

Ongoing Emergence: A Core Concept in Epigenetic Robotics

Christopher G. Prince, Nathan A. Helder, & George J. Hollich*

Department of Computer Science
University of Minnesota Duluth
Duluth, MN 55812 USA
chris@cprince.com, nhelder@nerp.net

*Purdue University
Department of Psychological Sciences
West Lafayette, IN 47907 USA
ghollich@psych.purdue.edu

Abstract

We propose *ongoing emergence* as a core concept in epigenetic robotics. Ongoing emergence refers to the continuous development and integration of new skills and is exhibited when six criteria are satisfied: (1) continuous skill acquisition, (2) incorporation of new skills with existing skills, (3) autonomous development of values and goals, (4) bootstrapping of initial skills, (5) stability of skills, and (6) reproducibility. In this paper we: (a) provide a conceptual synthesis of ongoing emergence based on previous theorizing, (b) review current research in epigenetic robotics in light of ongoing emergence, (c) provide prototypical examples of ongoing emergence from infant development, and (d) outline computational issues relevant to creating robots exhibiting ongoing emergence.

1. Introduction

Epigenetic robotics is a new field that focuses on modeling cognitive development and creating robots that show autonomous mental development (Lungarella, Metta, Pfeifer, & Sandini, 2003; Weng, McClelland, Pentland, Sporns, Stockman, Sur, & Thelen, 2001). For example, robots have been implemented that generate visual discrimination behavior using large-scale neural networks (Seth, McKinstry, Edelman, & Krichmar, 2004), that model early infant-caregiver interaction using behavioral rules (Breazeal & Scassellati, 2000), and that explore the knowledge needed by infants to succeed in perceptual object permanence experiments (Chen & Weng, 2004; Lovett & Scassellati, 2004; see also: Schlesinger & Casey, 2003). Given these and other diverse contributions to this new field (for a review, see Lungarella et al., 2003) it seems an opportune time to synthesize a few core concepts from this corpus of research.

In this paper, we distill one such core concept, *ongoing emergence*, which refers to the continuous development and integration of new skills. An agent exhibiting ongoing emergence, in a motivationally autonomous manner, will continue develop and refine its skills across development. This vision of open-ended development is evident in recent work. For example, in efforts to “allow a mobile robot to incrementally progress through levels of increasingly sophisticated behavior” (p. 1, ms., Blank, Kumar, Meeden, & Marshall, 2005), in

efforts to build robots that exhibit “new behavior, which in turn, becomes a precursor for successive stages of development” (p. 27, Grupen, 2003), and in efforts to achieve robots exhibiting a “successive emergence of behaviors in a developmental progression of increasing processing power and complexity” (p. 1, ms., Dominey & Boucher, 2005). Unfortunately, while humans clearly show such long-term progressions, epigenetic robots as yet do not—they are typically designed to achieve particular behaviors or to learn specific tasks.

To escape this impasse, we propose a theoretical framework for achieving ongoing emergence. To this end, in Section 2 we review previous theoretical conceptions regarding ongoing emergence and synthesize the current state of the art in terms of six criteria. Section 3 considers how current examples of robotic systems fare with respect to these criteria for ongoing emergence. In Section 4, we look to infant developmental research for examples of ongoing emergence. Section 5 outlines some computational issues for designing robots that exhibit ongoing emergence. We close with a discussion.

2. Conceptual Synthesis

2.1. Background

Blank et al. (2005) discuss the possibility that a robot can use a *developmental algorithm* to learn, via a process of self-exploration, its repertoire of behaviors and mental capabilities, instead of being preprogrammed with “the capabilities of a human body and human concepts” (p. 2, ms). Robots are proposed to discover even the most primitive behaviors through a process of exploration.

A possible benefit of providing such a developmental algorithm is avoiding specification of task-goals for the robot. Instead, “it is the goal of developmental robotics to explore the range of tasks that can be learned (or grown) by a robot, given a specific developmental algorithm and a control architecture” (p. 2, ms). These authors consider three mechanisms to be essential to developmental algorithms: abstraction, anticipation, and self-motivation. Abstractions are seen as necessary to focus the robot’s attention on relevant environmental features, given the “constant stream of perceptual information” (p. 2, ms.). Anticipations enable the robot to predict environmental change to “go beyond simple reflexive behavior to purposeful behavior” (p. 3, ms.). And self-motivation “push[es] the system toward further abstractions and

more complex anticipations” (p. 3, ms.). It is thought that a developmental algorithm incorporating these three mechanisms could be successively applied to move an agent from a discovery of initial behaviors (“reflexes”) to more complex behaviors.

Weng (Weng, 2004; Weng et al., 2001) also emphasizes the need for robots to autonomously generate their own task-specific representations in order to cope with dynamic, unknown, or uncontrolled environments. “A developmental program for robots must be able to generate automatically representations for unknown knowledge and skills” (Weng et al., 2001) so as to adapt to these environmental variations. An agent with the capacity to construct its own representations has the potential of understanding these representations. “Without understanding, an agent is not able to select rules when new situations arise, e.g. in uncontrolled environments” (p. 205, Weng, 2004). These processes are viewed as open-ended and cumulative. “A robot cannot learn complex skills successfully without first learning necessary simpler skills, e.g., without learning how to hold a pen, the robot will not be able to learn how to write” (Weng et al., 2001).

