

Heterochrony: It's (all) about time!

Matthew Schlesinger
Department of Psychology
Brain and Cognitive Sciences
Southern Illinois University
Carbondale, IL 62901
matthews@siu.edu

Abstract

Heterochrony—variation in the rate or timing of developmental processes or events over evolutionary time—plays an important role in the study of evolutionary developmental biology. I review the historical background of heterochrony, and highlight examples of how both physical structure and behavior are influenced by changes in the rate of development. I also describe neurogenetic models of evolution and development, which have been used to investigate heterochronic mechanisms of change. To help illustrate how epigenetic robotics complements the use of neurogenetic models, I propose two research questions for epigenetic robotics that focus on evolutionary changes in developmental timing.

1. Introduction

A fundamental goal of **epigenetic robotics** is the design and study of artificial organisms or agents that develop through interaction with their environment (e.g., Lungarella, Metta, Pfeifer, & Sandini, 2003; Weng, McClelland, Pentland, Sporns, Stockman, Sur, & Thelen, 2001). This work focuses on changes in physical structure or behavior that occur during the **ontogenetic** timescale, that is, the lifespan of a single individual. A closely-related field of research, **evolutionary robotics**, focuses instead on the **phylogenetic** timescale by investigating changes that occur over multiple generations in a simulated species (e.g., Husbands, Harvey, Cliff, & Miller, 1997; Nolfi & Floreano, 2001).

As a research strategy, it may seem expedient to treat the developmental and evolutionary timescales as two distinct or independent levels of change (analogous to how one robotics researcher studies object recognition while another focuses on reaching and grasping). Indeed, researchers employing either the epigenetic- or evolutionary-robotic approach often rely on different modeling paradigms and learning algorithms (e.g., reinforcement learning vs. genetic algorithms).

In the current paper, I argue instead for precisely the opposite strategy. In particular, I highlight the concept of **heterochrony**—systematic change in the rate or timing of a developmental stage or event that occurs over successive generations (Gould, 1977)—as a mechanism that not only integrates developmental and evolutionary timescales into a unifying framework, but that also reveals bi-directional influences between the two processes.

The rest of the paper is organized as follows. In

Section 2, I first provide a brief overview of the historical and conceptual background from which the concept of heterochrony has emerged. Next, Section 3 presents a survey of existing computational models that simulate heterochrony in a population of artificial organisms. Section 4 outlines future research questions, including the development of locomotion and tool-use. In the final section, I discuss the mutually-beneficial role that models of evolution and development can play in the ongoing dialog between robotics and developmental science.

2. Historical/conceptual background

Over evolutionary time, how do new structures or skills appear within a species? As the field of *evolutionary developmental biology* has taken shape, several different answers to this question have been proposed.

2.1 Recapitulation theory

One of the most well-known solutions was inspired by Ernst Haeckel, who compared the appearance of different animals during **embryogenesis** (i.e., growth during the embryonic period), and ultimately proposed that *ontogeny recapitulates phylogeny*. As Figure 1 illustrates, Haeckel noted the similarity in form of different species at corresponding points in embryonic development. For example, during phase I, the human embryo resembles the fish embryo; during phase II, the human embryo resembles the tortoise or chick embryo; and so on.




Figure 1: A copy of Haeckel's embryo drawings, used to highlight cross-species similarities that occur during embryogenesis.

According to the strong form of Haeckel's **recapitulation theory**, evolutionary advances occur by "adding on" to the adult forms of previous species. In other words, during embryonic growth, a given species quickly "passes through" (i.e., recapitulates) the stages of its ancestors, and then adds a new stage.

Although Haeckel's recapitulation theory was highly influential well into the 20th century, his account was ultimately rejected by evolutionary developmental biologists. In particular, as Gould (1977) notes, Haeckel's theory was based on a critical, but unfortunately incorrect premise: during embryonic growth, it is not the *adult ancestors* that are "recapitulated", but rather the *embryonic ancestors*.

