mutation rate may be low even though the phenotypic variation rate is high, so this problem does not arise.

Discussion

Two mechanisms have often been presented as required components to demonstrate the Baldwin effect, *genetic assimilation* and *canalization*. However, as highlighted in previous work (Mills and Watson 2005), only genetic assimilation is necessary. By canalisation we mean a reduction in phenotypic plasticity, whereas genetic assimilation occurs when a behaviour that was once acquired in the phenotype becomes specified in the genotype. The conceptual distinction is easily recognised by considering how the mean and variance of the distribution of phenotypes of an individual changes over evolutionary time: canalisation means that the variance in phenotypes reduces, genetic assimilation means that the mean phenotype is moved (but does not necessarily suggest that the width of that distribution might reduce). In (Mills and Watson 2005) we stress that many works have conflated these two concepts and that this confusion is in large part because they are difficult to disentangle in the particular model that Hinton and Nowlan provided.

The issue of canalisation impacts the saddle-crossing ability of the Baldwin effect considerably. Specifically, to the extent that a population canalises to one peak it will be unable to explore varied phenotypes that may find another, perhaps higher, peak. In the more general case, if there were a number of valleys which were individually crossable, if moving to the first peak requires canalising to that peak, then the first jump would prohibit the crossing of any further valleys. In the single-peaked landscape that

Hinton and Nowlan use, this issue cannot arise and canalising on a peak, even in the limit of removing all phenotypic variation, seems unproblematic. But in a multi-peaked landscape it is not at all clear how a population might avoid what might be termed ‘premature canalisation’. In contrast, in a ‘constant plasticity’ model such as we have used in our study, this problem is moot. If some mechanism for canalisation were to operate after one peak had been found, then to the degree that no plasticity remained, the Baldwin effect could provide no further valley crossing capability, since its fundamental requirement is for individuals to exhibit some phenotypic plasticity. But there is a further issue about the particular way that Hinton and Nowlan model canalisation that we would like to discuss.

A large factor contributing to the difficulty of finding a phenotype at the high-peak is the low probability of variation changing not only the correct number of loci, but ensuring those changes occur at the appropriate loci – changing those that need to be changed, leaving the other loci unchanged. The number of ways of choosing the correct k -loci in an n -locus problem increases with n factorial. In other words, for non-trivial sized problems, the probability of the arrangement of mutations occurring in exactly the right combination to jump a long distance is very small. Naturally, the probability of making a useful jump could be increased if somehow the variation mechanism knew which bits required modification. Hinton and Nowlan allow their individuals to adapt which alleles are variable and which are not (using the special ‘question mark’ alleles that indicate a locus of phenotypic plasticity, and a model that does not allow non-question-mark alleles to vary at all in an organism’s lifetime). It essentially guards the bits which are already correct from further modification, and in each learning trial, incurs a new random allele variation rate of exactly 1 in the loci which are not yet correct. Again, for the single-peaked landscape this seems unproblematic. The fitness rewards that are enjoyed by an individual that has canalised a locus with a correct allele are unambiguously valuable to that individual. However, in a multi-peaked landscape there are many canalisations that would increase the average fitness of an offspring at the low peak but only some of these are “correct” for the high peak. There is no means for the selective pressures that promotes canalisation of loci to distinguish between the 15 loci (to use our example) where the alleles of the two peaks are the same, and the other 5 loci that should remain plastic because the alleles of the two peaks disagree. Thus, we see no way to utilise Hinton and Nowlan’s particular model of canalisation appropriately in a multi-peaked scenario. Hence, although their mechanism of canalisation works well on a single-peaked landscape, where fitter alleles are unambiguously “correct”, and it would clearly be an advantage to saddle-crossing if a population could somehow identify which loci should remain plastic, it is not clear to us that this mechanism of canalisation makes sense in general.

¹ To consider performance from an engineering perspective, we find the learning population takes approx. $13.3 \times 200 \times 1024 = 2,723,840$ evaluations, which is significantly more than for the non-plastic case (approx. $2382 \times 200 = 476,400$ evaluations). However, this model is not intended to show any engineering advantage; from a biological viewpoint, generation time is approximately fixed. Thus the important comparison to make is upon the number of generations required.

