DOI: http://dx.doi.org/10.7551/978-0-262-31050-5-ch006

Heterochronous Neural Baldwinism

Keith L. Downing
Department of Computer Science
The Norwegian University of Science and Technology
Trondheim, Norway
keithd@idi.ntnu.no

Abstract

Neural Baldwinism concerns the Baldwin Effect in the evolution of brains and intelligence. The first phase of the Baldwin Effect (B.E.), wherein plasticity provides a selective advantage, is intuitive and commonplace in simulations of adaptive systems. However, the second (assimilation) phase often poses problems for Baldwinism in general, and this is particularly acute for biological neural networks, where a complex developmental process greatly confounds the mapping from genotype to functional phenotype: a brain whose synapses are tuned to perform particular tasks. Since a strong genotype-phenotype correlation is often viewed as a prerequisite to this second phase, the body's most plastic organ would appear to defy Baldwinism. However, a detailed examination of 3 key processes of neural adaptation blurs the distinction between classic developmental and learning stages of brain maturation, thus supporting a re-interpretation of Neural Baldwinism's phase II as a heterochronous shift of the bulk of these three adaptive processes from postnatal to prenatal stages. This article illustrates Heterochronous Neural Baldwinism (HNB) with artificial neural networks that evolve, develop and learn, and in which some degree of synaptic tuning shifts to the prenatal stage.

Introduction

The Baldwin Effect (B.E.) (Baldwin, 1896; Turney et al., 1997) concerns the ability of learning to accelerate evolution via a two-stage process. In phase I, individuals with phenotypic plasticity achieve higher fitness than those relying purely on innate skills. This moves the population distribution toward plastic individuals. In phase II, some of these learned skills become innate by chance mutations. This assimilation of plastic features into the genome and developmental process becomes more probable when the genotype-phenotype mapping is not overly complex (with correlations maintained between genotype and phenotype spaces); and selection pressure favors assimilation when a) the environment is reasonably static across the generations, and b) learning has a fitness cost (Mayley, 1996).

Although B.E. seems plausible for some phenotypic traits, such as the size of muscles and the efficacy of certain physical skills, its relationship to the evolution of intelligence is more tenuous, given contemporary understanding of the brain, neural development and synaptic change. If learning is generally equated with synaptic change, then how can the modification of a few of the (human) brain's 100 trillion synapses be assimilated into DNA consisting of approximately 25,000 genes? In general, the mapping from genotype to phenotype is highly *indirect*, and correlations between genotype space and neural network space seem highly unlikely, once again precluding the assimilation of specific synaptic change into the genome.

In search of a more plausible reconciliation of neural evolution and Baldwinism, we examine the mechanisms traditionally associated with neural-network development and learning; the border between the two seems fuzzy with respect to the creation of new neurons (a.k.a. neurogenesis) and synapses (a.k.a. synaptogenesis), along with the tuning of those synapses. For instance, neurogenesis and synaptogenesis are not restricted to early neuro-development, as once believed. Recent evidence (Shors, 2009) shows that neurons can be generated and inter-connected throughout life, depending upon an animal's mental (and physical) challenges. Also, a good deal of synaptic tuning has been shown to occur prenatally (Sanes et al., 2006). Thus, neuro- and synaptogenesis, along with synaptic tuning, can be shared between development and learning, with the genome brokering the actual division of labor. Furthermore, heterochronous shifts in these distributions - that transfer some of the burden from stages of life that are strongly influenced by the environment (a.k.a. nurture) to those strongly governed by the genes (a.k.a nature) - seem to support the assimilatory requirements of B.E. phase II.

Motivated by these biological findings and their implications for the B.E., we investigate models in which a) artificial neural networks (ANNs) evolve, develop and learn, b) the core processes of neurogenesis, synaptogenesis and synaptic tuning have varying levels of activity in early (developmental) and late (learning) stages of life, and c) these levels are determined by the genome. Then, by monitoring the evolving distribution of these 3 processes among development and learning, we observe this more flexible interpretation of Neural Baldwinism.

Related Work

Hinton and Nowlan's simulations - classic in their simplicity and elegance - first illustrated B.E (Hinton and Nowlan, 1987). They showed that early learning helped guide evolution toward a difficult goal (B.E. phase I), but as the population approached the target, the flexible portions of the phenotype became hard-wired to the correct values, thus jettisoning the (costly) learning capabilities (B.E. phase II). Their model involved simple bit-string genotypes, which doubled as phenotypes, so no development nor neural networks were involved, though they commented that the model could serve as a coarse abstraction for the evolution of neural networks.

