11

The Baldwin Effect in the Age of Computation

Ruben R. Puentedura

1 Introduction—The Baldwin Effect

Baldwin (1896) introduced the notion that animals might, via plastic learning within their lifetimes, affect their subsequent evolutionary path in areas related to this plastic learning. Baldwin's proposal was not constructed as an exemplar of Lamarckian thought, but was rather designed as a viable alternative to it. Thus, the plastic learning involved is not incorporated directly into the germ line (ă la Lamarck), but rather serves as an "in-life" crutch that allows animals that have learned a "good trick" to outcompete the others. While Baldwin clearly thought of his theory in terms of Mind as the source of plastic learning, it is generally applicable to scenarios invoking only a rudiment of neural thought—see, for instance Dennett (1995)—and in fact could be reasonably applicable to scenarios invoking no neural interaction at all, but only a modicum of in-lifetime behavioral plasticity. I will leave a more detailed description of the general mechanism of the Baldwin effect to other authors in this volume (see, for instance, Paul E. Griffiths's chapter in this volume). The specifics of one nonbiological incarnation of the Baldwin effect—the computational model developed by Hinton and Nowlan (1987)—will instead be the focus of this essay. The rationale for the interest in computational models such as this one is twofold: first, given the difficulty of determining unequivocally the presence of the Baldwin effect in biological systems (Waddington 1942), it is desirable to provide a model that can be investigated directly and might hence shed some light on the actual activities of its biological counterparts. Second, the Baldwin effect may in fact provide a tool that may prove of considerable use in computer modeling. Genetic algorithms on their own having proven in many cases somewhat disappointing,

a mechanism that would allow a neural network component to aid the genetic algorithm's evolutionary path would indeed be warmly welcomed.

2 The Hinton and Nowlan Model

The task defined in the model is the following: start out with a string containing twenty loci, all set to 1: this will provide a "fitness target." Next, create a population of 1000 strings (the "organisms"), all containing twenty loci, all set at random to either 1 or 0. Compare the strings to the fitness target—if any string matches the target exactly, assign it a fitness value of 20; otherwise, assign it a fitness of 1. Assign a probability of reproduction to the organisms based upon their fitness value (i.e., the higher the fitness value, the greater the reproductive probability), then create a new generation by random crossover among the fitter organisms. Repeat the process, until a majority of the population has found the target.

Clearly, the problem is insoluble in any reasonable fashion by this mechanism—the space to be explored yields only $1/2^{20}$ as the probability of a correct string match for any given randomly selected string, a negligible percentage. However, now assume that the "digital critter" is allowed to apply an in-lifetime "guessing approach" to some of the locations. The coding used by Hinton and Nowlan takes the form:

1: definitely set

0: definitely not set

?: available to guessing

In other words, locations with a ? in them allow for a set of lifetime guesses by the organism (with a lifetime set by Hinton and Nowlan at 1000 tries) at the correct fitness target. A possible coding gene for the organism then looks like:

1100?11?001??1?0?110

Hinton and Nowlan set the average distribution of 1s, 0s, and ?s in the starting population at 25 percent, 25 percent, and 50 percent, respectively. Furthermore, to go along with this "plastic genome," matching to the fitness target also needs to be adjusted. The fitness is now given by a function F, where

$$F = 1 + 19 \cdot n/1000$$

and n is the number of trials remaining (out of 1000) for a particular organism after it stumbles upon the correct result. As can be readily seen, organisms that start out with all twenty locations set to 1 have a fitness of 20, while those with any location set to 0 have a fitness of 1. Organisms with a number q of locations in their genome set to ? and all other locations set to 1 have a probability p of matching the string in any given round given by:

$$p = 1/2^q$$

Remarkably, with this change, the task goes from being impossible to being readily achievable in a few generations. The result of a standard run of this model is depicted in figure 11.1a—as can be seen, the solution is reached in roughly 15 generations.

