
Introduction to the Special Issue

Evolution, Learning, and Instinct: 100 Years of the Baldwin Effect

Peter Turney

Institute for Information Technology
National Research Council Canada
Ottawa, Ontario, Canada K1A 0R6
peter@ai.iit.nrc.ca

Darrell Whitley

Department of Computer Science
Colorado State University
Fort Collins, CO 80523
whitley@cs.colostate.edu

Russell W. Anderson

Research Scientist
Smith-Kettlewell Eye Research Institute
2232 Webster Street
San Francisco, CA 94115
rwa@milo.berkeley.edu

At the turn of the century, it was unclear whether Darwin's theory or Lamarck's better explained evolution. Lamarck believed in direct inheritance of characteristics acquired by individuals during their lifetime. Darwin proposed that natural selection coupled with diversity could largely explain evolution. Darwin himself believed that Lamarckian evolution might play a small role in life, but most Darwinians rejected Lamarckism. One potentially verifiable difference between the two theories was that Darwinians were committed to gradualism (evolution in tiny, incremental steps), whereas Lamarckians expected occasional rapid change. Lamarckians cited the gaps in the fossil record (which are now associated with punctuated equilibria) as supporting evidence.

Lamarckism was a viable theory until August Weismann's (1893) work was widely accepted. Weismann argued that higher organisms have two types of cells—germ cells that pass genetic information to offspring and somatic cells that have no direct role in reproduction. He argued that there is no way for information acquired by somatic cells to be transmitted to germ cells.

In the context of this debate, James Mark Baldwin (1896) proposed “a new factor in evolution,” whereby acquired characteristics could be indirectly inherited. Morgan (1896) and Osborn (1896) independently proposed similar ideas. The “new factor” was phenotypic plasticity: the ability of an organism to adapt to its environment during its lifetime. The ability to learn is the most obvious example of phenotypic plasticity, but other examples are the ability to tan with exposure to sun, to form a callus with exposure to abrasion, or to increase muscle strength with exercise. Baldwin (1896) pointed out that, among other things, the new factor could explain punctuated equilibria.

The Baldwin effect works in two steps. First, phenotypic plasticity allows an individual to adapt to a partially successful mutation, which might otherwise be useless to the individual. If

this mutation increases inclusive fitness, it will tend to proliferate in the population. However, phenotypic plasticity is typically costly for an individual. For example, learning requires energy and time, and it sometimes involves dangerous mistakes. Therefore there is a second step: given sufficient time, evolution may find a rigid mechanism that can replace the plastic mechanism. Thus a behavior that was once learned (the first step) may eventually become instinctive (the second step). On the surface, this looks the same as Lamarckian evolution, but there is no direct alteration of the genotype, based on the experience of the phenotype. This effect is similar to Waddington's (1942) "canalization."

The Baldwin effect came to the attention of computer scientists with the work of Hinton and Nowlan (1987). The Baldwin effect may arise in evolutionary computation when a genetic algorithm is used to evolve a population of individuals that also employ a local search algorithm. Local search is the computational analog of phenotypic plasticity in biological evolution. In computational terms, in the first step of the Baldwin effect, local search smooths the fitness landscape, which can facilitate evolutionary search. In the second step, as more optimal genotypes arise in the population, there is selective pressure for reduction in local search, driven by the intrinsic costs associated with the search.

The cost of local search may be expressed in various ways. For example, learning may only probabilistically increase fitness; even if what is learned increases fitness with high probability, there may still be a selective advantage for evolving an instinctual response. This assumes, of course, that the instinctual trait more reliably increases individual fitness than learned behaviors and that the environment favoring a particular behavior is stable. In a rapidly changing environment, learning can have obvious advantages over instinctual responses.

A recently published collection of papers on "adaptive individuals in evolving populations" (Belew & Mitchell, 1996) illustrates the growing interest of computer scientists in the implications of the Baldwin effect for evolutionary computation. Belew and Mitchell (1996) have gathered classical and recent papers from biology, psychology, and computer science, under the theme of interactions between phenotypic plasticity and evolution. The collection includes a reprint of Baldwin's (1896) article, and almost every paper mentions the Baldwin effect. Although Belew and Mitchell (1996) have collected more than 30 papers, this special issue of *Evolutionary Computation* shows that we are far from exhausting the research potential of the Baldwin effect.

