(In R. Brooks & P. Maes (eds.) 1994 *Artificial Life IV* Cambridge, MA: The MIT Press, 277-282)

Genes, Phenes and the Baldwin Effect: Learning and Evolution in a Simulated Population

by
Robert M. French and Adam Messinger
Willamette University,
Salem, OR 97301
(french@willamette.edu and amessing@willamette.edu)

Introduction

The view that evolution is *influenced by* acquired behaviors and traits is regarded by many as being uncomfortably close to the discredited Lamarckian contention that evolution *consists of* the inheritance of acquired behaviors and traits. It is perhaps for this reason more than any other that the evolutionary mechanism first proposed first by J. Mark Baldwin and Lloyd Morgan in 1896 [Baldwin, 1896; Morgan, 1896] is still veiled in controversy. This mechanism, known today simply as the Baldwin Effect, states that learned behavior and characteristics at the level of individuals can significantly affect evolution at the level of the species. Schull [1990] sums up the process as one in which "individual developmental responses will necessarily lead to directed and non-random evolutionary change." And while many evolutionary biologists accept the Baldwin Effect as a significant force in evolutionary change, the theory also has many detractors. For example, in a recent article, Piattelli-Palmarini [1990] writes, "One would have hoped that, in 1990, all talk of the Baldwin effect . . . would have been mercifully forgotten."

Parisi, Nolfi, and Cecconi [1990] give three further reasons that evolutionary biologists tend to dismiss the Baldwin Effect. The orthodoxy of evolutionary biologists, they claim, is strongly reductionist, "which implies that the causes and basic mechanisms of evolution are only to be found at the level of genetics." As a consequence, behavior and learning, both being highly holistic processes, have been largely ignored in attempting to understand evolutionary processes. Another reason for the lack of attention, according to these authors, is that evolutionary biologists feel "behavior and learning are the province not of biology but of psychology and ethology." And finally, until recently, there have been very few empirical studies of the Baldwin Effect in either real or simulated populations.

In this paper, clear evidence is presented that the Baldwin Effect can indeed significantly alter the course of Darwinian evolution at the level of the genotype. In other words, we show that plasticity at the phenotypic level can and does produce directed changes at the genotypic level. In addition, the amount of plasticity and the amount of benefit of the learned behavior are also demonstrated to be crucial to the size of the effect: either too little or too much and the effect disappears or is significantly reduced. Finally, we demonstrate certain conditions under which the Baldwin Effect is more powerful in sexually reproducing populations than in asexually reproducing ones. This research confirms and extends earlier experimental work done by Hinton & Nowlan [1987], Belew [1990], Parisi, Nolfi, & Cecconi [1992], and Fontanari & Meir [1990], among others.

We empirically tested various aspects of the Baldwin Effect on populations similar to those described in Bedau [1993] and Holland [1993]. We created a large population of agents with varying metabolic, feeding, locomotive and reproductive characteristics and allowed them to evolve in a world in which the amount and distribution of food varied over time. We then examined the effects of phenotypic plasticity on the evolution of the genotype. Three different areas of plasticity were considered — namely, more efficient metabolism, movement or reproduction.

Genes and Phenes: an overview of the simulation

Our simulated world consists of a population of agents whose genetic material consists of a fixed-length bitstring. Food is regularly, but randomly, added to this world in discrete piles of uniform depth, much as if someone were regularly throwing handfuls of food into the world at random locations. The size of the piles of food and the frequency with which they are added to the world can vary. Food is required for energy. "Energy taxes" are levied for movement, reproduction and existence.

We stipulate that one particular genotype (i.e., one particular pattern of bits) — the "Good Gene" or GG, for short — will result in a fitness-enhancing behavior or trait — the "Good Phene" or GP — at the phenotypic level. We assume that the closer an agent is to the Good Gene (in terms of some metric on gene space), the easier it will be for it to learn the Good Phene that will enhance survival and reproductive possibilities. Phenes can either be learned or be the direct product of the possession of a Good Gene. A “natural” Good Phene (one which is the direct result of having the Good Gene) and a “learned” Good Phene are indistinguishable at the level of the phenotype. Possessing the Good Gene necessarily implies that its associated Good Phene will be expressed.

In our simulations, the Good Phene can be one of three things: improved locomotive, metabolic or reproductive efficiency. In each case, these were implemented as a reduction in the default taxes on each of these activities. In contrast to Hinton & Nowlan [1987], our simulation incorporates no explicit "fitness function" [Holland, 1975] to make an *a priori* determination of how good each genotype is. In other models, an explicit fitness function is generally used to determine the future reproductive success of each genotype.

In our model, however, the fitness of a particular genotype is determined implicitly by how well it survives in the population over time. The success of a particular genotype, then, is a measure of the percentage of agents in the population that have it. We are particularly interested in the evolution of the percentage of organisms in the population with the Good Gene. Specifically, we compare how often and under what circumstances the Good Gene appears in a population with learning compared to a population without learning.