The interaction of canalisation with genetic assimilation, and in particular the model of canalisation that Hinton and Nowlan employ, has become quite embedded in how researchers view the Baldwin effect. In particular, the ‘two phase’ aspect of the Baldwin effect in their model – first the purging of incorrect alleles, then the continued replacement of plastic alleles with correct alleles – seems unnecessarily complicated to us. In the constant plasticity model, genetic assimilation occurs by the continued genetic change of genotypes toward phenotypes that are fit. This simplification allows us to exploit the intuitive notion of smoothing the fitness landscape as the only concept that needs to be understood in order to understand how lifetime learning can guide evolution, as we have shown in the above example.

In summary, the problem of how to escape local optima is fundamental for incremental improvement processes. In this paper we have verified that a population with learning can alleviate this problem, by means of the Baldwin effect. We have provided analysis of the modifications to the fitness landscape that learning grants, and this helps us to understand the capabilities and limitations of the effect, specifically when crossing fitness valleys. The analysis is, in principle, applicable to arbitrary fitness landscapes. Our study shows how the Baldwin effect can operate without canalisation and this aids significantly in simplifying understanding of how the Baldwin effect works by smoothing out of the fitness landscape.

References

- Baldwin, JM, 1896, “A New Factor in Evolution”, *American Naturalist*, 30, 441-451.
- Belew, R.K. and Mitchell, M., (1996). eds, *Adaptive Individuals in Evolving Populations: Models and Algorithms*, Massachusetts: Addison-Wesley.
- Conrad, M, 1990, “The Geometry of Evolution.” *BioSystems* 24: 61-81.
- Darwin, C, 1859, *On the Origin of Species by means of natural selection*, John Murray, London.
- Gould, SJ & Vrba, E, 1982, “Exaptation—a missing term in the science of form”, *Paleobiology*, 8: 4-15. Ithaca.
- Harvey I, 1993, “The puzzle of the persistent question marks: A case study of genetic drift”, Forrest S (Ed) *Proceedings of the Fifth International Conference on Genetic Algorithms*, California: Morgan Kaufmann
- Hinton, GE, & Nowlan, SJ, 1987, “How Learning Can Guide Evolution,” *Complex Systems*, 1, 495-502.
- Huynen, MA, Stadler, PF & Fontana, W, 1996, “Smoothness Within Ruggedness: The Role of Neutrality in Adaptation.” *Proc. Natl. Acad. Sci. (USA)*, 93: pp.397-401.
- Jansen, T, and Wegener, I, 1999, “On the Analysis of Evolutionary Algorithms - A Proof that Crossover Really Can Help”, Nesetrl J (Ed.), *Procs. 7th Annual European Symposium on Algorithms*, Springer, Berlin, Germany, pp 184-193.
- Kauffman, SA, 1993, *The Origins of Order*, Oxford University Press.
- Mayley, G, 2000. *Explorations into the interactions between learning and evolution using algorithms*”, DPhil Thesis, School of COGS, University of Sussex
- Mills, R. and Watson, R.A., 2005. Genetic assimilation and canalisation in the Baldwin effect. In *Proceedings of the European Conference on Artificial Intelligence*, 353-362, Springer-Verlag, New York
- Todd, P.M. and Miller, G.F, 1997. Biodiversity through Sexual Selection. In C.G. Langton and K. Shimohara (eds.), *Proceedings of the fifth International Workshop on the Synthesis and Simulation of Living Systems* (pp 289-299), Cambridge, MA
- Turney, P., Whitley, D., & Anderson, R. 1996. Evolution, learning, and instinct: 100 years of the Baldwin effect. *Evolutionary Computation*, 4(3), iv-viii.
- Watson, R.A., 2004, “A Simple Two-Module Problem to Exemplify Building-Block Assembly Under Crossover” in *Proc. Parallel Problem Solving from Nature*, Xin Yao et al. (eds), Springer-Verlag, pp160-169.
- Watson R.A., 2006, *Compositional Evolution: The impact of Sex, Symbiosis and Modularity on the Gradualist Framework of Evolution*, MIT Press, forthcoming.
- Watson, R.A., Reil, T, & Pollack JB 2000, “Mutualism, Parasitism, and Evolutionary Adaptation”, *Procs. of Artificial Life VII*, Bedau, M, McCaskill, J, Packard, N, Rasmussen, S (eds.). MIT Press, pp. 170-178.
- Wiles, J., Tonkes, B., and Watson, J. R. (2001). How learning can guide evolution in hierarchical modular tasks. In Moore, JD and Stenning, K, (eds.), *Proceedings of the 23rd Annual Conference of the Cognitive Science Society*, pp 1130-1135.
- Wright, S, 1935, “Evolution in populations in approximate equilibrium”, *Journal of Genetics* 30, 257-66.