

Genetic Assimilation and Canalisation in The Baldwin Effect

Rob Mills and Richard A. Watson

School of Electronics and Computer Science, University of Southampton, Southampton,
UK, SO17 1BJ

Email: rmm101@zepler.net, raw@ecs.soton.ac.uk

Abstract. The Baldwin Effect indicates that individually learned behaviours acquired during an organism's lifetime can influence the evolutionary path taken by a population, without any direct Lamarckian transfer of traits from phenotype to genotype. Several computational studies modelling this effect have included complications that restrict its applicability. Here we present a simplified model that is used to reveal the essential mechanisms and highlight several conceptual issues that have not been clearly defined in prior literature. In particular, we suggest that canalisation and genetic assimilation, often conflated in previous studies, are separate concepts and the former is actually not required for non-heritable phenotypic variation to guide genetic variation. Additionally, learning, often considered to be essential for the Baldwin Effect, can be replaced with a more general phenotypic plasticity model. These simplifications potentially permit the Baldwin Effect to operate in much more general circumstances.

1 Introduction

Our knowledge of modern genetics suggests that an organism's lifetime adaptations cannot influence the course of evolution because learned characteristics do not change one's own genes. In the late 19th century, Baldwin argued that although a direct effect of lifetime adaptation on genes is not possible, an indirect influence on the course of evolution is [1]. Subsequently his name has been associated with the impact that learning can have upon evolution. The 'Baldwin Effect' is based on two levels of search occurring: from generation to generation we have a slow genetic variation; and within each generation the variation due to lifetime learning is faster. The combination of the two search mechanisms allows the space to be explored more completely than it would be by genetic search alone; an in-depth search around the genetically specified position is performed by the lifetime plasticity, and the genetic starting points are selected for the lifetime phenotypes they enable. This can change the selection of genotypes, providing selective gradients where none was previously available, and in particular, if the discovery of fit phenotypes during lifetime plasticity is correlated well with the genetic closeness of those genotypes to fit configura-

tions then selection will guide evolution toward fit genotypes that may not have been discovered otherwise [2].

1.1 The Baldwin Effect

Controversy has surrounded the Baldwin Effect since Baldwin first proposed it in 1896. The hypothesis appears very similar to Lamarck's disproved beliefs that an acquired trait is directly inherited by offspring; as Turney put it, "*Baldwinian and Lamarckian evolution are virtually indistinguishable in their effect*" [3]. However, unlike Lamarckian evolution, the Baldwin Effect is compatible with genetics since it does not require the direct inheritance of acquired characteristics. There is perhaps a little irony here in the debate over Baldwin's 'organic selection' hypothesis since it was meant to replace Lamarck's failed theory [4-5].

In 1987 Hinton and Nowlan published the first computational model [6] to demonstrate the Baldwin Effect which provides excellent insight into how the effect works. They use a simple model where a population is given a 'needle on a plateau' problem with a single phenotype having increased fitness from an otherwise equally fit plateau. A string with 20 'genes' of 0's, 1's, and ?'s (in initial frequencies 0.25, 0.25, 0.5, respectively) is used to represent each individual's genotype; 0's and 1's represent genetically specified traits, and ?'s represent phenotypically plastic traits. A population of 1000 individuals are randomly initialised. The population is evolved using one-point crossover between two parents selected proportional to their fitness. Within each lifetime, each organism completes 1000 lifetime learning trials, where loci with ?'s in the genotype are randomly assigned a new allele of 0 or 1. The all 1's phenotype is the fitter combination; the fitness of an organism is proportional to the number of lifetime trials left after the all 1's phenotype is found. Thus, a genotype with no 0's may reach the peak in the fitness landscape, and the more 1's it has the more likely its phenotypes are to hit that peak. So when an individual finds this peak in a phenotypic trial, it obtains a greater fitness and begins to dominate the population quickly removing individuals containing 0's. Once there exists a number of individuals who can all reach the peak in their lifetime, the selection pressure shifts to differentiate between these individuals. An individual who is genetically closer to the peak will more reliably hit the peak during its lifetime. In this way, exhibition of the good trait in the genotype is selected for. In typical runs of the simulation after very few generations, the all 1's phenotype is found by lifetime learning, and the average number of 1's in the genotype increases rapidly. In subsequent generations the average number of 1's in the genotype increases slowly towards the fittest genotype.