In another seminal Baldwinian simulation, (Ackley and Littman, 1992) showed the B.E. in evolved pairs of interacting neural networks, one of which learned by back-propagation while the other evolved (but could not learn) to provide proper world-state evaluations to guide learning in the former network. Upon adding learning (in neural networks) to cellular encoding (CE) (Gruau and Whitley, 1993) observed the confounding effects of development (in CE) upon the B.E.. Later, (Downing, 2004) extended the Hinton and Nowlan model to include an abstract developmental process based on a Turing machine (TM) (whose specifications were encoded in the genome). Those experiments showed the scaffolding effect that development can manifest to reduce the learning burden and thus support B.E. phase II. That scaffolding effect is also evident in this article, but now with fully-functioning neural networks as the phenotype and developmental synaptic tuning replacing the TM.

Recently, (Paenke et al., 2009) proposed a mathematical framework to help quantify when, in fact, learning will accelerate evolution, while many important B.E. studies employ models other than neural networks (Suzuki and Arita, 2007; Bull, 1999; Mayley, 1996) to reveal critical relationships between fitness landscapes, epistasis, the genotype-phenotype mapping, and B.E. Of critical relevance to this article are Mayley's two key prerequisites for B.E. phase II: a) a strong correlation between genotype and phenotype space, and b) a significant learning cost.

Despite the obstacles to B.E. posed by neural development, a plausible reconciliation of the two involves a heterochronous shift of a significant degree of neurogenesis and synaptogenesis from postnatal (experience-driven) learning to the prenatal (gene-governed) phase of life, as shown in (Downing, 2010). In this article, we turn to the third key factor, synaptic tuning, and explore the degree to which it too can be assimilated into neural development.

Heterochronous Neural Baldwinism (HNB)

Figure 1 summarizes a few of the key processes involved in the mapping from genes (roughly 25,000 in humans) to the brain (containing around 100 billion neurons and 100 trillion synapses). The high degree of scrambling and elaboration of genetic information that occurs during neural development and tuning must clearly ruin any correlations between genotype and phenotype space, thus violating a key precondition for B.E. phase II (Mayley, 1996). Thus, a plausible model of Neural Baldwinism would seem to require a different perspective and/or abstraction level.

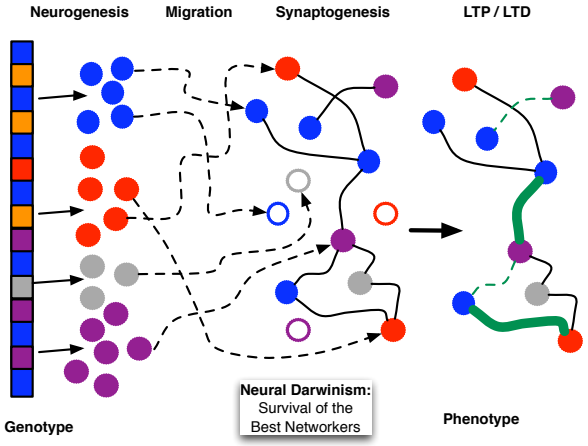

Figure 1: The complex gene-to-brain mapping.

A common *lock-step* scenario for brain formation and maturation consists of two clearly distinct phases: a) prenatal development wherein neurons are produced and linked together, and b) postnatal learning, wherein synaptic strengths are modified to enhance behavioral control. Though convenient for computational models and general explanations, this over-simplifies temporal relationships whose details may prove useful for understanding Neural Baldwinism. For example, many studies, summarized in (Sanes et al., 2006), find high levels of long-term potentiation (LTP) and long-term depression (LTD) - both forms of synaptic tuning - during prenatal development. In fact, the rates of LTP and LTD (i.e. learning rates) are actually very high during development and much lower during adult life. In addition, recent work (Shors, 2009) reveals that a) neurogenesis occurs throughout life, particularly in the dentate gyrus (DG) of the hippocampus, but b) those neurons only hook up to other neurons (and ultimately survive) if the organism subsequently performs cognitively-challenging tasks. Thus, although we can retain terminology that equates development with all prenatal brain formation, and learning with postnatal activity, the constituent processes of development and learning are clearly not mutually exclusive in this (more biologically realistic) *overlapping* model.

This new perspective motivates a reinterpretation of B.E. in neural networks. In a lock-step model, B.E. phase II entails converting synaptic-strength changes (i.e. classic learning) into genomic codes for controlling neurogenesis and synaptogenesis (i.e. classic development). This represents