Learning in the simulated world

Two factors will determine the extent to which the Good Gene will eventually proliferate in the population, namely:

- the difficulty across the population of learning the Good Phene (phenotypic plasticity);

- the amount of benefit conferred on an agent by having learned the Good Phene.

First, let us consider phenotypic plasticity. We assume that each agent is born with a genotype that is a certain hamming distance from the Good Gene. Based on that distance, the probability that an agent will *learn* the Good Phene associated with the Good Gene is determined by a phenotypic plasticity curve like the one shown in Figure 1. The x-axis represents the agent's normalized hamming distance from the Good Gene (i.e., the number of bits differing from the Good Phene divided by the total number of bits) and the y-axis shows the probability of learning the Good Phene. Therefore if the agent is genetically close to the Good Gene, it will stand a good chance of learning the corresponding Good Phene. As the normalized hamming distance from the Good Gene increases, the probability of learning the Good Phene falls off. The shape of the curve depends on the Good Phene. Phenes range from being very hard (or impossible) to learn to comparatively easy to learn. For example, consider a Good Phene that is relatively hard to learn across the entire population (ρ is small; see Figure 4). When the genotype of a particular organism differs from the underlying Good Gene by half of its bits , it will very little chance of learning the Good Phene. Whether or not the agent will actually learn the Good Phene is done stochastically based on the probability taken from its phenotypic plasticity curve.

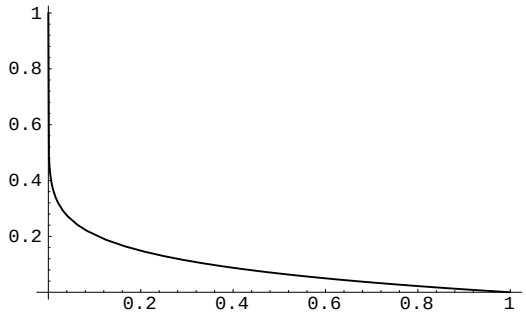

Figure 1: A phenotypic plasticity curve ($\rho = 0.1$) showing the probability of learning a Good Phene based on hamming distance from the Good Gene (if 20% of an agent's bits differ from the Good Gene, it will have a 15% chance of learning the Good Phene).

Some Good Phenes are clearly harder to learn than others. For example, the trait of possessing blue eyes has no phenotypic plasticity. Individuals who do not possess the precise genotype that codes for blue eyes will never be able to learn this characteristic.

Other traits and behaviors are easier to learn. These are represented by the family of phenotypic plasticity curves shown in Figure 2. Each curve represents a different amount of native plasticity for a particular Good Phene. In a population where the plasticity parameter ρ is very low, an agent has to have been born very close to the Good Gene in order to have a chance of learning the Good Phene. In essence, when ρ is very low, the Good Phene is almost never learned.

However, as ρ increases, it becomes possible for some members of the population to actually learn the Good Phene. For example, consider the ability to hum Middle C. Some people can do it perfectly with no learning at all. These people possess the Good Phene — in this case, perfect pitch — from birth. Presumably, there is something in their

genes that allows them to perform this task flawlessly. People who are genotypically a bit farther from the Good Gene may still be able to learn to hum Middle C, eventually, but will have to adopt a variety of clever strategies to do so. Most people, though, could try forever and still never be able to hum Middle C correctly. They are genetically too far from the Good Gene. This is a case of low, but non-zero, phenotypic plasticity.

Now, consider the ability to memorize the written word. Again, some people are naturally very good at this particular task. After reading a poem a few times, they can recite it without errors. For others, the task is much more difficult, requiring many hours and clever mnemonic strategies, but given enough time, they will eventually memorize succeed in memorizing the poem. A relatively small, but significant number of individuals in the population would never be able, try as they might, to learn the poem. This is an example of average phenotypic plasticity.

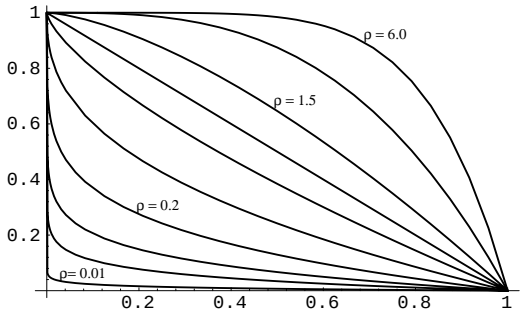

Figure 2: Curves showing varying degrees of phenotypic plasticity: $y=1-x^{\rho}$, where ρ is phenotypic plasticity associated with a particular Good Phene

Finally, certain behaviors are very plastic. For example, if the Good Phene consisted of "writing with your left hand" or “winking”, virtually everyone in the population could master this. The phenotypic plasticity, ρ, for this trait is very high.