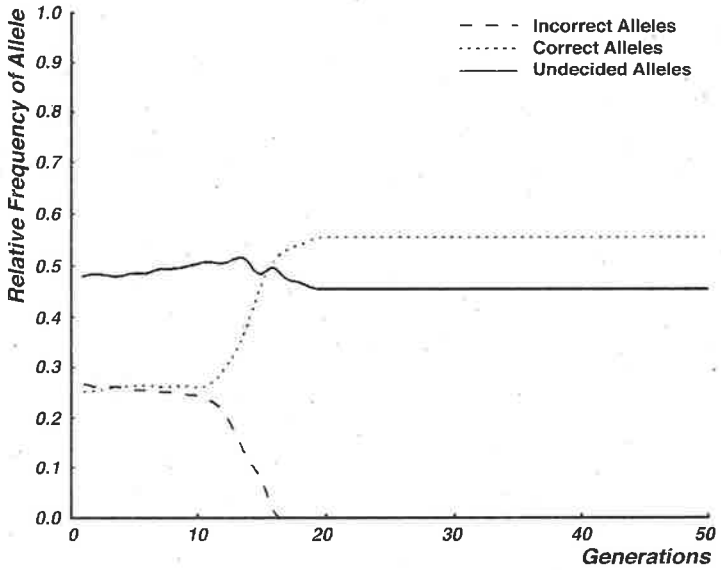

Figure 11.1a
The frequency in the population of "correct" (1) alleles rises sharply at around the 15th generation, indicating that a partial solution has been found; the frequency of "incorrect" (0) alleles drops sharply at this same point, while that of "undecided" alleles drops only slightly, indicating that the "correctness" is achieved primarily at the expense of the absolutely incorrect alleles, rather than the undecided ones. Note that: $\text{Freq}_1 + \text{Freq}_0 + \text{Freq}_? = 1$. (From Hinton and Nowlan 1987.)

3 What Does This Mean?

The Hinton and Nowlan model clearly establishes that, under some very specific and stringent limitations, the Baldwin effect can be observed in computational systems. The task, clearly impossible under the conditions originally stated for the system, becomes eminently possible. This is done via a "smoothing" of the evolutionary landscape: the plasticity transforms the single spike into a gentler curve, as shown in figure 11.1b.

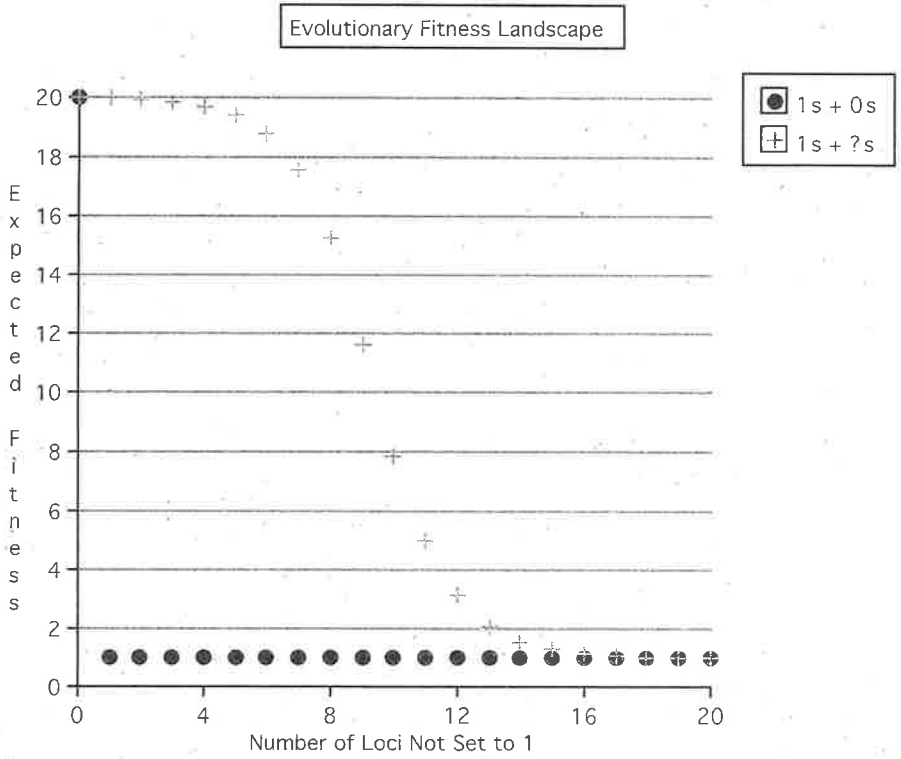

Figure 11.1b
The expected fitness for the nonguessing setup is a single sharp spike, corresponding to all 20 loci being correctly set. By contrast, the probabilistic element introduced by the guessing loci allows for reasonably high fitness results for up to about 10 loci being set to ?.

4 What Does It Not Mean?

Some aspects of the original model have been overinterpreted, and given far too much prominence. By a long shot, the worst of these is the question of what have been called the "persistent question marks" (Harvey 1993). This refers to the fact that, as can be seen in the diagram above, the fitness "levels off" before all the question marks have been replaced by ones. A common misreading of this aspect of the model is that this implies that the Baldwin effect cannot, in fact, yield the global optimum (i.e., a population with all genetic loci set to 1) for the problem, a misinterpretation further confused by some erroneous results of Belew (1989). As pointed out by Harvey (1993), the observed leveling off is nothing more than the result of "hitchhiking" in the genome with a limited population. Consider the fact that no mutation at all is allowed in the Hinton and Nowlan model; therefore, in a population where all 1000 members have a specific locus with ? in it, there is no possibility that it will be replaced by a 1. A simple probabilistic analysis reveals that this would not be the result of an aberrant run, and in fact runs such as this are guaranteed to be the norm (again, for the details, see Harvey 1993). Adding even a minute amount of mutational probability quickly rectifies the situation.

