Biologically Inspired Cognitive Architectures 18 (2016) 116–127

Contents lists available at ScienceDirect

Biologically Inspired Cognitive Architectures

journal homepage: www.elsevier.com/locate/bica

Research article

Desiderata for developmental cognitive architectures

David Vernon [a,*], Claes von Hofsten [b], Luciano Fadiga [c,d]

[a] Interaction Lab, School of Informatics, University of Skövde, Sweden
[b] Department of Psychology, University of Uppsala, Sweden
[c] Section of Human Physiology, University of Ferrara, Italy
[d] IIT@UniFe Center for Translational Neurophysiology, Istituto Italiano di Tecnologia, Ferrara, Italy

ARTICLE INFO

Article history:
Received 26 July 2016
Revised 4 October 2016
Accepted 5 October 2016

Keywords:
Autonomy
Cognitive architecture
Desiderata
Development
Ontogeny
Phylogeny

ABSTRACT

This paper complements Ron Sun's influential Desiderata for Cognitive Architectures by focussing on the desirable attributes of a biologically-inspired cognitive architecture for an agent with a capacity for autonomous development. Ten desiderata are identified, dealing with value systems & motives, embodiment, sensorimotor contingencies, perception, attention, prospective action, memory, learning, internal simulation, and constitutive autonomy. These desiderata are motivated by studies in developmental psychology, cognitive neuroscience, and enactive cognitive science. All ten focus on the ultimate aspects of cognitive development — why a feature is necessary and what it enables — rather on than the proximate mechanisms by which it can be realized. As such, the desiderata are for the most part neutral regarding the paradigm of cognitive science — cognitivist or emergent — that is adopted when designing a cognitive architecture. Where some element of a desideratum is specific to a particular paradigm, this is noted.

1. Introduction

In 2004, Ron Sun published an influential paper "Desiderata for Cognitive Architectures" in which he listed and discussed several attributes that a general cognitive architecture should exhibit (Sun, 2004). While the Desiderata emphasize the desirability of biological realism, the attribute of development[1] — exhibited by all altricial species and one of the chief characteristics of cognition in humans — is not the central focus of the article. That is not to say that the article underestimates the importance of development. On the contrary, it notes that some structures in a cognitive architecture result from ontogenetic development under the influence of physical, social, and cultural environments, and it states that an architecture should include only minimal structures and minimal learning mechanisms which should be capable of "bootstrapping all the way to a full-fledged cognitive model" (Sun, 2004). It notes that while human everyday activities are largely made of routines, i.e. habitual sequences of behavioral responses, these are the result of gradual trial-and-error adaptation and are subject to constant modification. It also emphasizes the importance of bottom-up learning. It discusses Anderson's proposal (Anderson, 1983) that skill development is characterized by the initial acquisition of declarative knowledge — "explicit verbal knowledge concerning a task" — and through practice the development of a set of implicit procedures that allow the task to be performed without using declarative knowledge. It also addresses Karmiloff-Smith's notion of "representational redescription", whereby during development low-level implicit representations are transformed into more abstract and explicit representations and are thus made more accessible (Karmiloff-Smith, 1992, 1994). Sun's own cognitive architecture CLARION (Sun, 2003, 2004) exhibits both of these explicit-to-implicit and implicit-to-explicit desiderata, among others.

Significantly, Sun remarks in the context of the desideratum of cognitive realism that cognitive architectures should "capture essential characteristics of human behavior and cognitive processes" and he suggests that "we may incorporate even more desiderata down the road". Since ontogeny is such an important aspect of human cognition, we wish to follow through on what Sun himself anticipated by identifying several desiderata for *developmental* cognitive architectures. We begin by defining exactly what is meant by the term cognitive architecture, recognizing that it has different connotations in different branches of cognitive science. We follow this with a brief summary of the main points

* Corresponding author.
E-mail address: david@vernon.eu (D. Vernon).

[1] Development, which involves various forms of learning, is a process by which an agent adapts over an extended timescale on the basis of interaction between a cognitive agent and its environment, including other cognitive agents. It results in an extended repertoire of skills and prospective control of those skills, as well as an ability to anticipate the outcome of the agent's own actions and the actions of other agents. It proceeds in phases, scaffolding skills developed in previous phases. Crucially, development is driven by internal motives. See Section 4 for more details.

http://dx.doi.org/10.1016/j.bica.2016.10.004
2212-683X/

of Sun's Desiderata. We then outline very briefly some of the key aspects of cognitive development, drawing on the body of knowledge in developmental psychology and cognitive neuroscience. Finally, we identify and discuss ten desiderata for developmental cognitive architectures. Our motivation in proposing them is to provide a minimal set of characteristics that can guide the design of either a developmental cognitive architecture or a developmental cognitive architecture schema. The focus here is on those aspects that form the essential basis for scaffolding subsequent development and other faculties (which may themselves be the result of development) will almost certainly have to be deployed to facilitate that scaffolding.

2. Cognitive architectures

Ron Sun defines a cognitive architecture as follows.

> "A cognitive *architecture* is the overall, essential structure and process of a broadly-scoped domain-generic computational cognitive model, used for broad, multiple-level, multiple-domain analysis of cognition and behavior" (Sun, 2004) (emphasis in the original).

In defining a cognitive architecture in this way, he builds on the work of Newell (1990). In fact, the term cognitive architecture can be traced to Newell's pioneering work in cognitivist cognitive science[2] where the term has a very specific meaning: a cognitive architecture represents any attempt to create a unified theory of cognition (Anderson et al., 2004; Newell, 1990). This is a theory that covers a broad range of cognitive issues, such as attention, memory, problem solving, decision making, and learning from several perspectives including psychology, neuroscience, and computer science. Allen Newell's and John Laird's Soar architecture (Laird, 2008, 2009, 2012; Laird, Newell, & Rosenbloom, 1987; Lehman, Laird, & Rosenbloom, 1998; Rosenbloom et al., 1993), John Anderson's ACT-R architecture (Anderson, 1996; Anderson et al., 2004), and Ron Sun's CLARION architecture (Sun, 2003, 2004) are typical candidate unified theories of cognition.

In the cognitivist paradigm, the focus in a cognitive architecture is on the aspects of cognition that are *constant over time* and that are *independent of the task*, i.e. unchanging from situation to situation (Gray, Young, & Kirschenbaum, 1997; Langley, 2005; Langley, Laird, & Rogers, 2009; Ritter & Young, 2001). The cognitive architecture determines the overall structure and organization of a cognitive system, including the component parts or modules, the relations between these modules, and the essential algorithmic and representational details within them. Again, in Sun's words:

> "A cognitive architecture provides a concrete framework for more detailed modeling of cognitive phenomena, through specifying the essential structures, division of modules, relations between modules, and a variety of other aspects" (Sun, 2004).

Thus, the architecture specifies the formalisms for knowledge representations and the types of memories used to store them, the processes that act upon that knowledge, and the learning mechanisms that acquire it.

A cognitive architecture plays an important role in computational modeling of cognition in that it makes explicit the set of assumptions upon which that cognitive model is founded. These assumptions are typically derived from several sources: biological or psychological data, philosophical arguments, or *ad hoc* working hypotheses inspired by work in different disciplines such as neurophysiology, psychology, or artificial intelligence. Once it has been created, a cognitive architecture also provides a framework for developing the ideas and assumptions encapsulated in the architecture (Sun, 2004).

Although the term cognitive architecture originated in cognitivist cognitive science, it has also been adopted in the emergent paradigm of cognitive science[3] where it has a sightly different meaning. Emergent approaches to cognition focus on the development of the agent from a primitive state to a fully cognitive state over an extended period of time. Here, the term cognitive architecture is used to denote the framework that facilitates that development. In this sense, an emergent cognitive architecture is essentially equivalent to the phylogenetic configuration of a newborn or newly-created cognitive agent: the initial state from which it subsequently develops (Vernon, 2014; Vernon et al., 2007). It is a way of dealing with the intrinsic complexity of a cognitive system by providing a structure within which to embed the mechanisms for perception, action, adaptation, anticipation, and motivation that enable the ontogenetic development over the system's life-time.

In contrast to the cognitivist stance, the presence of innate capabilities in an emergent system does not necessarily imply that the architecture is functionally modular, *i.e.* that the cognitive system is comprised of distinct modules, each one carrying out a specialized cognitive task. Heinz von Foerster argues that the constituents of a cognitive architecture cannot be separated into distinct functional components:

> "In the stream of cognitive processes one can conceptually isolate certain components, for instance (i) the faculty to perceive, (ii) the faculty to remember, and (iii) the faculty to infer. But if one wishes to isolate these faculties functionally or locally, one is doomed to fail. Consequently, if the mechanisms that are responsible for any of these faculties are to be discovered, then the totality of cognitive processes must be considered" (von Foerster, 2003, p. 105).

This quotation from von Foerster speaks powerfully in favor of the cognitive architecture endeavor because the cognitive architecture is an attempt to address and circumscribe that totality in a single well-specified framework. If modularity is present, it may be because the system develops this modularity through experience as part of its ontogenesis rather than being prefigured by the phylogeny of the system.

The cognitivist and emergent perspectives differ somewhat on the issue of innate structure. While in an emergent system the cognitive architecture *is* the innate structure, this is not necessarily so with a cognitivist system. Sun contends that "an innate structure can, but need not, be specified in an initial architecture" (Sun, 2004). Indeed, a cognitive architecture can be both more than and less than a set of innate structures: *more than* when it characterizes a mature developed (adult) cognitive model where the structures that are not innate have resulted from ontogenetic development, and *less than* when the innate structure is not required for the current model and can be specified later in a more detailed modeling.

