PALADYN Journal of Behavioral Robotics

Research Article · DOI: 10.2478/s13230-010-0016-y · JBR · 1(2) · 2010 · 89-98

Enaction as a Conceptual Framework for Developmental Cognitive Robotics

David Vernon*

Department of Robotics, Brain, and Cognitive Sciences, Italian Institute of Technology, Genoa, Italy

Received 9 June 2010

Accepted 27 August 2010

Abstract

This paper provides an accessible introduction to the cognitive systems paradigm of enaction and shows how it forms a practical framework for robotic systems that can develop cognitive abilities. The principal idea of enaction is that a cognitive system develops it own understanding of the world around it through its interactions with the environment. Thus, enaction entails that the cognitive system operates autonomously and that it generates its own models of how the world works. A discussion of the five key elements of enaction — autonomy, embodiment, emergence, experience, and sense-making — leads to a core set of functional, organizational, and developmental requirements which are then used in the design of a cognitive architecture for the iCub humanoid robot.

Keywords

enaction · enactive systems · cognition · autonomy · embodiment · emergence · experience · sense-making

1. Introduction

The field of cognitive robotics looks increasingly to developmental techniques to provide a way of acquiring knowledge and learning new skills. Often, we seek to apply our knowledge of development in natural cognitive systems, *i.e.* human infants, to the problem of creating artificial cognitive systems, typically in the guise of humanoid robots. In the following, we discuss *enaction*[1] as a conceptual framework which identifies and explains a specific developmental stance on cognitive systems that views them as emergent embodied systems which develop cognitive skills as a result of their action in the world. Enaction draws out explicitly the theoretical and practical consequences of adopting this stance and thereby provides a clear conceptual framework within which to position our insights from developmental psychology.

We begin by considering the operational characteristics of a cognitive system, focussing on the purpose of cognition rather than debating the relative merits of competing paradigms of cognition [1]. Of course, such a debate is important because it allows us to understand the preconditions for cognition so, once we have established the role cognition plays and see why it is important, we move on to elaborate on these pre-conditions and we introduce the underlying framework of enaction which we adopt as the basis for the research described in this paper.

By working through the implications of the enactive approach to cognition, the central role of development in cognition becomes clear, as do several other key issues such as the crucial role played by action, the inter-dependence of perception and action, and the consequent constructivist nature of the cognitive system's knowledge.

The framework of enaction provides a foundation for the identification of the phylogeny and the ontogeny of cognitive systems: their initial capabilities and subsequent development.

2. Cognition

Cognitive systems anticipate, assimilate, and adapt. In doing so, they learn and develop[2] [1]. Cognitive systems anticipate future events when selecting actions, they subsequently learn from what actually happens when they do act, and thereby they modify subsequent expectations and, in the process, they change how the world is perceived and what actions are possible. Cognitive systems do all of this autonomously. The adaptive, anticipatory, autonomous viewpoint reflects the position of Freeman and Núñez who, in their book *Reclaiming Cognition* [2], assert the primacy of action, intention, and emotion in cognition. In the past, however, cognition was viewed in a very different light as a symbol-processing module of mind concerned with rational planning and reasoning. Today, however, this is changing and even proponents of these early approaches now see a much tighter relationship between perception, action, and cognition (*e.g.* see [3, 4]).

---

*E-mail: david@vernon.eu

[1] *The enactive approach is quite subtle and it has its own well-established but idiosyncratic terminology. To make this paper as accessible as possible, the more subtle issues are dealt with in footnotes in order to keep the main narrative text as fluid as possible.*

[2] *The difference between learning and development is not always clear. Some researchers argue that development is simply a specific form of learning. Others view learning as a process for estimating or improving the parameter values that govern the behaviour of a known or operational model, with development being a process for generating or discovering the model itself. We adhere to the second view in this paper. We would also remark that learning arises as a consequence of the interaction between the cognitive agent and the world around him, whereas development arises from learning as a consequence of the interaction of the cognitive agent with itself. As we will see later on, this position is consistent with the enactive stance on cognition (see Section 3). Both learning and development are forms of adaptation.*

So, if cognitive systems anticipate, assimilate, and adapt, if they develop and learn, it is natural to ask is *why* do they do this?