A final comment about phenotypic plasticity. For the purpose of the simulations in this paper, the phenotypic plasticity of a particular Good Phene remains constant over time, but in real animals phenotypic plasticity seems to decrease with age. Consider, for example, humans’ ability to learn foreign languages without an accent. Most people are born with an extraordinary phenotypic plasticity in this area (i.e., with ρ very high) but, around the time of adolescence, this ability gradually seems to disappear. Once adolescence has passed, it is almost impossible to learn a new language without an accent [Grosjean, 1982, p. 300].

The second factor that plays a major role in the Baldwin Effect is the benefit to the organism of the Good Phene. If the advantage conferred by the Good Phene is too low, then one would expect little movement of the population towards an increased incidence of the corresponding Good Gene. On the other hand, if the advantage conferred is extremely high, then any organisms, however few, who manage to learn it will beat out all competitors in the survival game. The Good Gene, however rare to begin with, will propagate and will soon come to dominate the population. Consider an extreme example that illustrates this point: if the (Very) Good Phene were, say, "Energy taxes drop to zero", then, even with no learning, once this Good Gene succeeded in entering the population by a lucky mutation, it would stay in the population forever. This

would happen because the agent who happened to have been born with it would be able to survive and reproduce even in the complete absence of food.

The real evolutionary value of the Baldwin Effect is that it gives good — but not extraordinarily good — genes an improved chance of remaining in the population. Extremely good genes will, in general, stay in a population. The benefit they confer is simply so overwhelming that they do not get washed out. But the notion that an organism could stumble upon a single Extraordinary Gene that would produce an Extraordinary Phene smacks of saltationism [Mayr, 1988], whereby major, highly beneficial evolutionary adaptations — such as the development of wings from forelegs — suddenly appear in the population. If one accepts Darwinian gradualism, the value of the Baldwin Effect as a mechanism for protecting moderately good phenes becomes apparent.

Experimental results

In the following experiments, we will first demonstrate that, even with no explicit fitness function, the Baldwin Effect can significantly influence evolving populations. We compare "Baldwin" populations (i.e., ones in which learning takes places) to non-Baldwin populations by means of three different measures — namely:

- total population;
- number of agents in the evolving population who have the Good Phene (GP);
- number of agents in the evolving population who have the Good Gene (GG) that produced the original Good Phene.

The last measure — the genetic shift towards the Good Gene — is the most important for establishing that learning can have a significant effect on the genetic composition of the population. It turns out that the Baldwin Effect is most pronounced for phenes whose plasticity is neither too great nor too small.

In a later section of this paper, we will consider Good Phenes that confer more or less benefit to an organism. For the moment, though, we will hold this benefit constant and examine the effect of learning versus non-learning on the genetic makeup of the population.

Simulation 1: The Baldwin Effect

In the first experiment, using asexually reproducing agents, we chose a Good Gene whose naturally corresponding Good Phene decreased the organism's "energy tax" on movement. Any agent born with this particular Good Gene had phenotypic characteristics that allowed it to move around in its environment more efficiently than those that did not have it. We chose a phenotypic plasticity for this Good Phene of ρ of 0.1. In other words, this phene is neither particularly easy nor particularly difficult to learn (see Figure 1).

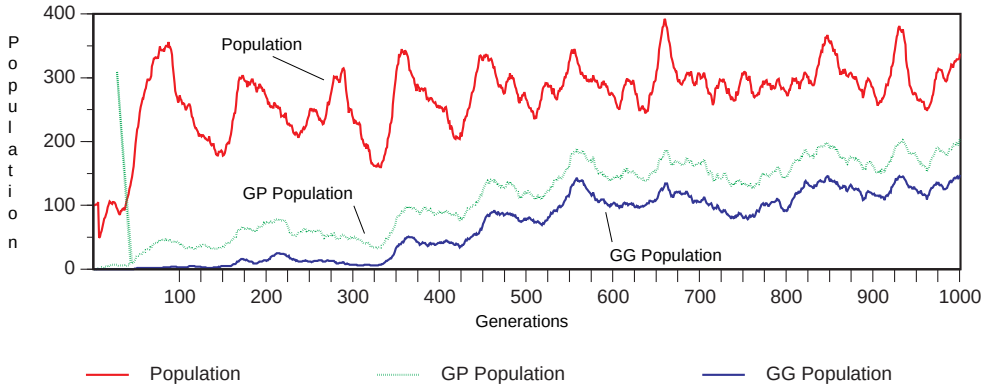

Figure 3: The effect of learning on the evolution of the genotype

Recall that the probability that a given agent will learn the Good Phene is determined by how far the agent is from the Good Gene (Figure 1). The graph in Figure 3 shows the evolution of a population in which learning of the Good Phene is occurring. The top line shows the total population. The middle line indicates the number of agents in the population who have the successfully learned the Good Phene. In other words, these agents possess the Good Phene either by dint of learning it or because they were born with the Good Gene. Finally, the lowest line represents the number of agents in the population who actually possess the Good Gene.