Cognitive architectures, developmental or otherwise, are designed with some purpose in mind, be it as a candidate unified theory of cognition (human and artificial) or as the basis for the development of an adaptive robot. In the latter case, the relative importance of social interaction will affect the nature of the design

[2] For an overview of the cognitivist and emergent paradigms of cognitive science see Vernon, Metta, and Sandini (2007) and Vernon (2014).

[3] See FootNote 2.

and the degree to which a focus is placed on development based on human-robot interaction. Many of the desiderata in this paper are based on the assumption that such interaction is desirable. This is certainly the case for a cognitive architecture that is a candidate unified theory of cognition and will also be the case for cognitive robots that are intended to work and play with people. Furthermore, these desiderata do not reflect the possibility of designing a cognitive architecture to exploit forms of advanced information processing that lie beyond the boundaries of biological cognition. On the contrary, much of the background research that forms the foundation for this paper was in the area of cognitive robotics where the interest was in (a) using robot as an apparatus for empirical investigations of human cognition in developmental psychology and neuroscience and (b) building artificial cognitive systems that can interact effectively with humans.

3. Sun's *Desiderata for Cognitive Architectures*

In his *Desiderata for Cognitive Architectures* (Sun, 2004), Sun identifies four desirable features of a cognitive architecture. These are

1. Ecological realism;
2. Bio-evolutionary realism;
3. Cognitive realism;
4. Eclecticism of methodologies and techniques.

The key idea behind *ecological realism* is that a cognitive architecture should focus on allowing the cognitive system to operate in its natural environment, engaging in "everyday activities" (Sun, 2004). This means it has to be able to deal with being embodied and the attendant natural constraints on its actions and perceptions. It also means that the architecture has to deal with many concurrent and often conflicting goals with many environmental contingencies. Since human intelligence evolved from the capabilities of earlier primates, *bio-evolutionary realism* asserts that a cognitive model of human intelligence should be reducible to a model of animal intelligence. *Cognitive realism* means that a cognitive architecture should capture the essential characteristics of human cognition as we understand them from the perspective of psychology, neuroscience, and philosophy. Finally, the design of a cognitive architecture should include prior perspectives and capabilities, i.e. it should exhibit *eclecticism of methodologies and techniques*: new models should draw on, subsume, or supersede older models.

Sun also elaborates on the behavioral and cognitive characteristics which should ideally be captured by a cognitive architecture and exhibited by a cognitive system.

From a behavioral perspective, a cognitive system should exhibit *reactivity*: it should act and react without employing excessively complicated conceptual representations and extensive computation devoted to working through alternative strategies. That is, the system should behave in a "direct and immediate" manner. Furthermore, a cognitive system should exhibit *sequentiality*: it should operate sequentially, one step at a time, in a temporally-extended sequence of actions. This leads naturally to the characteristic of *routineness*: a cognitive system should have gradually-learned routine behaviors that are subject to constant modification and are typically acquired through a process of trial-and-error adaptation.

From the perspective of cognitive characteristics, Sun suggests that a cognitive architecture should reflect the dichotomy of implicit and explicit processes. The explicit processes are accessible and precise whereas the implicit ones are inaccessible and imprecise. Furthermore, there should be *synergistic interaction*, i.e. there should be a synergy between these two types of process, one supplementing and one complementing the other in various ways. There are, for example, explicit and implicit learning processes and these interact. The most important type of learning in a cognitive architecture is what Sun refers to as *bottom-up learning* whereby implicit learning is followed by explicit learning. In Sun's own work (Sun, 2003, 2007), implicit processes operate on connectionist representations and explicit processes on symbolic representations. Significantly, they interact and cooperate both in action selection and in learning: the CLARION cognitive architecture is able to effect autonomous generation of explicit conceptual structures by exploiting implicit knowledge acquired by trial-and-error learning, and it can also effect top-down learning by integrating externally-provided knowledge in the form of explicit rule-based conceptual structures and assimilating these into the bottom level implicit representation.

Finally, Sun argues for a form of *modularity* in a cognitive architecture so that some cognitive faculties are specialized and separate, either as functionally-encapsulated modules or as physically — neurophysiologically — encapsulated modules.

4. Cognitive development

Cognitive development is a complex process that takes place over an extended period of time, especially during the first years of a cognitive agent's lifetime. It would be impossible to summarize all the facets of psychological development in a short paper such as this. Instead, we simply present some key principles of development and substantiate these by drawing on what is known of development in infants. We do this in three steps. First we note the importance of prospection, internal simulation, motives, and value systems. We then highlight some core cognitive abilities linked to development. Finally, we illustrate the developmental process by looking very briefly at the development of one of the hallmarks of cognitive behavior: an infant's capacity to cooperate and collaborate with others. Taken together, these considerations motivate the ten desiderata in the next section.

To begin with, we remark that in the emergent paradigm of cognitive science, which focusses on the autonomous nature of cognitive systems, development is the process whereby a cognitive agent (a) expands its repertoire of action capabilities and (b) extends the time horizon of its ability to anticipate events in its world, including the need to act, the outcome of selected actions, the intentions of other cognitive agents, and the outcome of their actions (Vernon, 2010; Vernon, von Hofsten, & Fadiga, 2010). Ultimately, development results in an embodied agent capable of adaptive anticipatory interaction with the world in which it is embedded. Development itself arises from dynamic interaction with the environment and especially with other cognitive agents. It results in new forms of action and predictive control of these actions: anticipating the outcome of actions as well as anticipating the need for them in the first place. Ultimately, through development, cognition becomes a process of effective action: action that contributes prospectively to the maintenance of the system's autonomy (Maturana & Varela, 1987).

4.1. The importance of prospection, internal simulation, motives, and value systems

To act effectively, an agent must be able to infer the occurrence of upcoming events. Here we see the importance of prospection and, in particular, the establishment of prospective control of these movements in the context of the goals of an action (Lintern, 1998; Reed, 1996). In this, it is the goal that is the focus: an action defines a movement made in order to achieve a goal. The goal, therefore, is the fundamental property of the action rather than the specific motoric details of how it is achieved.[4]

Development, particularly the prospective aspect, may be accelerated by internal simulation, i.e. mentally rehearsing the execution of actions and inferring the likely outcome of those actions (Hesslow, 2002, 2012).

While a cognitive agent's actions are guided by prospection and directed by goals, they are also triggered by affective motives (Núñez & Freeman, 1999). However, not only are actions modulated by motivation, but development itself is driven by these same motives. Significantly, these motives are not task-specific; instead they modulate the affective state of the system and, indirectly, the actions in which it engages (Oudeyer, Kaplan, & Hafner, 2007). In turn, these motivations reflect an inherited intrinsic value system that constrains behavior, with actions being selected based on the motivations that manifest the value system (Edelman, 2006) (e.g. actions that lead to the greatest decrease in the mean error rate of the system's predictive learning mechanism (Oudeyer et al., 2007)). From this perspective, the value system "mediates the saliency of environmental stimuli" yielding an intrinsic motivation system that signals the occurrence of important events and triggers the formation of goals which are then acted upon by a behavioral system (Merrick, 2010). However, value systems focus not only on behavioral autonomy but also on constitutive autonomy (Froese & Ziemke, 2009; Froese, Virgo, & Izquierdo, 2007; Vernon, 2016; Vernon, Lowe, Thill, & Ziemke, 2015b). That is, they include internal events in the state of the agent, specifically to maximize the potential for development of that system and thereby maintain and enhance the agent's autonomy. This mirrors the concept of a self-aware self-effecting (SASE) agent (Weng, 2002, 2004a, 2004b) that incorporates mechanisms for self-modification.

4.2. Core cognitive abilities in infants

The phylogenetic configuration of a cognitive system provides the core on which development builds. Young infants have two core knowledge systems that provide the basis for representing objects (including persons and places) and the concept of number (numerosities) (Spelke, 2000). While these core systems provide the foundation of cognitive flexibility, they are themselves limited in a number of ways: they are domain specific, task specific, and encapsulated, i.e. they operate relatively independently of one another.

Infants build representations of objects but only if they exhibit certain characteristics. Specifically, the entities that are considered to be object-like are complete, connected, solid bodies that maintain their identity over time, and persist through occlusion when they are hidden by other objects. Infants can keep track of multiple objects simultaneously but the number is limited to about three objects and this ability is tolerant to changes in object properties such as color, precise shape, and spatial location.

An important part of core knowledge has to do with people. Specifically, it has to do with interaction between infants and their carers, and the predisposition of human infants to interact with other humans. Infants are attracted by other people and are endowed with abilities to recognize them and their expressions, and to communicate with them. They develop an ability to perceive the goal-directedness of the actions of other people quite quickly. Similarly, young infants exhibit a preference for the motions produced by a moving person over other motions, so-called *biological motion*. The preference for biological motion by human infants aged four to six months was suggested in 1982 as evidence in support of the hypothesis that this is an intrinsic capacity of the human visual system (Fox & McDaniel, 1982). Recently, in 2008, it was shown by Francesca Simion and colleagues that even newborn babies are sensitive to biological motion (Simion, Regolin, & Bulf, 2008).