The view of cognition taken in this paper is that cognition is the process whereby an autonomous self-governing system acts effectively in the world in which it is embedded [5]. However, in natural systems, the latencies inherent in the neural processing of sense data are too great to allow effective action. This is one of the primary reasons a cognitive agent must anticipate future events: so that it can prepare the actions it may need to take. In addition, there are also limitations imposed by the environment and the cognitive system's body. To perform an action, one needs to have the relevant body part in a certain place at a certain time. In a dynamic environment that is constantly changing and with a body that takes time to move, this requires preparation and prediction. Furthermore, the world in which the agent is embedded is unconstrained and the sensory data which is available to the cognitive system is not only 'out-of-date' but it is also uncertain and incomplete. Consequently, it is not possible to encapsulate ***a priori*** all the knowledge required to deal successfully with the circumstances it will experience so that it must also be able to adapt, progressively increasing its space of possible actions as well as the time horizon of its prospective capabilities. It must do this, not as a reaction to external stimuli but as a self-generated process of proactive understanding.[3] This process is what we mean by development.

In summary, the position being set out in this paper is that (a) cognition is the process by which an autonomous self-governing agent acts effectively in the world in which it is embedded, that (b) the dual purpose of cognition is to increase the agent's repertoire of effective actions and its power to anticipate the need for and outcome of future actions, and that (c) development plays an essential role in the realization of these cognitive capabilities.

We will now introduce a framework which encapsulates all these considerations.

3. Enaction

There are many alternative perspectives on cognition: what it is, why it is necessary, and how it is achieved. We have already argued that cognition arises from an agent's need to compensate for latencies in neural processing by anticipating what may be about to happen and by preparing its actions accordingly. So we can agree fairly easily what cognition is — a process of anticipating events and acting appropriately and effectively[4] — and why it is necessary — to overcome the physical limitations of biological brains and the limitations of bodily movements operating in a dynamic environment. The difficulty arises when we consider how cognition is achieved. There are several competing theories of cognition,[5] each of which makes its own set of assumptions. Here, we wish to focus on one particularly important paradigm — enaction — and pick out its most salient aspects in order to provide a sound theoretical foundation for the role of development in cognition [5, 9–14]
.

The principal idea of enaction is that a cognitive system develops it own understanding of the world around it through its interactions with the environment. Thus, enaction entails that the cognitive system operates autonomously and that it generates its own models of how the world works. When dealing with enactive systems, there are five key elements to consider [15, 16]:

1. Autonomy
2. Embodiment
3. Emergence
4. Experience
5. Sense-making

We have already mentioned autonomy. It is the self-maintaining organizational characteristic of living creatures that enables them to use their own capacities to manage their interactions with the world, and with themselves, in order to remain viable [17, 18].[6] This simply means that the system is entirely self-governing and self-regulating: it is not controlled by any outside agency and this allows it to stand apart from the rest of the environment and operate independently of it. That's not to say that the system isn't influenced by the world around it, but rather that these influences are brought about through interactions that don't threaten the autonomous operation of the system.[7]

---

[3] *As we will see in the next section, this is often referred to as a process of 'sense-making'.*

[4] *Once you have the cognitive ability to anticipate events, you can use that ability in other ways, not just as a means to overcome the delays introduced by neural processing and as a means to react quickly to events in the world. For example, you can reverse the prediction process to try to explain how something might have happened or you can imagine alternative outcomes based on slightly different circumstances or associations. Thus, from a need to predict outcomes to deal effectively with dynamic events, cognition also brings with it an ability to construct explanations and imagine unforeseen events, both of which can be used to expand the agent's repertoire of actions and enhance its ability to interact effectively with the world around it, all the while maintaining its autonomy.*

[5] *Many of the assumptions made in various theories of cognition concern the workings of cognitive systems and how information is represented. Other assumptions concern some deep philosophical issues regarding ontology and epistemology: the nature of reality and how we come to know it. For this reason, we often speak of different paradigms of cognition rather than different theories of cognition since each posits an entirely different view of the world and how we apprehend it. A review of these paradigms and their relative merits is outside the scope of this paper but the interested reader might wish to refer to other sources, such as [1, 6–8], to get a sense of the diversity of approaches that exist.*

