Guiding or Hiding:
Explorations into the Effects of Learning on the Rate of Evolution.

Giles Mayley
School of Cognitive and Computing Sciences,
University of Sussex, Brighton, BN1 9QH, England.
`email: gilesm@cogs.susx.ac.uk`

Abstract

Individual lifetime learning can 'guide' an evolving population to areas of high fitness in genotype space through an evolutionary phenomenon known as the Baldwin effect (Baldwin, 1896; Hinton & Nowlan, 1987). It is the accepted wisdom that this guiding speeds up the rate of evolution. By highlighting another interaction between learning and evolution, that will be termed the *Hiding effect*, it will be argued here that this depends on the measure of evolutionary speed one adopts. The Hiding effect shows that learning can reduce the selection pressure between individuals by 'hiding' their genetic differences. There is thus a trade-off between the Baldwin effect and the Hiding effect to determine learning's influence on evolution and two factors that contribute to this trade-off, the cost of learning and landscape epistasis, are investigated experimentally.

1 Introduction

In recent years there has been a renewed interest in the Artificial Life and Adaptive Behaviour communities in an evolutionary phenomenon known as the Baldwin Effect (Baldwin, 1896; Hinton & Nowlan, 1987; Gruau & Whitley, 1993; Belew, 1989; Belew & Mitchell, 1996; Whitley, Scott Gordon, & Mathias, 1994). It has become the accepted wisdom that, through the Baldwin Effect, the inclusion of learning 'guides' an evolutionary system to fit solutions and therefore 'speeds up' the evolutionary process. The emphasis in this body of work has been on the rate at which fit phenotypes can be produced by a genetic algorithm (G.A.) in what are largely function optimisation experiments. However, there are many evolutionary scenarios that cannot be described as function optimisation; for example, the Species Adaptive Genetic Algorithms (SAGA) of Harvey (1992) or co-evolutionary systems that are prone to Red-Queen effects (Cliff & Miller, 1995). In these cases, measures of evolutionary speed, rate or progress that simply involve the number of generations until a fit phenotype is produced are meaningless since there will always be evolutionary driven movement in genotype space that may not lead to fitness increases.

The purpose of this paper is to show that the view concerning the speeding up of evolutionary progress afforded by the inclusion of learning depends on the measure of evolutionary speed one adopts. As an aid to this I will highlight a second evolutionary phenomenon which occurs when learning is included in an evolutionary system that, under some measures, slows down evolution. I will call this effect the 'Hiding effect'.

I will first describe the Baldwin effect and how it's effect on the dynamics of an evolving population on a fitness landscape is generally viewed. I will then describe the Hiding effect and how the evolutionary dynamics are different to the Baldwin effect. Section 4 discusses the trade-off between the two effects, suggesting conditions under which one would dominate the other. Section 5 describes the simulation experiments that were run to investigate the topics discussed, Section 6 defines the measures used to assess the rate of evolution and the results are presented in Section 7.

2 The Baldwin Effect

The Baldwin effect is an evolutionary phenomenon that has been discussed, on and off, for the past 100 years. Introduced independently by Baldwin (1896) and Lloyd Morgan (1896) it was a means of explaining cases of apparent inheritance of acquired characteristics without recourse to Lamarckianism. It was first tested empirically on *drosophila* by Waddington in the 1950's (Waddington, 1953, 1956) and was finally brought to the attention of the A-life/Adaptive Behaviour communities by Hinton and Nowlan (1987). The Baldwin effect deals with a specific interaction between the two adaptive processes of evolution and learning. Here, learning is taken to be any environmentally-driven phenotypic change that increases an individual's survival chances (fitness). The two processes are arranged thus: a population of individuals are evolving to perform a specific task (in artificial evolution with a G.A.). Their phenotypes are generated from their

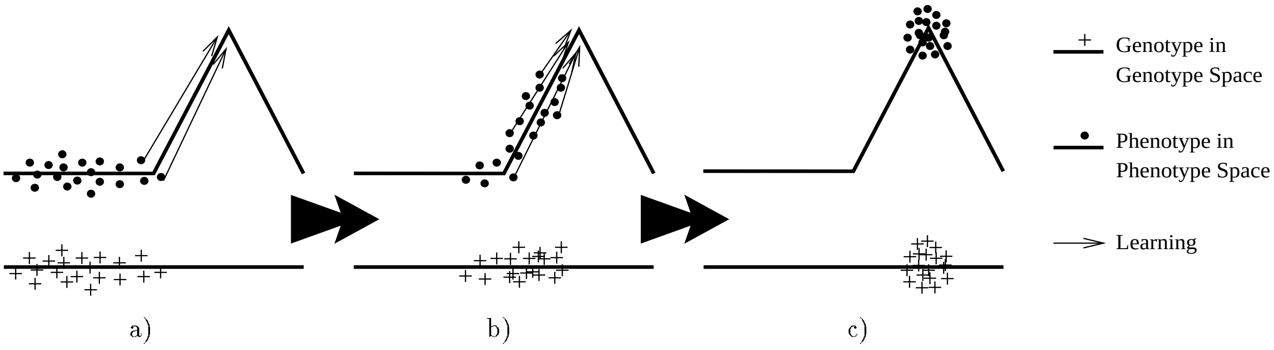

a) b) c)

Figure 1: *The Baldwin Effect: The one dimensional genotypes are translated to phenotypes which can then modify themselves in phenotype space. The heights of the phenotypes represent relative fitness. a) Some individuals in the population are able to achieve a fit trait through learning and will therefore be selected; b) Differing levels of learning among individuals in the population leads to the genetic assimilation of the trait. (Note: for clarity only a few arrows are shown — in fact the majority of phenotypes are learning); c) The trait is now innate. Learning has guided evolution to the fitness peak.*

genotypes at the beginning of a generation and are then modified by learning *on the same task* (Menczer & Belew, 1994) and given a fitness score. The fitness score of the *modified individual* is used for the purposes of selection among conspecifics but no information about the structure or behaviour of the modified individual is translated back into the genotype. Despite this non-Lamarckian framework, the learning process can still effect the course of evolution through the Baldwin effect.