When a comparable population without lifetime learning is simulated, the all 1's phenotype takes an unreasonable time to be found and there is thus no pressure to increase the number of 1's in the genotype. From this result we can see that the presence of lifetime learning can influence the selective pressure for genetic traits; a learned trait can be genetically assimilated without any direct genetic transfer from phenotype to genotype. A second effect, canalisation, is also exhibited by such mod-

els of the Baldwin Effect. Canalisation, or reduction in lifetime plasticity, is facilitated by means of reduction in numbers of ?'s – the allele representing that plasticity. The reduction of ?'s only begins to occur after all-1's phenotypes have been discovered and the 0's have been removed from the population. Selection favours those who find the all-1's phenotype more quickly over those that find it more slowly, and in the Hinton and Nowlan model, the only way to achieve this is a reduction in ?'s and hence canalisation. Indeed, it may be viewed that an individual's genetically specified traits are preferred in this model because it requires less lifetime learning – implying that canalisation is a necessary part of the Baldwin Effect. This is not correct. In general, a genotype may be selected for in the Baldwin Effect because it produces better phenotypes and, as we will show, this need not necessarily imply that it has less variation in phenotypes as it does in the Hinton and Nowlan model.

Several research papers have been inspired by the work of Hinton and Nowlan (H&N), on a variety of topics. H&N's model is analysed in [7-9], which all focus on the canalisation effect in the experiment; that is pressures for the reduction in plasticity following a learned behaviour becoming genetic. A different kind of evolutionary simulation is used in [10] to demonstrate the Baldwin Effect; they use a simulated physical world in which the population size is variable, and has food costs associated with reproduction, movement and metabolism. Fitter phenotypes benefit an organism in reduced costs for one of the actions. Watson and Pollack adapt H&N's model to demonstrate how symbiosis can produce organisms which would not have evolved without the support of a symbiont [11]. This work is extended to show that in a sparse ecosystem, when one of the symbionts can perform the task independently, the symbiont can offer no advantages and thus becomes a parasite [12]. Another study presents a cultural model with learned and inherited behaviours, using a physical world similar to [10], but with shared behaviours between phenotypes [13]. Results are similar to [11-12]; when behaviours are shared in abundance, assimilation advantage (i.e. selection for independence) reduces. Turney identifies confusion surrounding Baldwinian and Lamarckian evolution, and highlights that although the benefits of learning are demonstrated, often the costs of learning pass without acknowledgement. A series of experiments on drosophila demonstrate 'genetic assimilation' towards wings without cross-veins following a temperature shock [14].

However, the Baldwin Effect does not appeal to all researchers: [15] and [16] both write that the effect is not of much interest or importance in evolution and that Baldwin himself was not primarily interested in 'organic selection'; instead, social heredity or niche construction should be the subject of further study.

Some relevant summary points are made by Turney:

1. Lifetime learning smoothes the fitness landscape since the phenotypic exploration is local to its inherited genotype.
2. There are benefits to phenotypic rigidity: it is advantageous to evolve some rigid mechanism to replace learned mechanisms over time, since learning requires experimentation (for example, learning how to hunt could be dangerous).

Though both these points may often be true in natural populations, and Hinton and Nowlan's model includes both, they are quite separate issues. Reviewing the literature on the Baldwin Effect and phenotypic plasticity in general, a question is raised regarding the difference between genetic assimilation [14] and canalisation [17]. Only a subtle distinction exists and we find no previous model that attempts to show the Baldwin Effect without canalisation, i.e. reduction in plasticity, or discussion that identifies this distinction. Thus, we propose to use the words as follows: canalisation is a reduction in phenotypic plasticity; 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 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. Thus we suggest that genetic assimilation, the key aspect of the Baldwin Effect, does not logically require canalisation, i.e. the phenotype to reduce its level of plasticity.