Gruppen (2003) is similarly concerned with enabling robots to solve “open tasks in unstructured environments” (p. 2, ms.). The approach he advocates is to use “developmental processes [that] construct increasingly complex mental representations from a sequence of tractable incremental learning tasks” (p. 1, ms.). He proposes “computational mechanisms whereby a robot can acquire hierarchies of physical schemata” (Gruppen, 2003, p. 1, ms.). Physical schemata provide parameterized, and in that sense reusable, sensorimotor control knowledge.

Dominey and Boucher (2005) model linguistic grammar acquisition, based on visual and auditory pre-processing of sensory inputs and connectionist models. The authors use the developmental theory of Mandler (1999), who “suggested that the infant begins to construct meaning from ... scene[s] based on the extraction of perceptual primitives. From simple representations such as contact, support, and attachment ... the infant [may] construct progressively more elaborate representations of visuospatial meaning” (p. 244, Dominey & Boucher, 2005).

2.2. Synthesis

From this earlier thinking, we wish to synthesize a picture of what we refer to as ongoing emergence. We propose six defining criteria for ongoing emergence (see Table 1). Our first two criteria are: (1) An agent creates new skills by utilizing its current environmental resources, internal state, physical resources, and by integrating current skills from the agent’s repertoire, and (2) These new skills are incorporated into the agent’s existing skill repertoire and form the basis from which further development can proceed. By “skills” we include overt behaviors, perceptual abilities, and internal representational schemes.

These first two criteria express the notion that when we view agents as developing systems, with certain skills

in their repertoire, they have the potential to develop related skills. For example, under this view a developmental robot that can learn to kick a ball might then later develop skills for playing soccer. Ongoing emergence thus has the property of *developmental systematicity*¹. In developmental systematicity if an agent demonstrates skill *aRb*, then we also expect competence with directly related skills, *bRa* (i.e., systematicity; Fodor & Pylyshyn, 1988). Furthermore, we expect the emergence of developmentally related skills such as *f(a)* and *g(aRb)*, where *f(x)* and *g(y)* are developmental processes producing emergent skills in the agent’s repertoire over time. This process would in part be based on earlier skills (e.g., *x* and *y* in *f(x)* and *g(y)* above). For example, if a robot exhibits a range of object tracking behaviors (*aRb*, *bRa*) through the composition of blob tracking skill (*a*) and motion finding skill (*b*), and the robot is a developing agent, we would have further, developmental expectations about its future behaviors such as facial tracking and gaze following (e.g., *f(a)*, *g(aRb)*).

Another central notion in the work described in Section 2.1 is that of autonomy: avoidance of specification of task goals, autonomous generation of task-specific representations, and the ability to solve open tasks. We include this as a third criterion for ongoing emergence: (3) An agent that exhibits ongoing emergence autonomously develops adaptive skills on the basis of having its own values (e.g., see Sporns, 2005; Sporns & Alexander, 2002) and goals, with these values and goals being developed by the system in a manner similar to its skills. If an agent develops its own values and goals, it can use these for self-supervision and to determine the tasks that need to be solved. In brief, the agent needs some way to evaluate its own behaviors, and determine when a particular skill is useful². This is true in both the short and long term. For example, in the short-term, a robotic agent might tradeoff energy output for the gain of information, while long-term goals might include improving communication amongst the robot’s cohorts.

To these initial three criteria for ongoing emergence we add three additional criteria: (4) bootstrapping (when the system starts, some skills rapidly become available), (5) stability (skills persist over an interval of time), and (6) reproducibility (the same system started in similar initial states and in similar environments also displays

¹ We introduce the concept of developmental systematicity to avoid viewing behavior generation an infinite domain. This is analogous to the way that Fodor & Pylyshyn (1988) introduced systematicity to avoid viewing language generation as an infinite domain.

² We refrain from adopting the idea that skills that emerge in development should necessarily be more complex (i.e., be more powerful in some sense) than prior skills. From our view, this criterion is too strong for several reasons. First, strictly increasing adaptation is violated in some instances of child development (e.g., the “U” shaped curves of child performance on various tasks over time; see Siegler, 2004). Second, a view of strictly increasing complexity of skills does not allow for escape (“detours”) from local maxima, where behavior needs to get worse before it can get better. Third, strictly increasing skill complexity may remove the possibility of discovering simpler means to achieve the same (or similar) ends as existing skills—as in evolution, “different” is sometimes at least as good as “better.” Relatedly, a strict view of building complexity does not seem to allow for the loss (e.g., forgetting) of some skills over time.

similar ongoing emergence). We include bootstrapping as a criterion for ongoing emergence because it seems inevitable that a robot either needs to have some means of spontaneously developing its own set of initial skills or, more conventionally, will need to have some initial skills pre-programmed prior to its being turned on. While pre-programming of bootstrap skills is not consistent with the concept of skills being developed by the agent itself, we consider this an acceptable practice if for no other reason than keeping the scope of research projects tractable. However, in our view, the preferred method to establish bootstrapping skills is to represent those skills in the same manner as later emerging skills such that both the bootstrapping and developed skills comprise a uniform part of the agent's skill repertoire.