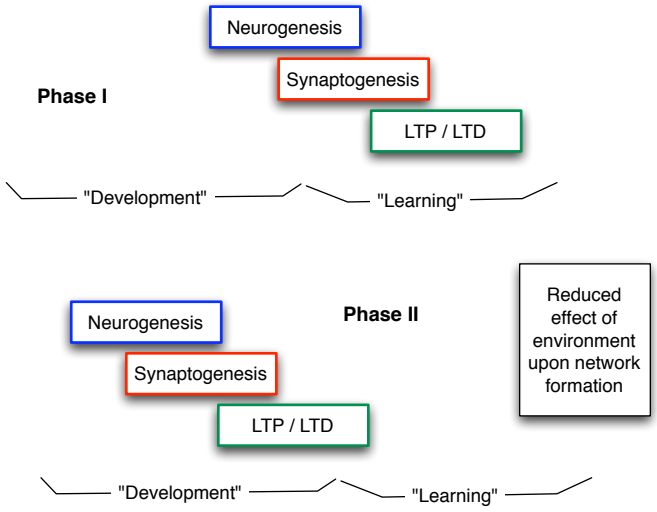

Figure 2: Overview of Heterochronous Neural Baldwinism, wherein the hallmark of phase II is the transfer of considerable neuro- and synaptogenesis, along with synaptic tuning, from postnatal to prenatal stages.

the reverse encoding of the results of one process into two dramatically different processes. However, in an overlapping model, the assimilation phase involves only quantitative, not qualitative, change.

Thus, B.E. Phase II may primarily constitute heterochrony: a change in the *onset*, *termination* and *rates* of neurogenesis, synaptogenesis and LTP/LTD across an organism's life stages, as shown in Figure 2. Under the view that adaptive changes in later life are predominantly governed by the environment, not the genome, a Baldwinian modification could simply be to move more of that adaptive change into earlier life stages, where genomic control may dominate.

For example, when the biochemical bases for LTP and LTD arose in evolution, both processes may have been very active throughout life, requiring constant environmental signaling to tune neural circuitry. However, over many (thousands of) generations, genomic changes could have arisen such that the early stages of development utilized neurogenesis, synaptogenesis and high LTP/LTD to form much of this circuitry with a minimum of environmental influence. Similarly, the rates of neurogenesis and synaptogenesis could have originally been much less variable throughout life, but evolution has gradually found genomes coding for an acceleration of these processes in early development; and thus, more of these activities became governed by genomic rather than environmental factors. The neural plasticity that remains in today's adult genomes (of any species), may represent that flexibility which evolution found optimal with respect to factors such as a) the coding limits of the genome, b) constraints of the animal's brain and body, and c) earth's environment and the rates of change associated with it.

This research illustrates Heterochronous Neural Baldwinism using ANNs which a) evolve using genetic algorithms, b) learn via backpropagation, and c) employ a complex developmental procedure for tuning synaptic weights prior to backpropagation. Importantly, evolution controls both the extent of backpropagation and the detailed nature of developmental tuning.

In this model, phase II of the Baldwin Effect is evident when the genome transfers a significant level of synaptic tuning from postnatal to prenatal stages, i.e., from a life stage when the environment governs a good deal of neural activity to a stage when the genome holds more control. Hence, although this model employs a complex, correlation-destroying mapping from genes to adaptive (developmental and learning) parameters to synaptic weights, the genome still possesses the ability to evolve the recipe for a developmental procedure that can reduce the burden of postnatal adaptation and thereby increase phenotypic fitness.

Developmental Synaptic Tuning (DST)

The DST model introduces a biologically-inspired mechanism for modifying connection weights prior to exposure to the environment (i.e., training cases). This mechanism abstracts from neurological studies showing that spontaneous waves of neural activity, modulated by cyclic-AMP (cAMP) concentrations, lead to early synaptic tuning during development, prior to the exposure to normal sensory inputs. This has been shown to play an important role in the binocular segregation of connections from the retina to the lateral geniculate nucleus (LGN) (Stellwagen and Shatz, 2002), while others (McNaughton et al., 2006) postulate similar wave-induced synaptic tuning in the hippocampus, and a variety of evidence, summarized in (Sanes et al., 2006), indicates both a) the presence of these waves throughout the brain during neural development, and b) their instructive role in synaptic formation and tuning.

These waves promote neural firing such that neurons in adjacent regions that happen to fire simultaneously (due to stimulation from their respective activation waves) will have their synaptic connections modified, typically by Hebbian means. Thus, early chemical waves strongly influence the patterning of neuronal connections, prior to the molding effects of normal sensory stimuli.

A comprehensive model of this phenomena would include the chemical and physical bases of reaction-diffusion processes, a reasonably straightforward but computationally-intensive endeavor. Fortunately, compositional pattern-producing networks (CPPNs) (Stanley, 2007) provide an efficient alternative for abstractly modelling any number of natural pattern-generating processes.

Composite Pattern-Producing Networks (CPPNs)

As shown on top of Figure 3, a CPPN (Stanley, 2007) resembles a neural network, but with each node housing one of a number of alternative activation functions, as opposed to the standard sigmoids, step functions and hyperbolic tangents of ANN nodes. For example, the CPPN may include Gaussians, absolute values, and sine waves (as well as the common ANN activation functions). Each CPPN connection includes a weight, and all nodes compute the sum of their weighted inputs, which serves as input to the activation function, whose result becomes the node's output.