Note that H&N's model shows both a change in the mean and the variance of phenotypes, i.e. genetic assimilation and canalisation. In the following section we provide a simplified adaptation of H&N's model to show just genetic assimilation without canalisation in order to separate these concepts and illustrate that the Baldwin Effect (being essentially genetic assimilation) does not require canalisation. Since an organism with plasticity (in a particular activity) will still benefit from having that trait genetically well-adapted; it has more chance of an appropriate phenotype occurring than if the genotype is poorly adapted. We will see this demonstrated in the first experiment in section 3.

2 Constant Plasticity Model

The work of Hinton and Nowlan as described above provides a simple model which demonstrates the Baldwin Effect. However, some features of the model are not required to show the Baldwin Effect, and here we present a simpler model. This simplification aims to reduce the assumptions required and realise a model which could be applied to a more general set of cases, and also to assist understanding, specifically to separate the concepts of genetic assimilation and canalisation. The main issues addressed by these simplifications concern a learning model that recognises successful phenotypes, the conflation of canalisation and genetic assimilation, and the mismatch of genetic and phenotypic variation spaces.

H&N use a lifetime plasticity model that involves recognising when a good phenotype is discovered (which may be called learning). If the mechanism of the Baldwin Effect derives simply from smoothing the fitness landscape, as Turney suggests, then a simpler more direct model of lifetime plasticity should suffice such as the average fitness of random phenotypic variants. To remove the assumption of learning phenotypes, we evaluate fitness as the mean fitness of the lifetime phenotypes, rather than

the number of trials remaining after the phenotypic solution is first found. This means that the organisms do not have to recognise their own success (as is implicit in H&N's model).

As established in section 1, H&N's model contains both genetic assimilation and canalisation. We propose that this canalisation is an unnecessary element to display the Baldwin Effect, and as such a significant change to the model is required. In our constant plasticity model there is no designation of particular traits that are phenotypically plastic, i.e. we do not use '?' alleles. Instead lifetime plasticity varies any trait with equal probability, using non-heritable mutation-like variations applied to the genotypic trait specifications in each lifetime trial. This models a constant plasticity level which separates the effects of lifetime plasticity from canalisation.

It has been suggested that in order to enable genetic variation to follow lifetime variation it is desirable to have genetic variations and lifetime plasticity using the same or correlated variation operators [2]. Thus, we can simplify the model further by facilitating genotypic variation using the same method as the phenotypic variation in lifetime trials, i.e. by spontaneous point mutation, instead of using sexual recombination.

These three key changes to H&N's model produce one which is considerably simplified: a population of binary strings is reproduced with fitness proportionate selection and mutational variation. The fitness of each individual is the average fitness of the phenotypes it produces during its lifetime and the fitness of each phenotype is F_{\max} if it is all 1's and 1 otherwise. Each phenotype is a random mutation-like variation of the genetically specified traits (rather than a random filling-in of a variable number of '?') This model is described in more detail below:

- 1) Initialise population of P individuals,
- 2) for each generation
 - a) for all individuals, i:
 - i) for each lifetime trial, 1 to L:
 - (1) generate phenotype by adding random variation
 - (2) evaluate fitness
 - ii) calculate mean fitness across all phenotypes for individual i
 - b) select P parents with probability proportional to fitness
 - c) generate offspring by mutating the parents genotype

The number of mutations in each lifetime trial (and each genetic reproduction) is governed by an exponentially decaying distribution in which it is most likely for no mutations to occur, but some small probability of a large mutation count exists. The probability of a given phenotypic trial (or evolutionary step) having k mutations is given by $P(m=k) = \exp[-ak/N] \cdot (1-\exp[-a/N])$, where N is the number of traits, and a is a rate parameter. At each loci marked for mutation, a new random allele is produced with equal probability of 0 and 1. This is used in order to allow the possibility

of a large number of mutations, whilst maintaining the condition that the most likely phenotype will be identical to its genotype (this is not true of the usual Poisson distribution of mutation counts which results from common mutational models).

3 Simulation Experiments

We simulate populations in a variety of configurations firstly to demonstrate that a population with a constant lifetime plasticity level can manifest the Baldwin Effect (and that populations without lifetime plasticity do not show the effect), and secondly to aid explanations which distinguish genetic assimilation from canalisation.