Intentions and emotions are displayed by elaborate and specific movements, gestures, and sounds that become important to perceive and control. Some of these abilities are already present in newborn infants and reflect their preparedness for social interaction. Young infants are very attracted by people, especially to the sounds, movements, and features of the human face, and they look longer at a face that directs the eyes straight at them than at one that looks to the side. Teresa Farroni and colleagues (Farroni, Csibra, Simion, & Johnson, 2002) describe an experiment that demonstrates that, from birth, human infants prefer to look at faces that engage them in mutual gaze, i.e. to interact with people that make direct eye contact. Human infants also engage in some social interaction and turn-taking that among other things is expressed in their imitation of facial gestures: the two classic papers on imitation of facial features by neonates, and imitation generally, are the result of work by Meltzoff and Moore (1977, 1997).

The ability to spatially reorient and navigate, often taken as a typical cognitive ability, is also subject to development. While adults solve reorientation tasks by combining non-geometric information (e.g. color) with geometric information, young children rely only on geometry even when nongeometric information is available (Hermer & Spelke, 1996). This suggests that children exploit a core cognitive process for representing the shape of the environment and for using this representation to determine their own position within it. Similarly, there is evidence that navigation is based on representations that are momentary rather than enduring, egocentric rather than geocentric, and limited in the information they capture about the environment. Wang and Spelke (2002) discuss the characteristic dependence by humans on momentary, egocentric, and informationally-limited cues in spatial representation, contrasting it with the more widely-held and more intuitive assumption that it is based on enduring, geocentric "cognitive maps". When navigating, children and adults base their turning decisions on local, view-dependent, and geometry-based representations. They navigate by forming and updating a dynamic representation of their relationship to the environment. This capacity for path integration, whereby an agent navigates from point to point, cumulatively basing the next step on the previous ones, has been found to be one of the primary forms of navigation in insects, birds, and mammals. Like other animals, humans can return to the origin of a path and travel to familiar locations along novel paths, reorienting by recognizing landmarks rather than by forming global representations of scenes.

4.3. Development of the ability to help others and collaborate with them

The phased aspect of development is particularly relevant in the manner in which infants and children come to understand the intentions of others and to help them achieve their goals.[5] It takes several years for human infants to develop the requisite abilities.

[4] The importance of the goal in an action is exemplified by mirror neurons (Gallese, Fadiga, Fogassi, & Rizzolatti, 1996; Rizzolatti, Fadiga, Gallese, & Fogassi, 1996) which are activated both when an action is performed and when the same or similar action is observed being performed by another agent. These neurons are specific to the goal of the action and not the mechanics of carrying it out. For more details on the goal-orientation of actions, see von Hofsten (2004).

[5] This brief overview of the development of the ability to engage in instrumental helping and to collaborate with others follows closely the summary derived in Vernon (2014). A more detailed account can be found in Vernon et al. (2010).

During the first year of life the progressive acquisition of motor skills facilitates the development of an ability to understand the intentions of other agents, initially by anticipating the goal of simple movements and eventually understanding more complex goals. During this period, the ability to infer what another agent is focussing their attention on and the ability to interpret emotional expressions begins to improve substantially.

Around 14–18 months of age children begin to exhibit instrumental helping behavior, i.e. they display spontaneous, unrewarded helping behaviors when another person is unable to achieve his goal (Warneken & Tomasello, 2009). This is a critical stage in the development of a capacity for collaborative behavior, a process that progresses past three and four years of age.

Around 2 years of age children start to solve simple cooperation tasks together with adults (Warneken, Chen, & Tomasello, 2006). This phase of development sees the beginning of shared intentionality where a child and an adult form a shared goal and both engage in joint activity. It also seems that children seem to be motivated not just by the goal but by the cooperation itself, i.e. the social aspect of the interaction. The ability to cooperate with peers and become a social partner in joint activities develops over the second and third years of life as social understanding increases (Brownell, Ramani, & Zerwas, 2006).

More complex collaboration, which necessitates the sharing of intentions and joint coordination of actions, appears at about three years of age when children master more difficult cooperation tasks such as those involving complementary roles for the two partners in a collaborative task (Meyer, Bekkering, Paulus, & Hunnius, 2010). At three years of age, children begin to develop the ability to cooperate by coordinating two complementary actions. By three-and-a-half years of age children quickly master the task, can deal effectively with the roles in the task being reversed, and can even teach new partners (Ashley & Tomasello, 1998).

The motives which drive instrumental helping are simpler than those of collaborative behaviors: they are based on wanting to see the goal completed or wanting to perceive pleasure in the human at being able to complete it. In this case, the motivational focus is solely on the needs of the second agent and the needs of the first agent do not enter into the equation. The motives underlying collaborative behavior are more complicated. In this case, the intentions and the goals have to be shared and the motivational focus is on the needs of both agents.

5. Desiderata for developmental cognitive architectures

We come now to the main objective of this short paper: to complement Ron Sun's desiderata with others that focus explicitly on endowing a cognitive architecture with a capacity for development that is driven by both exploratory and social motives, as espoused by the two pioneers in studying child development, Piaget (1936) and Vygotsky (1978). These are derived in part from insights drawn from research in developmental psychology and cognitive neuroscience, encapsulated in a roadmap for cognitive development in humanoid robots (Vernon et al., 2010), and from the increasingly important enactive stance on cognitive science (Stewart, Gapenne, & Di Paolo, 2010), encapsulated in a study of enaction as a framework for development in cognitive robotics (Vernon, 2010). Together, they lead to ten desirable design goals — desiderata — for developmental cognitive architectures.

5.1. Desideratum 1: Value systems and motives

The development of a cognitive agent is dependent on value systems and motives which determine the goals of actions and provide the drive for achieving them (Merrick, 2010; Oudeyer et al., 2007). Thus, *a cognitive architecture capable of development needs to have a value system that guides the selection of actions and drives development.*

There are two broad classes of motive: exploratory motives and social motives, broadly reflecting the psychology of development espoused by Piaget and Vygotsky, respectively (Lindblom, 2015; Lindblom & Ziemke, 2003; Piaget, 1954; Vygotsky, 1978). Both forms of motive function from birth and provide the driving force for action throughout life.

There are at least two exploratory motives, one to do with the discovery of novelty and regularity in the world and the other to do with finding out about the potential of one's own action capabilities. Infants are visually attracted to new objects and events but after a while they cease to be attracted. Infants also have a strong motivation to discover what they can do with objects, especially with respect to their own sensorimotor capabilities and the particular characteristics of their embodiment. Effectively, infants have a strong motivation to discover the affordances of objects around them and it isn't necessarily success at achieving task-specific goals that drives development in infants but rather the discovery of new modes of interaction with the world in which the infant is embedded: the acquisition of a new way of doing something through exploration (von Hofsten, 2003, 2004).

The social motive focusses on finding comfort, security, and satisfaction through interaction with others, allowing the agent to learn new skills and acquire knowledge about the world from the experience of others. It is manifest from birth in the tendency to fixate social stimuli, imitate basic gestures, and engage in social interaction. The social motive is so important that it has been suggested that without it a person will stop developing altogether. Social motives also include a strong need to belong, a drive for self-preservation, and the need for cognitive consistency with other (Forgas, Kipling, & Laham, 2005).

The cognitive architecture's value system should manifest both exploratory and social motives. The exploratory motives should facilitate the discovery of novelty and regularities in both the environment and the cognitive agent's action capabilities. The social motives should exhibit a positive bias toward fixation on social stimuli, imitation of basic gestures, and engagement in social interaction.

Both exploratory and social motives reflect the agent's interaction with its environment and its behavioral autonomy (Froese et al., 2007; Vernon et al., 2015b). However, the value systems that drive development are relevant not just to the processes of behavioral autonomy but also to those of constitutive autonomy, a concept addressed in the last desideratum. Arguably, the value systems and motives that drive development can support both constitutive autonomy and behavioral autonomy (Vernon, 2016).

For a discussion of the importance of prospective goal-directed action in cognitive development and the role of exploratory and social motives in this developmental process, see von Hofsten (2009) and von Hofsten (2013). For examples of value systems in self-motivated cognitive robots, see Merrick (2010), Oudeyer et al. (2007) and Huang and Weng (2002, 2007).

5.2. Desideratum 2: Physical embodiment

The concept of embodiment is complex and nuanced, with several stances on what it means to be embodied (Chrisley & Ziemke, 2002; Ziemke, 2003), what embodied cognition entails (Anderson, 2003; Shapiro, 2011), and the key role embodiment plays in cognitive development (Thelen, 1995; Thelen & Smith, 1998, 2003). Given development's focus on expansion of action capabilities and extension of the time-horizon of anticipatory capabilities, *an agent that is capable of development should have a variety of physical sensory and motor interfaces to allow the system to act on the world* *and perceive the effects of these actions.* The richer the sensorimotor interface, the richer the model of the world the agent can construct as it makes sense of the environment with which it is interacting and, hence, the greater the potential for development.

The embodied cognition thesis states that "Many features of cognition ...are deeply dependent upon characteristics of the physical body of an agent, such that the agent's beyond-the-brain body plays a significant causal role, or physically constitutive role, in that agent's cognitive processing" (Wilson & Foglia, 2011). Variants include the conceptualization hypothesis, the constitution hypothesis, and the replacement hypothesis (Shapiro, 2011) which state, respectively: (a) that the characteristics of an agent's body determine the concepts an organism can acquire, (b) that the body plays a constitutive rather than a supportive role in cognitive processing and is itself an integral part of cognition, and (c) that the need for representations and representational processes is removed because an agent's body is engaged in real-time interaction with its environment.