Is it a coincidence, then, that as a species takes shape during embryogenesis, it *just happens* to resemble the embryos of its ancestors? According to Gould, this similarity is neither a coincidence nor epiphenomenal, but instead reflects a set of universal morphological principles, which are exploited by fish, amphibians, reptiles, and mammals as early growth occurs.

2.2 Heterochrony

Despite the demise of recapitulation theory, two key ideas have been taken from Haeckel, reshaped, and incorporated into a more comprehensive theoretical framework. First, developmental **acceleration** refers to an increase in the rate of a developmental process over successive generations. Second, **terminal addition** refers to the appearance of a new stage or period of development that is added or appended to an ancestor's development sequence.

Figure 2 illustrates the complementary processes of acceleration and terminal addition (the diagram is borrowed from Gould, 1977, Figure 26). The horizontal arrow indicates the direction of developmental time (e.g., a progression of stages) within a single individual, which is composed of a juvenile period followed by adulthood. Similarly, the vertical arrow indicates the direction of evolutionary time (e.g., successive generations in a hypothetical species).

In the first generation (i.e., G_1), there are two developmental stages (S_1 and S_2), which (for simplicity) are assumed to last equally long. In G_2 , these two stages are compressed (i.e., the developmental rate is accelerated), and a new, third stage (S_3) is added to the end of the developmental sequence. In a similar manner,

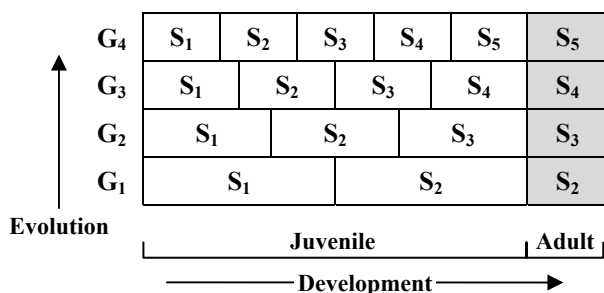


Figure 2: Schematic representation of acceleration and terminal addition. The horizontal arrow indicates developmental time (S_i = stages within an individual), while the vertical arrow indicates evolutionary time (G_i = generations).

Figure 2 illustrates the process of acceleration and terminal addition over successive generations.

It should be noted that terminal addition can occur *without* developmental acceleration, that is, by preserving the ancestral rate of development, and then adding a new stage after the ancestral stages have completed. However, the cost for this strategy is a necessary lengthening of the organism's lifespan, which accumulates over generations and leads to the "paradox of infinite length" (Gould, 1977).

In more general terms, acceleration and terminal addition are two related examples of **heterochrony**, in which the timing (i.e., the onset, offset, or rate) of a developmental stage, process, or event is modified over successive generations (see Alberch, Gould, Oster, & Wake, 1979, for a detailed quantitative model that describes six major categories of heterochrony). Thus, as Figure 2 illustrates, accelerating development results in the earlier emergence of stages, structures, or behaviors in a descendant, relative to their ancestors.

In contrast, Figure 3 illustrates the effect of developmental **retardation**. By slowing development over evolutionary time, the extent or endpoint of development is shortened (i.e., fewer stages are reached) while those stages that occur become progressively longer.

An important consequence of retardation, as Figure 3 highlights, is that features or characteristics that were present in immature or juvenile organisms in ancestors become retained during adulthood in descendants. This phenomenon is referred to as **neoteny**. At first glance, neoteny may appear to be a maladaptive or regressive trend. Indeed, how might it be adaptive for evolution to "go backwards", that is, for development to be truncated or shortened, so that the adults of a species come to resemble or behave like the young of their ancestors?

A familiar example of neoteny is the domestication of dogs, who are descended from wolves (Morell, 1997). In this case, wolves were selectively bred with the goal of slowing development so that adult "dogs" are effectively large, immature wolves that retain both the physical and behavioral characteristics of puppies.