Referring to Figure 1, we can see that within a population of non-identical individuals, some will be 'closer' in genotype space than others to areas of increased fitness. If these individuals are able to learn the appropriate phenotypic trait that corresponds to the increased fitness they will, therefore, be selected for (Figure 1a). The 'centre of gravity' of the population thus moves in genotype space towards the area of increased fitness (Figure 1b) and the initially learned trait may become innate in subsequent generations (Figure 1c) in a process known as *genetic assimilation* (Waddington, 1942). It can be said that learning is 'guiding' evolution (Hinton & Nowlan, 1987; Maynard Smith, 1987). This process is subject to a few conditions, though. First, once the entire population is learning the fitter trait, there must be further selection pressure for the reduction in the level of learning to occur. This is supplied by a cost for learning. That is, there are certain evolutionary costs to be paid by a learning individual over a non-learning one for displaying a particular phenotypic trait (see Johnston (1982) for biological review and Turney (1996), Anderson (1995), Mayley (1996a, 1996b) for examples from the A-Life/S.A.B. literature). Assuming these costs become reduced the 'more innate' a trait becomes, then they provide the selection pressure for the reduction in the level of learning that is necessary for genetic assimilation to occur. Another condition that has been discussed in (Mayley, 1996b) is that the genotypic space and the phenotypic space must have the property of *neighbourhood correlation* with respect to the genetic operators and the learning rule for genetic assimilation of a learned trait to be guaranteed. That is, the small changes in the phenotype that are a result of learning must correlate to possible small changes in the genotype using the genetic operators available.

Now that I have described the Baldwin effect, we can look at it in terms of its influence on the speed or rate of evolution. In Figure 1a, those individuals that are learning the fitter trait will achieve an increased fitness score. An experimenter monitoring the average fitness of the population will notice an increase in this measure. However, at this stage, there has been no movement of the population in genotype space and therefore no evolution. Learning is increasing the fitness scores additively. It is in the transition from Figure 1a to Figure 1b that a learning-directed movement in the population's position in genotype space occurs and so it is only at this stage that we can say learning is guiding evolution. With respect to the Baldwin effect, the question is: How can we be sure that evolution and learning are not producing a fitness increase in a purely additive fashion rather than through learning guiding evolution? That is, in comparing a graph of average fitness over generations in a learning and evolution experiment with that from an evolution alone experiment, if we were to remove that aspect of the fitness scores that was attributable to learning, how can we be sure we wouldn't have identical graphs? Evolution implies genotypic change, and an increased level of average population fitness in an experiment combining learning and evolution is not necessarily a good indication of increased evolutionary progress. So what is? Well, this will be dealt with in Section 6 but for now I just want to emphasise that the Baldwin effect implies that learning produces (directed) population movement through *genotype* space that would not be there if evolution were applied alone (as was most definitely the case in Hinton and Nowlan (1987)).

3 The Hiding Effect

This phenomenon, that I will call the Hiding effect of learning on evolution, is by no means a new idea. It has been described in several reviews of the effects of learning on evolution in the biology literature (Johnston, 1982; Gordon, 1992) and is mentioned in the Artificial Life/Adaptive Behaviour literature in the introduction to Belew and Mitchell (1996). However, unlike the Baldwin effect, I have not see any investigation into the nature or conditions of this effect.

The Hiding effect can be considered as one of the costs or selective disadvantages of learning (intro. Belew & Mitchell, 1996). It occurs like this: Members of an evolving population with different genotypes are selected for according to their phenotypic traits. The differences between the genotypes produce differences between their associated phenotypes that allow selection to get a hold and make a discrimination. Therefore, individuals with differing genotypes that learn to perform the same trait, or modify their phenotypes so that they are the same, reduce selection's ability to discriminate between them. Genetic differences are hidden from selection by learning. This is illustrated in Figure 2. Conceptually, each genotype in genotype space, represented by the black crosses in Figure 2, can be thought of as having an innate phenotype in phenotypic space, represented by the black blobs. The distances between the blobs indicate differences in the innately specified trait and translate directly to their associated genetic differences. If these individuals were unable to learn, then selection would have these phenotypic differences to work on and, because of the fitness differentials indicated by the height of the phenotypic landscape and the fact that the phenotypic differences are a good indicator of genetic differences, evolution would move the population towards areas of increased fitness. However, with each individual possessing the ability to learn, as indicated by the arrows, (i.e. change their phenotype to one situated in an area of increased fitness in phenotype space) then the differences between the individuals' phenotypes are reduced. Since their phenotypes are now identical, each member of the population will achieve the same fitness and so the selection pressure between them is reduced by the learning over a non-learning population in the same position. We would expect this decrease in the selection pressure to lead to a reduction in the rate of movement of the population through genotype space, i.e. the rate of evolution would be reduced by the inclusion of learning in this case.

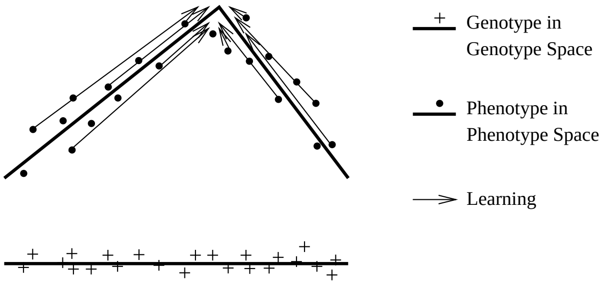

Figure 2: *The Hiding Effect: The one dimensional genotypes are translated to phenotypes which can then modify themselves in phenotype space. (Note: again, only a few arrows are shown). The heights of the phenotypes represent relative fitness. The differences between individual's genotypes are hidden to selection by their phenotypes learning the same trait.*