The evolutionary algorithm used follows Hinton and Nowlan's model, with the modifications as detailed in section 2. The parameter values can be summarised as follows:

Population size, P :	2^{10}
Lifetime trials, L :	2^{10}
Number of traits, N :	20
Fitness of best phenotype, F_{\max} :	1000
Phenotypic mutation parameter, α_p :	5
Genotypic mutation parameter, α_g :	9.1

For the non-plastic experiment, all parameters are as above except that there are no phenotypic trials, i.e. the evaluated phenotype is identical to the genotype. A parallel pair of experiments using Hinton and Nowlan's model are also run; in the non-learning H&N population the '?'s are not used and all variation is performed by sexual crossover in the genotype. Figure 1 shows one typical run of the constant plasticity model, and one run of H&N's model is depicted in figure 2. Due to the stochastic nature of the effect, a single typical run is more informative than mean values over many runs.

Figures 1(a) and 2(a) show data which reveal the mean composition of the population. By calculating the mean distance from the consensus phenotype across all phenotypic trials in each individual and taking the mean of these, we provide a measure of the mean number of phenotypic variations per individual, which is shown by the dotted line. This is a more suitable measure of phenotypic variations than variance of the number of 1's in the phenotypes of an individual, since the value is not dependent on the position in the phenotype space, whereas the variance of the number of 1's is dependent on the mean number of 1's of that particular phenotypic distribution. The dot-dashed line shows the mean number of 1's in the non-plastic population; 100 generations are shown but the behaviour continues similarly for several hundred generations; all non-plastic runs take longer than all plastic runs, with a mean of greater than 1600 generations to find the fitter solution (although some fortunate runs do succeed without lifetime plasticity in much less than this). Figures 1(b) and 2(b)

show the progression of distributions of phenotypes across the population through the experiments with plasticity or learning.




Fig. 1. Typical behaviour of the constant plasticity model

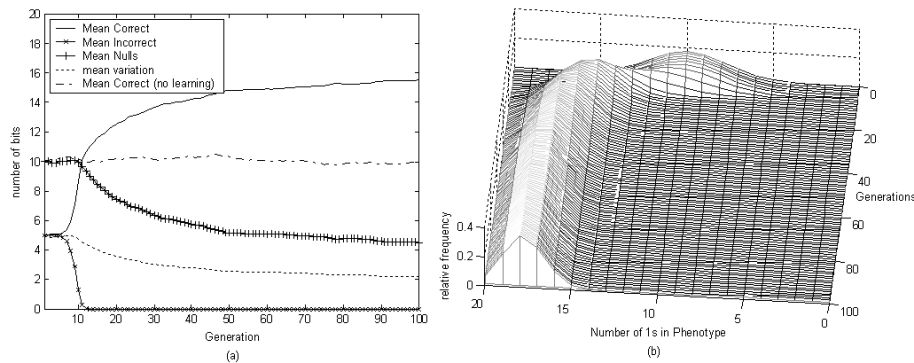


Fig. 2. Typical behaviour of Hinton and Nowlan's model

4 Discussion

Here we have presented a model which shows that although already simple, H&N's model has elements unnecessary to manifest the Baldwin Effect. Specifically, the results for our model show that lifetime plasticity can accelerate evolutionary search without: 1) the need for subsequent reduction in lifetime plasticity, 2) a learning model that recognises success, or 3) sexual recombination.

If we first consider the results in figure 1(a), it takes approximately 60 generations before the selective pressure favouring genotypes with more 1's overcomes stochastic effects (genetic drift) and is able to affect genotype frequencies. When this happens the population finds the fittest genotype very shortly afterwards. The transition is really too rapid to see that genotypes with more than 50% 1's *but not all 1's* increase in frequency before the genotypes with all 1's appear, but we can see that this pres-

sure is present because in the case without plasticity the mean number of 1's in the genotype does not increase at all. Part of the reason for the rapidity of this transition, as compared to the H&N model, is that since genetic crossover is not used, the population does not have to wait for two fitter individuals to 'find' each other before the fitter genotype can proliferate. The fittest genotype dominates the population which is indicated by the mean number of 1's in the genotype being very close to 20. The mean number of ones in the phenotype is lower since mutations still occur (in both the phenotype and the genotype) after genetic assimilation has taken place.