In this population it is apparent that, by as little as 500 generations, nearly half of the population has learned the Good Phene. This causes an overall population increase compared to an equivalent population in which none of the members have the Good Phene (Figure 4). There is a consistent increase in total population directly related to the benefit the Good Phene confers.

The number of agents possessing the Good Gene gradually increases but will always remain significantly below the number of agents possessing the Good Phene because all that is required for improved survival is the Good *Phene* — however it was come by — and not necessarily the Good Gene. Individuals who did not obtain the Good Phene genetically and are not able to acquire it through learning will gradually be eliminated from the population by their fitter GP competitors.

Eventually virtually *all organisms* in the population will possess the Good Phene. At this point, the genotype will cease to evolve significantly. Darwinian selective pressure ceases because all of the organisms are now equally fit, at least with respect to the advantages conferred by the Good Phene. Nature cares only that the organism has a Good Phene and is not concerned with where it came from. In this simulation, by 5,000 generations, nearly all members of the population have acquired the Good Phene, and approximately 70% possess the Good Gene. Even after an additional 5,000 generations, the number of Good Genes in the population had not changed significantly. Once the entire population has become a GP-population, then the difference between the total population and the GG-population is directly proportional to phenotypic plasticity.

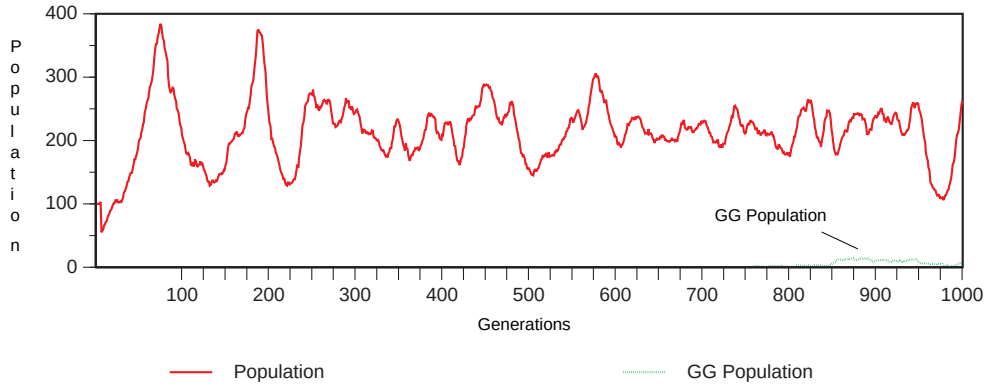

Figure 4: The same population with no phenotypic learning
(notice the lack of change of the genotype)

Finally, in an identical non-learning population (Figure 4), the number of agents possessing the Good Gene remains insignificant. Clearly, the ability to learn the Good Phene has a significant influence on the evolution of the genotype, thus demonstrating the Baldwin Effect.

Simulation 2: When phenotypic plasticity is too high or too low

We have already seen in Figure 4 that, without phenotypic plasticity (i.e., when no learning is possible), the genotype of the population does not evolve towards the Good Gene. As phenotypic plasticity increases, there is a corresponding increase in genotypic movement towards GG. But as phenotypic plasticity grows, the number of agents with the Good Gene actually begins to decrease and, for very high plasticity, there is virtually no trend at all towards a GG population (Figure 5). As in the first simulation, once the

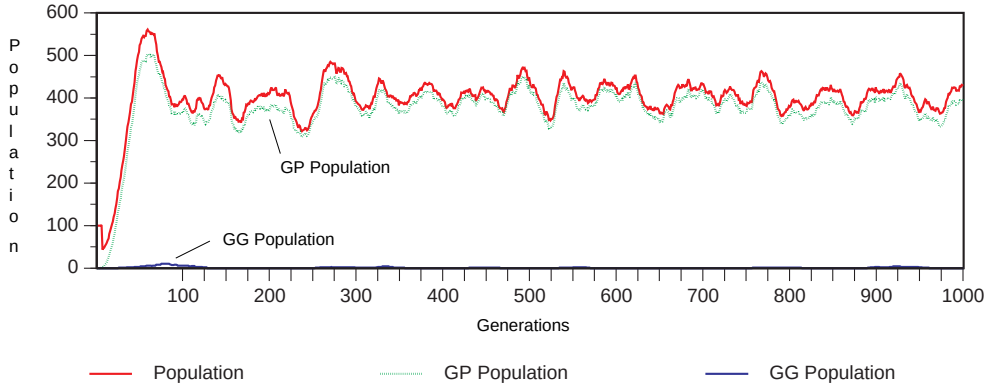

Figure 5: The Baldwin Effect ceases for high phenotypic plasticity
(Note that the GG population remains insignificant)

entire population possesses the Good Phene, the percentage of GG agents in the population stops increasing. If phenotypic plasticity is too high, the entire population will quickly acquire the Good Phene, effectively bringing to an end any further reason for the

genotypic profile of the population to change. The topographic plot in Figure 6 shows how different phenotypic plasticities affect the percentage of agents in the population with the Good Gene. From this graph it can been seen that the Baldwin Effect disappears for high and low phenotypic plasticities.