4 The Baldwin/Hiding Trade-off

The previous two sections have described how evolution can be both speeded up and slowed down by the inclusion of learning. There must, therefore be some sort of trade-off between the Baldwin effect and the Hiding effect or a set of conditions that determine whether one dominates the other in the determination of the actual rate of evolution. This section will describe two possible conditions which will then be tested in the experiments presented in Section 5. The first factor that influences the trade-off between the Baldwin effect and the Hiding effect is the cost of learning. It has already been stated in Section 2 that the ability to learn a specific trait can cost an individual over a non-learning conspecific. This cost is what supplies the selection pressure for the genetic assimilation of a learned trait, as in Figure 1b. It was also stated that the Hiding effect itself can be considered a learning cost or a selective disadvantage. However, the Hiding effect is of a fundamentally different nature to the costs which provide the selection pressure for genetic assimilation. Instead of being paid by the individual relative to its conspecifics, it is borne by the population as a whole[1]. It is the slowing-down of the rate of evolution that is the penalty that the Hiding effect bestows on the whole population relative to a non-learning one (or indeed a different population that is competing for resources may be able to 'out-evolve' the population that is experiencing the Hiding effect). Although the Hiding effect is itself a selective disadvantage against learning, for it to be sustained and genetic assimilation of the learned trait prevented, the individual-specific learning costs must be minimal. That is, to show this effect, we want to avoid the situation in Figure 1b in which the costs provide the selection pressure for the population to reduce the level of learning and climb to the area of

---

[1] In the classification of the various costs that have to be paid in an evolutionary system for learning, presented by Mayley (1996b), the Hiding effect falls into category 4 — individual non-specific costs.

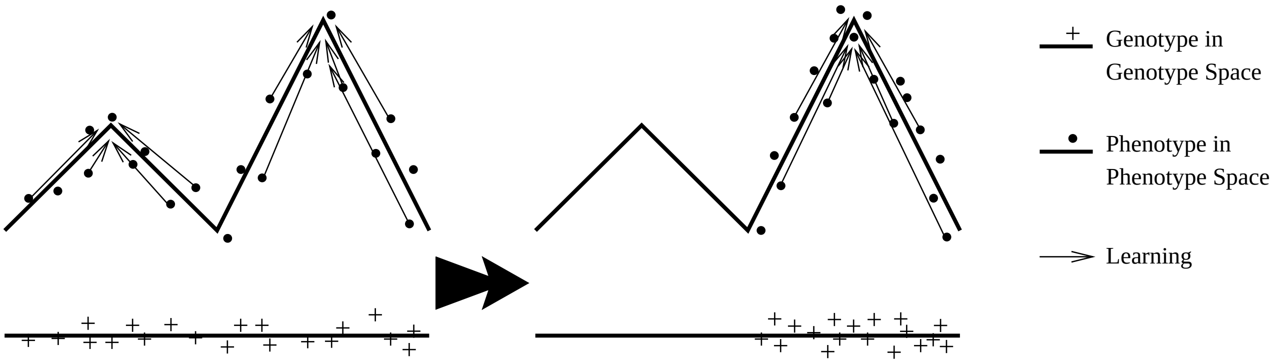

a) b)

Figure 3: *The Trade-off Between the Baldwin Effect and the Hiding Effect Brought about by Epistasis: The one dimensional genotypes are translated to phenotypes which can then modify themselves in phenotype space. The heights of the phenotypes represent relative fitness. a) Members of the population are learning their local fitness maxima. (Note: again, only a few arrows are shown); b) Differences in the heights of the local fitness maxima lead to the selection of individuals that achieve the highest maximum. This maximum is then hidden from selection by the individuals' learning.*

increased fitness in genotype space. In natural systems it is very difficult for organisms to avoid the costs of learning since their fitness is effectively continuously assessed throughout their lifetime and any mistakes made or time wasted whilst learning will cost them over a non-learning individual. However, it is possible to avoid individual specific costs in artificial evolution by awarding individuals their fitness depending on their performance achieved at the end of any learning trials as in Chalmers (1990) (see Mayley (1996a) for greater discussion of this point).

The second condition that influences the trade-off between the Baldwin effect and the Hiding effect is epistasis. We can see this if we refer to Figure 3. Members of a learning population that are spread over several peaks will each be able to learn their local peak (Figure 3a). Those that learn the highest peak will be selected for and so come to dominate the population (Figure 3b). Thus we have a directed movement of the population in genotype space towards an area of high fitness that is driven by traits that have been learnt by the individuals during their lifetime (Baldwin effect). However, once all the members of the population are learning the high peak, the genetic differences between them are no longer available for selection and, in the absence of any individual-specific learning cost, the learned trait(s) cannot be genetically assimilated (Hiding effect). This illustrates the trade-off between the Baldwin effect and the Hiding effect and we can have situations where there is a mixture of the two phenomenon. We would therefore expect the Hiding effect to dominate in evolutionary situations where there is low individual-specific learning costs and low epistasis but as epistasis increases the Baldwin effect gains more and more influence and as costs increase it should come to dominate the Hiding effect.

5 Simulation Experiments

A set of experiments is described that make explicit the cost of learning and landscape epistasis as parameters. Section 6 will describe the techniques used to monitor the rates of evolution in these experiments.

5.1 Genotypes, Phenotypes and Fitness

Each individual has a genotype and a phenotype that both consist of a binary string of length N. The goals of evolution and learning are to produce phenotypes with the highest fitness according to Kauffman's NK fitness landscape model (Kauffman, 1993). It is expected that the majority of the readership is familiar with NK fitness landscapes but for those who are not, here is the briefest of descriptions. N refers to the length of a binary string and the value of K sets the level of epistasis by determining the dependence the partial fitness of a bit at location n has on the bits in a neighbourhood of K other locations. The neighbourhood may be the K locations nearest to n in the string or a set of K locations randomly picked from anywhere on the string.