The CPPNs in this research have no explicit layered organization (other than pre-defined input and output nodes), so any node can send outputs to any other node; and all nodes (except the inputs) can receive weighted outputs. At each timestep, the nodes undergo asynchronous activation, wherein each node simply sums the weighted outputs in its input buffer and feeds that sum to its activation function to produce an output value, which is immediately propagated to the input buffers of all post-synaptic neighbors. After a user-determined number of update rounds, the CPPN's outputs are gathered from the output nodes.

By sending Cartesian coordinates through a CPPN and using the output value to encode pixel color or intensity, the CPPN can generate pictures (Stanley, 2007). Similarly, by adding the time step as input, the CPPN can produce a time series of patterns, depicting a dynamic structure such as an activation wave, as shown at the bottom of Figure 3.

In the DST model, an ANN's genome encodes various parameters for each of its layers, such as the number of initial neurons. In addition, it can include a set of CPPN genes for any layer such that the decoded CPPN can be used to generate activation waves during development.

Neural layers are modeled as 2d surfaces, where each neuron (n) has a center coordinate, (x_n,y_n). During development, to compute the wave-induced activity of n at time t, simply input x_n, y_n and t to the layer's CPPN and interpret the output value as a local activation.

When adjacent layers in an ANN include CPPNs, each can be run to produce activation patterns. As shown in Figure 4, when neurons j and k in adjacent layers (J and K) have correlated wave-induced activation, Hebbian-based synaptic tuning on the j-k connection provides an early bias of the network. When the activation waves fortuitously reflect some aspect of the sensory world to which the organism will eventually be exposed, this preliminary synaptic tuning should provide a useful *head start* for the ANN and agent.

Hence, by including CPPN parameters with the other layer-specific genes in an evolving ANN, any pair of interconnected layers with CPPN-based developmental stimulation can achieve an evolving prenatal bias of its weights.

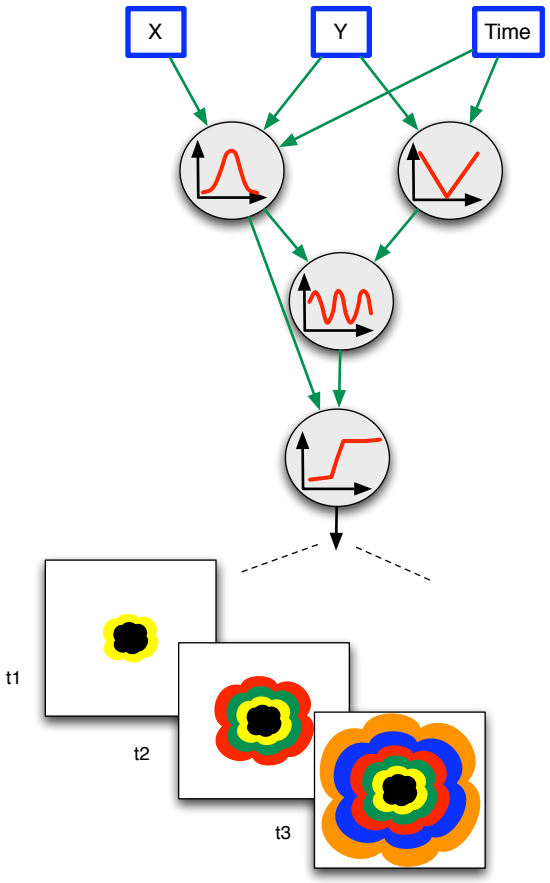

Figure 3: A CPPN, when provided with Cartesian coordinates and time as inputs, produces abstract temporal activation patterns.

Evolving CPPNs

CPPNs, as defined in (Stanley, 2007) are evolved via the NEAT system (Stanley and Miikkulainen, 2002), which has the advantage of supporting gradual complexification but which is a rather direct encoding, with one gene required for each node and weight. This work employs a more generative CPPN encoding to reduce the need for individual weight genes and achieve a bit more biological plausibility.

These CPPNs evolve via a simple bit-vector chromosome consisting of multiple segments, one for each input, internal and output node in the network. Each segment consists of 5 genes that encode the: a) activation function, b) afferent connection tag, c) efferent connection tag, d) afferent weight tag, and e) efferent weight tag.