5 Critiques and Responses

Interestingly, some possible critiques of the model have merited little or no serious commentary in the literature. In the remainder of this paper, I will deal with three of these problems: the coupled genome problem, the autonomous learning function problem, and the teleological problem.

6 The Coupled Genome Problem

The coding used by Hinton and Nowlan suffers from one significant problem in terms of its applicability to the broader biological picture: the allele that codes for plasticity (?) is a replacement for the allele that codes for a fixed location (1 or 0). In other words, one necessarily replaces the other. This is a weak aspect of the model: it militates against the likelihood that

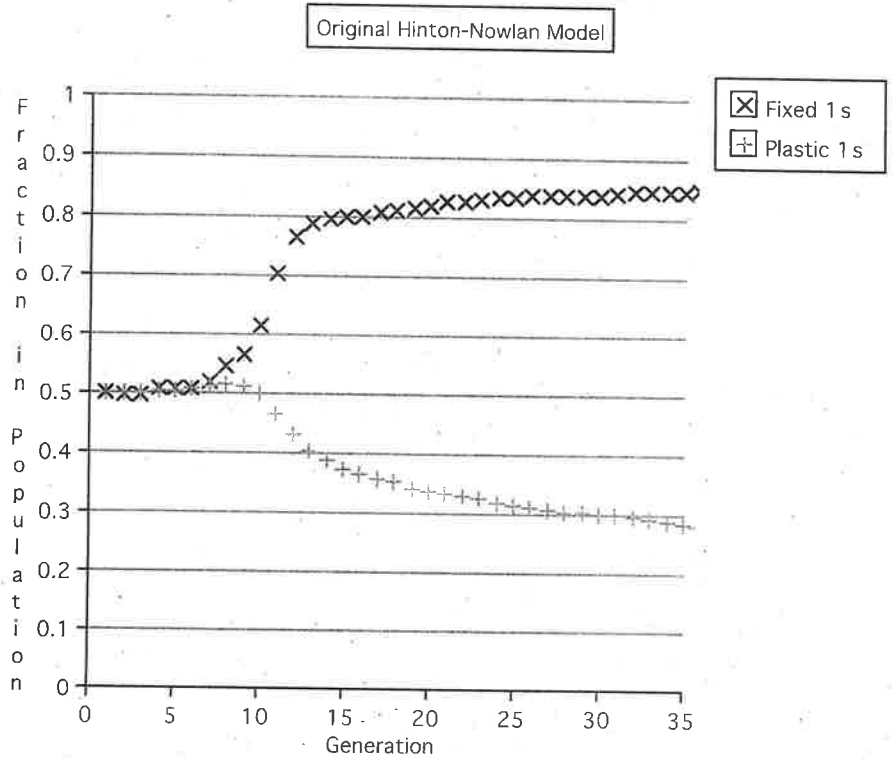

Figure 11.2a
The curves shown here are akin to those displayed above in the original Hinton and Nowlan experiment; for exact interpretation purposes, it should be noted that some of the "fixed" 1s here are overlaid with "plastic" 1s. The exact frequencies observed in this run are somewhat different from those obtained by Hinton and Nowlan, but fully compatible with those subsequently obtained by Harvey (1993).

this would be a mechanism of significant biological import—consider that this would mean that the key gene for the "good trick" learning potential in the organism's lifetime, and the "hard-coded" version would have to be one and the same. Fortunately, this is an easily soluble problem. We adopt the following "two gene" model:

Fixed location gene: 20 sites, all 1 or 0

Plasticity gene: 20 sites, all 1 (plastic) or 0 (fixed)