If the cognitive architecture is to adhere to the embodied cognition thesis or any of its three variants *the system's physical structure, its kinematics, and its dynamics should form a constitutive element of the cognitive process, and changes in morphology should involve matching changes in perceptual processing to improve the control of specific actions.* For example, we can recognize objects based on the nature of the actions we can perform on it in the absence of visual information about its shape (Campanella, Sandini, & Morone, 2011). Also, attention is selectively modulated in motor-dependent manner with some objects rather than others being more apparent depending on the motor system; for example, the ability to detect an object is enhanced when the appearance and features of the object match the configuration of the agent's hands as it prepares to grasp an object (Craighero, Fadiga, Rizzolatti, & Umiltà, 1999).

5.3. Desideratum 3: Sensorimotor contingencies

Linking exploratory motives and embodiment, *the cognitive architecture should ideally have innate mechanisms for the agent to learn sensorimotor contingencies (SMC)*, i.e. the relation between the actions that the agent performs and the change it experiences in its sensed data because of those actions (O'Regan & Noë, 2001). As the agent develops, it masters its sensorimotor contingency knowledge. This knowledge is practical, not propositional, and the agent's perceptions of its world arises from its experiences as it interacts with its environment. Several developmental mechanisms have been proposed by which this can be accomplished. These include motor babbling, whereby the agent's motor system is randomly activated and associated with the sensory consequences (Cangelosi & Schlesinger, 2015), and goal babbling, which effects a more efficient strategy than exhaustive search of the space of motor skills, bootstrapping coordination skills and learning inverse kinematic mapping by repetitively trying to accomplish multiple goals related to that skill (Rolf & Steil, 2012; Rolf, Steil, & Gienger, 2010).

The agent's sensorimotor contingency mechanism also allows it to develop the ability to distinguish between sensed data associated with its own body and data associated with other objects and agents in its environment. Furthermore, it provides the basis for the development of the ability to perceive the affordance of objects, a concept due to the influential psychologist Gibson (1977). An affordance refers to the potential use to which an object can be put, as perceived by an observer of the object. This perceived potential depends on the skills the observer possesses. Thus, the affordance is dependent both on the object itself and the perception and action capabilities of the observing agent. See Barsingerhorn, Zaal, Smith, and Pepping (2012) and Jamone et al. (2016) for surveys of work on affordances and the role of affordance in cognitive systems.

5.4. Desideratum 4: Perception

The ability to interpret sensorimotor data is essential to any cognitive agent and there are many perceptual processes involved. Here, however, we focus on those processes that are particularly relevant to the agent's capacity for development. Bearing in mind what we said in the brief outline of cognitive development in humans above, we can highlight the following desirable processes.

First, regarding the perception of objecthood, *it is desirable that the cognitive architecture have a capacity for unsupervised perceptual categorization*, i.e. the capacity to organize unlabeled sensory signals of all modalities into categories without prior knowledge or external instruction. When these categories exhibit the five characteristics of inner unity, a persistent outer boundary, cohesive and distinct motion, relatively constant size and shape when in motion, a change in the behavior or motion of one or both of the objects when contact occurs with another object, they are perceived as objects. The perception of motion is fundamental because it introduces time into perception and makes it possible to predict future events. It develops quite early as shown by for instance infants catching moving objects. Infants catch moving objects as soon as they grasp stationary ones (von Hofsten, 1980).

Second, to facilitate social interaction, *the cognitive architecture should be able to distinguish objects that exhibit biological motion, to localize faces and key features of faces, such as eyes and mouth, and to distinguish facial expressions*. For auditory perception, *the cognitive architecture should have a capacity to identify and localize human voices and to distinguish between different tones of voice.*

Thirdly, *the cognitive architecture should have a capacity to extract geometric relationships between objects and build geometric representations of shape.*

5.5. Desideratum 5: Attention

Spatial and selective attention make perception more effective. Consequently, *the cognitive architecture should have an attentional mechanism that has certain innate characteristics for facilitating cognitive development*. Given the prospective nature of action selection and action execution, the attention process should fixate on the goal of an action, not on its constituent movements. Given the importance of social interaction in development, it should pay preferential attention to biological motion rather than non-biological mechanical motion, it should be attracted to people and especially to their faces, sounds, and movements, and it should gaze longer when a person looks directly at it. Furthermore, the preparation of a motor program in readiness to act in some spatial regions should predispose the perceptual system to process stimuli coming from that region (spatial attention) (Craighero, Nascimben, & Fadiga, 2004) and the preparation of a motor program in readiness to act on specific objects should predispose the perceptual system focus attention on those objects (selective attention) (Craighero et al., 1999).

Depending on the eventual purpose of the cognitive architecture (see final paragraph of Section 2), it may or may not be desirable to include attentional features that are characteristic of human attention, such as distractibility (Forster & Lavie, 2015; Kelley, Rees, & Lavie, 2013) or attentional blocking and habituation (Billing & Balkenius, 2014).

5.6. Desideratum 6: Prospective action

Evidence from developmental psychology clearly shows that infants perform goal-directed movements from the very beginning of life and goal-directed action plays a central role in infant motor development (Rolf & Steil, 2012; Rolf et al., 2010): "Before infants master reaching, they spend hours and hours trying to get the hand to an object in spite of the fact that they will fail, at least to begin with" (von Hofsten, 2004). The capacity to develop an understanding of sensorimotor contingencies though goal babbling is an important mechanism in the development of goal-directed action (see Desideratum 3 Sensorimotor Contingencies above). Consequently, *a cognitive architecture capable of development should facilitate actions that are directed by goals, guided by prospection, and triggered by motives*. Again, see von Hofsten (2009) and von Hofsten (2013) for a discussion of the importance of prospective goal-directed action in cognitive development and the role of exploratory and social motives in this developmental process.

This desideratum dovetails with several others. The selection of action should be modulated by affective motivation mechanisms (Desideratum 1, Value Systems and Motives) as well as rational analysis. The cognitive architecture should have a capacity to learn through exploration the actions an object affords, i.e. its affordances (Desideratum 3, Sensorimotor Contingencies). The motor system should have a capacity to encode a repertoire of goal-oriented actions and the goal of the action should be provided by a value system (Desideratum 1, Value Systems and Motives) modulated by internal perceptuomotor simulation (Desideratum 9, Internal Simulation) and associated effector-specific percepts (Desideratum 5, Attention).

5.7. Desideratum 7: Declarative and procedural memory

Memory serves to preserve what has been achieved through learning and development, ensuring that, when a cognitive system adapts to new circumstances, it does not lose its ability to act effectively in situations to which it had adapted previously. However, memory has another role: to anticipate the future. It forms the basis for one of the central pillars of cognitive capacity, i.e. the ability to simulate internally the outcomes of possible actions and select the one that seems most appropriate for the current situation (Desideratum 9, Internal Simulation). Viewed in this light, memory can be seen as a mechanism that allows a cognitive agent *to prepare to act*, overcoming through anticipation the inherent limiting immediacy of its perceptual capabilities.

Given that the focus of development on expansion of action capabilities and extension of the time-horizon of anticipatory capabilities, manifested by an ability to anticipate the need to act, the outcome of selected actions, the intentions of other cognitive agents, and the outcome of their actions, it is safe to say that *a cognitive architecture capable of development must have memory*. We can distinguish different forms memory on the basis of the nature of what is remembered and the type of access we have to it (Squire, 2004; Wood, Baxter, & Belpaeme, 2012). It can be categorized as either declarative or procedural, depending on whether it captures knowledge of things — facts — or actions — skills, sometimes characterized as memory of knowledge and know-how: "knowing that" and "knowing how" (Ryle, 1949).

Two different types of declarative memory can be distinguished. These are *episodic memory* and *semantic memory*. Episodic memory, a term coined by Endel Tulving in 1972 (Tulving, 1972, 1984), plays a key role in the anticipatory aspect of cognition. It refers to specific instances in the agent's experience while semantic memory refers to general knowledge about the agent's world which may be independent of the agent's specific experiences. Episodic memory has an explicit spatial and temporal context: what happened, where it happened, and when it happened. Crucially, episodic memory is a *constructive* process (Seligman, Railton, Baumeister, & Sripada, 2013). Each time an event is assimilated into episodic memory, past episodes are reconstructed. However, they are reconstructed a little differently each time. This constructive characteristic is central to development and is related to the role that episodic memory plays in the process of internal simulation that forms the basis of prospection. In contrast, semantic memory "is the memory necessary for the use of language. It is a mental thesaurus, organized knowledge a person possesses about words and other verbal symbols, their meaning and referents, about relations among them, and about rules, formulas, and algorithms for the manipulation of the symbols, concepts, and relations" (Tulving, 1972, 1984).

In addition to the need for a developmental cognitive architecture to have both declarative and procedural memory, in order to facilitate sensorimotor contingencies (Desiderata 3) and internal simulation (Desideratum 9), it is desirable that, based on the established mutual connection of perception and action in the brain, *the cognitive architecture should have some mechanism that facilitates constructive association of procedural memory of actions and their outcomes with episodic memories*, possibly with joint episodic-procedural memory (Vernon, Beetz, & Sandini, 2015a).

5.8. Desideratum 8: Multiple modes of learning

Reflecting the existence of implicit and explicit processes in a cognitive architecture, Sun's Desiderata notes the need for a cognitive architecture to have the capacity for implicit learning, explicit learning, bottom-up learning (implicit learning first, explicit learning later), top-down learning (explicit learning first, implicit learning later), and parallel learning (implicit and explicit learning simultaneously). The adaptation inherent in development is dependent on this capacity for learning. Here, we wish to amplify Sun's point by emphasizing that *a developmental cognitive architecture to have at least three different modes of learning: supervised learning, reinforcement learning, and unsupervised learning* (Doya, 1999, 2000).