Stability of skills is in part a practical matter: in order for a skill to be measured (i.e., by researchers), it must exist for a measurable duration. In terms of the robot, however, stability may be more than merely a practical matter in that, in order for ongoing emergence to be achieved, a certain degree of skill stability over time will likely prove necessary. If the behaviors exhibited by the robot are merely transient then those behaviors may not contribute to the basis for the acquisition of new skills.

The reproducibility criterion asks the question: Under what starting-state and environmental conditions does a given developmental algorithm produce an ongoing emergence of behavior? We presume that if a developmental algorithm is well-understood, then the starting-state and conditions under which it produces ongoing emergence will also be well-understood. These conditions may be limited (e.g., to specific values of initial variables), but once known can reproducibly generate ongoing emergence.

3. Current Research & Ongoing Emergence

In this section, we review examples of epigenetic robotic research (see Table 2) in light of our criteria for ongoing emergence. Our selection of these particular papers is not intended to reflect some *a priori* sense that they have achieved ongoing emergence. Rather they simply reflect our subjective impression of good illustrative examples of research in this area.

Several lines of research satisfy Criterion 1 (new skill acquisition). For example, a swinging behavior emerges in the robot of Berthouze and Lungarella (2004), and the skill of tracking a face view to an object emerges in the robot of Nagai et al. (2003). To varied extents, some research has also satisfied Criteria 3 through 6. Criterion 3 (autonomy of goals and values) is satisfied to some degree by the robot of Seth et al. (2004), and also the work of Kaplan and Oudeyer (2003). Seth et al. (2004) utilize a value system in the Darwin VIII robot to signal the occurrence of salient sensory events. Initially, Darwin VIII's value system was activated by sounds detected by the robot's auditory system, but through learning became activated by particular visual stimulus attributes. Criterion 4 (bootstrapping) is satisfied by the Dominguez and Jacobs (2003) system in that the system

uses progressive changes in visual acuity to increase its binocular disparity sensitivity. Criterion 5 (stability) appears satisfied, for example, by the Lungarella and Berthouze (2002) system in that the robots' swinging behavior reaches stable states. Criterion 6 (reproducibility) is satisfied by studies which replicated their robots' behavior, perhaps under varied conditions. For example, Chen and Weng's (2004) experiments were replicated with 12 separate robot "subjects" (the same robot and algorithms, but with different environmental conditions).

To give a more extended example of how these criteria can be applied, we consider the work of Nagai et al. (2003). These authors modeled joint visual attention behavior using a robot. Joint attention occurs when individuals both look at the same object, and may include knowledge of shared attentional states (e.g., Carpenter, Nagell, & Tomasello, 1998). The Nagai et al. (2003) robot learned to track the face view of a person to the object the person was looking at. Learning started with the robot (a) knowing how to visually detect faces and salient objects, (b) knowing to switch its gaze from a face to an object when both the object and face were in its field of view, and (c) having a predefined transition function (a sigmoidal) to switch from how salient objects were found—either directly in its visual field, or indirectly by first looking at a face³. Initially the robot did not know how much to turn its head based on a particular view of a face to find the object that the person was looking at, and learning acquired this skill. The transition function enabled this skill to be gradually applied.

The Nagai et al. (2003) robot has some behaviors that are programmed into the system (e.g., bootstrapping, Criterion 4). A new behavior is constructed from the initial behaviors (e.g., visually detecting faces and objects) and environmental interaction—i.e., the robot learns to track faces to the objects that they are looking at (Criterion 1 is satisfied). However, once the new behavior has emerged, there is no further potential for development. That is, the new skill was not incorporated into the system in such a way that it contributed to the basis for further skill development. Thus, Criterion 2 is not satisfied. Reproducibility (Criterion 6) was demonstrated in the system by conducting experimental runs with 1, 3, 5, or 10 objects in which the emergent behavior was maintained. In summary, while some of the criteria for ongoing emergence are satisfied, the behavior of the Nagai et al. (2003) robot seems best classified as demonstrating emergence and not ongoing emergence.

Notably absent in this review of current work is full evidence for Criterion 2 (incorporation of new skills with existing skills so that those new skills can be used as part of the basis for further development). We have yet to find examples of robots exhibiting this ability (but we hope to be corrected on this point!). This leads us to view current

³ The use of this sigmoidal seems unnecessary, but in our view should not be viewed as a shortcoming of this research. In principle, the authors could have used a method of self-supervised learning to transition between modes: when the robot was sufficiently able to predict the amount of head turn required for accurately turning to face an object, it could have then begun utilizing its self-generated head turn.

examples of epigenetic robots (e.g., see Table 2) as demonstrating emergence, but not ongoing emergence.

4. Human Infant Developmental Examples

In contrast to the state of the art in epigenetic robotics, human infants clearly exhibit ongoing emergence. Development is an unending process that continually produces new skills by making use of currently available skills, environmental conditions, and other resources. In this section, we provide prototypical examples of ongoing emergence in infants from three developmental areas: walking, language, and visual object skills.