In contrast to dogs, who are neotenuous by design, there are several lines of evidence which suggest that slowed or delayed development in humans is in fact an evolutionary adaptation. First, Bjorklund (1997) notes that while humans are **altricial** (i.e., developmentally immature) at birth relative to other species, the human infant's

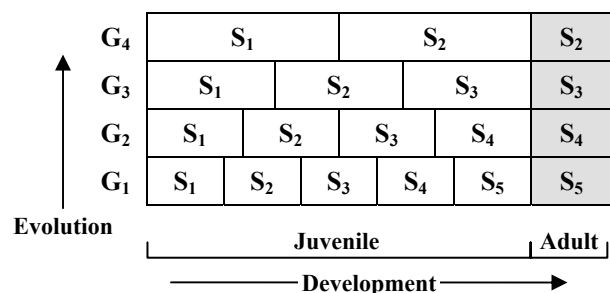


Figure 3: Schematic representation of retardation and neoteny. The horizontal and vertical dimensions are the same as Figure 2. In contrast to terminal addition, during neoteny, ancestral juvenile traits are retained in adulthood.

prolonged period of slow development helps to promote both cognitive and social development. In particular, Bjorklund suggests that developmental immaturity aids: (1) visual development, by limiting both the rate and amount of information that is attended to and stored in short-term memory (see Elman, 1993, and Schlesinger, Parisi, & Langer, 2000, for comparable accounts of language acquisition and motor skill development, respectively), (2) cognitive development, by motivating curiosity, exploration, and play, and (3) social development, by virtue of the need for protection and supervision from adult caretakers.

A second, related line of evidence is provided by Gould (1977), who suggests that human evolution reflects a gradual retardation of our primate ancestor's morphological development. This view is supported by a comparison of human and chimpanzee skulls (e.g., Mitteroecker, Gunz, Bernhard, Schaefer, & Bookstein, 2004). As Figure 4 illustrates, during the fetal period, human and chimpanzee skulls are remarkably similar. In contrast, by adulthood the chimpanzee skull has undergone major qualitative change, while the human skull has changed less dramatically. Gould uses these data to argue that by slowing the rate of skeletal growth in our primate ancestors, the hominid brain case has progressively enlarged, creating the opportunity for developmental increases in brain size.

Although comparisons among living animals and their ancestors are a fundamental aspect of paleontology, these comparisons are limited not only by gaps in the fossil record, but perhaps more importantly, by the fact that each species represents a “natural experiment” that does not develop or evolve under precisely controlled conditions (e.g., physical environment, historical time period, etc.).

Fortunately, other research methods are available that complement morphological comparisons. For example, molecular-genetic methods offer a measure of genetic “relatedness” across species, and in particular, an estimate of the time points at which related species diverge from a common ancestor. Another key element of molecular genetics is the distinction between **structural genes**—which “encode” morphological structure—and **regulator genes**—which modulate the timing and activity of structural genes.

Indeed, the concept of a regulator gene that directs the timing of developmental events has been explicitly incorporated and investigated in computational models of heterochrony. Perhaps not surprisingly, these models span a wide array of biological phenomena, including molecular and cellular mechanisms, morphology and growth patterns, neurophysiology, behavior, and population dynamics.

3. Neurogenetic models of heterochrony

Of particular relevance to epigenetic robotics are **neurogenetic** or **cellular encoding** models (e.g., Cangelosi, Parisi, & Nolfi, 1994; Gruau & Whitley, 1993; Kitano, 1990), which combine genetic algorithms (GAs)