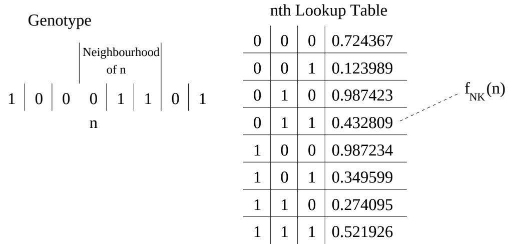

nth Lookup Table

0	0	0	0.724367
0	0	1	0.123989
0	1	0	0.987423
0	1	1	0.432809
1	0	0	0.987234
1	0	1	0.349599
1	1	0	0.274095
1	1	1	0.521926

Figure 4: *Illustration of the calculation of $f_{NK}(n)$ on Kauffman's NK fitness model with $N = 8$, $K = 2$.*

A series of N lookup tables are generated, one for each location. Each table has 2^{K+1} random entries in the interval (0,1). The fitness, F_{NK}, of a particular string is calculated by $F_{NK} = \frac{1}{N}\sum_{n=1}^{N} f_{NK}(n)$ where the partial fitness $f_{NK}(n)$ is obtained from the nth lookup table using the values of the bits in location n and its neighbourhood as the lookup key (Figure 4). Thus when $K = 0$, each location contributes independently to the overall fitness of the string and the landscape is smooth; when $K = (N - 1)$ the fitness landscape is maximally rugged.

5.2 Learning

Lifetime learning is implemented as a steepest-ascent hill-climb by the phenotype on the NK landscape. The bits in an individual's phenotype are initially set to be the same as those in its genotype, giving the individual a definite innate phenotype and a starting point from which to begin its learning search. The fitness of this innate ***phenotype*** is tested using the NK fitness model. Learning then proceeds as a set of learning trials. Each bit is flipped, the fitness of the resulting phenotype tested and the bit flipped back again. When all N bits have been tested, the bit that produced the best fitness increase is flipped permanently. This process of testing each bit and selecting the best, called a learning operation, continues until no further fitness increase is achieved. We thus have the phenotype performing a local hill-climb to its nearest local optima, in accordance that with the view that learning is a local search whilst evolution is a population based one. To make the learning cost an explicit parameter, it is incorporated into the fitness function (Section 5.3).

5.3 Experiment Overview

A population of random genotypes of size 100 is generated with $N = 20$. Each individual's genotype is copied to its phenotype and learning takes place as described in Section 5.2. The fitness is evaluated using the equation:

$$F(g_i) = F_{NK}(p'_i) - cx \tag{1}$$

where $F(g_i)$ is the fitness awarded to genotype g_i, $F_{NK}(p'_i)$ is the NK fitness of the best string found by learning, c is the cost of each learning operation and x is the number of learning operations it took find p'_i. The cx term in Equation 1 is used to provide a cost for learning. The individual is penalised a fixed amount for each learning operation that actually takes place. Once the fitnesses of all the individuals in the population have been evaluated, they are bred to form the next generation using fitness-proportional selection with 0.99 × the fitness of the worst as the base, a crossover probability of 0.7 and an average mutation rate of 0.3 bits per genotype ($0.3/N$ per bit).

6 Measuring the Rate of Evolution

In the above discussion we have been considering the influence of learning on the rate of evolution, but what exactly is meant by this term?

It was stated in Section 2 that monitoring the average fitness of the population in a non-Lamarckian framework could lead to a false idea of the rate of evolution. The average fitness of the population is an indication of the rate at which the system is able to produce fit phenotypes. If the experimenter's motivation is an engineering one, where the requirement is to produce a good, working solution to a problem, then this is sufficient. However, that is not the motivation here. We are interested in the effect learning has on the movement of the population through genotype space: The Baldwin effect promotes that movement; the Hiding effect suppresses it. A similar situation was encountered by Cliff and Miller (1995) in which problems of assessing continued evolutionary activity in co-evolutionary scenarios were considered. Here, tracking average fitness is meaningless since the fitness of one population depends on the other. Cliff and Miller used image-processing techniques to look at patterns of genetic change and persistence over generations in bitmap images of elite and consensus sequences[2]. Although giving a good indicator of genetic activity, their results were largely qualitative making checking the rate of evolution in any given system difficult. Bedau and Packard (1991) monitored evolutionary activity in a different way. All of the phenotypic traits of the members of the population were monitored for their lifetime usage and therefore their contribution to fitness[3] could be individually assessed. The genes that directly coded for those traits that persisted population-wise over generations were thus considered evolutionarily favourable and the rate at which new, favourable genes were created was considered a measure of evolutionary activity. The problem with the approach as a general method of measuring evolutionary activity is that once the direct mapping between genotype and phenotypic traits is lost (e.g. some sort of morphogenesis scheme is used to construct a phenotype from a genotype) then the decision as to which gene contributes to which phenotypic trait is a hard one. More specifically, when dealing with learning and evolution, if the phenotype changes during its lifetime through learning then the decision is impossible.

One of the biggest problems encountered by both Cliff and Miller (1995) and Bedau and Packard (1991) in tracking evolutionary movement in genotype space is that of distinguishing between useful, directed population movement and random drift. Harvey and Thomp-

---

[2] Other techniques were also described that only apply to co-evolutionary scenarios.

[3] Bedau and Packard actually used an energetic model where an individual's fitness was implicit but that distinction is not important here.

son (1996) have developed methods to try and accomplish this and the techniques described here draw on this work. I will describe a collection of measures of evolutionary activity that together are sufficient to give us a picture of the different effects that learning has on evolution when influenced by the Baldwin effect and the Hiding effect.

Centroid Movement: The centroid of a population at any given generation is its center of gravity in genotype space. Each genotype is treated as a vector of 0.0's and 1.0's and the centroid is average of these vectors: $C_i = \frac{1}{P}\sum_{j=1}^{p} g_{ij}$. Where C_i is the centroid at generation i; P is the population size and g_{ij} is the jth genotypic vector at generation i. We can thus plot the distance moved by the centroid from generation $i-1$ to i: $|C_i - C_{i-1}|$.

Centroid Direction Correlation (C.D.C.): A population under no selection pressure that is performing a random walk will show a large centroid movement each generation. We want to capture a more directed form of movement; a movement that is sustained over generations as new genetic material takes over the population. This is done by taking the cosine of the angle between the vectors of the motion of the centroid between two generations:

$$C.D.C._i = \frac{(C_i - C_{i-1}) \cdot (C_{i-1} - C_{i-2})}{|C_i - C_{i-1}||C_{i-1} - C_{i-2}|} \qquad (2)$$

Thus, when the centroid moves in a similar direction through genotype space over two generations, this measure will be close to 1.0; in opposite directions it will be close to -1.0 and orthogonally, close to 0.0 (most likely with random movement in high dimensional space).