In supervised learning, the agent is provided with examples of what it needs to learn and it can determine an error value between the correct answer and its estimate of the correct answer. These errors are vector values: they show the magnitude of the error as well as the direction in which it needs to adjust its estimate in order to reduce the error next time it makes that estimate. In reinforcement learning, a reward signal is provided at each step of the learning process. In this case, the teaching signals don't contain any directional information. They are simply scalar values called rewards or reinforcement signals. Unsupervised learning operates with no teaching signals, just a stream of input data. Here, the goal of learning is to uncover the statistical regularity in this input stream and, in particular, to find some mapping between the input data and the learned output that reflects the underlying order in the input data.

As you would expect, different types of development require different learning mechanisms and the human brain is good at all three types of learning, with different regions being specialized for the different types (Doya, 1999). These regions and the learning processes are also interdependent (McClelland, NcNaughton, & O'Reilly, 1995). Innate behaviors are honed through continuous knowledge-free reinforcement-like learning while new skills develop through a different form of learning, driven by spontaneous unsupervised play and exploration which is not directly reinforced. On the other hand, imitative learning and learning by instruction makes use of supervised learning (Argall, Chernova, Veloso, & Browning, 2009).

These different forms of learning support development that arises through exploratory and social motives (Desideratum 1, Value Systems and Motives). For example, unsupervised learning and reinforcement learning supports motor and goal babbling (Desideratum 3, Sensorimotor Contingencies) while supervised learning and reinforcement learning support teaching and learning by imitation, a key mechanism in neonatal development (Billard, 2002; Dautenhahn & Billard, 1999; Meltzoff, 2002; Meltzoff & Moore, 1977, 1997).

5.9. Desideratum 9: Internal simulation

We have already mentioned that a cognitive agent's capacity for prospection relies on mentally simulating future possibilities (Seligman et al., 2013) and mentally rehearsing the execution of actions and the likely outcome of those actions (Hesslow, 2002, 2012). Through episodic memory, it provides the key anticipatory capacity of a cognitive agent for self-projection: the ability to shift perspective from itself in the here-and-now and to take an alternative perspective through the mental construction of an imagined alternative perspective (Schacter, Addis, & Buckner, 2008). Internal simulation (or emulation) (Grush, 2004; Svensson, Lindblom, & Ziemke, 2007) also features in desiderata for sensorimotor contingencies, prospective action, and declarative and procedural memory. Thus, while essential for cognition, internal simulation is also important for development because it provides support for these other developmental desiderata.

The literature on internal simulation is extensive (see Vernon, 2014 for an overview) but it is instructive to highlight a few key points to reinforce the assertion that *a cognitive architecture that is capable of development requires some mechanism to effect internal simulation*.

While self-projection is effected through internal simulation based on episodic memory, action-directed internal simulation involves three different types of anticipation: implicit, internal, and external (Svensson, Morse, & Ziemke, 2009). Implicit anticipation concerns the prediction of motor commands from perceptions (which may have been simulated in a previous phase of internal simulation). Internal anticipation concerns the prediction of the proprioceptive consequences of carrying out an action, i.e. the effect of an action on the agent's own body. External anticipation concerns the prediction of the consequences for external objects and other agents of carrying out an action. Implicit anticipation selects some motor activity (possibly covert, i.e. simulated) to be carried out based on an association between stimulus and actions; internal and external anticipation then predict the consequences of that action. Collectively, they simulate actions and the effects of actions.

Covert action involves what is referred to as motor imagery and simulation of perception is often referred to as visual imagery. Reflecting the interdependence of perception and action, covert action often has elements of both motor and visual imagery and, *vice versa*, the simulation of perception often has elements of motor imagery. Visual and motor imagery are sometimes referred to collectively as mental imagery. Indeed, mental imagery can be viewed as a synonym for internal simulation: "all imagery is mental emulation" (Moulton & Kosslyn, 2009), p. 1276. For a computational perspective on mental imagery, specifically in the context of cognitive architectures, see Wintermute (2012).

While internal simulation supports development, it too is subject to development in the sense that the internal models have to be bootstrapped in some manner. One possible way this might happen is addressed by the *inception of simulation hypothesis* which asserts that internal simulations in young infants are formed by re-enacting sensory-motor experiences in dreams (Thill & Svensson, 2011). The models that arise from these simulations are validated while awake and subsequently refined. As the child develops, the simulations become more accurate and reliable, adjustments are needed less and less, and internal simulation can be used increasingly in everyday cognitive activities. Experiments using a robot simulator showed that robot "dreams" can lead to faster development of improved internal simulation during waking behavior (Svensson, Thill, & Ziemke, 2013).

5.10. Desideratum 10: Constitutive autonomy

Autonomy is an obscure concept (Boden, 2008) that can be labeled in more than twenty different ways (Vernon, 2014). However, one can distinguish two types of particular relevance to cognitive architectures: behavioral autonomy and constitutive autonomy (Barandiaran & Moreno, 2008; Froese et al., 2007; Vernon et al., 2015b). Behavioral autonomy focusses on the external characteristics of the system: the extent to which the agent sets its own goals and its robustness and flexibility in dealing with an uncertain and possibly precarious environment. On the other hand, constitutive autonomy focusses on the internal organization and the organizational processes that keep the system viable and maintain itself as an identifiable autonomous entity through on-going processes of self-construction and self-repair. Agents that are constitutively autonomous can make different levels of contribution to the maintenance of their autonomy, making them less or more effective in dealing with the uncertainty and precariousness of the environment in which the system is embedded and in which it has to survive. Behavioral and constitutive autonomy are linked: an agent can't deal with uncertainty and danger if it is not organizationally — constitutively — equipped to do so (Vernon, 2016). The constitutive-behavioral distinction is sometimes cast as a difference between constitutive processes and interactive processes (Froese & Ziemke, 2009). As noted already, constitutive processes deal with the agent itself, its organization, and its maintenance as an agent through on-going processes of self-construction and self-repair. On the other hand, interactive processes deal with the interaction of the agent with its environment. Both processes play complementary roles in autonomous operation of the agents.

All nine desiderata above focus on development in the context of behavioral autonomy. However, *if the cognitive architecture is to adhere to the tenets of the emergent paradigm of cognitive science generally, and enactive cognitive science, in particular, it should also be constitutively autonomous*. This entails that it should have a range of processes of self-maintenance (Bickhard, 2000) effecting both homeostatic (Damasio, 2003; Damasio & Carvalho, 2013; Morse, Lowe, & Ziemke, 2008; Ziemke & Lowe, 2009) and allostatic (i.e. adaptive and predictive) (Schulkin, 2011; Sterling, 2004, 2012) regulation that maintain the agent's organizational parameters within operational bounds. It also means that *the developmental processes should operate autonomously so that changes are not a deterministic consequence of external stimuli but result from an internal process of generative model construction, driven by value systems that promote self-organization*. These changes should serve to modify the system's organization so that its dynamics are altered to effect (a) the expansion in the space of viable actions and (b) the extension of the time horizon of the system's anticipatory capability. That is, the processes that serve to achieve constitutive autonomy should also serve to achieve behavioral autonomy (Vernon, 2016).

6. Discussion

To check the consistency of these ten desiderata, it is instructive to consider how they are reflected in the design principles proposed by others in the field. First, we consider Jeffrey Krichmar's five design principles for developmental artificial brain-based devices (Krichmar & Edelman, 2005; Krichmar & Edelman, 2006; Krichmar & Reeke, 2005), all of which are directly applicable to cognitive architectures.

The first principle is that the design (i.e. the architecture) should address the dynamics of the neural elements in different regions of the brain, the structure of these regions, and especially the connectivity and interaction between these regions. While this principle does not speak directly to development *per se* it does mirror von Foerster's caveat above and it reflects the need to view a cognitive architecture not just as a set of interconnected modules but as network of flows among constituent elements with an appropriate balance between localized functional responsibility and functionality that arises from distributed dynamics. For example, Desideratum No. 9, Internal Simulation, might not be effected by a some cohesive encapsulated component but through the cognitive architecture as a whole by inhibiting the execution of prepared motor commands.

The second principle is that the architecture should support perceptual categorization, i.e. the organization of unlabeled sensory signals of all modalities into categories without prior knowledge or external instruction. From a developmental perspective, this means the architecture should have a capacity for model generation as well as model fitting. The distinction between model generation and model fitting in cognitive systems is also emphasized by John Weng (Weng, 2004a). Model generation is tightly linked to unsupervised learning in Desideratum No. 8, internal simulation in Desideratum No. 9, and sensorimotor contingencies in Desideratum No. 3.

Krichmar's third principle is that a developmental system should have a physical instantiation, i.e. it should be embodied, so that it is tightly coupled with its own morphology and so that it can explore its environment. This directly mirrors Desideratum No. 2, Physical Embodiment.

The fourth principle is that the system should engage in some behavioral task and, consequently, it should have some minimal set of innate behaviors or reflexes in order to explore and survive in its initial environmental niche. From this minimum set, the system can learn and adapt so that it improves its behavior over time. This principle is reflected in Desiderata No. 1, Value Systems and Motives, No. 4, Perception, and No. 5, Attention.