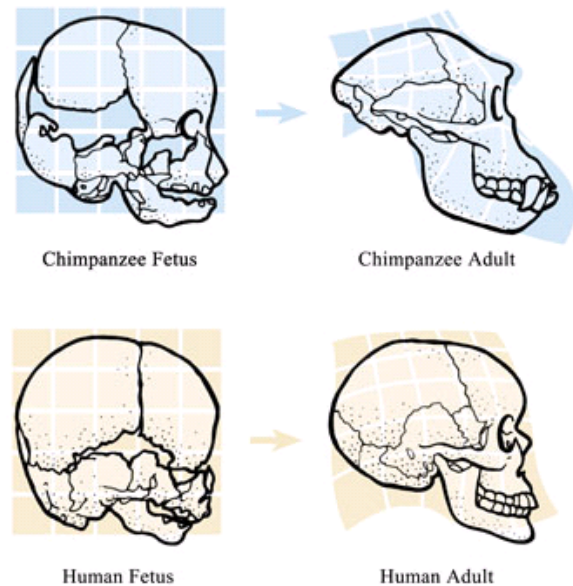


Figure 4: Comparison of skull growth between the fetal period and adulthood, in chimpanzees and humans.

and artificial neural networks (ANNs) into a unified computational framework. While neurogenetic models resemble conventional GAs and ANNs, they also possess two unique features: (1) a genome (i.e., structural genes) that determines the initial topology of a 2D neural network in a population of artificial networks, and (2) a set of growth rules that allow the networks to develop novel structures through the combined influence of both external (i.e., interaction with the environment) and internal (i.e., regulator genes) factors (see Stanley & Miikkulainen, 2003, for a systematic comparison of grammatical and cell chemistry approaches to neurogenetic modeling).

From a computational perspective, neurogenetic models raise a number of interesting questions. Does adding a developmental process to a GA influence the evolutionary pattern? Do ANNs that evolve *and* develop outperform those that only evolve? How much of the network's structure and/or behavior is encoded in the genes (i.e., innately determined), either at “birth” or through maturation? Alternatively, which aspects of neural structure and/or behavior are determined by the developmental process?

In general, combining evolutionary and developmental mechanisms in a neurogenetic model offers several benefits, including: (1) accelerated learning on a number of basic computational tasks (e.g., Boolean parity problems), (2) improved performance over the use of evolutionary learning alone, (3) greater scalability to large-scale networks and tasks domains, and (4) the creation of task-specific or modular sub-networks that can be copied and reused (Bongard & Pfeifer, 2001; Downing, 2004; Gruau, 1994; Gruau & Whitley, 1993; Matos, Suzuki, & Arita, 2007).

3.1 The Baldwin Effect

The latter question—to what degree are neural structures or behaviors genetically encoded versus determined by development—is related to the **Baldwin Effect** (Baldwin,

1896). Baldwin proposed a bi-directional interaction between evolution and development, in which (1) evolution initially guides development by the selection of individuals who are “good learners”, and then subsequently (2) the skills or behaviors that are re-learned by each generation gradually become genetically encoded (Downing, 2004; Hinton & Nowlan, 1987).

Several important findings have emerged from neurogenetic models of the Baldwin Effect. First, rather than viewing evolution and development as competing for control of a given structure or behavior, it may be more advantageous to view the two levels (mechanisms, timescales, etc.) of change as coordinating or cooperating. For example, Downing (2004) contrasts **genetic blueprints** with **developmental recipes**: blueprints are relatively direct mappings from genotype to phenotype, with little or no influence from the environment (i.e., **canalization**), while recipes are comparatively flexible strategies or plans that rely on environmental structure or experience to guide them.

By systematically varying task size and complexity, Downing (2004) shows that smaller tasks, as well as those with repetitive elements, are often dominated by blueprints (i.e., genetic control). Alternatively, on large tasks that lack repetitive sub-patterns, cooperative strategies become more likely, in which blueprints and recipes emerge in parallel and serve as bootstraps for each other.

A second finding is that genetic encoding tends to occur more often when the environment is predictable or stationary. Alternatively, in dynamic environments in which optimal behaviors (i.e., reward or objective functions) vary from generation to generation, genotypes tend to encode the capacity for efficient learning, rather than specific input-output behavior patterns (e.g., Nolfi & Parisi, 1997).

Finally, a third finding is that learning costs constrain the relative influences of evolutionary and developmental processes (e.g., Mayley, 1996). For example, Munroe and Cangelosi (2002) show that when learning costs are high (e.g., non-adaptive exploratory behaviors reduce fitness), genetic control of behavior predominates. Conversely, developmental processes tend to guide behavioral change when organisms are not penalized for exploring.