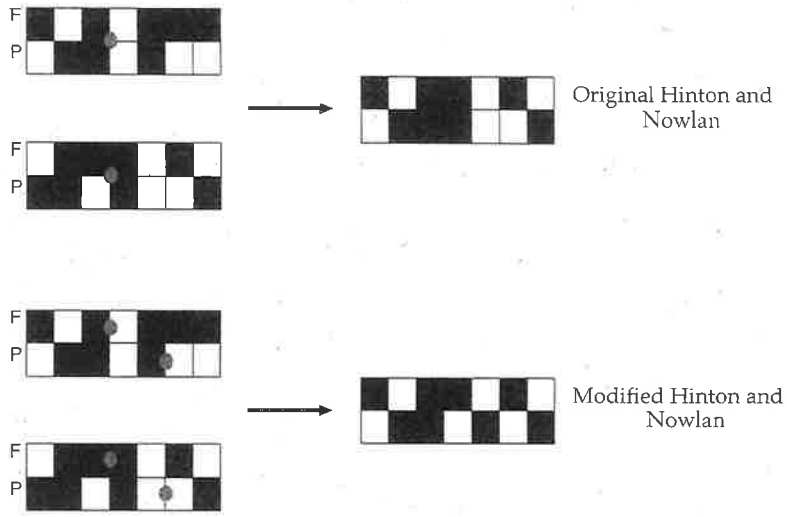

Figure 11.2b
The dots represent the crossover point for the creation of a "child" gene from two parents; in this diagram we assume the first part of the child gene comes from the top "parent," the second part from the bottom one. In the original Hinton and Nowlan model, the crossover point for both the fixed (F) and plastic (P) genes is forced to be the same, since plastic genes are not autonomous from fixed genes; in the modified model, this restriction is lifted.

This new coding can easily be mapped onto the old Hinton and Nowlan model:

Fixed: 1101101...

Plastic: 1001110...

Hinton and Nowlan: ?10???1...

Note that, when the crossover location in both genes is forced to be the same, and the percentage of ones in each gene is set to 50 percent, we recover the original Hinton-Nowlan model. A run with this constraint is illustrated in figure 11.2a. However, if we loosen this constraint—in other words, if we allow the two crossover points to vary independently, we obtain a new model, not subject to the preceding critique. Again, figure 11.2b should help clarify this point. It is gratifying to note that the results for a typical run for this model are akin to those already noted for the original Hinton and Nowlan model, as shown in figure 11.2c.

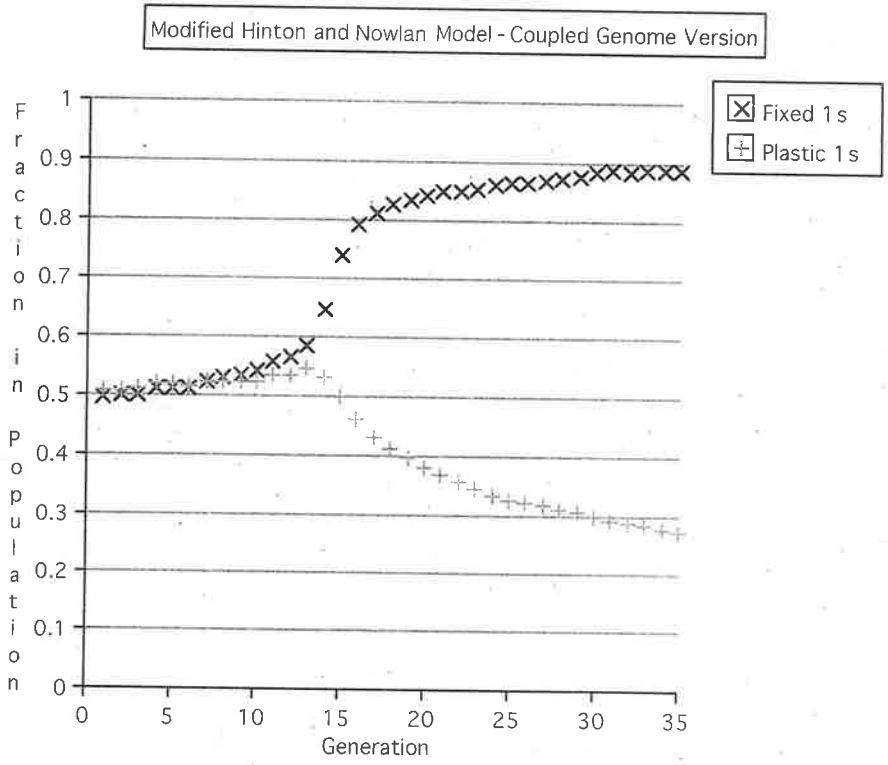

Figure 11.2c
While the onset of the sharp rise in the fraction of fixed 1s in the population is delayed (by about 5 generations) from that observed in the original Hinton and Nowlan model, the overall curve profiles, and hence conclusions, are essentially similar.