4.1 Emergence of Walking

As any one-year-old would acknowledge, walking is more difficult than it may first appear. To properly walk, children must achieve the right mix of balance, head control, and coordinated oscillation of the limbs that have thousands of muscle fibers and billions of nerves as well as their own length, mass and transitory inertia. The degrees of freedom for the task are enormous (Bernstein, 1967).

Further complicating this process, children grow physically. They begin life top-heavy, which makes stabilizing this system all the more difficult. Their growth is also erratic and dramatic—children can go up to 63 days with no measurable change in height and then suddenly grow up to 2.5cm in a single night (Lampl, Veldhuis, & Johnson, 1992). As if this wasn't difficult enough, children must learn to navigate different slopes and uneven terrain, perhaps while carrying objects (Adolph & Avolio, 2000). Yet somehow nearly all children learn to walk, and continue to walk, despite the complexity of achieving the right mix of skills, changes in body morphology, and varying situations.

Current theory (Thelen, 1995) views this process as a dynamic self-organizing system in which integration of diverse skills plays a key role. Because the world, the task, and even children's bodies are constantly changing, each component is constantly being weighted differently, as dictated by the interaction of the nervous system and the environment. For example, while all infants possess a stepping and a kicking reflex at birth, the stepping reflex disappears after a few months. Why? In short, babies don't have the strength to keep up this reflex as they grow heavier—even though the nervous system is still sending the signals. Stepping and, by extension, walking must wait for stronger muscles to grow before infants can take their first steps. If one makes the task easier, by supporting the infants (on a treadmill or underwater), even newborns can walk (Thelen & Fisher, 1982). In contrast, if one makes the task harder by placing weights on older infants, their walking again approximates that of younger infants (Thelen & Fisher, 1982). Thus, it is the dynamic interaction between current skills, the state of the system, and the environment that allows for walking behavior to self-organize into coherent patterns across changes in morphology and task. This illustrates Criterion

1 for ongoing emergence—namely that skills are created through the integration of resources including environment and existing skills.

4.2 Emergence of Language

While purely physical skills like walking show ongoing emergence, skills that are more cognitive also require the use and integration of multiple developing skills. For example, word learning can be seen as the product of social skills (e.g., sensitivity to eye gaze), domain-general skills (e.g., sensitivity to statistical structure such as invariances), and linguistic skills (e.g., a bias toward labeling objects based on shape).

Just as in walking, the weight placed on each of these skills likely changes with time and situation. In the beginning, infants may depend on a range of perceptual biases and statistical relations to establish the meaning of each new word (Hollich, Hirsh-Pasek, & Golinkoff, 2000). However, as more words are learned, children use their knowledge of known words to help them learn additional words. This illustrates another property (Criterion 2) of systems exhibiting ongoing emergence: incorporation of new skills into a skill repertoire. For example, Smith (1999) provides evidence that infants may notice how particular types of words get extended (e.g., nouns are generalized, a.k.a. extended, to different objects on the basis of shape). Infants then use this knowledge to guide their own extensions of novel words. When told a U-shaped object is a “dax,” infants will spontaneously extend that word to other U-shaped objects. Even so the system is flexible—infants will not extend a word based on shape if the object happens to have eyes (or even shoes), suggesting that children have noticed that living creatures can often change their overall shape in ways that static objects do not.

Related to the use of multiple skills, the more skills that an agent can bring to bear, the more fault-resistant and flexible the system. Loss of one skill does not cause the system to fail entirely, and the interaction among skills insures that children can successfully acquire a language under extremely impoverished conditions. For example, even deaf children growing up in an area without exposure to any fully formed language will create their own language (Sengas, 1995). With both biological and robotic systems, more pathways to success imply greater adaptivity and increased likelihood of organism survival. With regular upheaval at the neurological and muscular levels, it is no wonder that developmental architectures are massively fault tolerant with multiple, redundant skills. Thus, the self-organization of new skills combined with an increasing skill repertoire can lead to a process of ongoing emergence.

4.3. Object Skill Developments

At the same time that human infants are developing the walking and language acquisition skills discussed above, they also show an ongoing emergence of physical and mental capabilities related to visual objects. Starting from birth, infants are able to extract information about object

size and shape (Slater & Morison, 1985), remember objects over time (Slater, Morison, & Rose, 1982), and perceive similarities and differences between visual stimuli (Slater, Morison, & Rose, 1984). Newborn infants are also able to track a moving target with eye and head movements (albeit in a jerky fashion, e.g., Aslin, 1981; von Hofsten 1982), and can recognize the constancy of an object's identity across transformations in orientation and movement (Slater, Morison, Town & Rose, 1985).

While constituting a perhaps surprisingly robust set of initial skills, developing and incorporating these skills into more complex behaviors takes time. For example, it is not until 4 months of age that an infant's muscular control and object understanding have matured to the point of allowing an infant to successfully reach for and grasp an object (e.g., von Hofsten, 1989). Also at 4 months, infants begin to perceive (measured via looking-time) a partially occluded object as a single unified object (Kellman & Spelke, 1983; Johnson & Nanez, 1995). However, it is not until about 6 months of age that infants combine their object tracking skills, their understanding of object unity, and their reaching skills to reach for an object that has been partially obscured from view by an occluding object (von Hofsten & Lindhagen, 1979; Shinskey & Munakata, 2001).