3.2 Developmental timing

As I noted above, a unique feature of neurogenetic models is the use of a genotype that is comprised of both structural and regulator genes. In this section, I highlight three recent neurogenetic models that employ this feature to investigate heterochrony and the role of developmental timing.

First, Cangelosi (1999) simulates a population of organisms that compete on a foraging task. The genotype of each organism includes a **genetic regulatory network** (GRN) that specifies a complex set of interactions among regulator genes. In particular, the GRN modulates the physical growth and experience-dependent changes of each organism’s neural network (e.g., cell duplication and migration, synaptogenesis, etc.).

As part of his analysis, Cangelosi systematically

compares not only the structures of the ANNs that emerge over evolution, but also how these structures vary in their developmental pattern during successive generations. This analysis reveals a number of heterochronic changes—both accelerations and retardations—that modify neural growth patterns and structure, and consequently, lead to adaptive changes in behavior. A particularly interesting example involves a simulation run in which both acceleration and retardation occur (in the growth of sensory neuron synapses), giving rise to a new behavior: in the presence of “poisonous” food, these organisms quickly learn to avoid approaching and eating the food, in contrast to their ancestors, who lacked the ability to inhibit approach.

Second, Bongard (2002) investigates an ambitious neurogenetic model in which the GRN not only encodes neural structure, but also body morphology in a population of artificial organisms. Each generation of organisms first undergoes a phase of developmental growth, including both **neurogenesis** (i.e., formation of the neural network) and **morphogenesis** (i.e., growth of a multi-segmented body). The growth phase is followed by an evaluation phase, in which individuals compete on a locomotion task.

After evolving several successful populations, Bongard (2002) describes a “heterochronic analysis” in which regulator genes are selectively modified and resulting changes in developmental pattern are identified. A key finding from this analysis is that while regulator genes play an essential role during both neurogenesis and morphogenesis, there is an evolutionary trend toward modularity. In particular, one segment of the GRN evolves into a set of regulator genes that shape neurogenesis, while another segment of the GRN regulates morphogenesis. This effectively allows the GRN to explore novel body configurations while preserving adaptive neural structures, and vice versa.

A third neurogenetic model is proposed by Matos, Suzuki, and Arita (2007). What is unique about this study is *not* the model—Matos et al. essentially reuse the neurogenetic model proposed by Gruau (1994)—but how the data are analyzed. In particular, Matos et al. (2007) use the quantitative analysis methods proposed by Alberch et al (1979) to evaluate their findings.

Matos et al. first estimate the growth of several traits, including the topology of the ANNs, and then apply Alberch’s framework to analyze the onset, offset, and growth rates of these traits. Interestingly, a variety of different heterochronic patterns are observed, including both neoteny (i.e., retardation) and **hypermorphosis** (i.e., extended growth due to a delayed offset of neurogenesis). In comparing these two forms of heterochrony, Matos et al. report that while both neoteny and hypermorphosis result in increased fitness, hypermorphosis is more common across simulation runs, and has a significantly larger effect on fitness.

4. Heterochronic robotics? Future research questions

In contrast to epigenetic and evolutionary robotics, which

have become well-established fields in the last several years, the design and study of neurogenetic models is still a comparatively new area of exploratory research. There remain a number of important theories, concepts, and empirical phenomena in evolutionary developmental biology that have not been fully explored from the computational perspective.

In this section, I highlight two potential research questions that the epigenetic-robotic approach may be well-suited to address. Of course, investigating either of these questions will likely require expanding and redefining the boundaries of epigenetic robotics, to include perspectives and techniques used by other areas of machine learning and computational modeling.

4.1 Attachment and social development

According to Bowlby's attachment theory (1969), infants develop a variety of behaviors during the first year of life that serve to help monitor and maintain proximity to caretakers. These behaviors include crying, smiling, and vocalizing, as well as locomotion. Bowlby proposes that these behaviors are part of an attachment system that has evolved in primates to promote the safety and survival of infants.