Krichmar's fifth and final principle is that a developmental system should have a means to adapt (cf. Desiderata No. 8, Multiple Modes of Learning and No. 9, Internal Simulation). This entails the presence of a value system, i.e. a set of motives that guide or govern its development (cf. Desideratum No. 1, Value Systems and Motives). These should be non-specific modulatory signals (in the sense that the signals don't specify what actions to take) that bias the dynamics of the system so that the global needs of the system are satisfied: in effect, so that the system's autonomy is preserved or enhanced (cf. Desideratum No. 10, Constitutive Autonomy).

The ten desiderata are also reflected in several of Pfeifer's and Bongard's eight design principles for intelligent agents, all eight focussing on the importance of embodiment and interaction in cognition (Pfeifer & Bongard, 2007). Krichmar applies these eight principles to biologically inspired cognitive architectures (Krichmar, 2012), as follows. In each case, we identify the corresponding desiderata.

Principle 1 — The Three-Constituents Principle — is that an intelligent agent should have (a) a defined ecological niche, (b) a defined behavior, and (c) an agent design. This principle does not speak directly to development, per se, but neverthless (a) and (b) are reflected in Desiderata No. 1 Value Systems and Motives whereas (c) is essentially the same as the overall goal of the exercise: to identify the constituents of a cognitive architecture. Krichmar notes that while many cognitive robots are built to test theories of biological cognition and are therefore tested in conditions that are similar to the heavily restrictive environments of these cognitive science experiments, on the basis that conditions must be similar if one it to make meaningful claims about the theory, it is also essential that these cognitive robots can make a successful transition to the real world outside the laboratory, just as humans and animal do.

Principle 2 — The Complete Agent — is that one must consider the complete agent behaving in the real world when designing the agent. This focus on a balanced blend of brain, body, and environment — a embodied brain embedded in an environment, matching morphology and ecological niche — mirrors Desiderata No. 2, Physical Embodiment.

Principle 3 — Cheap Design — is that the construction and design of agents that are built to exploit properties of the ecological niche will be much easier or cheaper. This off-loading of cognitive process to the embodiment, exploiting the well-tuned morphology, kinematics, and dynamics to accomplish the necessary control rather than having it effected by a central processing is a key aspect of the embodiment thesis as noted in Desideratum No. 2 Physical Embodiment.

Principle 4 — Redundancy — asserts that agents should be designed such that different subsystems function on the basis of different processes but that there should be an overlap of functionality between subsystems. Krichmar adopts the term degeneracy to better reflect the ability of differently structured subsystems to accomplish the same function. Although no desideratum corresponds directly to this principle, the emphasis in Desideratum No. 10, Constitutive Autonomy, on self-organization and self-maintenance to promote continued behavioral autonomy is relevant to the manner in which this degeneracy is managed.

Principle 5 — Sensory-Motor Coordination — highlights the advantages of a cognitive robot embedded in its environment being able to induce sensory stimulation through coordinated motor-based actions. This greatly simplifies the processing required to interpret the sensory information and make sense of the environment. This principle is reflected in Desideratum No. 3, Sensorimotor Contingencies.

Principle 6 — Ecological Balance — emphasizes the need to evenly balance the complexity of the agent's sensory, motor, and neural systems, and, especially, that the agent's morphology should balance the nature and complexity of its environment and ecological niche. Krichmar makes the point that the main functions of the central nervous system (and, implicitly, cognition) are predicting and planning for future events, leaving the agent's body itself to take responsibility for interaction with the environment without direct neural control. This principle is reflected in Desideratum No. 2, Physical Embodiment.

Principle 7 — Parallel, Loosely Coupled Processes — asserts that intelligence is emergent from a large number of parallel processes which are coordinated through embodied interaction with the environment. Consequently, cognitive architectures should be designed as complete systems comprising loosely coupled processes, operating asynchronously and concurrently, reflecting the manner in which events in the environment occur. This principle reflects von Foerster's caveat but is not explicitly addressed by the ten desiderata.

Principle 8 — Value — states that intelligent agents are equipped with a value system that encapsulates what is good and bad for the agent. Krichmar emphasizes that this value system is innate and forms the basis for the agent's learned ability to predict the way the environment operates to the agent's advantage. He notes too the importance of linking this value system to the state and operation of the agent's body. This principle directly is reflected in Desideratum No. 1, Value Systems and Motives.

At the outset, we noted our motivation in proposing these desiderata is to provide a minimal set of characteristics that can guide the design of either a developmental cognitive architecture or a developmental cognitive architecture schema. We do not claim that this list is exhaustive. There are certainly other characteristics and mechanisms required for development. Language, for example, is an important aspect of human interaction and plays a crucial roles in later phases of cognitive development. At the same time, there is also the open question as to whether or not language is innate and present at birth. If it is, then certainly a capability for language development would merit being included in these desiderata. Perhaps, though, there is another way of viewing this question. We know that Broca's area might realize a supramodal or polymodal sensorimotor representation of hierarchical syntactic structures and, as such, it appears that this machinery is a multi-purpose syntactic processor, ready-for-language, but not specific-for-language. Thus, rather than language being a desirable characteristic of a developmental cognitive architecture, on the contrary, one might argue for the desirability of a mechanism that facilitates the realization of such hierarchical syntactic structures and the development of language, among other things.

7. Conclusions

The ambitions of this short article are relatively modest. We wish to draw attention to the pivotal role played by development in cognitive agents in transforming cognitive *capacity* into cognitive *ability* over time. In particular, we seek to highlight some of the key elements that should, ideally, be present in a biologically-inspired cognitive architecture if it is to facilitate such development. In doing this, we have focussed on the ultimate aspects of cognitive development — why that feature is necessary and what it enables — rather on than the proximate mechanisms by which they can be realized (Scott-Phillips, Dickins, & West, 2011). As such, the desiderata are for the most part neutral regarding the paradigm of cognitive science — cognitivist or emergent — that is adopted when designing a cognitive architecture. From an ultimate perspective, it taken as a given that the development facilitated by the cognitive architecture involves interaction between the resultant cognitive agent (e.g. a cognitive robot) and humans, and furthermore that the developing cognitive agent will be capable of meaningful social interaction, with all that entails for engagement, cooperation, collaboration, trust, and safety.

Acknowledgements

This work was supported by the European Commission (Grant agreement No: 688441; Action RockEU2).