[6] *Froese and Ziemke [18] refine the concept of autonomy further, introducing the term* constitutive autonomy *to denote a particular form of autonomy which results from* organizational closure, *Maturana's and Varela's generalization of the self-organizational process of* autopoiesis *(literally, self-production) [5, 11, 12] to include any system that actively generates and sustains its existence and systemic identity under "precarious conditions",* i.e. *conditions which would cause the system to cease to exist and its identity to disappear. Froese and Ziemke argue that constitutive autonomy entails the emergence of an identity and is necessary for intrinsic teleology (*i.e. *internally-generated intentionality) and sense-making. They also note that adaptivity — an implicit capacity of constitutive autonomy to actively regulate sensorimotor interaction in relation to some viability constraint — is also required for sense-making. Ziemke and Lowe have deployed constitutive autonomy in their Cognitive-Affective Architecture Schematic in which different levels of cognitive function and behavioural complexity are associated with, and are brought about by, different levels of emotion, each linked to affective homeostatic processes ranging from reflexes right through to internal simulation [19]. Autopoiesis is a special type of autonomous self-organization: an autopoietic system is a homeostatic system (*i.e. *self-regulating system) but one in which the regulation applies not to some system parameter but to the organization of the system itself. Thus, autopoiesis is a process whereby a system reinstantiates itself through a network of relations between its component entities. Autopoiesis was originally introduced in the context of cellular dynamics but its central tenets have been generalized to embrace other autonomous systems; see Footnote 8 for more details.*

[7] *When an influence on a system isn't directly controlling it but nonetheless has some impact on the behaviour of the system, we refer to it as a* perturbation.

The second element of enaction is the idea of embodiment. For our purposes here, embodiment means that the system must exist in the world as a physical entity which can interact directly with the environment. This means the system can act on things in the world around it and they, in turn, can act on the system. These things can be inanimate objects or animate agents, cognitive or not. As it happens, there are some subtle distinctions which can be made about different types of embodiment and we will return to this topic later in Section 4.

The element of emergence refers to the manner in which cognition arises in the system. Specifically, it refers to the laws and mechanisms which govern the behaviour of the component parts of the system. In an emergent system, the behaviour we call cognition arises from the dynamic interplay between the components and the laws and mechanisms we mentioned govern only the behaviour of the component parts; they don't specify the behaviour of the interplay between the components. Thus, behaviour emerges indirectly because of the internal dynamics. Crucially, these internal dynamics must maintain the autonomy of the system and, as we will see shortly, they also condition the experiences of the system through their embodiment in a specific structure.

Experience is the fourth element of enaction. This is nothing more than the cognitive system's history of interaction with the world around it: the actions it takes in the environment and the actions arising in the environment which impinge on the cognitive system. These interactions don't control the system (otherwise it wouldn't be autonomous) but they do trigger changes in the state of the system. The changes that can be triggered are *structurally determined*: they depend on the system structure, *i.e.* the embodiment of the self-organizational principles that make the system autonomous.[8] This structure is also referred to as the system's phylogeny: the innate capabilities of an autonomous system with which it is equipped at the outset and which form the basis for its continued existence. The experience of the systems — its history of interactions — involving *structural coupling* between the system and its environment is referred to as its ontogeny.

Finally, we come to the fifth and, arguably, the most important element of enaction: sense-making. This term refers to the relationship between the knowledge encapsulated by a cognitive system and the interactions which gave rise to it. In particular, it refers to the idea that this emergent knowledge is generated by the system itself and that it captures some regularity or lawfulness in the interactions of the system, *i.e.* its experience. However, the sense it makes is dependent on the way in which it can interact: its own actions and its perceptions of the environment's action on it. Since these perceptions and actions are the result of an emergent dynamic process that is first and foremost concerned with maintaining the autonomy and operational identity of the system, these perceptions and actions are unique to the system itself and the resultant knowledge makes sense only insofar as it contributes to the maintenance of the system's autonomy. This ties in neatly with our view of cognition, the role of which is to anticipate events and increase the space of actions in which a system can engage. By making sense of its experience, the cognitive system is constructing a model that has some predictive value, exactly because it captures some regularity or lawfulness in its interactions. This self-generated model of the system's experience lends the system greater flexibility in how it interacts in the future. In other words, it endows the system with a larger repertoire of possible actions that allow richer interactions, increased perceptual capacity, and the possibility of constructing even better models that encapsulate knowlege with even greater predictive power. And so it goes, in a virtuous circle. Note that this sense-making and the resultant knowledge says nothing at all about what is really out there in the environment; it doesn't have to. All it has to do is make sense for the continued existence and autonomy of the cognitive system. Sense-making is actually the source of the term enaction. In making sense of its experience, the cognitive system is somehow bringing out through its actions — enacting — what is important for the continued existence of the system. This enaction is effected by the system as it is embedded in its environment, but as an autonomous entity distinct from the environment, through an emergent process of making sense of its experience. This sense-making is, in fact, cognition [18].