Support for Bowlby's theory is provided by Parisi, Cecconi, and Cerini (1995), who simulated a population of artificial organisms on a foraging task. A somewhat expected result from this study is that when parents find food and are given the option to either eat or share with their offspring, a tendency to share rapidly evolves (this of course benefits the offspring, which then increases the chances of survival). Surprisingly, however, Parisi et al. also note that as parents evolve the tendency to share, the offspring **co-evolve** the tendency to follow their parents (i.e., maintain proximity) during foraging!

An open question is how the evolution of proximity-seeking (and other components of the attachment system) in humans is influenced by the historical trend toward neoteny. For example, one hypothesis is that in a relatively **precocious** (i.e., developmentally-accelerated) species, the infant rapidly acquires basic survival skill, which then diminishes the need for parental support, and consequently, attachment with a caretaker. Alternatively, in an altricial (i.e., developmentally-retarded) species, a prolonged period of helplessness and dependence on a caretaker is likely to promote attachment, and in particular, a number of related social skills (e.g., reading the caretaker's emotional state).

This question can be addressed in a relatively straightforward manner, by combining neurogenetic and epigenetic-robotic approaches. First, a neurogenetic model can be used to evolve a population of artificial organisms. Second, these organisms can be designed to allow for any number of social-cognitive skills and behaviors, including **face recognition**, **gaze-following**, **affect-sharing**, **social referencing**, **theory of mind**, and so on (note that each of these is an established research topic in both developmental science and epigenetic robotics). Finally, by manipulating the developmental rates of neurogenesis and morphogenesis in the population, the effects of developmental acceleration and

retardation on attachment and social development can be systematically investigated.

4.2 Locomotion, grasping, and tool-use

A second question focuses on the development and evolution of sensorimotor skill. This question is inspired by comparative research across a wide variety of primates, including human infants, chimps, gorillas, and monkeys (e.g., Antinucci, 1989; Vauclair, 1984).

These studies employ a two-step research strategy. First, a behavioral assessment method is used to identify the developmental pattern or sequence that occurs in each species. For example, a common assessment method is Uzgiris and Hunt's (1975) ordinal scales, which are a comprehensive series of behavioral tasks that are used to measure sensorimotor development (e.g., hand-eye coordination, object manipulation, tool-use, etc.). Second, after the developmental profile (i.e., onset, rate, and extent of each stage) is generated for each species, the profiles are normalized (i.e., corrected for differences in life expectancy across species), and compared side-by-side.

A consistent finding from these comparisons is that human infants develop farther (i.e., reach higher stages of sensorimotor development) than non-human primates. Perhaps more importantly, however, when the cross-species comparison is limited to the initial stages that all primates share in common, human infants consistently have *the slowest rate of sensorimotor development*.

An intriguing hypothesis is that the onset of self-produced locomotion (i.e., crawling and walking) is a rate-limiting factor on the development of grasping and tool-use (e.g., Antinucci, 1989; Vauclair, 1984). In particular, when macaques, chimpanzees, gorillas, and human infants are compared, the first group to develop locomotion is macaques (at 7 days), while human infants are the last (at 255 days).

According to this hypothesis, because human infants are "flat on their backs" for more time than non-human primates, they may become more skilled at grasping, object manipulation, and ultimately, the use of intermediaries (e.g., strings, sticks, etc.) as a means to reach their goals. Alternatively, because other primates quickly acquire the develop the ability to locomote, they spend less time learning to manipulate nearby objects, as well as using intermediaries to reach distant objects.

As before, this hypothesis can be tested by simulating a population of artificial organisms that develop the capacity to **reach**, **grasp**, and **manipulate objects**, as well as **locomote**. It should be acknowledged, however, that while each of these skills is an ongoing area of study in epigenetic robotics, none of them are trivial tasks. Nevertheless, since it is the relative timing of the skills that is critical, rather than the realism and complexity of the skills as they are modeled *per se*, a pragmatic research strategy would be to design the model so that each behavior is appropriately scaled (i.e., pre-structured, modularized, etc.). The focal question, then, is whether evolutionary delays in the onset of locomotion will result in developmental advances in sensorimotor skill.