7 The Autonomous Learning Function Problem

It can be argued that the Hinton and Nowlan model may have made the search task of its critter a little too easy by making the learning function too autonomous. In other words, the fact that an independent coin toss occurs for each ?-locus, regardless of the "hard-coded" aspects of other gene sites, indicates that the hard-coding has no effect at all on the plastic behavior—again, not that likely an assumption if we would like to apply our results to the biological domain. Fortunately, the modification from the preceding

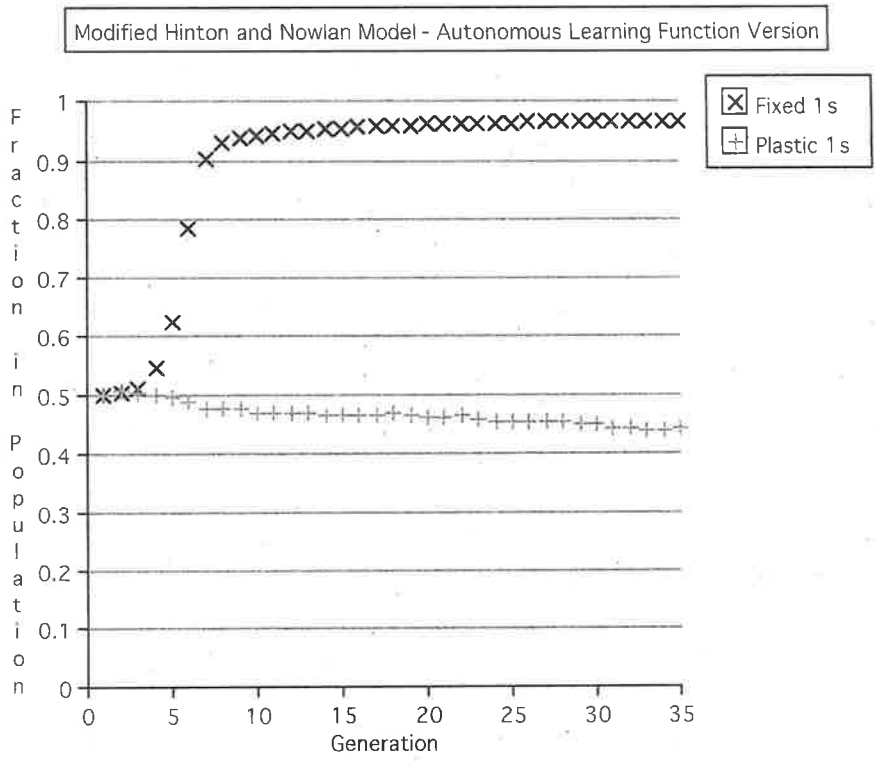

Figure 11.3
Again, the sharp rise in the population of fixed 1s is observed. Two interesting features differ from the Hinton and Nowlan model: first, the onset of the sharp rise in the population of fixed 1s occurs about 5 generations earlier than in the original model, and reaches higher values earlier. Second, the fraction of plastic 1s drops off at a slower rate than in the original model. These two observations, taken together, indicate that there is a higher evolutionary premium for having "good" underlying fixed behavior in this model, but less of a penalty for retaining plasticity.

step allows us to sidestep this difficulty. We now make the dependence of the fixed/plastic genes to be probabilistic, rather than deterministic (in fashion akin to Turney 1997). That is, loci in the plastic gene will not determine uniquely the resulting behavior, but will rather interact probabilistically with the underlying fixed gene according to the following table:

If locus in plastic gene is off -> use the unmodified original fixed gene when determining fitness.

If locus in plastic gene is on -> 50 percent of the time "flip a coin," 50 percent of the time use the unmodified original fixed gene when determining fitness.

It is rewarding to report that the results of the basic model survive this modification, as can be seen in the sample run shown in figure 11.3.

8 The Teleological Problem

Even after the preceding modifications have been incorporated into it, it can be argued that the Hinton and Nowlan model is unrealistic in a different fashion. To a very large extent, one of the key aspects of the problem at hand is the length of the string to be determined. By hard coding the length of their original gene at 20 loci—identical to the length of the string to be found—Hinton and Nowlan essentially predefined this aspect of the question in a fashion that can only be called teleological. Putting the question in biological terms, how would the organism conveniently know beforehand that the length of its plastic and fixed genes should be "just right" for a question it has not even encountered yet? Fortunately, the preceding modifications can be adapted to deal with this question. We now allow our dual genes to now vary independently in length such that they can both be too short (in which case the organism will always fail to find an answer), just right, or too long (in which case the organism may find an answer, but it will carry "junk" unused baggage in its genome). Yet again, the basic model proves fundamentally stable. While there are limits to how much this particular variation may be pushed—fixed genes of length one are clearly outside the bound of anything that could reasonably work—it is nonetheless gratifying to report that the basic model survives even this. As can be seen in the sample run in figure 11.4, a 20 percent variation still allows for reasonable goal-achievement by our digital critters.