The first is simply an index into a list of possible activation functions (identity, sine, absolute value, gaussian and sigmoid), while the afferent tags for node N help determine a) which nodes can send input to N, and b) the weights on those incoming arcs. Similarly, the efferent tags influence a)

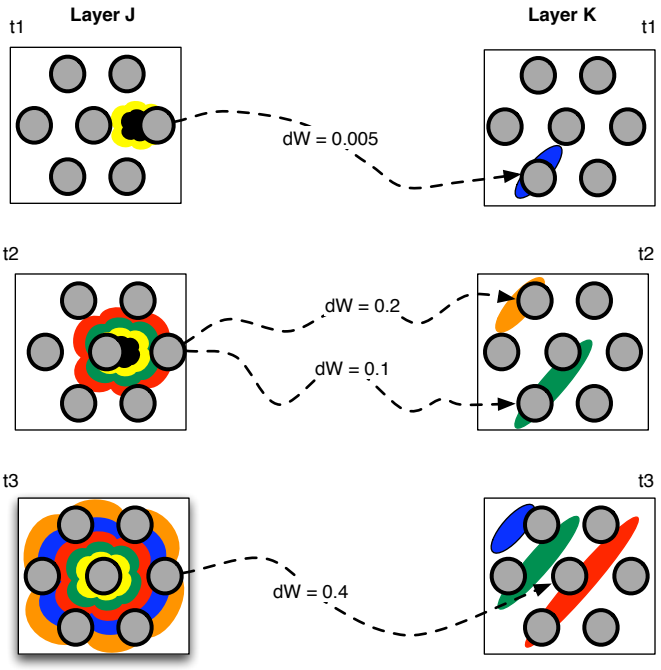

Figure 4: CPPN-generated activation patterns stimulate adjacent neural layers, leading to correlation-based weight changes (dW) to synapses between co-active neurons.

the nodes to which N can send its output, and b) the weights on those outgoing arcs. The two afferent (efferent) tags constitute the afferent (efferent) *mask* of each node.

More specifically, if the afferent connection tag of node N matches the efferent connection tag of node M (above a user-defined match threshold, e.g. 0.75), then M will send an excitatory connection to N. Conversely, if the match is very poor, and thus below a similar threshold, e.g. 0.25, then M will send an inhibitory connection to N. For medium-strength tag matches, no connection between M and N is created. Then, the strength of an excitatory or inhibitory arc is positively correlated with the matching degree of M's efferent-weight- and N's afferent-weight tags.

Our CPPNs have a pre-defined number of input and output nodes, but the efferent and afferent masks (along with the activation functions) of output nodes can evolve, as can the efferent masks of each input node. Since recurrent links are permitted in CPPNs, both types of masks are relevant for output nodes, while input nodes are strictly entry ports for external data. During CPPN configuration, once all connections are determined, nodes that either form no connections or, more generally, do not lie along at least one pathway from some input to some output node, are pruned.

Combining DST with Backpropagation

To investigate Heterochronous Neural Baldwinism using DST, we employ a 3-layered (input, hidden, output) ANN, with 9 linear input, 9 sigmoidal output, and a variable number of sigmoidal hidden nodes (coded in the genome). Standard backpropagation (BP) accounts for all of the learning, while CPPN-generated waves handle developmental tuning of connections between the hidden and output neurons. The complete DST-BP process is summarized in Figure 5.

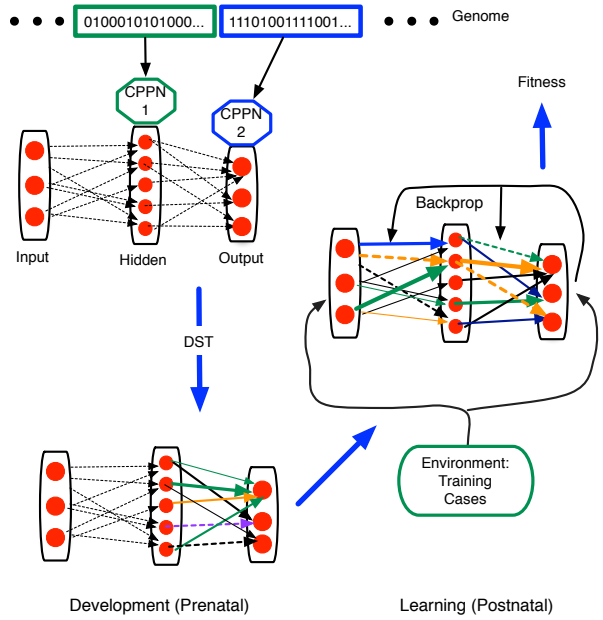

Figure 5: Key stages of the DST-BP process: (Upper left) Translation of chromosomal segments to CPPNs for 2 layers; (Lower left) DST provides an initial bias to hidden-to-output connections; (Right) BP learning further tunes all synapses to capture the training set, with fitness based on training error, test error, and BP learning effort.

The Genetic Algorithm

In the spirit of Hinton and Nowlan's original work, we use relatively small populations evolved over relatively few generations to solve simple problems, primarily as a proof of concept. The key GA parameters include a population size of 20, full-generational replacement with a rank-based selection mechanism and elitism of two individuals, a crossover rate of 0.8 and a mutation rate of 0.05 per bit.