Also in the realm of visual-object skills is object permanence, which relates to the child's understanding that an object continues to exist even when the object cannot be seen. It has been shown that 3.5-month-old infants show recognition of an impossible object event (i.e., a violation of object permanence), such as a drawbridge closing despite a solid object appearing to have been blocking its path (Baillargeon, 1987, 1993, 1995). However, this sort of "perceptual object permanence" is not manifested as a behavior indicating an understanding of physical (i.e., more conventional) object permanence until much later, when 8- to 10-month-old infants will begin to search for an object that has been hidden from view (Piaget, 1954). Still, infant searching behavior at this age is not free of difficulties and is subject to the "A-not-B error" (the infant searches for a hidden object at location A when the object was initially uncovered at location A but subsequently hidden at location B). Infants persevere in this error until roughly 12 months of age (at which time infants will correctly search for the hidden object at location B; e.g., see Wellman, Cross, & Bartsch, 1986; Newcombe & Huttenlocher, 2000).

In developing from initial skills of being able to identify and track objects (birth), to perceptually distinguishing impossible object events (3.5 months), to being able to maintain perception of object unity despite an occlusion (4 months), to successfully reaching for an object (4 months), to successfully reaching for an object despite an occlusion (6 months), to searching for a hidden object (8-10 months), to searching for a hidden object without displaying the A-not-B error (12 months), infants demonstrate an ongoing emergence of behavior. Changes occurring in the visual, conceptual, and motor systems of the infants interact to produce unique, observable behaviors at multiple points along the developmental path of these visual object skills, with each developed skill

being incorporated and providing a contributing factor to the emergence of subsequent skills.

5. Designing For Ongoing Emergence

Past a theoretical understanding of ongoing emergence, our most burning question was well-expressed by one of the anonymous referees of this paper: How can we design robots so that the behaviors exhibited by the robot continue to be adaptive and open to further development throughout their duration of use (e.g., either as models of infants, or deployed in some industrial environment)? That is, how do we design robots that exhibit ongoing emergence? Our thinking here divides broadly into two possibilities. The first possibility we address is that of designing robots that exhibit ongoing emergence where the bootstrapping components (see Criterion 4) of the system are not generated by ongoing emergence. Effectively, this corresponds to basing the design of the robots on existing research (e.g., the robotic systems in Table 2). The second possibility we address is that of designing robots that exhibit ongoing emergence where the bootstrapping components themselves are generated by processes of ongoing emergence. This corresponds to discovering a different way of approaching the design of the initial components of a robotic system.

5.1. Bootstrapping Ongoing Emergence Without Primitive Ongoing Emergence

Ongoing emergence in humans results in part from the dynamic integration of multiple skills with the environment (i.e., Criterion 1). One way to achieve an analog of this in robots may be to add an integration layer on top of an existing system or systems (see Table 2), providing soft-assembly of the component skills. For example, we might combine robotic behaviors across several systems, such as the perceptual object permanence behavior of Cheng and Weng (2004), the joint attention of Nagai, et al. (2003) and the social skills of Breazeal and Scassellati (2000). For this integration layer to satisfy Criterion 1, it would be appropriate for these skills, in interaction with the environment, to produce new, adaptive, emergent skill(s). For example, given the integration of the above three prior research projects, the integrated system might express surprise towards a caregiver when an object permanence situation was violated.

An approach that might be useful for this integration layer was given by Cheng, Nagakubo, and Kuniyoshi (2001). These authors proposed an integration mechanism to combine components in a humanoid robotic system, involving integrating the results of various component mechanism, which themselves show adaptation over time. Combining components involves weighting the components for their relative contributions, and such contributions may vary according to factors such as learning and context. The authors use a sensory-cue competition approach to integration, and generating motor outputs. They define the motor output of a robot as the vector $U_i(t)$, expressed by equation [1]. $U_i(t)$

gives the output for motor sub-system i (e.g., a head control motor), at time t .

$$U_i(t) = \frac{\sum_k \alpha_k(t) a_k(t) v_k(t)}{\sum_k a_k(t)} \quad [1]$$

In equation [1], $v_k(t)$ is a vector giving the current sensory input from sensory subsystem k (e.g., a joint angle or a camera) at time t , $a_k(t)$ is a measure of the reliability (confidence) of sensory subsystem k (a scalar), and $\alpha_k(t)$, defines the strength (e.g., priority) of a particular sensory input (also a scalar).

In a perceptual context, the weighted component integration (or *democratic integration*) algorithm forwarded by Triesch and von der Malsburg (2001) offers similar ideas, and presents a more detailed investigation of the integration concept than Cheng et al. (2001). In Triesch and von der Malsburg (2001), a group of perceptual components such as motion, color, and shape detection are adapted both in terms of the weighting of the components and in terms of prototypes for the perceptual components.