5. Summary

As I noted in the introduction, researchers who adopt the epigenetic-robotic approach are committed to the idea that development is the result of the interaction between biological and environmental influences. In the current paper, I have argued for an expansion of the notion of “biological” to include not only an organism’s physical body and nervous system, but also its genes and the cumulative experience of its evolutionary past.

I began by reviewing the concept of heterochrony, which plays an important role in the field of evolutionary developmental biology. Speeding up or slowing down the developmental process is one of the fundamental strategies that evolution uses to create new structures and behaviors. In particular, the human species may have evolved through the progressive slowing of development during our ancestors’ juvenile period. While this slower rate of development comes with a cost (e.g., increased parental investment of time and energy), there are numerous benefits that offset the cost: a larger, more complex brain; more advanced perceptual, cognitive, and social development; language; and so on.

In addition to the wide array of computational tools and models that are available to study heterochrony, a relatively new approach is offered by neurogenetic models that use a GA to simulate the evolution and development of ANNs. Neurogenetic models have not only provided an ideal research tool for identifying interactions between evolution and development, but also for investigating heterochrony in particular.

In the previous section, I outlined two specific research questions that can be addressed by employing a neurogenetic model within the epigenetic-robotic framework. It is important to stress that while these questions could be studied using a neurogenetic model alone, epigenetic robotics adds two essential elements. First, while traditional machine-learning approaches tend to focus on performance-related outcomes, a unique feature of epigenetic robotics is a view that recognizes the importance of studying both developmental *outcomes* and developmental *patterns*.

Second, epigenetic robotics also provides a rich platform for studying behavior as it unfolds in realtime. As Lungarella et al. (2003) note in their review, most epigenetic-robotic research is concerned with “competence in interacting with the local environment—in particular, basic visuo-motor competencies such as saccading, gaze fixation, joint attention, hand-eye coordination and visually-guided reaching” (pp. 175-176). This emphasis on the development of sensorimotor skill is also a unique strength of epigenetic robotics, insofar as it not only gives the neurogenetic model a body to control and learn from, but also a physical (and perhaps also social) environment to interact with, and a diverse range of behavioral skills to develop.

I conclude by noting that, on occasion, developmental researchers may resort to describing an innate behavior as “hard-wired”. The analogue of this in epigenetic robotics is a behavior that is “hand-built” (hand-coded, etc.). To the degree that this phrase is a promissory note—that is, that someday the researcher will return to the problem and

explain how the capacity “gets into the genes” in the first place—the use of neurogenetic models and the study of heterochrony offer the means to move beyond promissory notes and actually answer the question.

References

- Alberch, P., Gould, S.J., Oster, G.F., & Wake, D.B. (1979) Size and shape in ontogeny and phylogeny. *Paleobiology*, 5, 296–317.
- Antinucci, F. (1989). Systematic comparison of early sensorimotor development. In F. Antinucci (Ed.), *Cognitive structure and development in nonhuman primates* (pp. 67-85). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Baldwin, J.M. (1896). A new factor in evolution. *American Naturalist*, 30, 441.
- Bjorklund, D.F. (1997). The role of immaturity in human development. *Psychological Bulletin*, 122, 153-169.
- Bongard, J.C. (2002) Evolving modular genetic regulatory networks. *Evolutionary Computation*, 2, 1872-1877.
- Bongard, J.C., & Pfeifer, R. (2001) Repeated structure and dissociation of genotypic and phenotypic complexity in artificial ontogeny. In L. Spector et al. (Eds.), *Proceedings of the Genetic and Evolutionary Computation Conference* (pp. 829-836). San Francisco, CA: Morgan Kaufmann.
- Bowlby, J. (1969). *Attachment*. New York: Basic Books, Inc.
- Cangelosi, A. (1999). Heterochrony and adaptation in developing neural networks. In W. Banzhaf et al. (Eds.), *Proceedings of the Genetic and Evolutionary Computation Conference* (pp 1241–1248). San Francisco, CA: Morgan Kaufmann.
- Cangelosi A., Parisi D., & Nolfi S. (1994). Cell division and migration in a 'genotype' for neural networks. *Network: Computation in Neural Systems*, 5, 497-515.
- Downing, K.L. (2004). Development and the Baldwin Effect. *Artificial Life*, 10, 39-63.
- Elman, J.L. (1993). Learning and development in neural networks: The importance of starting small. *Cognition*, 48, 71-99.
- Gould, S.J. (1977). *Ontogeny and phylogeny*. Cambridge, MA: Harvard University Press.
- Gruau, F. (1994). Automatic definition of modular neural networks. *Adaptive Behavior*, 3, 151-183.
- Gruau, F., & Whitley, D. (1993). Adding learning to the cellular development of neural networks: Evolution