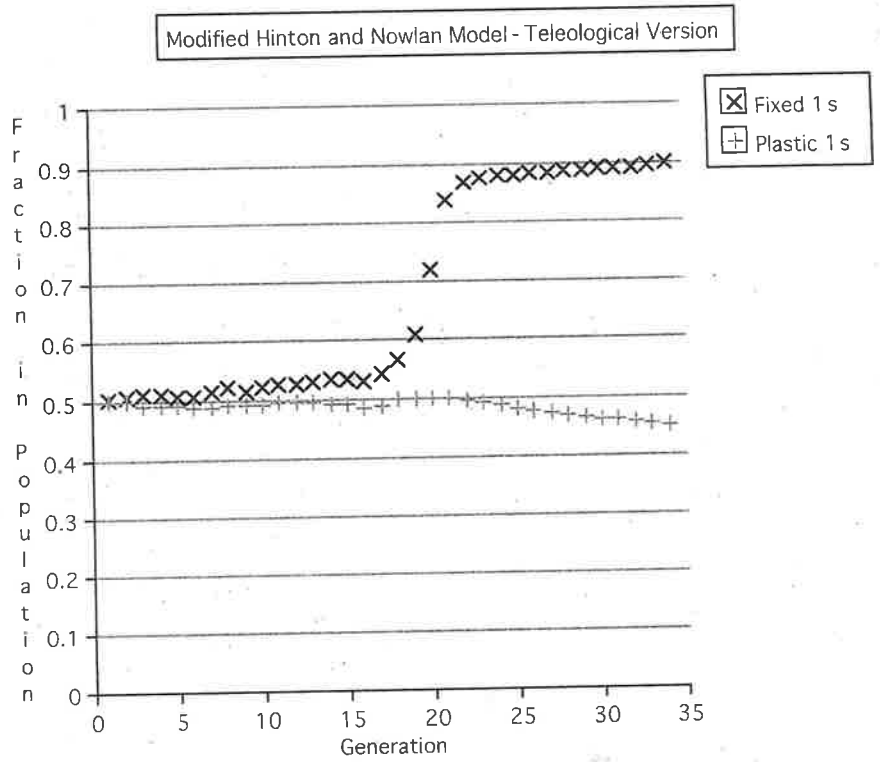

Figure 11.4
The familiar sharp rise is still present here, albeit with an onset delayed by about 10 generations, and a slower drop-off in the plasticity fraction. This latter point may be interpreted as indicating an even more significant role for plasticity in providing a good search path in the context of the increased difficulty of this problem.

9 Issues Still Ahead

From my expansion of the original Hinton and Nowlan model, it can be seen that many basic questions can be successfully answered without doing violence to the results or simplicity of the model. Simple extensions and uses of the model in computational questions are therefore both safe and desirable. However, in terms of what the results mean for biological systems, I believe a fundamental question is yet to be answered. Thus far, the model has operated on a one-dimensional landscape, with strict homology between the plastic and fixed universes. Furthermore, the distance between genotype

and phenotype in the model is minimal. In a biological system, this twin situation is much less likely to arise. It may be indeed at the heart of Waddington's canalization phenomena (Waddington 1942), but probably not elsewhere. In particular, it would be desirable to have a multidimensional problem landscape that would allow for the fixed/plastic aspects to scan directions of exploration that are not strictly collinear, like the one shown in figure 11.5a. A few models have been proposed that attempt to deal with this, but unfortunately it is not clear that any of them has quite risen to the challenge. The work done, among others, by Ackley and Littman (1992), Parisi, Nolfi, and Cecconi (1992), Menczer and Belew (1996), and Nolfi, Elman, and Parisi (1994) sets its digital critters the task of surviving in a two-dimensional landscape, riddled with sources of food and poison. In each "animat" a set of plastic neurons dedicated to the predictive evaluation/determination of actions on the landscape based on a set of sensory inputs is coupled to a genetically determined set of parameters, describable as tendencies/goals (see fig. 11.5b). Reproductive fitness is not determined by a fixed fitness function, but rather is determined by the "energy level" of the animat and its survival capacity. These models generally show poor results when animats lack the capacity to learn plastically in their lifetime, far better results when animats can learn plastically, but lack any form of genetic evolution, and (sometimes) very modest improvements over this latter category when learning is coupled to evolution (see fig. 11.5c). From these results, coupled to the greater observed variability in the learning gene when the third case is compared to the second, many authors have designated this a "better" example of the Baldwin effect than the Hinton and Nowlan model (see, for instance, Mitchell 1996). However, as Harvey (1997) points out, there exist significant flaws in considering this as a Baldwin effect model: two separate tasks (plastic and fixed) are required; the plastic and fixed tasks are completely uncorrelated, and the plastic "good trick" is not assimilated into the genome; at-birth performance of the animat is not improved; too much learning causes the effect to disappear. Harvey claims a different effect, which he terms "Another New Factor (ANF)," is responsible for the observed results here, one which he attributes to the recovery of weight perturbations in a neural network via the learning of uncorrelated tasks. This latter element of Harvey's interpretation has recently been