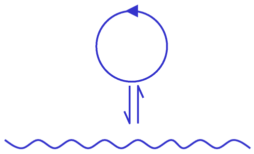

Figure 1. *Maturana and Varela's ideograms to denote structurally-determined autopoietic and organizationally-closed systems. The arrow circle denotes the autonomy, self-organization, and self-production of the system, the rippled line the environment, and the bi-directional half-lines the mutual perturbation — structural coupling — between the two.*

The founders of the enactive approach, Maturana and Varela, introduced a diagrammatic way of conveying the self-organizing and self-maintaining autonomous nature of an enactive system, perturbing and being perturbed by its environment [5]: see figure 1. The arrowed circle denotes the autonomy and self-organization of the system, the rippled line the environment, and the bi-directional half-arrows the mutual perturbation.

3.1. Enaction and Development

So what has all this to do with development? As we noted earlier in Footnote 2 of Section 2, our position in this paper is that learning arises as a consequence of the interaction between the cognitive agent and the world around it, whereas development arises from learning through a process of interaction of the cognitive agent with itself. We remarked above that the process of sense-making forms a virtuous circle in that the self-generated model of the system's experience provides a larger repertoire of possible actions, richer interactions, increased perceptual capacity, and potentially better self-generated models, and so on. Recall also our earlier remarks that the cognitive system's knowledge is represented by the state of the system. When this state is embodied in the system's central nervous system, the system has much greater plasticity in two senses: (a) the nervous system can

---

[8] *The founders of the enactive approach use the term* structural determination *to denote the dependence of a system's space of viable environmentally-triggered changes on the structure and its internal dynamics [5, 20]. The interactions of this structurally-determined system with the environment in which it is embedded are referred to as* structural coupling: *a process of mutual perturbations of the system and environment that facilitate the on-going operational identity of the system and its autonomous self-maintenance. Furthermore, the process of structural coupling produces a congruence between the system and its environment. For this reason, we say that the system and the environment are* co-determined. *The concepts of structural determination structural coupling of autopoietic systems [5] are similar to Kelso's circular causality of action and perception [21] and the organizational principles inherent in Bickhard's self-maintenant systems [22]. The concept of enactive development has its roots in the structural coupling of organizationally-closed systems which have a central nervous system and is mirrored in Bickhard's concept of recursive self-maintenance [22].*

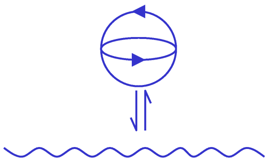

Figure 2. *Maturana and Varela's ideograms to denote structurally-determined autopoietic and operationally-closed systems. The diagram (denotes an organizationally-closed autonomous system with a central nervous system. This system is capable of development by means of self-modification of its nervous system, so that it can accommodate a much larger space of effective system action.*

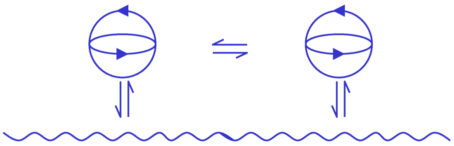

Figure 3. *Maturana and Varela's ideogram to denote the development engendered by interaction between cognitive systems.*

accommodate a much larger space of possible associations between system-environment interactions, and (b) it can accommodate a much larger space of potential actions. Consequently, the process of cognition involves the system modifying its own state, specifically its central nervous system, as it enhances its predictive capacity and its action capabilities. This is exactly what we mean by development. This generative (*i.e.* self-constructed) autonomous learning and development is one of the hallmarks of the enactive approach [15, 18].

Development, then, is identically the cognitive process of establishing and enlarging the possible space of mutually-consistent couplings in which a system can engage (or, perhaps more appropriately, which it can withstand without compromising its autonomy). The space of perceptual possibilities is predicated not on an absolute objective environment, but on the space of possible actions that the system can engage in whilst still maintaining the consistency of the coupling with the environment. These environmental perturbations don't control the system since they are not components of the system (and, by definition, don't play a part in the self-organization) but they do play a part in the ontogenetic development of the system. Through this ontogenetic development, the cognitive system develops its own epistemology, *i.e.* its own system-specific history- and context-dependent knowledge of its world, knowledge that has meaning exactly because it captures the consistency and invariance that emerges from the dynamic self-organization in the face of environmental coupling. Again, it comes down to the preservation of autonomy, but this time doing so in an every increasing space of autonomy-preserving couplings.