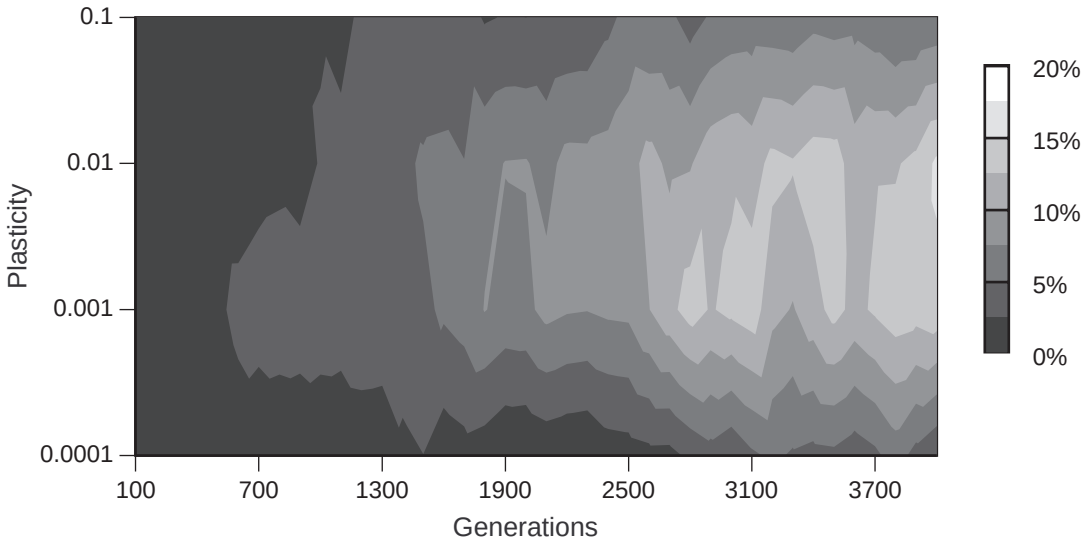

Figure 6: The effects of differing phenotypic plasticities on the evolution of the Good Gene in the population (lighter areas indicate higher GG percentages)

Simulation 3: The Baldwin Effect and the Quality of the Good Phene

The extent to which the Baldwin Effect is also operative is affected by the selective advantage conferred by the Good Phene. If the Good Phene confers little or no selective advantage, then the question of its acquisition is of little importance. GP organisms will not survive better than non-GP organisms and, consequently, there will be little or no increase of GG organisms in the population.

On the other hand, if the Good Phene is extremely good, thereby conferring a very large selective advantage, those individuals who somehow manage to acquire it will have a much better chance of surviving. These organisms will almost invariably out-compete their non-GP rivals and reproduce more successfully (this is especially true in asexual populations). Even if phenotypic plasticity is extremely low (or even zero), the progeny of a single organism that happened to stumble onto the (Very) Good Gene through random mutation would stand a good chance of going on to dominate the population.

In Figure 7 below the "benefit factor" of the Good Phene is set very high and phenotypic plasticity is set very low. This figure shows how the total number of individuals with the Good Gene evolves depending on the phenotypic plasticity and quality of the Good Phene.

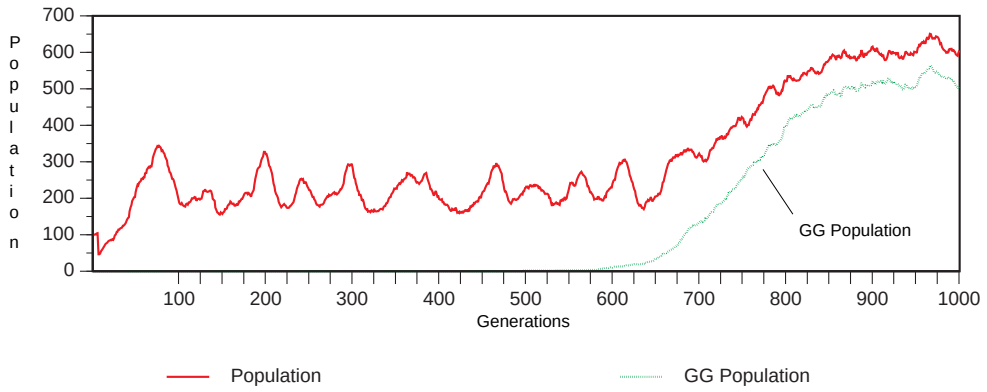

Figure 7
Effects on genotypic evolution of a *very* Good Phene

Sexual populations

So far, all of our simulations have used asexual populations. Menczer & Parisi [1992] have demonstrated adaptive advantages in sexual over asexual reproduction. The advantage of genetic crossover has been pointed out on many occasions by Holland [1975, 1986], Belew [1990], and others. In addition, parthenogenic insects such as aphids and certain species of flies actually use both forms of reproduction, depending on environmental conditions. As long as there is abundant food, the organisms reproduce asexually — which is far faster, thus producing more abundant offspring [Gould, 1977]. As the food supply dwindles, the insects switch to sexual reproduction. Sexual reproduction, while possibly not as reproductively efficient in terms of sheer numbers, ensures broader genetic diversity, which means that the population as a whole will be less likely to succumb to a single highly detrimental environmental change.