Unfortunately, the integration mechanisms proposed by both Cheng et al. (2001) and Triesch and von der Malsburg (2001) do not focus on or provide specific means of incorporating skills that result from the process of integration—leaving Criterion 2 unsatisfied. Two additional computational mechanisms would seem needed past Cheng et al. (2001) and Triesch and von der Malsburg (2001) in order to provide skill incorporation. First, the system needs a (at least implicit) means of determining that a soft-assembled skill is re-occurring. That is, the system needs a way to determine when that skill should be considered “stable” (Criterion 5). Second, once stable, the skill needs to be represented in a manner similar to the existing skills. This last part, at least in terms of this present scenario seems particularly difficult. We have been working from the design premise of utilizing current results from epigenetic robotics to form the bootstrapping components of a robot, intended to display ongoing emergence. But, in this case, there is no particular means to add to this static collection of bootstrapping skills. Presumably, a computational mechanism would be needed to learn the aspects of the new, now-stable soft-assembled skill. This new skill would, after this learning, be part of the repertoire of the system and with the other skills would form the basis for further development (i.e., it would then be termed *incorporated*; Criterion 2).

5.2. Bootstrapping Ongoing Emergence With Primitive Ongoing Emergence

A possible limitation of adopting the strategy proposed above is that one may miss common underlying mechanisms that helped create the individual skills in the first place. That is, in the ideal case, the goal would be to create a robot that exhibits ongoing emergence, where the bootstrapping primitives themselves are emergent. Thus,

in this ideal case the bootstrapping primitives are generated by the same processes that underlie subsequent skill development.

This presents a rather different problem than in Section 5.1. On the one hand, a robot that has a pre-programmed set of behaviors can presumably exhibit those behaviors (e.g., in a sequence, or through a blending of behaviors, such as shown in Breazeal, Buchsbaum, Gray, Gatenby, & Blumberg, 2005), but is in need of mechanisms to incorporate stabilized soft-assembled behaviors into its skill repertoire. On the other hand, a robot without a pre-programmed set of behaviors, in addition to needing mechanisms to provide ongoing emergence itself, is in need of an initial set of skills—it needs initial means of perceiving, representing, and behaving.

One conceptual way that such initial—and emergent—skills might be created is through self-exploration. A number of authors in epigenetic robotics have suggested the need for some form of “self” in these robotic systems. For example, Weng (2004) proposes that developing robots must be SASE—Self-Aware and Self-Effecting agents, Blank et al. (2005) talk about self-motivation and self-organization, and Steels (2004) suggests that robotic agents should self-regulate their build-up of skills and knowledge as a way to increase their rate of learning. In the present context of bootstrapping a developing agent without pre-programmed skills, self-exploration could be used to facilitate differentiation between self and other (e.g., see Michel, Gold, & Scassellati, 2004), which is important because such a developing agent would likely need to figure out what parts of its “environment” are part of the agent (e.g., its own limbs) versus part of the external world. We hypothesize that the basic properties of ongoing emergence (i.e., Criterion 1 through 6) could also provide the basis for these self-other discrimination skills, and hence can provide the means to bootstrap the skills of a developing robot.

6. Discussion

The foregoing has been a largely theoretical discussion of ongoing emergence. Ongoing emergence describes, in brief, behavioral growth in humans and (hopefully, in the future) in robots. If we have achieved our goal, this paper will stimulate further theoretical and empirical research towards these ends. We hope that this is but one of many entries to follow in the continuing discussion of behavioral growth in robots. In closing this paper, we want to argue for a relation between ongoing emergence and theorizing in cognitive science, we discuss additional means by which ongoing emergence may be achieved incrementally in epigenetic robotics research, and we close with a view to the future.

In Section 5 we raised a distinction between using pre-programmed initial skills (Section 5.1) and not using pre-programmed initial skills but instead relying purely on the properties of ongoing emergence (Section 5.2). In the pre-programmed initial skills case, we take this to be analogous to Fodor’s concept of modularity (Fodor,

1983). We take this position because the amount of interaction between the components will be limited and because the components show limited development. This provides another way to view the Section 5 alternatives: ongoing emergence through separate modules versus ongoing emergence through “modules” that develop.

It seems crucial to establish methodological ways to achieve research progress in ongoing emergence. While we have implicitly offered some ideas to this end in the body of the paper, three further ideas come to mind. First, it seems conceptually possible that ongoing emergence could be exhibited strictly within particular domains. For example, a robot might exhibit ongoing emergence only in its language and communication skills, or only in its object manipulation skills. Second, it also seems conceptually possible that ongoing emergence may be achieved in a primarily perceptual manner. We feel justified in part for this statement by the productivity of psychological methods with infants that have focused largely on the development of perceptual knowledge (e.g., Baillargeon, 1995; Hollich et al., 2000). Third, a potentially useful research step towards a full sense of ongoing emergence may be a linear emergence of a limited number of skills. In this case, a *single* skill would emerge, that skill would then be incorporated into the robot's existing skill repertoire, and then this new pool of skills would be used to develop *one* additional skill.

In closing, we recollect the statements of György Gergely, in his invited address at EpiRob 2003. György suggested that “recent research in epigenetic robotics has been strongly preoccupied with and [has] made significant advances towards modeling the ‘lower level’ mechanisms and ‘bottom-up’ processes involved in systems of action perception and production and the ways in which these systems [may be] inherently interrelated” (p. 192, Gergely, 2003). Clearly, with goals including modeling cognitive development, epigenetic robotics should not be limited to modeling ‘lower level’ mechanisms. But, how do we make progress? In the terms of this paper, we advocate directly tackling the challenge of ongoing emergence, and in particular our Criterion 2 (incorporation of skills) seems in most need of further research. If cognitive skills arise out of ongoing emergence, then if we achieve robots with ongoing emergence, there is a good chance that those robots will have instantiated models of cognitive skills.