Much of the computational study of the Baldwin effect has emphasized the first step—the synergy between learning and evolution. The second step of the Baldwin effect is known as genetic assimilation: Plastic mechanisms are assimilated by the genotype; learned behaviors become instinctive. Genetic assimilation is facilitated by using a fitness function that incorporates the cost of local search and by allowing the genotype to have some control over the amount of local search. Many hybrids of genetic algorithms and local search take advantage of the first part of the Baldwin effect (synergy), but they miss the advantage of the second part (assimilation), because they do not take deliberate measures to encourage genetic assimilation.

If there is no cost for learning, genetic assimilation will halt, with a residual amount of unassimilated learning, when the population has evolved to the point where every individual can reliably learn to reach the optimum fitness level. However, if the fitness function incorporates the cost of learning, then genetic assimilation will continue until learning is completely assimilated (ignoring genetic drift). For complete genetic assimilation, it must also be possible to evolve a genome that encodes a phenotype that behaves optimally without learning. That is, such a genome must be in the space of possible

genomes, and there must be some reasonable chance that evolutionary search can discover the genome.

Mayley (this issue) explores two criteria for genetic assimilation under the Baldwin effect: the evolutionary cost of learning and the existence of a neighborhood correlation relationship between genotypic and phenotypic space. Rates of convergence of simple evolution and evolution with learning are compared under varying degrees of evolutionary costs. The cost of learning is assumed to be proportional to the distance between the innate phenotype and the learned phenotype. He emphasizes that the intrinsic costs associated with learning drive the process of genetic assimilation. He then analyzes the role of correlations between genotype-phenotype and genotype-learned-phenotype mappings. He demonstrates that high correlations between genotypic and phenotypic space can create favorable conditions for genetic assimilation to occur. He concludes that these two conditions (high relative cost for learning and neighborhood correlation) are necessary conditions for genetic assimilation to occur.

Harvey (this issue) analyzes the effect of learning on evolution under conditions where there is no correlation between the task learned during an individual's lifetime and evolutionary fitness. Under certain conditions, lifetime learning on one task paradoxically improves performance on unrelated evolutionary tasks (Parisi, Nolfi, & Cecconi, 1992). Harvey constructs a minimal model consisting of evolving weight vectors and shows how excursions in phenotypic space due to learning can lead to recovery of performance that has been degraded by genomic perturbations introduced during reproduction (e.g., mutation). Harvey presents a geometric analysis to calculate the likelihood of benefit from learning as a function of mutation, learning investment, and genotypic dimension. This "recovery effect" is presented as a special case of the "relearning effect" described by neural network researchers (Hinton & Sejnowski, 1986; Harvey, 1996).

For Harvey, the defining characteristic of the Baldwin effect is genetic assimilation. Since the "recovery effect" does not involve genetic assimilation, Harvey argues that it constitutes "another new factor in evolution," orthogonal to the Baldwin effect. This claim may provoke some controversy, since some researchers have a more general interpretation of the Baldwin effect. It could be argued that learning on "unrelated tasks" has the same net result as learning in general—maintaining species polymorphism. Enhanced polymorphism facilitates evolution by increasing the likelihood of bringing together beneficial mutations and is a necessary component of the Baldwin effect (Anderson, 1995). However, Harvey has clearly identified an important and distinct route whereby learning affects genetic diversity.

In biological organisms, learning can be driven by pleasure and pain. Turing (1950) argued that artificial intelligence researchers would be wise to build a pain-pleasure mechanism into their software. Most research in reinforcement learning examines how to learn from a reinforcement signal but does not consider the origin of the signal. Batali and Grundy (this issue) call the pain-pleasure mechanism the "motivation system," and they investigate how motivation systems might evolve. Batali and Grundy show that interaction between a learning system and a motivation system can be much more complex and interesting than one might assume. The motivation system can evolve to encode regularities in the individual's evolutionary environment, which can simplify the learning task. On reflection, we can see that the motivation systems of biological organisms have a kind of "wisdom," which we tend to overlook.