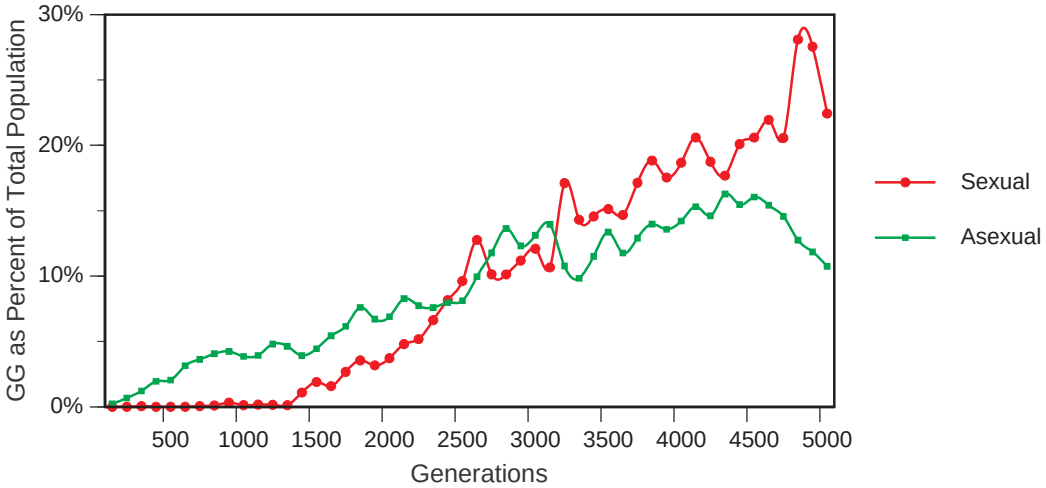

Figure 8: A comparison of evolving GG percentages in asexual and sexual populations

We also ran the simulation with sexually reproducing agents. These sexual populations ultimately always had higher concentrations of Good Genes than asexual

ones. For equivalent total population sizes, the GG-population in the sexually reproducing population increased considerably faster than the GG-population in a asexual population of the same size (Figure 8). This result was not entirely unexpected in light of predictions about the behavior of genetic algorithms [Holland, 1975].

Justification of the learning procedure

In Hinton & Nowlan's [1987] classic simulation of the Baldwin Effect, the genotype for each organism in the population consists of a bitstring of twenty 0's, 1's and ?'s ("undetermined" alleles). A "phenotypic copy" (instantiated as the weights of a network in their model) of this bitstring is made on which "phenotypic learning" is done. Learning is done by randomly assigning 1's or 0's to all of the question marks approximately 1000 times, each time checking to see if an all-1's phenotype has been hit upon. If the all-1's phenotype is discovered during the 1000 learning trials, the original genotype is then assigned a high explicit "fitness" making it much more likely to be chosen as a parent for mating and crossover in the next time cycle.

Our model differs in certain significant respects from Hinton & Nowlan's model. Phenotypic learning, in particular, is done differently in our simulation. We start from the assumption that certain Phenes, as we have called them, are harder to learn than others *in general*, i.e., for all organisms in the population. Across the human population, for example, it is harder to learn perfect pitch than it is to learn to wink. In our model, differing degrees of phenotypic plasticity express those differences. With respect to learning a *particular* Phene, say, winking, certain *individuals* have a considerably harder time than others. This is a function of the organism's genetic hamming distance from the phene in question.

In Hinton and Nowlan's model, phenotypic plasticity is fixed; there is no mechanism for varying the difficulty of learning according to the Good Phene under consideration. There is only one Good Phene and the difficulty of learning it depends solely on how far an organism's phenotype is from it. If there were a second Good Phene, there seems to be no simple mechanism in their model for making it easier (or harder) to learn than the first one.

In addition, in Hinton & Nowlan, an explicit function is used to measure the fitness for reproduction of an individual. This function is implicit in our simulation, arising from the interaction of the agents in the world. While the Good Phene does confer an efficiency advantage in our model, it does not *expressly* add to the fitness of the individual. In fact, in certain simulations when the reproduction tax was too low, the agents rapidly overpopulated the world and actually became extinct due to starvation.

By varying phenotypic plasticity and benefit to the organism we are able to more completely characterize the Baldwin Effect.