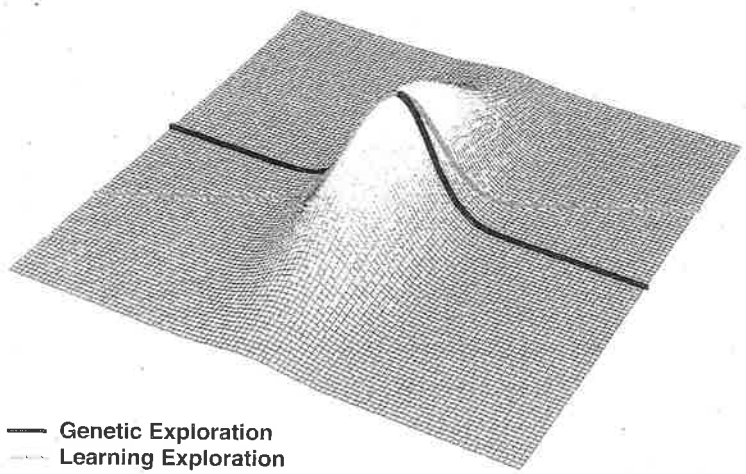

Figure 11.5a
The diagram illustrates a two-dimensional analogue of figure 11.1b above. The dark line indicates the more sharply spiked path available to genetic exploration, while the white line indicates the more softly curved path accessible to learning exploration. The problem to be solved is assumed to be essentially the same, but the exploratory approaches are partially different.

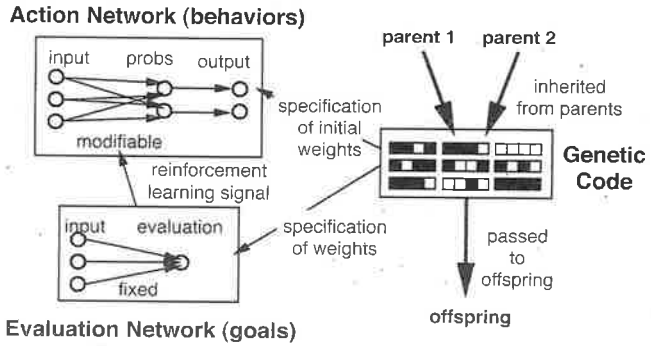

Figure 11.5b
A schematic diagram of the Ackley-Littman model. The "action network" is modifiable within the lifetime of the animat, and hence corresponds to the learning component, while the "evaluation network" is modifiable only genetically via reproduction. (From Ackley and Littman 1992.)

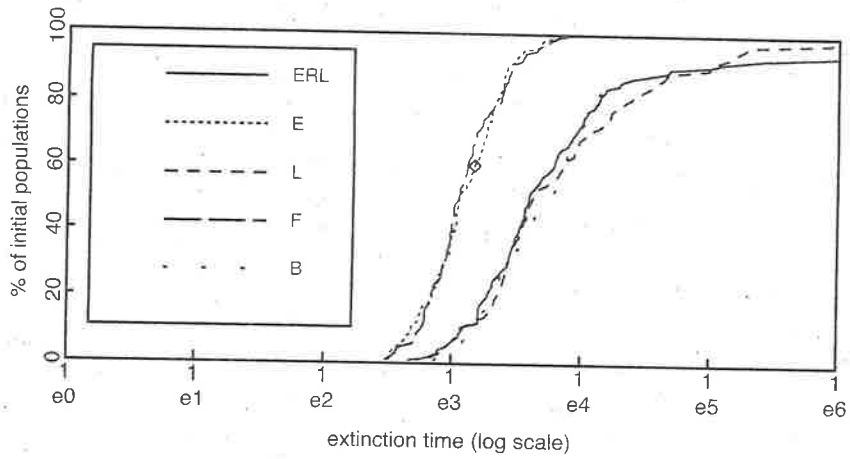

ERL = Evolution Plus Learning
E = Evolution Only
L = Learning Only
F = Fixed Random Weights
B = Random (Brownian) Movements

Figure 11.5c
Ackley and Littman's results. The sooner a population goes extinct, the less fit it is. Note that there is roughly an order of magnitude in difference between the evolution-only results and the learning-only and learning plus evolution results, but only very small differences between these latter two. (From Ackley and Littman 1992.)

challenged by Nolfi (2000), although his results in no way reestablish the observed effect as being Baldwinian in nature.

Thus, the question of a generalized computational Baldwin effect is still very much open. Further work in this field will require careful model definition, in order to avoid having the peculiarities of specific neural network setups (such as are encountered in the training of feedforward neural nets) swamp out any potential observation of the Baldwin effect. A paper by Giles Mayley (1997) has two important implications in terms of any future research in this area. First, any future model must clearly establish the role of the cost of learning within it as a source of evolutionary pressure for genetic assimilation. Second, the existence of correlation between learned/innate distances in phenotypic space and learned/innate distances in genotypic space must also be clearly parametrized, since it defines the path by

which assimilation might occur. It is clear that we have so far only scratched the surface of the understanding that the age of computation might bring to the Baldwin effect.