Batali and Grundy systematically explore the evolution of six different neural network architectures in three different simulated worlds. Three of the network architectures evolve innate behaviors, whereas the other three evolve both a learning system and a motivation

system. This research significantly extends Littman's (1996) paradigm of "evolutionary reinforcement learning." It appears that the learning–motivation system architectures are particularly advantageous when a complex environment can be decomposed into a relatively simple physics function (simple state transition rules) and a relatively simple fitness function.

Bala, De Jong, Huang, Vafaie, and Wechsler (this issue) present what we believe is the first published evidence of the Baldwin effect in a practical, concrete, realistic application. Their application is recognition of visual concepts. Their algorithm is a hybrid of a genetic algorithm and a decision tree induction algorithm. They perform experiments with satellite and facial image data, and the results appear to support the hypothesis that their algorithm manifests both aspects of the Baldwin effect: synergy between evolution and local search (decision tree induction) and genetic assimilation of learned behavior. This is also evidence for our belief that the Baldwin effect has many practical applications in computational problem solving.

Turney (this issue) points out that machine learning algorithms have a bias. A bias is any factor inherent in the learning algorithm that causes it to prefer one result over another; bias is independent of the training data. This notion also extends to optimization routines: given a set of points that have been sampled from the search space, different algorithms may have very different probability distributions when selecting the next point in the search space to sample. Thus, all search algorithms have a bias.

Turney associates a "strong bias" with "instinct" and a "weak bias" with learning. The Baldwin effect then involves a shift from a weak bias to a strong bias that occurs over time. In Turney's experiments, a bias factor is included as part of the genetic encoding. The bias factor determines to what degree an individual depends on learning versus genetically determined behavior. The experiments show that a weak bias is preferred during early stages of learning and evolution, but as the genetic code for instinctual behavior becomes more reliable, the bias becomes stronger. This is not only consistent with other studies of the Baldwin effect, it also results in effective methods for dynamically adjusting the bias of algorithms for machine learning.

As Harvey (this issue) remarks, the Baldwin effect "can be subtle and often counterintuitive." The Baldwin effect is a rich and fertile area for research in evolutionary computation, as can be seen from the papers in this issue. We would like to thank the authors and the referees for their concerted efforts to make this an outstanding special issue.

Acknowledgments

Russell W. Anderson was supported by NSF grant number IBN-9511365 (awarded to Edward L. Keller). Darrell Whitley was supported by NSF grants IRI-9312748 and IRI-9503366.

References

- Anderson, R. W. (1995). Learning and evolution: A quantitative genetics approach. *Journal of Theoretical Biology*, 175, 89–101.
- Baldwin, J. M. (1896). A new factor in evolution. *American Naturalist*, 30, 441–451.
- Below, R. K. & Mitchell, M. (Eds.) (1996). *Adaptive individuals in evolving populations: Models and algorithms*. Reading, MA: Addison-Wesley.
- Harvey, I. (1996). Relearning and evolution in neural networks. *Adaptive Behavior*, 4(1), 79–82.
- Hinton, G. E., & Sejnowski, T. (1986). Learning and relearning in Boltzmann machines. In D. E. Rumelhart, J. L. McClelland, & the PDP Research Group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition*. Volume 1: Foundations (pp. 282–317). Cambridge, MA: MIT Press.

Introduction to the Special Issue

- Hinton, G. E., & Nowlan, S. J. (1987). How learning can guide evolution. *Complex Systems*, 1, 495–502.
- Littman, M. (1996). Simulations combining evolution and learning. In R. K. Belew & M. Mitchell (Eds.), *Adaptive individuals in evolving populations: Models and algorithms* (pp. 465–477). Reading, MA: Addison-Wesley.
- Morgan, C. L. (1896). On modification and variation. *Science*, 4, 733–740.
- Osborn, H. F. (1896). Ontogenic and phylogenic variation. *Science*, 4, 786–789.
- Parisi, D., Nolfi, S., & Cecconi, F. (1992). Learning, behavior and evolution. In F. J. Varela & P. Bourguine (Eds.), *Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life* (pp. 207–216). Cambridge, MA: MIT Press.
- Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, LIX, 433–460.
- Waddington, C. H. (1942). Canalization of development and the inheritance of acquired characters. *Nature*, 150, 563–565.
- Weismann, A. (1893). *The germ-plasm: A theory of heredity*. New York: Scribners.