A final remark

In regards to our strict separation of the *ability* to learn the Good Phene from the underlying genetics, the following objection might be raised: Isn't any ability to modify one's phenotypical characteristics through learning *itself* determined by the underlying genotype? And, if so, doesn't the exclusion of "learning" genes from the underlying genotype constitute a sort of legerdemain contrived to allow us to pull the Baldwin Effect

out of our simulation hat? In other words, would the Baldwin Effect go away if a big enough genotype were considered, one that included the genes that coded for motivation, for learning, for general intelligence, for the amount of fast-twitch muscle fiber, etc? The answer is, of course, yes, but a rather hollow yes. This genetic reductionism, wherein *everything* is reduced to the level of genes, is certainly correct but lacking in high-level explanatory power, rather like describing King John's signature on the Magna Carta in terms of molecules of ink on molecules of paper.

The fact is that we *could have* added "learning genes" to our base genotype and the effects demonstrated in this paper would no doubt have remained unchanged. This is because, even though an organism's ability to learn is, admittedly, derived from its genes (like everything else), general-purpose learning abilities are assumed to be the product of a very large, highly robust set of genes. This set of "learning genes", while obviously not identical in each individual of the population, is assumed to be similar across the entire population and determines how difficult it will be, on average, to learn a particular Phene. These learning genes determine the shape of the *population-wide* phenotypic plasticity curve for any Good Phene.

In our attempt to better isolate the Baldwin Effect, we elected not to include any specific set of genes in the base genotypes of the agents that coded for learning. Isolating the genes that coded for one specific behavior allowed us to give an improved characterization of certain important aspects of the Baldwin Effect.

Conclusion

We have empirically demonstrated a pronounced Baldwin Effect in a simulated population of naturally evolving agents. In other words, the ability to learn at the phenotypic level had a significant effect on the genotypic evolution of the population. In addition, certain factors, in particular, the amount of phenotypic plasticity and the benefit associated with the learned phenotypic behavior or characteristic, have a significant influence on the amount of genotypic change produced. It also appears that excessively high or low levels of phenotypic plasticity have the same effect — namely, they are significantly less successful at promoting genotypic evolution than moderate levels of plasticity. Finally, we have shown more pronounced Baldwin Effect in sexually reproducing organisms than in otherwise comparable asexually reproducing organisms.

Acknowledgments

The authors would like to thank Mark Bedau and Melanie Mitchell for their suggestions and assistance with this paper.

Appendix: Details of the Simulation

Bugs!, the simulation described in this paper was written in Objective C on a NeXT computer. The code is available through the Willamette University cognitive science anonymous ftp. (ftp to willamette.edu then cd to pub/cogsci and get baldwin.bugs.Z)

Bugs! is a discrete time simulation of free-acting agents (bugs) in a two-dimensional matrix arranged to form a torus. Each cell in the 100x100 matrix contains at most one

agent. The world is randomly seeded with agents having randomized genotypes and with food piles of unifrom size. At each timestep more food piles are randomly placed in the world.

These agents are similar to the "tropic bugs" described in [Packard, 1989]. They contain an internal food counter which must remain positive in order for the agent to remain alive. The agents are taxed for movement, metabolism and reproduction by decrementing their internal food counter. All bugs can sense the amount of food in the four cells immediately surrounding them and to move up this gradient. At each timestep, after checking their internal food, all of the agents do the following, in order.

Metabolize: The "existence tax" is subtracted from the internal food count.
Eat: If the cell they occupy contains food, they eat up to a prespecified amount.
Reproduce: A variable is set that determines whether reproduction will be asexual, sexual, or both. Reproduction is only permitted only if an agent's internal food counter is above a pre-specified reproduction threshold.
Move: Agents move in the direction with the greatest amount of food. If all cells have the same amount of food, the agent will move in the same direction as last turn. Agents may move only move one space per time step.

An agent's genetic code consists of a fixed-length gene whose alleles are either 1's or 0's. As in Hinton & Nowlan [1987], there is a single bit pattern that has been determined to be better than the others — we call this pattern the Good Gene. Any individual with this gene will benefit from certain advantages phenotypic traits or behaviors — the Good Phene. Agents who do not have the Good Phene as a direct result of having the Good Gene have, nonetheless, a certain probability of learning it, depending on how far their genotype is from the Good Gene. The calculation of this probability is performed at the time of birth. The genotype of a particular organism does not change over its lifetime.

Reproduction:
Asexual: After checking that its internal food exceeds the reproduction threshold, the agent spawns a new agent containing a copy of its own genetic material. This copy is subject to mutation, the probability of which is set by a mutation rate parameter that specifies the probability that a given bit will be copied incorrectly. A pre-specified amount of food, corresponding to the reproduction tax, is taken from the internal food count of the parent and is given to the child.
Sexual: In addition to checking its internal food count, the agent checks its 4 space surround to see if another agent is present to mate with. If this is the case the agent "mates" with the new agent and a new agent is "born" with a combination of the two agents genotypes. The child's gene is the produced by crossing over the two parent genes at a point randomly chosen along the length of the gene. Because a small amount of mutation is required to prevent the gene pool from possible stagnation [Holland, 1975], a mutation parameter is included. The reproduction tax is only levied against the current individual and not its mate. As in asexual reproduction, the amount of this tax is given to the new agent as internal food.