The GA chromosome for DST-BP involves 9 basic developmental and learning parameters, while two CPPNs (for the hidden and output layers) require 14, 5-part genes apiece (encoding activation functions and masks). The 3 developmental parameters are: 1) initial hidden-layer size ($H_{init} \in [1, 10]$), 2) developmental tuning rate ($D_r \in [0, 1]$), and 3)

activation wave steps ($D_s \in [0, 5]$). The 6 learning parameters are 3 each for the input-hidden layer connections, and the hidden-output layer links. The former 3 are: a) learning rate, $L_{r1} \in [0, 1]$, b) learning epochs, $L_{e1} \in [1, 10]$, and c) momentum, $M_{e1} \in [0, 0.2]$, while the latter 3 are: d) $L_{r2} \in [0, 1]$, e) $L_{e2} \in [1, 10]$, and f) $M_{e2} \in [0, 0.2]$.

Fitness Testing

Data sets for backpropagation learning are generated by one-dimensional cellular automata (CA), run for 10 timesteps, with each consecutive pair of 9-cell states constituting a training case, i.e., $s_t \rightarrow s_{t+1}$,where s_t, the CA state at time t, is loaded onto the input neurons, with the target output being s_{t+1}. The key point is that if a bit is on (or off) in s_t, this influences the chances of it and other neighboring bits being on or off in s_{t+1}, thus adding some structure to the data.

Fitness stems from both training and testing error, with the former consisting of the average error per output neuron, per training case, per epoch; while the latter is per neuron per case for a single epoch, without learning. Thus, the abilities to a) *quickly* reduce error during training, and b) *eventually* reduce that error, are independently assessed.

In all of the runs reported below, each individual undergoes 5 independent rounds of fitness testing, wherein a different set of random initial weights (in the range [-0.2 0.2]) are assigned. The error terms E_{train} and E_{test} denote averages over these 5 rounds, and all runs employ a *tuning tax*, $\Theta = 0.04$. The fitness function of equation 1 accounts for both error terms along with the learning effort:

$$f = e^{-(E_{train}+E_{test}+\Theta L_e L_r)} \quad (1)$$

where $L_e = min(L_{e1}, L_{e2})$ (since only this minimum of the two values of epochs is actually performed), and $L_r = \frac{L_{r1}+L_{r2}}{2}$.

The Developmental Contribution

In each of the runs below, the contribution of developmental synaptic tuning to error reduction is estimated by a simple test, performed only on the best-of-generation individuals. First, the weights of the ANN are randomly initialized, before sending the entire training set through the network, but without learning. The average error, per output node, per training case, is then compared to the average error in a second test, wherein the same ANN, with the same initial weights, also undergoes the developmental tuning encoded in the genome. Differences in these two error terms gives a rough indication of the contribution of DST to error reduction, and thus to fitness. It is important to note that the fitness function does not explicitly reward this contribution. Its effect is only indirect, via the reduction in learning effort afforded by DST.

In the runs below, a typical training error prior to any tuning (developmental or learning) is 0.35 to 0.45, while the improvement typically varies from 0 to 0.05 (e.g. 0.45 - 0.40 = 0.05).

Results

For each of the runs presented below, 5 properties of the best-of-generation individual are plotted: fitness, learning effort, training error, test error, and developmental contribution to error reduction. To easily display each value on the same linear plot, the two error values and the developmental contribution are multiplied by 10.

Figure 6 illustrates a sample 300-generation run using DST-BP and a CA-generated data set. HNB phase I involves a gradual increase in learning effort over approximately 80 generations, during which fitness rises and error falls. Phase II begins thereafter, with learning effort dropping in several discrete steps over the final 220 generations, although the lowest learning-effort value appears evolutionarily unstable. In this plot, the rise of fitness during phase II is clearly evident. The developmental wave functions for the best-fit individual of the 300th generation are shown in Figure 7.

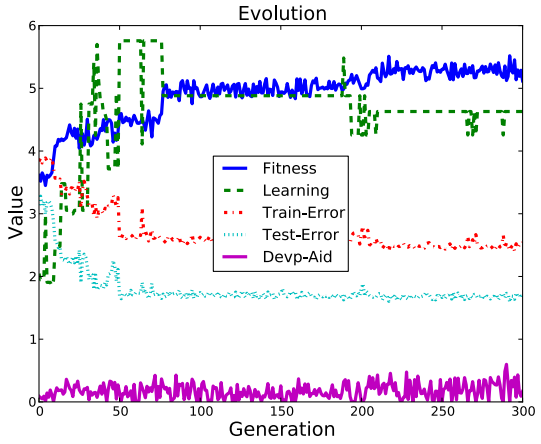

Figure 6: Time series of fitness, learning effort, two error terms, and developmental enhancement for each best-of-generation individual in a 300-generation DST-BP run.