Acknowledgments

My thanks go to Bruce Weber and David Depew for inviting me to participate in the current volume. Also, my thanks go to Stacy D. Evans for her helpful suggestions and sharp eye when reading early drafts of this essay—any remaining awkward spots are, needless to say, my responsibility alone. Finally, I would like to acknowledge the generous support of my colleagues and Bennington College.

References

Ackley, D. and M. Littman (1992). Interaction between learning and evolution. In C. Langton (ed.), *Artificial Life II*. Reading, Mass.: Addison-Wesley.

Baldwin, J. M. (1896). A new factor in evolution. *American Naturalist* 30: 441–451, 536–553.

Belew, R. K. (1989). When both individuals and populations search. In J. D. Schaffer (ed.), *Proceedings of the Third International Conference on Genetic Algorithms*. San Mateo, Calif.: Morgan Kaufmann.

Dennett, D. C. (1995). *Darwin's Dangerous Idea*. New York: Simon and Schuster.

Griffiths, P. E. (2002). Beyond the Baldwin effect: James Mark Baldwin's social heredity, epigenetic inheritance, and niche-construction. In D. J. Depew and B. H. Weber (eds.), *Evolution and Learning*. Cambridge, Mass.: MIT Press.

Harvey, I. (1993). The puzzle of the persistent question marks: A case study of genetic drift. In S. Forrest (ed.), *Proceedings of the Fifth International Conference on Genetic Algorithms, ICGA-93*. San Mateo, Calif.: Morgan Kaufmann.

Harvey, I. (1997). Is there another new factor in evolution? *Evolutionary Computation* 4 (3): 313–329.

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

Mayley, G. (1997). Landscapes, learning costs, and genetic assimilation. *Evolutionary Computation* 4(3): 213–234.

Menczer F. and R. K. Belew. (1996). Latent energy environments. In R. K. Belew and M. Mitchell (eds.), *Adaptive Individuals in Evolving Populations: Models and Algorithms*, SFI Studies in the Sciences of Complexity vol. 23. Reading, Mass.: Addison-Wesley.

Mitchell, M. (1996). *An Introduction to Genetic Algorithms*. Cambridge, Mass.: MIT Press.

Nolfi, S. (2000). How learning and evolution interact: The case of a learning task which differs from the evolutionary task. *Adaptive Behavior* 7(2): 231–236.

Nolfi, S., J. L. Elman, and D. Parisi. (1994). Learning and evolution in neural networks. *Adaptive Behavior* 3(1): 5–28.

Parisi, D., N. Nolfi, and F. Cecconi. (1992). Learning, behavior, and evolution. In C. Langton (ed.), *Artificial Life*. Reading, Mass.: Addison-Wesley.

Turney, P. (1997). How to shift bias: Lessons from the Baldwin effect. *Evolutionary Computation* 4(3): 271–295.

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

12

The Role of Predator-Induced Polyphenism in the Evolution of Cognition: A Baldwinian Speculation

Scott F. Gilbert

1 Development and the Baldwin Effect

Developmental biology is not often called upon to relate mind and body. However, the Baldwin effect almost demands that it be considered. The Baldwin effect (Baldwin 1896) was formulated to explain how psychological characteristics that made an individual more fit in a particular environment could be fixed in the genome. One way of restating this (see Marcos 2000) is to say that the greater the ability of an individual to adapt to external conditions, the greater its fitness (i.e., the production of progeny). This ability, which was originally a physiological response to particular conditions, will eventually be inherited even if the original initiating conditions are no longer present. Whereas Baldwin believed that a single mutation could transfer the inducing signal from the environment to the genotype, Ivan Schmalhausen and Conrad Waddington found that the transfer of competence from an environmental inducer to an internal inducer could also arise through the cryptic variation already present in the population. Waddington (1953, 1956) called this transfer "genetic assimilation."

The first tenet of the Baldwin effect and genetic assimilation is phenotypic plasticity. This idea that environment can induce phenotypic variation is now very well established. Reaction norms, dietary polyphenisms, seasonal polyphenisms, and predator-induced polyphenisms have become more familiar to biologists, especially as life history strategies research has begun to enter developmental biology (Gilbert 1997, 2001; Schlicting and Pigliucci 1998; Tollrian and Harvell 1998). Moreover, environmentally regulated gene expression has now been demonstrated. Physical stress, for instance, has been known to effect bone density, both positively and negatively. It is