Learning: Whether or not a given agent learns an adaptive behavior or trait depends on

the phenotypic plasticity curve and the agent's genotypic hamming distance from the so-called Good Gene. When a bug is born, its hamming distance from the Good Gene is calculated and this distance establishes the probability that the bug will phenotypically learn the Good Phene during its lifetime (Figures 1 and 2). This probability is used at the time of birth in a stochastic dermination of whether or not the agent learns the Good Phene.

If an agent acquires the Good Phene, either through birth or learning, then one of the three taxes (i.e., movement, reproduction, or metabolism) is reduced by dividing it by a factor called the GPBenefit factor. Currently only the rate of taxation is controlled by this mechanism, however method of reproduction has been considered as a possible addition.

Initial Settings: A number of parameters must be set before the simulation will work. Below is a list of parameter values for all of the simulations in the experimental section of this paper.

GeneLength:	10,12	StartingPopulation:	100
MetaTax (metabolism tax):	5	MoveTax (movement tax):	15
RepoTax (reproduction tax):	150	RepoThreshold (reproduction threshold):	200
StartingInternalFood:	100	Mouthful (max. bitesize per timestep):	100
StartingNumberFoodPiles:	30	NumberFoodPilesTurn (piles added per timestep):	4
FoodPileDepth (food per pile):	20		
MutationRate:	50		
AsexualReproduction:	0-1 (Off/On)	SexualReproduction:	0-1 (Off/On)
SexualMutation (sex. reprod. mutation rate):	0 (Off)		
GP Plasticity (phen. plasticity):	0-10		
GPBenefit (amount of tax decrease):	2,15		
GPType (type of Good Phene):	1,2, or 3 (Move. tax reduction)		

References

Ackley, D. H. and Littman, M. L. Interactions between learning and evolution. In C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen (eds.) *Artificial Life II*, pp. 487-507, Reading, MA, 1992, Addison-Wesley.

Baldwin, J. M. A new factor in evolution. *American Naturalist*, 30:441-451, 1896.

Bedau, M. A. and Packard, N. H. Measurement of evolutionary activity, teleology, and life. In C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen (eds.) *Artificial Life II*, pp. 487-507, Reading, MA, 1992, Addison-Wesley.

Belew, R. K. Evolution, learning, and culture: Computational metaphors for adaptive algorithms. *Complex Systems*, 4:11-49, 1990.

Dennett, D. C., *Consciousness Explained.* Boston, MA, 1991, Little, Brown & Company,

pp. 184-193

Fontanari, J. F. and Meir, R. The effect of learning on the evolution of asexual populations. *Complex Systems*, 4:401-414, 1990.

Gould, S. J. *Ever Since Darwin*. New York, NY, 1977, pp. 91-96.

Grosjean, Francois. *Life With Two Languages.* Cambridge, MA, 1982, Harvard University Press

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

Holland, J. H. *Adaptation in natural and artificial systems.* Cambridge, MA: MIT Press. Second edition (First edition: 1975)

Holland, J. H. Escaping brittleness: The possibilities of general-purpose learning algorithms applied to parallel rule-based systems. In R. S. Michalski, J. G. Carbonell, and T. M. Mitchell (eds.) *Machine Learning II*, pp. 593-623, San Mateo, CA, 1986, Morgan Kaufmann.

Holland, J. H. Echoing emergence: Objectives, rough definitions, and speculations for Echo-class models. Technical Report 93-04-023, Santa Fe Institute, 1993

Mayr, E. *Towards a New Philosophy of Biology*. Havard University Press, Cambridge, MA, 1988.

Morgan, L. On modi variation. variation. *Science* 4:733-40 (1896).

Menczer, F., and Parisi, D. A model for the emergence of sex in evolving networks: Adaptive advantage or random drift? In F. J. Varela and P. Bourgine (eds.) *Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*, Cambridge, MA, 1992, MIT Press/Bradford Books.S

Packard, N. H. Intrinsic adaptation in a simple model for evolution. In C. G. Langton (ed.) *Artificial Life*, pp. 141-155, Reading, MA, 1989, Addison-Wesley.

Parisi, D., Nolfi, S., and Cecconi, F. Learning, behavior, and evolution. In *Proceedings of the First European Conference on Artificial Life*, Cambridge, MA, 1992, MIT Press/Bradford Books.

Piattelli-Palmarini, M. Which came first, the egg-problem or the hen-solution? *Behavior and Brain Sciences* 13:1 pp. 84-86, 1990

Schull, J. Are species intelligent? *Behavior and Brain Sciences* 13:1 pp. 63-75, 1990.