This process of development is achieved through self-modification by virtue of the presence of a central nervous system: not only does environment perturb the system (and *vice versa*) but the system also perturbs itself and the central nervous system adapts as a result. Consequently, the system can develop to accommodate a much larger space of effective system action. This is captured in a second ideogram of Maturana and Varela (see Figure 2) which adds a second arrow circle to the ideogram to depict the process of development through self-perturbation and self-modification. In essence, development *is* autonomous self-modification and requires the existence of a viable phylogeny, including a nervous system, and a suitable ontogeny.

3.2. Enaction and Knowledge

Let us now move on to discuss in a little more detail the nature of the knowledge that an enactive cognitive system constructs. This knowledge is built on sensorimotor associations, achieved initially by exploration of what the world offers. However, this is only the beginning. The enactive system uses the knowledge gained to form new knowledge which is then subjected to empirical validation to see whether or not it is warranted (we, as enactive beings, imagine many things but not everything we imagine is plausible or corresponds well with reality). One of the key issues in cognition is the importance of internal simulation in accelerating the scaffolding of this early developmentally-acquired sensorimotor knowledge to provide a means to predict future events, to reconstruct or explain observed events (constructing a causal chain leading to that event), or to imagine new events [23–25]. Naturally, there is a need to focus on (re-)grounding predicted, explained, or imagined events in experience so that the system can *do* something new and interact with the environment in a new way. If the cognitive system wishes or needs to share this knowledge with other cognitive systems or communicate with other cognitive systems, it will only be possible if they have shared a common history of experiences and if they have a similar phylogeny and a compatible ontogeny. In other words, the meaning of the knowledge that is shared is negotiated and agreed by consensus through interaction.

When there are two or more cognitive agents involved, interaction is a shared activity in which the actions of each agent influence the actions of the other agents engaged in the same interaction, resulting in a mutually constructed pattern of shared behaviour [26]. Again, Maturana and Varela introduce a succinct diagrammatic way of of conveying this coupling between cognitive agent and the development it engenders [11]: see Figure 3. Such mutually-constructed patterns of complementary behaviour is also emphasized in Clark's notion of joint action [27]. Thus, explicit meaning is not necessary for anything to be communicated in an interaction, it is simply important that the agents are mutually engaged in a sequence of actions. Meaning emerges through shared consensual experience mediated by interation. The research programme encapsulated in this roadmap is based on this foundational principle of interaction. The developmental progress of imitation follows tightly that of the development of other interactive and communicative skills, such as joint attention, turn taking and language [28–30]. Imitation is one of the key stages in the development of more advanced cognitive capabilities.

3.3. Phylogeny and Ontogeny: The Complementarity of Structural Determination and Development

Let us summarize: enaction entails two complementary processes: (a) phylogenetically-dependent structural determination, *i.e.* the preservation of autonomy by a process of self-organization which determines the relevance and meaning of the system's interactions, and (b) ontogenesis, *i.e.* the increase in the system's predictive capacity and the enlargement of its action repertoire through a process of generative model construction by which the system develops its understanding of the world in which is it embedded. Ontogenesis results in development: the generation of new couplings effected by the self-modification of the system's own state, specifically its central nervous system. This complementarity of structural determination — phylogeny — and development — ontogeny — is crucial. Cognition is the result of a developmental process through which the system becomes progressively more skilled and acquires the ability to understand events, contexts, and actions, initially dealing with immediate situations and increasingly acquiring a predictive or prospective capability. Prediction, or anticipation, is one of the two hallmarks of cognition, the second being the ability to learn new knowledge by making sense of its interactions with the world around it and, in the process, enlarging its repertoire of effective actions. Both anticipation and sense-making are the direct result of the developmental process. This dependency on exploration and development is one of the reasons why the artificial cognitive system requires a rich sensory-motor interface with the environment.