- and the Baldwin Effect. *Evolutionary Computation*, 1, 213-233.
- Hinton, G.E., & Nowlan, S.J. (1987). How learning can guide evolution. *Complex Systems*, 1, 495-502.
- Husbands, P., Harvey, I., Cliff, D., & Miller, G. (1997). Artificial evolution: A new path for artificial intelligence? *Brain and Cognition*, 34, 130-159.
- Kitano, H. (1990). Designing neural networks using genetic algorithms with graph generation system. *Complex Systems*, 4, 461-476.
- Lungarella, M., Metta, G., Pfeifer, R., & Sandini, G. (2003). Developmental robotics: A survey. *Connection Science*, 15, 151-190.
- Matos, A, Suzuki, R., & Arita, T. (2007). Heterochrony and evolvability in neural network development. *Artificial Life Robotics*, 11, 175-182.
- Mayley, G. (1996). Landscapes, learning costs, and genetic assimilation. *Evolutionary Computation*, 4, 213-234.
- Mitteroecker, P., Gunz, P., Bernhard, M., Schaefer, K., & Bookstein, F.L. (2004). Comparison of cranial ontogenetic trajectories among great apes and humans. *Journal of Human Evolution*, 46, 679-698.
- Morell, V. (1997). The origin of dogs: Running with the wolves. *Science*, 13, 1647-1648.
- Munroe, S., & Cangelosi, A. (2002). Learning and the evolution of language: The role of cultural variation and learning costs in the Baldwin effect. *Artificial Life*, 8, 311-330.
- Nolfi, S., & Floreano, D. (2001). *Evolutionary robotics. The biology, intelligence, and technology of self-organizing machines*. Cambridge, MA: MIT Press.
- Nolfi, S., & Parisi, D. (1997). Learning to adapt to changing environments in evolving neural networks. *Adaptive Behavior*, 5, 99-105.
- Parisi D., Cecconi F., Cerini A. (1995). Kin-directed altruism and attachment behaviour in an evolving population of neural networks. In N. Gilbert & R. Conte (Eds.), *Artificial societies. The computer simulation of social life* (pp. 238-251). London: UCL Press.
- Schlesinger, M., Parisi, D., & Langer, J. (2000). Learning to reach by constraining the movement search space. *Developmental Science*, 3, 67-80.
- Stanley, K.O., & Miikkulainen, R. (2003). A taxonomy for artificial embryogeny. *Artificial Life*, 9, 93-130.
- Uzgiris, I, & Hunt, McV. (1975). *Assessment in infancy: Ordinal scales of development*. Urbana, IL: University of Illinois Press.
- Vauclair, J. (1984). Phylogenetic approach to object manipulation in human and ape infants. *Human Development*, 27, 312-328.
- Weng, J., McClelland, J., Pentland, A., Sporns, O., Stockman, I., Sur, M., & Thelen, E. (2001). Artificial intelligence: Autonomous mental development by robots and animals. *Science*, 291, 599-600.