Principal Component Analysis: Harvey and Thompson (1996) use a technique to visualise the movement of their evolving population through genotype space that reduces the dimensionality of their genotypes from 1800 to 2. This is done by projecting the centroids of the population over generations onto the first and second principal components of the trace they made through genotype space. The direct genotype to phenotype mapping and the one to one relationship between evolutionary task and the learning task in the particular experiments above, allow us to go further. As well as projecting the genotypic centroids onto the principal components, we can also project the phenotypic centroids. This will effectively lay phenotype space on top of genotype space, viewed from the angle of most movement in genotype space, and should allow us to see if learning is indeed guiding evolution.

Innate Phenotype Fitness: Since the above experiment was designed such that each individual has an innate phenotype on which learning acts, we can remove the components of the fitness scores that are attributable to learning. We can then plot the average of this across the population; it should indicate whether the population has moved to areas in genotype space that are innately fitter.

7 Results

The experiment was run with parameter values $c = 0.00$ (individual-specific cost-free learning), $c = 0.03$ (costly learning), and $K = 0$ (non-epistatic), $K = 5$ (epistatic). The results presented in this section are the average of 50 runs of the simulation unless otherwise stated.

7.1 Cost-Free Learning

Figure 5 shows the results for when $c = 0.00$. In the case where $K = 0$, the first thing to note is that the average fitness graph and the fitness of the best individual (Figure 5a) are constant and equal throughout the evolutionary run (the two graphs are laid over each other so there appears to be only one line). All members of the population have been able to find the globally optimal phenotype from the first generation, regardless of their genetic make-up. This continues for the entire simulation run. Because their phenotypes are all achieving the same score, there is no selection between the genotypes, and the average fitness scored by the innate phenotypes remains at 0.5, the average score that a randomly generated string would achieve on an NK fitness landscape. That is, there is no selection pressure for the genotypes to move to areas of increased fitness in genotype space that are achievable without learning. Each individual obtains it's fitness score purely through learning. With $c = 0.00$ and $K = 0$ we are seeing the Hiding effect at its most prominent: Each individual obtains its fitness score purely through learning and there is no evolution.

However, this is not to say that there is no movement of the genotypes in genotype space at all, as we can see if we look at Figure 5b, the graph of distance moved by the centroid of the population each generation. We notice that there is in fact a consistent level of movement throughout the simulation run (average over generations = 0.208). Since there is no evolution, i.e. no selection resulting in directed population movement, this motion is due to random drift[4]. Confirmation of this is given by the C.D.C. graph (Figure 5c). Remembering that this is a measure of directed centroid movement over generations and will be positive for a correlated direction, negative for anti-correlated and zero for uncorrelated, we can see that during this experiment there was very little directed movement within genotype space (average over generations = -0.0064). We can get strong qual-

---

[4] The level in Figure 5b is the same as that produced in a separate experiment where members of population were each awarded a random fitness with all other experimental conditions identical (not shown).

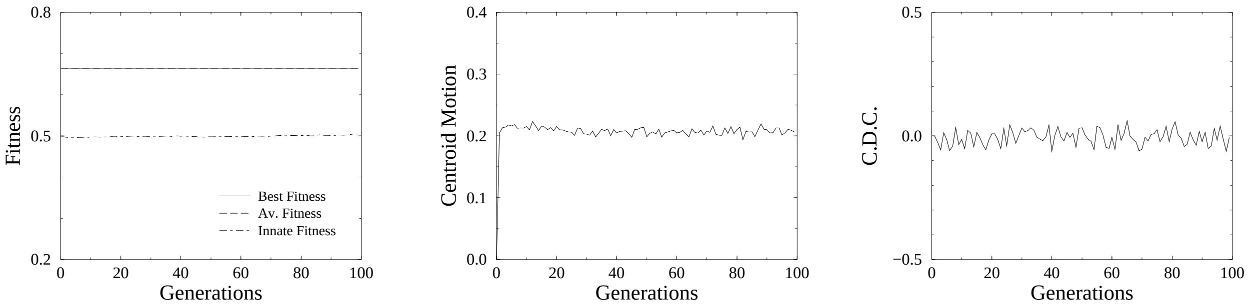

a) b) c)

$c = 0.00$, $K = 0$.

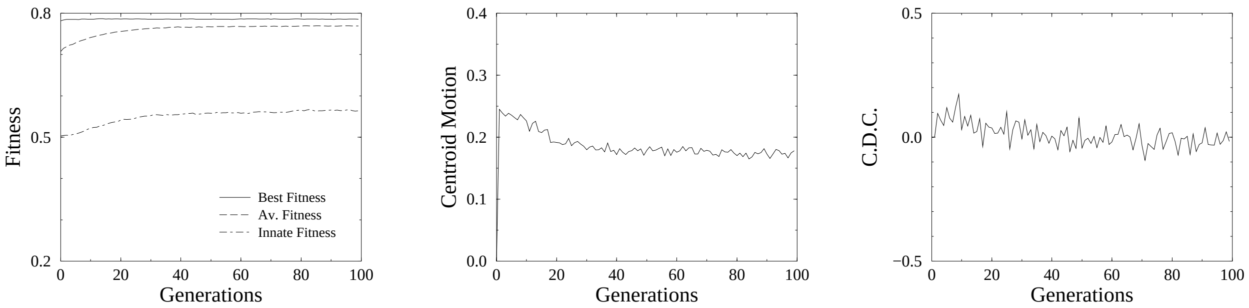

d) e) f)

$c = 0.00$, $K = 5$.