4. Embodiment: the Requirements and Consequences of Action

Cognitive systems as described above are intrinsically embodied and embedded in their environment in a situated historical developmental context [31]. Furthermore, as we have already noted, the system's physical embodiment plays a direct constitutive role in the cognitive process [32–34].[9] But what exactly is it to be embodied? One form of embodiment, and clearly the type envisaged by proponents of the enactive systems approach to cognition, is a physically-active body capable of moving in space, manipulating its environment, altering the state of the environment, and experiencing the physical forces associated with that manipulation [35]. But there are other forms of embodiment. Ziemke introduced a framework to characterise five different types of embodiment [36, 37]:

1. *Structural coupling* between agent and environment in the sense that a system can be perturbed by its environment and can in turn perturb its environment.
2. *Historical embodiment* as a result of a history of structural coupling;
3. *Physical embodiment* in a structure that is capable of forcible action;
4. *Organismoid embodiment*, *i.e.* organism-like bodily form (*e.g.* humanoid or rat-like robots);
5. *Organismic embodiment* of autopoietic living systems.

These five types are increasingly more restrictive. Structural coupling entails only that the system can influence and be influenced by the physical world. Historical embodiment adds the incorporation of a history of structural coupling to this level of physical interaction so that past interactions shape the embodiment. Physical embodiment is most closely allied to conventional robot systems, with organismoid embodiment adding the constraint that the robot morphology is modelled on specific natural species or some feature of natural species. Organismic embodiment corresponds to living beings.

To repeat again, the fundamental idea underpinning embodiment is that the morphology of the system is actually a key component of the systems dynamics. The morphology of the cognitive system not only matters, it is a constitutive part of the cognitive process and cognitive development depends on and is shaped by the form of the embodiment. There is, however, an important consequence of this. In a system that only satisfies the minimal requirements of embodiment, there is no guarantee that the resultant cognitive behaviour will be in any way consistent with human models or preconceptions of cognitive behaviour. Of course, this may be quite acceptable, as long as the system performs its task adequately. However, if we want to ensure compatibility with human cognition, then we have to admit the stronger version of embodiment and adopt a domain of discourse that is the same as the one in which we live: one that involves physical movement, forcible manipulation, and exploration, and perhaps even human form [38]. Why? Because when two cognitive systems interact or couple, the shared consensus of meaning — the systems' common epistemology — will only be semantically similar (have similar meaning) if the experiences of the two systems are compatible: phylogenetically, ontogenetically, and morphologically consistent [5]. Consequently, the approach to cognition we are advocating here requires that the cognitive system be embodied in a very specific sense: that it should lie in the organismoid space of embodied cognitive systems and, further still, that it should lie in the humanoid subspace of the organismoid space.

Apart from the morphology and phylogeny of the cognitive system, this also has strong implications for the development of the cognitive system. Specifically, the ontogeny of the system must follow the development of natural (human) systems.[10] This development follows a general path that begins with actions that are immediate and have minimal prospection, and progresses to much more complex actions that bring forth much more prospective cognitive capabilities. This involves the development of perception-action coordination, beginning with head-eye-hand coordination, progressing through manual and bi-manual manipulation, and extending to more prospective couplings involving inter-agent interaction, imitation, and (gestural) communication.[11] This development occurs in both the innate skills with which phylogeny equips the system and in the acquisition of new skills that are acquired as part of the ontgenetic development of the systems. As we have noted already, it is the ontogenetic development which provides for the greater

---

[9] *This distinguishes the approach from earlier cognitivist approaches in which cognition comprises computational operations defined over symbolic representations and in which these computational operations are not tied to any given instantiation. Although any computational system requires some physical realisation to effect its computations, the underlying computational model is independent of the physical platform on which it is implemented. This independence of model and instantiation is referred to as functionalism [2]. For this reason, it has also been noted that cognitivism exhibits a form of mind-body dualism [31, 35].*

[10] *In this discussion on the need for ontogeny to follow the development of natural human systems, we have focussed on the development of the couplings between system and environment and have said nothing about the role of physical development. Although we have emphasized from the point of view of phylogeny the importance of physical embodiment and human morphology to cognitive development, it remains to be seen to what degree humanoid cognitive development does or does not require bodily development. It will require some substantial advances in current technology to allow this aspect of ontogeny to be investigated. This issue reoccurs later in the paper in Section 6 on the challenges associated with modelling homeostasis and self-modification, specifically in Footnote 14 which discusses the common assumption in cognitive robotics that the physical system does not degenerate or develop.*

[11] *Communication in general, and especially language-based communication, is important in the development of prospective cognition with long time horizons, such as those involved in deliberation and reasoning. Gestural communication such as pointing is linked to the early development of language in human infants [39].*

prospective abilities of cognitive systems.