References

Anderson, J. R. (1983). *The architecture of cognition*. Cambridge MA: Harvard University Press.
Anderson, J. R. (1996). Act: A simple theory of complex cognition. *American Psychologist, 51*, 355–365.
Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review, 111*, 1036–1060.
Anderson, M. L. (2003). Embodied cognition: A field guide. *Artificial Intelligence, 149*, 91–130.
Argall, B. D., Chernova, S., Veloso, M., & Browning, B. (2009). A survey of robot learning from demonstration. *Robotics and Autonomous Systems, 57*, 469–483.
Ashley, J., & Tomasello, M. (1998). Cooperative problem solving and teaching in preschoolers. *Social Development, 7*, 143–163.
Barandiaran, X., & Moreno, A. (2008). Adaptivity: From metabolism to behavior. *Adaptive Behavior, 16*, 325–344.
Barsingerhorn, A. D., Zaal, F. T. J. M., Smith, J., & Pepping, G. -J. (2012). On the possibilities for action: The past, present and future of affordance research. AVANT, III.
Bickhard, M. H. (2000). Autonomy, function, and representation. *Communication and Control—Artificial Intelligence, 17*, 111–131.
Billard, A. (2002). Imitation. In M. A. Arbib (Ed.), *The handbook of brain theory and neural networks* (pp. 566–569). Cambridge, MA: MIT Press.
Billing, E., & Balkenius, C. (2014). Modeling the interplay between conditioning and attention in a humanoid robot: Habituation and attentional blocking. In *Proceedings of the 4th international conference on development and learning and on epigenetic robotics (IEEE ICDL-EPIROB)* (pp. 41–47).
Boden, M. A. (2008). Autonomy: What is it? *BioSystems, 91*, 305–308.
Brownell, C. A., Ramani, G. B., & Zerwas, S. (2006). Becoming a social partner with peers: cooperation and social understanding in one- and two-year-olds. *Child Development, 77*, 803–821.
Campanella, F., Sandini, G., & Morone, M. C. (2011). Visual information gleaned by observing grasping movement in allocentric and egocentric perspectives. *Proceedings of the Royal Society B, 278*, 2142–-2149.
Cangelosi, A., & Schlesinger, M. (2015). *Developmental robotics: From babies to robots*. Cambridge MA: MIT Press.
Chrisley, R., & Ziemke, T. (2002). Embodiment. In *Encyclopedia of cognitive science* (pp. 1102–1108). Macmillan.
Craighero, L., Fadiga, L., Rizzolatti, G., & Umiltà, C. A. (1999). Movement for perception: A motor-visual attentional effect. *Journal of Experimental Psychology: Human Perception and Performance, 25*, 1673–1692.
Craighero, L., Nascimben, M., & Fadiga, L. (2004). Eye position affects orienting of visuospatial attention. *Current Biology, 14*, 331–333.
Damasio, A., & Carvalho, G. B. (2013). The nature of feelings: Evolutionary and neurobiological origins. *Nature Reviews Neuroscience, 14*, 143–152.
Damasio, A. R. (2003). *Looking for Spinoza: Joy, sorrow and the feeling brain*. Orlando, Florida: Harcourt.
Dautenhahn, K., & Billard, A. (1999). Studying robot social cognition within a developmental psychology framework. In *Proceedings of Eurobot 99: Third European workshop on advanced mobile robots* (pp. 187–194). Switzerland.
Doya, K. (1999). What are the computations of the cerebellum, the basal ganglia and the cerebral cortex? *Neural Networks, 12*, 961–974.
Doya, K. (2000). Complementary roles of basal ganglia and cerebellum in learning and motor control. *Current Opinion in Neurobiology, 10*, 732–739.
Edelman, G. M. (2006). *Second nature: Brain science and human knowledge*. New Haven and London: Yale University Press.
Farroni, T., Csibra, G., Simion, F., & Johnson, M. H. (2002). Eye contact detection in humans from birth. *Proceeding of the National Academy of Sciences (PNAS), 99*, 9602–9605.
von Foerster, H. (2003). *Understanding: Essays on cybernetics and cognition*. New York: Springer.
Forgas, J. P., Kipling, K. D., & Laham, S. M. (Eds.). (2005). *Social motivation*. Cambridge University Press.
Forster, S., & Lavie, N. (2015). Establishing the attention-distractibility trait. *Psychological Science*, 1–10.
Fox, R., & McDaniel, C. (1982). The perception of biological motion by human infants. *Science, 218*, 486–487.
Froese, T., Virgo, N., & Izquierdo, E. (2007). Autonomy: A review and a reappraisal. In F. Almeida e Costa, L. Rocha, E. Costa, I. Harvey, & A. Coutinho (Eds.). *Proceedings of the 9th European conference on artificial life: Advances in artificial life* (Vol. 4648, pp. 455–465). Berlin Heidelberg: Springer. http://dx.doi.org/10.1007/978-3-540-74913-4_46.
Froese, T., & Ziemke, T. (2009). Enactive artificial intelligence: Investigating the systemic organization of life and mind. *Artificial Intelligence, 173*, 466–500.
Gallese, V., Fadiga, L., Fogassi, L., & Rizzolatti, G. (1996). Action recognition in the premotor cortex. *Brain, 119*, 593–609.
Gibson, J. J. (1977). The theory of affordances. In R. Shaw & J. Bransford (Eds.), *Perceiving, acting and knowing: Toward an ecological psychology* (pp. 67–82). Lawrence Erlbaum.
Gray, W. D., Young, R. M., & Kirschenbaum, S. S. (1997). Introduction to this special issue on cognitive architectures and human-computer interaction. *Human-Computer Interaction, 12*, 301–309.
Grush, R. (2004). The emulation theory of representation: Motor control, imagery, and perception. *Behavioral and Brain Sciences, 27*, 377–442.
Hermer, L., & Spelke, E. S. (1996). Modularity and development: The case of spatial reorientation. *Cognition, 61*.
Hesslow, G. (2002). Conscious thought as simulation of behaviour and perception. *Trends in Cognitive Sciences, 6*, 242–247.
Hesslow, G. (2012). The current status of the simulation theory of cognition. *Brain Research, 1428*, 71–79.
von Hofsten, C. (1980). Predictive reaching for moving objects by human infants. *Journal of Experimental Child Psychology, 30*, 369–382.
von Hofsten, C. (2003). On the development of perception and action. In J. Valsiner & K. J. Connolly (Eds.), *Handbook of developmental psychology* (pp. 114–140). London: Sage.
von Hofsten, C. (2004). An action perspective on motor development. *Trends in Cognitive Sciences, 8*, 266–272.
von Hofsten, C. (2009). Action, the foundation for cognitive development. *Scandinavian Journal of Psychology, 50*, 617–623.
von Hofsten, C. (2013). *Action science: The emergence of a new discipline. Chapter: Action in infancy: A foundation for cognitive development*. MIT Press.
Huang, X., & Weng, J. (2002). Novelty and reinforcement learning in the value system of developmental robots. In *Proc. second international workshop on epigenetic robotics: Modeling cognitive development in robotic systems (EPIROB '02)* (pp. 47–55). Edinburgh, Scotland.
Huang, X., & Weng, J. (2007). Inherent value systems for autonomous mental development. *International Journal of Humanoid Robotics, 4*, 407–-433.
Jamone, L., Ugur, E., Cangelosi, A., Fadiga, L., Bernardino, A., Piater, J., et al. (2016). Affordances in psychology, neuroscience and robotics: A survey. *IEEE Transactions on Cognitive and Developmental Systems, PP*. 1–1.
Karmiloff-Smith, A. (1992). *Beyond modularity: A developmental perspective on cognitive science*. Cambridge, MA: MIT Press.
Karmiloff-Smith, A. (1994). Precis of beyond modularity: A developmental perspective on cognitive science. *Behavioral and Brain Sciences, 17*, 693–745.

Kelley, T. A., Rees, G., & Lavie, N. (2013). The impact of distractor congruency on stimulus processing in retinotopic visual cortex. *NeuroImage, 81*, 158–163.

Krichmar, J. L. (2012). Design principles for biologically inspired cognitive architectures. *Biologically Inspired Cognitive Architectures, 1*, 73–81.

Krichmar, J. L., & Edelman, G. M. (2005). Brain-based devices for the study of nervous systems and the development of intelligent machines. *Artificial Life, 11*, 63–77.

Krichmar, J. L., & Edelman, G. M. (2006). Principles underlying the construction of brain-based devices. In T. Kovacs & J. A. R. Marshall (Eds.), *Proceedings of AISB '06 – Adaptation in artificial and biological systems* (pp. 37–42). Bristol: University of Bristol volume 2 of Symposium on Grand Challenge 5: Architecture of Brain and Mind.

Krichmar, J. L., & Reeke, G. N. (2005). The Darwin brain-based automata: Synthetic neural models and real-world devices. In G. N. Reeke, R. R. Poznanski, K. A. Lindsay, J. R. Rosenberg, & O. Sporns (Eds.), *Modelling in the neurosciences: From biological systems to neuromimetic robotics* (pp. 613–638). Boca Raton: Taylor and Francis.

Laird, J. E. (2008). Extending the soar cognitive architecture. In *Proceedings of the first conference on artificial general intelligence* (pp. 224–235). Amsterdam, The Netherlands: IOS Press.

Laird, J. E. (2009). Towards cognitive robotics. In G. R. Gerhart, D. W. Gage, & C. M. Shoemaker (Eds.). *Proceedings of the SPIE — Unmanned systems technology XI* (Vol. 7332). , pp. 73320Z–73320Z-11.

Laird, J. E. (2012). *The soar cognitive architecture*. Cambridge, MA: MIT Press.

Laird, J. E., Newell, A., & Rosenbloom, P. S. (1987). Soar: An architecture for general intelligence. *Artificial Intelligence, 33*.

Langley, P. (2005). An adaptive architecture for physical agents. In *IEEE/WIC/ACM international conference on intelligent agent technology* (pp. 18–25). Compiegne, France: IEEE Computer Society Press.

Langley, P., Laird, J. E., & Rogers, S. (2009). Cognitive architectures: Research issues and challenges. *Cognitive Systems Research, 10*, 141–160.

Lehman, J. F., Laird, J. E., & Rosenbloom, P. S. (1998). A gentle introduction to soar, an architecture for human cognition. In S. Sternberg & D. Scarborough (Eds.), *Invitation to cognitive science. Methods, models, and conceptual issues* (Vol. 4). Cambridge, MA: MIT Press.

Lindblom, J. (2015). *Embodied social cognition volume 26 of cognitive systems monographs (COSMOS)*. Berlin: Springer.

Lindblom, J., & Ziemke, T. (2003). Social situatedness of natural and artificial intelligence: Vygotsky and beyond. *Adaptive Behavior, 11*, 79–96.

Lintern, G. (1998). Encountering the world: Toward an ecological psychology by Edward S. Reed. *Complexity, 3*, 61–63.

Maturana, H., & Varela, F. (1987). *The tree of knowledge — The biological roots of human understanding*. Boston & London: New Science Library.

McClelland, J. L., NcNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review, 102*, 419–457.

Meltzoff, A. N. (2002). The elements of a developmental theory of imitation. In A. N. Meltzoff & W. Prinz (Eds.), *The imitative mind: Development, evolution, and brain bases* (pp. 19–41). Cambridge: Cambridge University Press.

Meltzoff, A. N., & Moore, M. K. (1977). Imitation of facial and manual gestures by human neonates. *Science, 198*, 75–78.

Meltzoff, A. N., & Moore, M. K. (1997). Explaining facial imitation: A theoretical model. *Early Development and Parenting, 6*, 179–192.

Merrick, K. E. (2010). A comparative study of value systems for self-motivated exploration and learning by robots. *IEEE Transactions on Autonomous Mental Development, 2*, 119–131.

Meyer, M., Bekkering, H., Paulus, M., & Hunnius, S. (2010). Joint action coordination in 2 1/2- and 3-year-old children. *Frontiers in Human Neuroscience, 4*, 1–7.

Morse, A., Lowe, R., & Ziemke, T. (2008). Towards an enactive cognitive architecture. In *Proceedings of the first international conference on cognitive systems*. Karlsruhe, Germany.

Moulton, S. T., & Kosslyn, S. M. (2009). Imagining predictions: Mental imagery as mental emulation. *Philosophical Transactions of the Royal Society B, 364*, 1273–1280.

Newell, A. (1990). *Unified theories of cognition*. Cambridge MA: Harvard University Press.

Núñez, R., & Freeman, W. J. (1999). *Reclaiming cognition — The primacy of action, intention and emotion*. Thorverton, UK: Imprint Academic.