Acknowledgements

The authors thank Lakshmi Gogate, Rich Maclin, and Eric Mislivec for discussions and comments. A discussion with John Weng started this paper. CGP thanks his Spring05 CS5541 class for discussions about some of the papers reviewed here. We thank the EpiRob05 anonymous referees for constructive feedback.

References

Adolph, K. E., & Avolio, A. M. (2000). Walking infants adapt locomotion to changing body dimensions. *Journal of*

- Experimental Psychology: Human Perception & Performance*, 26, 1148-1166.
- Aslin, R. N. (1981). Development of smooth pursuit in human infants. In D. Fisher, R. A. Monty, & J. W. Senders (Eds.), *Eye movements: Cognition and visual perception*. Hillsdale, NJ: Erlbaum.
- Baillargeon, R. (1987). Object permanence in 3 1/2- and 4 1/2-month old infants. *Developmental Psychology*, 23, 655-664.
- Baillargeon, R. (1993). The object concept revisited: New directions in the investigation of infants' physical knowledge. In C. Granrud (Ed.), *Visual Perception and Cognition in Infancy* (pp. 265-315). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.
- Baillargeon, R., (1995). A model of physical reasoning in infancy. In C. Rovee-Collier & L. P. Lipsitt (Eds.) *Advances in Infancy Research*, Vol. 9 (pp. 305-371). Norwood NJ: Ablex.
- Bernstein, N. (1967). *The Coordination and Regulation of Movements*. London: Pergamo.
- Berthouze, L., & Lungarella, M. (2004). Motor skill acquisition under environmental perturbations: On the necessity of alternate freezing and freeing of degrees of freedom. *Adaptive Behavior*, 12, 47-63.
- Blank, D., Kumar, D., Meeden, L., & Marshall, J. (2005). Bringing up robot: Fundamental mechanisms for creating a self-motivated, self-organizing architecture. *Cybernetics and Systems*, 36, 125-150.
- Breazeal, C., Buchsbaum, D., Gray, J., Gatenby, D., & Blumberg, B. (2005). Learning from and about others: Towards using imitation to bootstrap the social understanding of others by robots. *Artificial Life*, 11, 31-62.
- Breazeal, C. & Scassellati, B. (2000). Infant-like social interactions between a robot and a human caregiver. *Adaptive Behavior*, 8, 49-74.
- Carpenter, M., Nagell, K. & Tomasello, M. (1998). Social cognition, joint attention, and communicative competence from 9 to 15 months of age. *Monographs of the Society for Research in Child Development*, Serial No. 255, 63 (4).
- Chen, Y. & Weng, J. (2004). Developmental learning: A case study in understanding “object permanence.” In *Proceedings of the Fourth International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems*. Lund, Sweden: Lund University Cognitive Studies.
- Cheng, G. & Nagakubo, A., & Kuniyoshi, Y. (2001). Continuous humanoid interaction: An integrated perspective—gaining adaptivity, redundancy, flexibility—in one. *Robotics and Autonomous Systems*, 37, 161-183.
- Dominguez, M. & Jacobs, R. A. (2003). Developmental constraints aid the acquisition of binocular disparity sensitivities. *Neural Computation*, 15, 161-182.
- Dominey, P. F. & Boucher, J.-F. (2005). Developmental stages of perception and language acquisition in a perceptually grounded robot. *Cognitive Systems Research*, 6, 243-259.
- Fodor, J. A. (1983). *The Modularity of Mind*. Cambridge, MA: MIT Press.
- Fodor, J. & Pylyshyn, Z. (1988). Connectionism and cognitive architecture. *Cognition*, 28, 3-71.
- Gergely, G. (2003). What should a robot learn from an infant? Mechanisms of action interpretation and observational learning in infancy. *Connection Science*, 15, 191-209.