Results for non-plastic populations in figures 1(a) and 2(a) (the dot-dashed lines), show that both the non-plastic cases have approximately constant values for mean number of 1's in the genotype. In 30 additional runs this continued for a mean of over 1600 generations for the CP model, and a mean of over 3000 generations for the H&N model.

The dotted line in figures 1(a) and 2(a) provides a measure of the variation occurring in the phenotype from its genotype. This mean value of variation is unsurprisingly constant in the constant plasticity model, both before and after the fitter genotype has dominated the population. However, this is not the case in H&N's model; here the variation decreases through the course of the experiment. Specifically, the decrease in this variation is proportional to the mean number of '?'s present in the population's genotypes; the '?'s represent plasticity (or learning ability) so the relationship should be expected. More information is purveyed in figures 1(b) and 2(b) which depict the phenotypic distribution for these two experiments. Figure 1(b) shows a change in the mean of this distribution from approximately 10 to approximately 17 around generations 60-65; this is the point of genetic assimilation. (The mean values are also plotted on the dashed line in figure 1(a)). However, the width (variance) of this distribution is constant both before and after genetic assimilation has occurred in our model. The distribution progression shown in figure 2(b) has a different behaviour. The mean number of this distribution is again initially at 10 bits; as the 0's are purged and replaced with 1's the mean moves upwards; the distribution width is constant as the number of '?'s is also approximately constant. However, as the selection pressure begins to replace '?'s with 1's, both the mean and variance of the phenotypic distribution are changed. The mean is increased further as more 1's exist; since they replace '?'s the learning ability is reduced which directly reduces the variance. Thus, figure 2(b) shows first genetic assimilation – the shift in mean, followed by canalisation – the narrowing of the ridge, or reduction in learning ability. Since the mean of both distributions moves towards 20 1's, variation can only reduce the number of 1's present, so the shape of the distribution becomes skewed; this is an unfortunate artefact of compressing 20-dimensional data to a single axis. The variance is thus also plotted on the dotted line in figures 1(a) and 2(a), which shows these trends more clearly. We have demonstrated the cases of genetic assimilation (a change in mean) and canalisation (a reduction in variance) together as per H&N's model, as well as genetic assimilation independently in our simplified model. It is easy to imagine a third situation in which canalisation occurs without genetic assimilation; the variation of phenotypes would reduce about the mean, but the mean would be unaffected.

As already mentioned, a single typical run for each model is shown; however one point to note is that the specific number of generations required for the fitter genotype to be found in the constant plasticity model varies from run to run more than the number of generations required before the θ 's begin to be purged in H&N's model. This may be because incorrect alleles (θ 's) in the H&N model occur with half the probability that they do in the CP model. This relates to another issue with the CP model. The mechanism governing phenotypic variation in the CP model allows variations to occur on any trait whereas phenotypic variation in the H&N model only allows variation to occur on specified plastic traits. Since the plastic traits are exactly the traits that are not yet correct this gives the H&N model a distinct advantage with respect to the probability that mutations will occur in the 'correct' loci for an individual to gain fitness. Arguably, this restricts the ability of genetic assimilation without canalisation, but it also causes us to question the validity of the assumptions used in the H&N model.

5 Conclusion

The Baldwin Effect is investigated and a key ambiguity in current literature between genetic assimilation and canalisation is identified. A new model is presented which is simpler than Hinton and Nowlan's landmark model, yet still manifests the Baldwin Effect. Specifically, this new model does not use canalisation, individuals do not have to recognise their own success (i.e. cognitive learning is not required, only some form of phenotypic plasticity), and it unifies the genetic and phenotypic variation mechanisms. This shows that canalisation and learning, generally considered to be intrinsic features of the Baldwin Effect, are in fact not necessary to show that non-heritable phenotypic variation can guide genetic variation. Providing a simpler model assists us in revealing the essential mechanisms involved. These simplifications also widen the scope in which the Baldwin Effect can be applied by reducing the assumptions necessary for the effect: it may still guide the course of evolution in situations where a mechanism for canalisation is unavailable, in organisms or systems which are adaptable but not intelligently so, and in asexually reproducing populations.