O'Regan, J. K., & Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences, 24*, 883–-917.

Oudeyer, P., Kaplan, F., & Hafner, V. (2007). Intrinsic motivation systems for autonomous mental development. *IEEE Transactions on Evolutionary Computation, 11*, 265–286.

Pfeifer, R., & Bongard, J. (2007). *How the body shapes the way we think: A new view of intelligence*. Cambridge, MA: MIT Press.

Piaget, J. (1936). *The origin of intelligence in children*. London: International University Press, Inc. and Routledge & Kegan Paul Ltd.. Originally published in French in 1936.

Piaget, J. (1954). *The construction of reality in the child*. New York: Basic Books.

Reed, E. S. (1996). *Encountering the world: Towards an ecological psychology*. New York: Oxford University Press.

Ritter, F. E., & Young, R. M. (2001). Introduction to this special issue on using cognitive models to improve interface design. *International Journal of Human-Computer Studies, 55*, 1–14.

Rizzolatti, G., Fadiga, L., Gallese, V., & Fogassi, L. (1996). Premotor cortex and the recognition of motor actions. *Cognitive Brain Research, 3*, 131–141.

Rolf, M., & Steil, J. J. (2012). Goal babbling: A new concept for early sensorimotor exploration. In E. Ugur, Y. Nagai, E. Oztop, & M. Asada (Eds.), *Proceedings of Humanoids 2012 workshop on developmental robotics: Can developmental robotics yield human-like cognitive abilities?* (pp. 40–43). Osaka, Japan.

Rolf, M., Steil, J. J., & Gienger, M. (2010). Goal babbling permits direct learning of inverse kinematics. *IEEE Transactions on Autonomous Mental Development, 2*, 216–229.

Rosenbloom, P., Laird, J., & Newell, A. (Eds.). (1993). *The soar papers: Research on integrated intelligence*. Cambridge, Massachusetts: MIT Press.

Ryle, G. (1949). *The concept of mind*. London: Hutchinson's University Library.

Schacter, D. L., Addis, D. R., & Buckner, R. L. (2008). Episodic simulation of future events: Concepts, data, and applications. *Annals of the New York Academy of Sciences, 1124*, 39–60.

Schulkin, J. (2011). Social allostasis: Anticipatory regulation of the internal milieu. *Frontiers in Evolutionary Neuroscience, 2*, 1–15.

Scott-Phillips, T. C., Dickins, T. E., & West, S. A. (2011). Evolutionary theory and the ultimate-proximate distinction in the human behavioural sciences. *Perspectives on Psychological Science, 6*, 38–47.

Seligman, M. E. P., Railton, P., Baumeister, R. F., & Sripada, C. (2013). Navigating into the future or driven by the past. *Perspectives on Psychological Science, 8*, 119–141.

Shapiro, L. (2011). *Embodied cognition*. Routledge.

Simion, F., Regolin, L., & Bulf, H. (2008). A predisposition for biological motion in the newborn baby. *Proceeding of the National Academy of Sciences (PNAS), 105*, 809–813.

Spelke, E. S. (2000). Core knowledge. *American Psychologist*, 1233–1243.

Squire, L. (2004). Memory systems of the brain: A brief history and current perspective. *Neurobiology of Learning and Memory, 82*, 171–177.

Sterling, P. (2004). Principles of allostasis. In J. Schulkin (Ed.), *Allostasis, homeostasis, and the costs of adaptation* (pp. 17–64). Cambridge, England: Cambridge University Press.

Sterling, P. (2012). Allostasis: A model of predictive regulation. *Physiology and Behaviour, 106*, 5–15.

Stewart, J., Gapenne, O., & Di Paolo, E. A. (2010). *Enaction: Toward a new paradigm for cognitive science*. Cambridge MA: MIT Press.

Sun, R. (2003). A tutorial on CLARION 5.0. <http://www.cogsci.rpi.edu/rsun/sun.tutorial.pdf>.

Sun, R. (2004). Desiderata for cognitive architectures. *Philosophical Psychology, 17*, 341–373.

Sun, R. (2007). The importance of cognitive architectures: An analysis based on clarion. *Journal of Experimental & Theoretical Artificial Intelligence, 19*, 159–193.

Svensson, H., Lindblom, J., & Ziemke, T. (2007). Making sense of embodied cognition: Simulation theories of shared neural mechanisms for sensorimotor and cognitive processes. In T. Ziemke, J. Zlatev, & R. M. Frank (Eds.), *Body, language and mind* (pp. 241–269). Berlin: Mouton.

Svensson, H., Morse, A. F., & Ziemke, T. (2009). Representation as internal simulation: A minimalistic robotic model. In N. Taatgen & H. van Rijn (Eds.), *Proceedings of the thirty-first annual conference of the cognitive science society* (pp. 2890–2895). Austin, TX: Cognitive Science Society.

Svensson, H., Thill, S., & Ziemke, T. (2013). Dreaming of electric sheep? Exploring the functions of dream-like mechanisms in the development of mental imagery simulations. *Adaptive Behavior, 21*, 222–238.

Thelen, E. (1995). Time-scale dynamics and the development of embodied cognition. In R. F. Port & T. van Gelder (Eds.), *Mind as motion – Explorations in the dynamics of cognition* (pp. 69–100). Cambridge, Massachusetts: Bradford Books, MIT Press.

Thelen, E., & Smith, L. B. (1998). *A dynamic systems approach to the development of cognition and action* (3rd ed.). Cambridge, Massachusetts: MIT Press.

Thelen, E., & Smith, L. B. (2003). Development as a dynamic system. *Trends in Cognitive Sciences, 7*, 343–348.

Thill, S., & Svensson, H. (2011). The inception of simulation: A hypothesis for the role of dreams in young children. In L. Carlson, C. Hoelscher, & T. F. Shipley (Eds.), *Proceedings of the thirty-third annual conference of the cognitive science society* (pp. 231–236). Austin, TX: Cognitive Science Society.

Tulving, E. (1972). Episodic and semantic memory. In E. Tulving & W. Donaldson (Eds.), *Organization of memory* (pp. 381–403). New York: Academic Press.

Tulving, E. (1984). Précis of *elements of episodic memory*. *Behavioral and Brain Sciences, 7*, 223–268.

Vernon, D. (2010). Enaction as a conceptual framework for development in cognitive robotics. *Paladyn Journal of Behavioral Robotics, 1*, 89–98.

Vernon, D. (2014). *Artificial cognitive systems — A primer*. Cambridge, MA: MIT Press.

Vernon, D. (2016). Reconciling constitutive and behavioural autonomy: The challenge of modelling development in enactive cognition. *Intellectica: The Journal of the French Association for Cognitive Research, 65*, 63–79.

Vernon, D., Beetz, M., & Sandini, G. (2015a). Prospection in cognitive robotics: The case for joint episodic-procedural memory. *Frontiers in Robotics and AI, 2*, 1–14.

Vernon, D., von Hofsten, C., & Fadiga, L. (2010). *A roadmap for cognitive development in humanoid robots volume 11 of cognitive systems monographs (COSMOS)*. Berlin: Springer.

Vernon, D., Lowe, R., Thill, S., & Ziemke, T. (2015b). Embodied cognition and circular causality: On the role of constitutive autonomy in the reciprocal coupling of perception and action. *Frontiers in Psychology, 6*, 1–13.

Vernon, D., Metta, G., & Sandini, G. (2007). A survey of artificial cognitive systems: Implications for the autonomous development of mental capabilities in computational agents. *IEEE Transactions on Evolutionary Computation, 11*, 151–180.

Vygotsky, L. (1978). *Mind in society: The development of higher psychological processes*. Cambridge, MA: Harvard University Press.
Wang, R. F., & Spelke, E. S. (2002). Human spatial representation: Insights from animals. *Trends in Cognitive Sciences, 6*, 376–382.
Warneken, F., Chen, F., & Tomasello, M. (2006). Cooperative activities in young children and chimpanzees. *Child Development, 77*, 640–663.
Warneken, F., & Tomasello, M. (2009). The roots of human altruism. *British Journal of Psychology, 100*, 455–471.
Weng, J. (2002). A theory for mentally developing robots. In *Proceedings of the 2nd international conference on development and learning (ICDL 2002)*. IEEE Computer Society.
Weng, J. (2004a). Developmental robotics: Theory and experiments. *International Journal of Humanoid Robotics, 1*, 199–236.
Weng, J. (2004b). A theory of developmental architecture. In *Proceedings of the 3rd international conference on development and learning (ICDL 2004)*. La Jolla.
Wilson, R. A., & Foglia, L. (2011). Embodied cognition. In E. N. Zalta (Ed.), *The stanford encyclopedia of philosophy*. .
Wintermute, S. (2012). Imagery in cognitive architecture: Representation and control at multiple levels of abstraction. *Cognitive Systems Research*, 1–29.
Wood, R., Baxter, P., & Belpaeme, T. (2012). A review of long-term memory in natural and synthetic systems. *Adaptive Behavior, 20*, 81–103.
Ziemke, T. (2003). What's that thing called embodiment? In R. Alterman & D. Kirsh (Eds.), *Proceedings of the 25th annual conference of the cognitive science society Lund University cognitive studies* (pp. 1134–1139). Mahwah, NJ: Lawrence Erlbaum.
Ziemke, T., & Lowe, R. (2009). On the role of emotion in embodied cognitive architectures: From organisms to robots. *Cognition and Computation, 1*, 104–117.