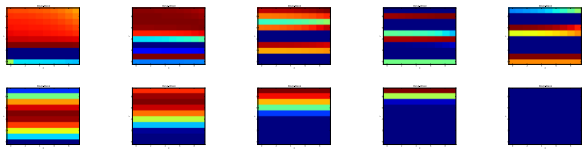

Figure 7: The 5-step, CPPN-generated, developmental activation waves for hidden (top row) and output (bottom row) layers for an evolved, 3-layered feed-forward network using backpropagation learning. Red (blue) indicates maximal (minimal) stimulation.

To test the generality of HNB using DST-BP, a series of 20 independent runs are performed, each using a different CA-generated dataset. The best-of-generation averages, shown in Figure 8, display the general Baldwin Effect in that learning initially boosts fitness, but then plasticity decreases while the error terms remain lower (than the early-generation values), and fitness gradually increases.

Heterochronous Neural Baldwinism is evidenced by the bottom line in the figure, which shows a gradual rise in the contribution of development to error reduction. This rise is barely perceptible in the figure, but a comparison of the first 50 averages (over 20 runs) to the last 50 shows a statistically significant difference (p = .0005) in a single-tailed Student-t test: the averages are 0.020 and 0.027 for the first and last quarter, respectively. This indicates that the developmental contribution to output error reduction, though small, does allow learning effort to decrease.

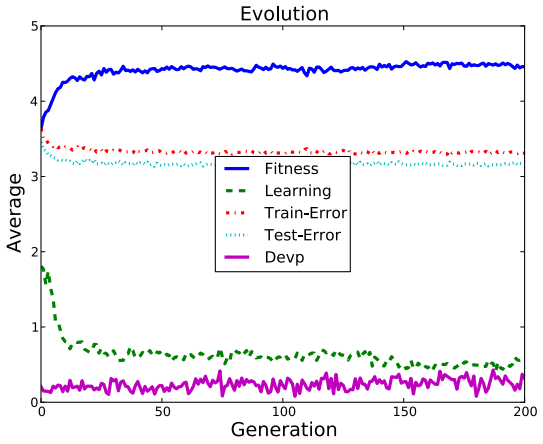

Figure 8: Averages (over 20 independent runs using different CA-generated datasets) of fitness, learning effort, two error terms, and developmental enhancement (devp) for best-of-generation individuals in a population of 20 DST-BP networks.

In another set of 20 runs, the datasets consist of randomly-generated sparse patterns, with exactly 3 ones and 6 zeros in each. In contrast to the CA-generated datasets, these have no spatial relationship between the on (1) bits of the input and output/target patterns, thus making it harder for evolution to find helpful developmental schemes. However, HNB occurs in these runs (not shown) as well, with gradually declining learning effort and gradually increasing (and statistically significant, p = .0005) developmental contribution and fitness.

Further evidence for the contribution of development appears in Figure 9, which displays 20-run averages for scenarios using CA-generated datasets, but no DST. Notice that the learning effort rises and remains high throughout the 200 generations. Without developmental assistance, learning must remain elevated to keep the error levels in check. This continuously-high learning cost keeps fitness levels in Figure 9 well below those of Figure 8.

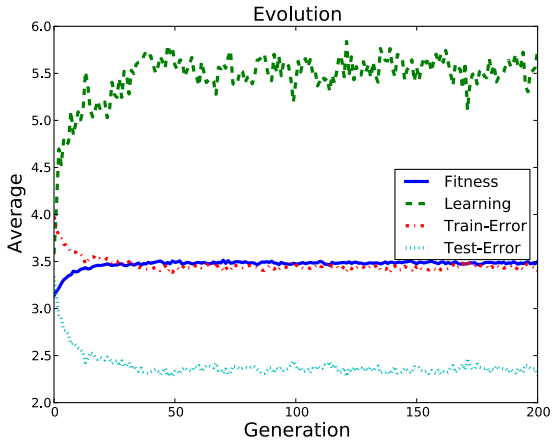

Figure 9: Averages (over 20 independent runs using different CA-generated datasets) of fitness, learning effort, and two error terms, for best-of-generation individuals in a population of 20 DST-BP networks, where DST is silenced.

Discussion

The DST-BP model gives preliminary evidence of the Baldwin effect in neural networks that undergo a developmental process. The validity of these results hinges on a quantitative (rather than qualitative) interpretation of the key differences between development and learning in neural systems. Namely, cross-generational changes in the pre- and postnatal **rates** of neurogenesis, synaptogenesis and synaptic tuning can transfer adaptive effort from learning to development, with the latter more closely governed by the genome and less by environmental factors. Thus, B.E. phase 2 transfers a portion of brain formation backwards, from learning to development, and thus to a stage where it is more strongly affected by the genome - and can thus lay claim to being *more innate* than characteristics acquired later in life, when environmental influences typically play a more decisive role.

As in (Downing, 2004), the effects of learning are not reverse-encoded into the genome, but strong learning *buys evolutionary time* until proper developmental scaffolding reduces the overall postnatal adaptive costs, thereby raising fitness to peak levels. In this work, scaffolding involves CPPN-generated activation waves, and the ensuing prenatal synaptic tuning, which provides the postnatal phase with a synaptic matrix that is already partially biased toward the environment (i.e. training set).