Figure 5: *Results from experimental runs with* $c = 0.00$ *and* $K = 0$, $K = 5$: *a), d) Fitness of the best member in the population, average fitness of members in the population and average innate fitness of members in the population, all over generations, plotted on the same axes; b), e) Distance moved by the centroid of the population in genotype space; c), f) Centroid Direction Correlation — indication of the correlation in the direction of movement of the genotypic centroid each generation.*

itative confirmation of this picture from the Principal Component Analysis shown in Figure 7a. The crosses represent the projection of the genotypic centroids on the first two principal components of the genotypic centroids over generations, the circles, the learnt phenotypic centroids from the same perspective and the lines connect genotypic centroids with phenotypic centroids from the same generation, representing learning. All the phenoytpic centroids are positioned on top of each other at the top of the Figure. This is because the phenotypes all learn the global optima throughout the evolutionary run. However, the genotypic centroids move around genotype space in a random way — there is no selection pressure for them to move in any particular direction.

Moving on to the case where $c = 0.00$ and $K = 5$ — a more epistatic fitness landscape — we can see a marked difference in the fitness graphs in Figure 5d when compared with 5a. First of all, the best fitness remains high for the entire run as in the case with $K = 0$ but the average fitness is no longer constant over generations. It starts off high, well above the 0.5 random average, and climbs steadily in the first 40 generations or so, until it is near the best fitness. We can see why this happens from the other three lines. The population starts off spread out all over genotype space. Their phenotypes all learn their local fitness peaks, which will vary in size, and the ones that achieve the higher fitness scores will be selected for. The phenotypic selection has the effect of moving the population in *genotype space* so that they are near the area corresponding to increased fitness in phenotype space. The graph of innate fitness thus starts off at the random average of 0.5 and increases steadily in the same period to reach a level of 0.56 as the population moves to the fitter area. Figure 5e show the magnitude of this movement in genotype space indicated by the increased level (above 0.208) of this line in the first 20 generations or so. Confirmation that this is *directed* movement is given in Figure 5f where the C.D.C. is beginning to show consistent levels above zero in the same period. The fact that the innate fitness graph doesn't climb any further indicates that there is still a large level of learning occurring once the population has settled around a peak; the peak now being hidden from selection by that learning.

The P.C.A. in figure 7b shows the movement in genotype space in this case. The genotypic centroids start off in the bottom righthand corner of the Figure with the centroids of their learned phenotypes to the northwest of them. The population then moves in genotype

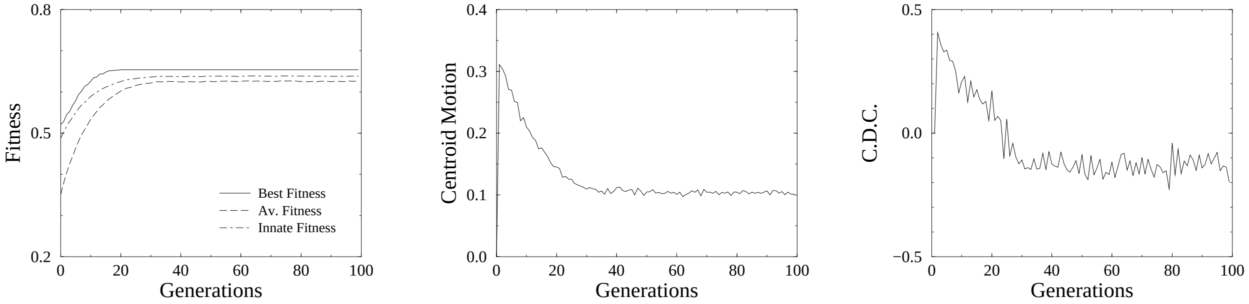

a) b) c)

$c = 0.03$, $K = 0$.

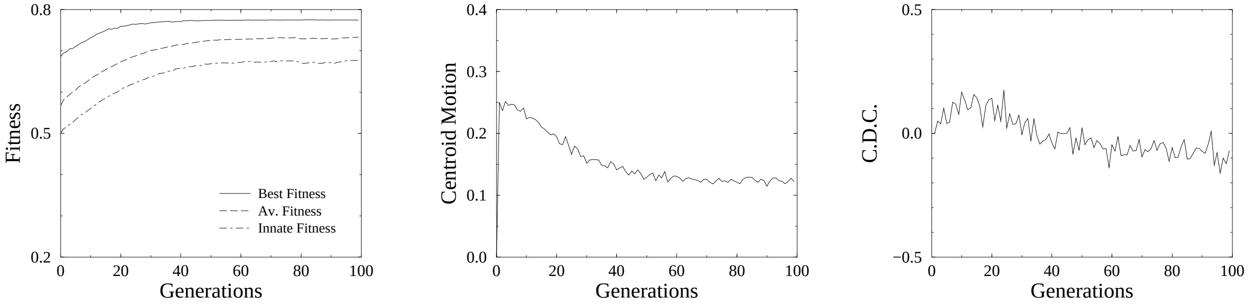

d) e) f)

$c = 0.03$, $K = 5$.

Figure 6: *Results from experimental runs with $c = 0.03$ and $K = 0$, $K = 5$: a), d) Fitness of the best member in the population, average fitness of members in the population and average innate fitness of members in the population, all over generations, plotted on the same axes; b), e) Distance moved by the centroid of the population in genotype space; c), f) Centroid Direction Correlation — indication of the correlation in the direction of movement of the genotypic centroid each generation.*

space in this same direction, following the phenotypes. In the latter generations, the genotypic centroids never quite 'catch up' with the phenotypic centroids indicating that full genetic assimilation has not occurred. The population has moved to an area in genotype space that innately produces fitter that random phenotyes, through fitness scores that were achieved through learning. It is fairly clear that the situation described in Section 4 has occurred here in which there has been a trade-off between the Baldwin effect and the Hiding effect.