Simulated experiments demonstrate the difference between the genetic assimilation and canalisation components which are often unnecessarily conflated, by consideration of the mean and variance of the distribution of phenotypes produced in a population. The conditions can be summarised as follows: canalisation occurs when there is a reduction in the variance of the phenotypes (but not necessarily a movement in the mean phenotype), whereas genetic assimilation produces a movement in the mean phenotype (but does not necessitate a reduction in the variance of phenotypes).

References

1. Baldwin, J.M.: *A New Factor In Evolution*, American Naturalist 30 (1896) 441-457, 536-554
2. Mayley, G.: *Explorations into the interactions between learning and evolution using algorithms*, DPhil Thesis, School of Cognitive and Computing Sciences, University of Sussex (2000)
3. Turney, P.: *Myths and Legends of the Baldwin Effect*, Proceedings of the 13th International Conference on Machine Learning (1996) 135-142
4. Richards, R.J.: *Darwin and the Emergence of Evolutionary Theories of Mind and Behavior*, Chicago: University of Chicago Press (1987)
5. Dennet, D.C.: *Darwin's Dangerous Idea: Evolution and the Meanings of Life*, London, Penguin (1995)
6. Hinton, G.E., Nowlan S.J., *How learning can guide evolution*, Adaptive individuals in evolving populations: models and algorithms, Addison-Wesley Longman (1987) 447-454
7. Belew, R.K.: *When both individuals and populations search*, In Schaffer, J.D. (ed.) Proceedings of the Third International Conference on Genetic Algorithms, San Mateo, California: Morgan Kaufmann (1989)
8. Harvey I.: *The puzzle of the persistent question marks: A case study of genetic drift*. In S. Forrest (editor) Proceedings of the Fifth International Conference on Genetic Algorithms, ICGA-93, California: Morgan Kaufmann (1993)
9. Wiles, J., Schultz, R., Hallinan, J., Bolland, S., Tonkes, B.: *Probing the persistent question marks*, In L. Spector, E. Goodman, A. Wu, W. B. Langdon, H.-M. Voigt, M. Gen, S. Sen, M. Dorigo, S. Pezeshk, M. Garzon & E. Burke (Eds), Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2001). San Francisco, CA: Morgan Kaufmann Publishers (2001) 710-717
10. French R.M., Messinger A.: *Genes, phenes and the Baldwin Effect: Learning and evolution in a simulated population*, Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems, Brooks, R. A. and Maes, P. (Eds.) MIT (1994) 277-282
11. Watson, R.A., Pollack, J.B.: *How Symbiosis can guide Evolution*, Fifth European Conference on Artificial Life. Dario Floreano, Jean-Daniel Nicoud, Francesco Mondada, eds. Springer (1999)
12. Watson, R.A., Reil, T., Pollack, J.B.: *Mutualism, Parasitism, and Evolutionary Adaptation*, Proceedings of Artificial Life VII, Bedau, M, McCaskill, J, Packard, N, Rasmussen, S (eds.) (2000)
13. Federici, D.: *Culture and the Baldwin Effect*, Lecture Notes in Computer Science 2801 (2003) 309-318
14. Waddington, C.H.: *Genetic Assimilation of an Acquired Character*, Evolution 4 (1953) 118-126
15. Simpson G.G.: *The Baldwin Effect*, Evolution 7 (1953) 110-117
16. Griffiths P.E.: *Beyond the Baldwin Effect: James Mark Baldwin's 'social heredity', epigenetic inheritance and niche-construction*, in Learning, Meaning and Emergence: Possible Baldwinian Mechanisms in the Co-Evolution of Mind and Language, Weber, B. and Depew, D. (eds), (2003)
17. Gibson, G., Wagner, G.: *Canalization in evolutionary genetics: a stabilizing theory?*, Bio-Essays 22:, John Wiley (2000) 372-380