The DST-BP model embodies a complex mapping from genes to synaptic weights (via CPPNs and DST), which

might seem to preclude B.E. phase II. However, the genome possesses enough flexibility to evolve developmental scaffolding capable of reducing postnatal adaptive demands. Thus, although the *results* of postnatal synaptic tuning do not become innate, their *attainment* becomes easier due to evolved innate processes.

The DST model says little about the potential applicability of CPPNs as activity-wave generators for evolving ANNs, and it seems impractical to go to such lengths to produce functioning ANNs for complex tasks. But in support of HNB, the CPPN provides an appropriate abstraction (over complex reaction-diffusion interactions) for generating activity waves whose biological counterparts do appear to play an important role in early neural circuit formation. The DST-BP model indicates that these activity waves may also help explain Neural Baldwinism, though a more convincing argument would involve a learning mechanism of greater biological realism than backpropagation. To this end, we have combined DST with Hebbian learning in simple two-layered ANNs. This produces a more dramatic Baldwin Effect than the DST-BP runs, both in terms of a greater learning declines and very sizeable developmental contributions to error reduction (often over 30%). However, the Hebbian model was only able to learn very simple training sets involving very sparse input and output vectors.

Despite these relatively weak results, ALife researchers should profit from this article's primary insight: the assimilatory phase of the Baldwin effect, when viewed as a quantitative rather than qualitative shift in activity from later to earlier stages of life, does reduce the general complexity (and near impossibility) of the transfer of the fruits of neural plasticity to the realm of genetic control.

As elaborated by several influential biologists (West-Eberhard, 2003; Kirschner and Gerhart, 2005), the interactions between evolution, development and learning are intricate and multifaceted. Though often intimidating, these complex relationships may in fact open the gate for interpretations of B.E., such as HNB, that can enhance its general plausibility with respect to the evolution of intelligence.

References

Ackley, D. H. and Littman, M. L. (1992). Interactions between learning and evolution. In Langton, C. G., Taylor, C., Farmer, J. D., and Rasmussen, S., editors, *Artificial Life II*, pages 487–509, Reading, Massachusetts. Addison-Wesley.

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

Bull, L. (1999). On the Baldwin Effect. *Artificial Life*, 5(3):465–480.

Downing, K. L. (2004). Development and the Baldwin Effect. *Artificial Life*, 10(1):39–63.

Downing, K. L. (2010). The Baldwin effect in developing neural networks. In *Proceedings of the 12th Genetic and Evolutionary Computation Conference*, pages 555–562, Portland, Oregon. ACM Press.

Gruau, F. and Whitley, D. (1993). Adding learning to the cellular development of neural networks. *Evolutionary Computation*, 1(3):213–233.

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

Kirschner, M. W. and Gerhart, J. C. (2005). *The Plausibility of Life: Resolving Darwin's Dilemma*. Yale University Press, New Haven, CN.

Mayley, G. (1996). Landscapes, learning costs and genetic assimilation. *Evolutionary Computation*, 4(3):213–234.

McNaughton, B., Battaglia, L., Jensen, O., Moser, E. I., and Moser, M. B. (2006). Path integration and the neural basis of the 'cognitive map'. *Nature Reviews Neuroscience*, 7(8):663–678.

Paenke, I., Kawecki, T., and Sendhoff, B. (2009). The influence of learning on evolution: a mathematical framework. *Artificial Life*, 15(2):227–245.

Sanes, D., Reh, T., and Harris, W. (2006). *Development of the Nervous System*. Elsevier Academic Press, Burlington, MA.

Shors, T. J. (2009). Saving new brain cells. *Scientific American*, 300(3):40–48.

Stanley, K. (2007). Compositional pattern producing networks: a novel abstraction of development. *Genetic Programming and Evolvable Machines: Special Issue on Developmental Systems*, 8(2):131–162.

Stanley, K. and Miikkulainen, R. (2002). Evolving neural networks through augmenting topologies. *Evolutionary Computation*, 10(2):99–127.

Stellwagen, D. and Shatz, C. (2002). An instructive role for retinal waves in the development of retinogeniculate connectivity. *Neuron*, 33(1):357–367.

Suzuki, R. and Arita, T. (2007). The dynamic changes of roles of learning through the Baldwin Effect. *Artificial Life*, 13(1):31–43.

Turney, P., Whitley, L. D., and Anderson, R. W. (1997). Introduction to the special issue: Evolution, learning, and instinct: 100 years of the Baldwin Effect. *Evolutionary Computation*, 4(3):iv–viii.

West-Eberhard, M. J. (2003). *Developmental Plasticity and Evolution*. Oxford University Press, New York, NY.