7.2 Costly Learning

We now turn to Figure 6 which shows the results from simulation runs where learning costs the individuals 0.03 fitness points for each learning operation performed. Firstly, when $K = 0$, it is noticeable from Figure 6a that both the best fitness and the average fitness start off appreciably lower than in the cost-free cases and then climb up to a similar value after about 30 generations. The reason for these lower levels of fitness is that, although the individuals are learning the global optima, they are being penalised for that learning. The increasing fitness levels over generations then comes about through the selection pressure to reduce these costs which can only be done by the population moving in genotype space towards the area that produces innate, globally-optimal phenotypes to reduce the level of learning. This is shown by the innate fitness graph in Figure 6a. It climbs to a level similar to the other two lines indicating that the majority of an individual's fitness is a result of the position of its innate phenotype in phenotypic space and only minimally as a result of learning. This shows strongly that the global optima has been genetically assimilated[5]. Figure 6b and c show the strong movement across genotype space as this assimilation takes place. In the early generations the levels are well above their baselines of 0.208 and 0.0 respectively, but it is interesting to note what happens to them after the genotypes reach the global optimum at about generation 30 — Figure 6b falls well

---

[5] One point to note in passing with these graphs is that the innate fitness graph is higher that the average fitness graph. This is because the individuals are being penalised heavily for their learning but, because we are imposing learning on them, can only evolve to areas in genotype space that require less learning rather than to not learn at all. This is of little consequence to the arguments presented here since we are dealing only with learning individuals.

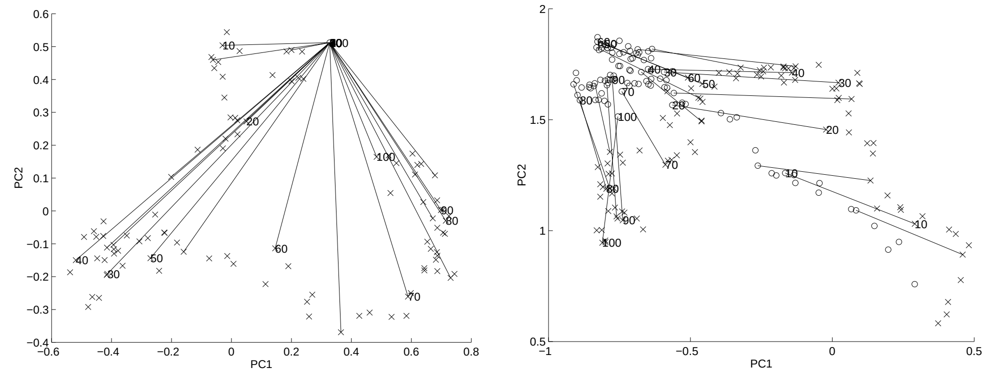

a) $c = 0.00$, $k = 0$ b) $c = 0.00$, $k = 5$

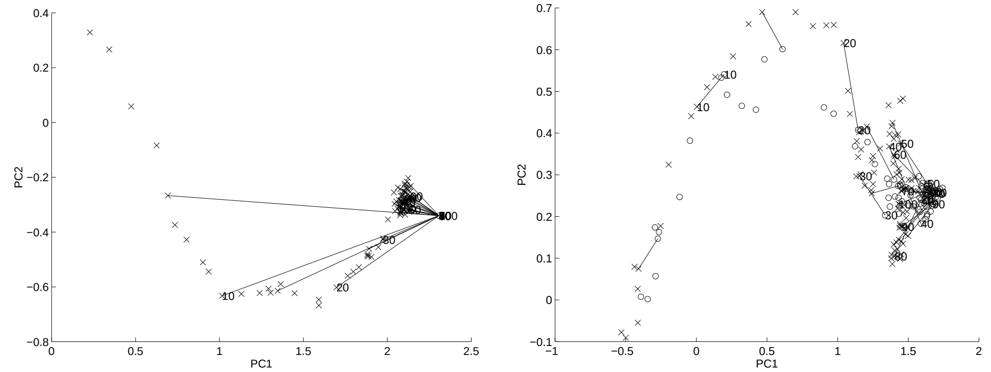

c) $c = 0.03$, $k = 0$ d) $c = 0.03$, $k = 5$

Figure 7: *Projection of the genotypic centroids (crosses) and the phenotypic centroids (circles) on the first two principal components of the genotypic centroids over generations. The lines connect genotypic and phenotypic centroids from the same generation to represent learning (for clarity they are only shown for every fifth generation) and the numbered labels indicate generation (every 10). These plots are each from a single simulation run with the parameter settings as shown.*

below 0.208 and the C.D.C. becomes negative. This is because the population is converged on the global optimum with little selection pressure in any direction. Mutation will cause individuals to fall off the optimum and any imbalance in the direction that individuals fall will cause the centroid to slowly shift away from the optimum. The fallen individuals will then have to learn the optimal phenotype and therefore incur a cost of learning. These individuals will then be selected against and the centroid will move back towards the optimum — the opposite direction from which it moved previously. We therefore get small, frequent movements of the centroid that are negatively correlated with the previous direction.

The P.C.A. for costly learning on a non-epistatic landscape is shown in Figure 7c and is rather different from the cost-free case in Figure 7a. Once again, all the centroids of the learned phenotypes are situated on the global optimum (to the right of the Figure) throughout the evolutionary run but in the first 30 generations the genotypic centroids quickly arrive in the same area as the learning becomes genetically assimilated, in contrast with Figure 7a. Subsequent to that, the genotypic centroids remain very close to the optimum, confirming the picture of a population continually falling off the peak and then being pushed back.

We look now at costly learning with $K = 5$. In this case the fitness, centroid movement and C.D.C. graphs indicate a less severe but more prolonged movement of the centroid across genotype space than in the $K = 0$ case. This is because there are more fitness peaks across phenotype space for the learning individuals to explore, reducing the selection pressure between them. This is born out by the P.C.A. shown in Figure 7d where the phenotypic centroids progress across the projection in contrast to Figure 7c where they were firmly fixed to the global optimum. The genotypic centroids follow the phenotypic centroids in their associated space in a similar manner to Figure 7b but at a faster rate because of the selection pressure of reducing the cost of learning. Notably also is that, in the later generations, the genotypic centroids are packed more tightly around the area of genotype space that innately corresponds to the fitness peak in phenotype space that the population eventually settles on. This is because the selection pressure from the cost of learning overcomes the Hiding effect that was seen in Figure 7b.