- Grupe, R. (2003). A developmental organization for robot behavior. In *Proceedings of the Third International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems*. Lund, Sweden: Lund University Cognitive Studies.
- Hofsten, C. von (1982). Eye-hand coordination in newborns. *Developmental Psychology*, 18, 450-461.
- Hofsten, C. von (1989). Transition mechanisms in sensorimotor development. In A. De Ribaupierre (Ed.), *Transition Mechanisms in Child Development: The Longitudinal Perspective* (pp. 233-258). New York, NY: Cambridge University Press.
- Hofsten, C. von, & Lindhagen, K. (1979). Observations on the development of reaching for moving objects. *Journal of Experimental Child Psychology*, 28, 158-173.
- Hollich, G., Hirsh-Pasek, K., & Golinkoff, R. (2000). Breaking the language barrier: An emergentist coalition model of word learning. *Monographs of the Society for Research in Child Development*, 65 (3, Serial No. 262).
- Johnson, S. P., & Nanez, J. E., Sr. (1995). Young infants' perception of object unity in two-dimensional displays. *Infant Behavior and Development*, 18, 133-143.
- Kaplan, F. & Oudeyer, P.-Y. (2003). Motivational principles for visual know-how development. In *Proceedings of the 3rd International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems*. Lund University Cognitive Studies.
- Kellman, P. J., & Spelke, E. S. (1983). Perception of partly occluded objects in infancy. *Cognitive Psychology*, 15, 483-524.
- Lampl, M., Veldhuis, J. D., & Johnson, M. L. (1992). Saltation and stasis: a model of human growth. *Science*, 258, 801-803.
- Lovett, A., & Scassellati, B. (2004). Using a robot to reexamine looking time experiments. In *Proceedings of the Third International Conference on Development and Learning*.
- Lungarella, M., & Berthouze, L. (2002). On the interplay between morphological, neural and environmental dynamics: A robotic case-study. *Adaptive Behavior*, 10, 223-241.
- Lungarella, M., Metta, G., Pfeifer, R., & Sandini, G. (2003). Developmental robotics: A survey. *Connection Science*, 15, 151-190.
- Nagai, Y., Hosoda, K., Morita, A., & Asada, M. (2003). A constructive model for the development of joint attention. *Connection Science*, 15, 211-229.
- Newcombe, N. & Huttenlocher, J. (2000). *Making Space: The Development of Spatial Representation and Reasoning*. Cambridge, MA: MIT Press.
- Mandler, J. (1999). Preverbal representations and language. In P. Bloom, M. A. Peterson, L. Nadel, & M. F. Garrett (Eds.), *Language and Space* (pp. 365-384). MIT Press.
- Michel, P., Gold, K., & Scassellati, B. (2004). Motion-based robotic self-recognition. Presented at the 2004 IEEE/RSJ International Conference on Intelligent Robots and Systems.
- Piaget, J. (1954). *The Construction of Reality in the Child*. New York: Basic Books.
- Schlesinger, M., & Casey, P. (2003). Where infants look when impossible things happen: Simulating and testing a gaze-direction model. *Connection Science*, 15, 271-280.
- Sengas, A. (1995) The development of Nicaraguan Sign Language via the language acquisition process. In D. MacLaughlin & S. McEwen (Eds.) *Proceedings of The 19th Boston University Conference on Language Development*. Boston: Casadilla Press.
- Siegler, R. S. (2004). U-shaped interest in u-shaped development – and what it means. *Journal of Cognition and Development*, 5, 1-10.
- Seth, A. K., McKinsty, J. L., Edelman, G. M., & Krichmar, J. L. (2004). Visual binding through reentrant connectivity and dynamic synchronization in a brain-based device. *Cerebral Cortex*, 14, 1185-1199.
- Shinsky, J. L., & Munakata, Y. (2001). Detecting transparent barriers: Clear evidence against the means-end deficit account of search failures. *Infancy*, 2, 395-404.
- Slater, A., & Morison, V. (1985). Shape constancy and slant perception at birth. *Perception*, 14, 337-344.
- Slater, A., Morison, V., & Rose, D. (1982). Visual memory at birth. *British Journal of Developmental Psychology*, 73, 5 19-525.
- Slater, A., Morison, V., & Rose, D. (1984). Newborn infants' perception of similarities and differences between two- and three-dimensional stimuli. *British Journal of Developmental Psychology*, 2, 287-294.
- Slater, A., Morison, V., Town, C., & Rose, D. (1985). Movement perception and identity constancy in the new-born baby. *British Journal of Developmental Psychology*, 3, 21 1- 220.
- Smith, L. B. (1999). Children's noun learning: How general learning processes make specialized learning mechanisms. In B. MacWhinney (Ed.), *The Emergence of Language* (pp. 277-304). Mahwah, NJ: Erlbaum.
- Sporns, O., & Alexander, W .H. (2002). Neuromodulation and plasticity in an autonomous robot. *Neural Networks*, 15, 761-774.
- Sporns, O. (2005). Dialog: What are the potential principles for the autonomous development of value and motivation? *AMD Newsletter*, 2(1).
- Steels, L. (2004). The autotelic principle. In Fumiya, I., Pfeifer, R., Steels, L., Kuniyoshi, K. (Eds.), *Embodied Artificial Intelligence, Lecture Notes in AI*, Vol. 3139 (pp. 231-242). Springer Verlag.
- Thelen, E. (1995). Motor development: A new synthesis. *American Psychologist*, 50, 79-95.
- Thelen, E., & Fisher, D. M. (1982). Newborn stepping: An explanation for a "disappearing reflex." *Developmental Psychology*, 18, 760-775.
- Triesch, J. & von der Malsburg, C. (2001). Democratic integration: Self-organized integration of adaptive cues. *Neural Computation*, 13, 2049-2074.
- Wellman, H. M., Cross, D., & Bartsch, K. (1986). Infant search and object permanence: A meta-analysis of the A-Not-B error. *Monographs of the Society for Research in Child Development*, 51 (3, Serial No. 214).
- Weng, J., McClelland, J., Pentland, A., Sporns, O., Stockman, I., Sur, M., & Thelen, E. (2001). Autonomous mental development by robots and animals. *Science*, 291, 599-600.
- Weng, J. (2004). Developmental robotics: Theory and experiments. *International Journal of Humanoid Robotics*, 1, 199-236.

Tables: See <http://www.cprince.com/PubRes/EpiRob05/tables>