8 Conclusion

I started with a discussion of the Baldwin effect and stated that the accepted wisdom is that the inclusion of learning in a evolutionary scenario speeded up the rate of evolution. It was noted that this is usually assumed from an increased production of fit phenotypes as indicated by an average or best fitness measure. I then suggested that these measures may not be appropriate, highlighting the Hiding effect as a phenomenon that slowed the progress of an evolving population through genotype space, whilst still producing fit phenotypes. Factors governing whether the Baldwin effect or the Hiding effect dominates in any given evolutionary scenario were suggested and it was shown experimentally that the Hiding effect dominates when epistasis and the cost of learning are low, the Baldwin effect dominates when the cost of learning is high, but when the costs are low and the epistasis is high there is a mixture of the two.

One note concerning C.D.C. measurement of evolutionary rate used here: On the NK fitness landscapes it was a good indicator of directed genotypic movement. However, on more realistic landscapes in which there is significant levels of 'junk D.N.A.' and neutral networks, hitchhiking may lead to some problems.

Acknowledgments

I would like to thank Dave Cliff, Inman Harvey, Phil Husbands and Adrian Thompson for their kind help.

References

Anderson, R. (1995). Learning and evolution: A quantitative genetics approach. *Journal of Theoretical Biology*, *175*, 89–101.

Baldwin, J. (1896). A new factor in evolution. *The American Naturalist*, *30*, 441–451.

Bedau, M., & Packard, N. (1991). Measurement of evolutionary activity, teleology, and life. In Langton, C., Taylor, C., Farmer, J., & Rasmussen, S. (Eds.), *Artificial Life II: Proceedings of the Workshop on Artificial Life.* Addison-Wesley.

Belew, R., & Mitchell, M. (Eds.). (1996). *Adaptive Individuals in Evolving Populations: Models snd Algorithms.* Addison-Wesley.

Belew, R. K. (1989). When both individuals and populations search: Adding simple learning to the genetic algorithm. In Schaffer, J. D. (Ed.), *Proceedings of the Third International Conference on Genetic Algorithms*, pp. 34–41 San Matteo, CA. Morgan Kauffmann.

Chalmers, D. J. (1990). The evolution of learning: An experiment in genetic connectionism. In Touretzky, D. S., Elman, J. L., Sejnowski, T. J., & Hinton, G. E. (Eds.), *Proceedings of the 1990 Connectionist Models Summer School* San Matteo, CA. Morgan Kaufmann, San Matteo, CA.

Cliff, D., & Miller, G. (1995). Tracking the Red Queen: Measurements of adaptive progress in co-evolutionary simulations. In Moran, F., Moreno, A., Merelo, J., & Cachon, P. (Eds.), *Advances in Artificial Life: Proceedings of the Third European Conference on Artificial Life (ECAL95)*, pp. 200–218. Springer Verlag.

Gordon, D. (1992). Phenotypic plasticity. In Fox Feller, E., & Lloyd, E. (Eds.), *keywords in Evolutionary Biology.* Harvard University Press.

Gruau, F., & Whitley, D. (1993). The cellular developmental of neural networks: the interaction of learning and evolution. Research report 93-04, Laboratoire de l'Informatique du Parallélisme, Ecole Normale Supérieure de Lyon.

Harvey, I. (1992). Species adaptation genetic algorithms: A basis for a continuing SAGA. In Varela, F., & Bourgine, P. (Eds.), *Towards a Practice of Autonomous Systems: Proceedings of first European Conference on Artificial Life.* MIT Press/Bradford Books.

Harvey, I., & Thompson, A. (1996). Through the labyrinth, evolution finds a way: A silicon ridge.. In *Proceedings of The First International Conference on Evolvable Systems: from Biology to Hardware (ICES96).* Springer-Verlag.

Hinton, G. E., & Nowlan, S. J. (1987). How learning can guide evolution. *Complex Systems*, *1*, 495 – 502.

Johnston, T. (1982). Selective costs and benefits in the evolution of learning. In Rosenblatt, J., Hinde, R., Beer, C., & Busnel, M. (Eds.), *Advances in the Study of Behavior.*, Vol. 12, pp. 65–106. Academic Press.

Kauffman, S. (1993). *The Origins of Order: Self-Organization and Selection in Evolution.* Oxford University Press.

Lloyd Morgan, C. (1896). On modification and variation. *Science*, *4*, 733–740.

Mayley, G. (1996a). The evolutionary cost of learning. In Maes, P., Mataric, M., Meyer, J., Pollack, J., & Wilson, S. (Eds.), *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, pp. 458–467. MIT Press.

Mayley, G. (1996b). Landscapes, learning costs and genetic assimilation. *Evolution, Learning and Instinct: 100 Years of the Baldwin Effect. A Special Edition of Evolutionary Computation*, *4*(3), 213–234.

Maynard Smith, J. (1987). When learning guides evolution. *Nature*, *329*, 761–762.

Menczer, F., & Belew, R. (1994). Evolving sensors in environments of controlled complexity. In Brooks, R., & Maes, P. (Eds.), *Artificial Life IV*, pp. 210–221.

Turney, P. (1996). Myths and legends of the Baldwin Effect. In Fogarty, T., & Venturini, G. (Eds.), *Proceedings of the ICML-96 (13th International Conference on Machine Learning)*, pp. 135–142 Bari, Italy. Workshop proceedings not published.

Waddington, C. (1942). Canalization of development and the inheritance of acquired characters. *Nature*, *150*, 563–565.

Waddington, C. (1953). Genetic assimilation of an acquired character. *Evolution*, *7*, 118–126.

Waddington, C. (1956). Genetic assimilation of the *bithorax* phenotype. *Evolution*, *10*, 1–13.

Whitley, D., Scott Gordon, V., & Mathias, K. (1994). Lamarckian evolution, the Baldwin Effect and functional optimization. In Davidor, Y., Schwefel, H. P., & Manner, R. (Eds.), *Parallel Problem Solving from Nature-PPSN III*, pp. 6–15. Springer-Verlag.