5. Enaction as a Practical Framework

The title of this paper refers to enaction as a conceptual framework for developmental cognitive robotics. However, it is also a practical framework in the sense that it provides well-defined constraints and requirements, as well as theoretical insights, for the realization of cognition in a physical robot. These constraints have been encapsulated in nine guidelines that form part of a larger set of forty-three guidelines[12] comprising a research roadmap for the development of cognitive capabilities in humanoid robots [40]. These guidelines have, in turn, been used in the design of a cognitive architecture for the iCub humanoid robot.[13] This cognitive architecture, together with the physical robot, constitute the phylogeny of the iCub and they provide the platform for the development of cognitive abilities through subsequent ontogenesis. In the following, we will identify these guidelines and summarize how they influenced the design of the iCub cognitive architecture.

The principles of enaction discussed in this paper yield the following nine guidelines.

1. The system should incorporate a rich array of physical sensory and motor interfaces which allow the system to act on the world and perceive the effects of these actions.
2. The system should exhibit structural determination: that is, the system should have a range of autonomy-preserving processes of homeostasis that maintain the system's operational identity and thereby determine the meaning of the system's interactions.
3. The system requires a humanoid morphology if it is to construct an understanding of its environment that is compatible with that of human cognitive agents.
4. The system must support developmental processes that modify the system's structure so that its dynamics of interaction are altered to effect
    - an increase in the space of viable actions, and
    - an extension of the time horizon of the system's anticipatory capability.
5. The system should operate autonomously so that developmental changes are not a deterministic reaction to an external stimulus but result from an internal process of generative model construction.
6. Development must be driven by internally-generated social and exploratory motives which enable the discovery of novelty and regularities in the world and the potential of the system's own actions.
7. The system should incorporate processes for the generation of knowledge effected by learning affordances whereby the perception of an object is interpreted as affording the opportunity for the system to act on it in a specific way with a specific outcome.
8. The system should incorporate processes of internal simulation to scaffold knowledge and to facilitate prediction of future events, explanation of observed events, and the imagination of new events.
9. The system should also incorporate processes for grounding internal simulations in actions to establish by observation their validity.

The remainder of this section provides a very brief outline of the iCub cognitive architecture and discusses the degree to which the enaction guidelines have been followed. The next section considers the challenges posed by a complete implementation of these guidelines. Our goal in this is not to describe the iCub cognitive architecture in any depth but to demonstrate the practical nature of the guidelines and the requirements that results from adoption of the enactive paradigm.

The iCub cognitive architecture focusses on several key capabilities. Gaze control, reaching, and locomotion constitute the initial simple goal-directed actions. Episodic and procedural memories are included to effect a simplified version of internal simulation in order to provide capabilities for prediction and reconstruction, as well as generative model construction bootstrapped by learned affordances.

In addition, motivations encapsulated in the system's affective state are made explicit so that they address curiosity and experimentation, both explorative motives, triggered by exogenous and endogenous factors, respectively. This distinction between the exogenous and the endogenous is reflected by the need to include an attention system to incorporate both factors.

A very simple process of homeostatic self-regulation governed by the affective state provides elementary action selection. Finally, all the various components of the cognitive architecture operate concurrently so that a sequence of states representing cognitive behaviour emerges from the interaction of many separate parallel processes rather than being dictated by some state-machine as in the case of most cognitive architectures. This preliminary cognitive architecture comprises thirteen components, as follows (refer also to Figure 4).

1. Exogenous Salience
2. Endogenous Salience
3. Egosphere
4. Attention Selection
5. Episodic Memory
6. Procedural Memory
7. Affective State
8. Action Selection
9. Gaze Control

---

[12] *The rest of the forty-three guidelines were derived from studies of developmental psychology, neurophysiology, and computational modelling [40].*

[13] *The iCub is an open-systems 53 degree-of-freedom humanoid robot [41, 42]. Measuring 1m tall, it is approximately the same size as a three or four year-old child although it weighs 22kg. It can crawl on all fours, its hands allow dexterous manipulation, and its head and eyes are fully articulated. The 53 degrees of freedom comprise six for the head, seven for each arm, nine for each hand, three for the waist, and six for each leg. Joint angles are sensed using a custom-designed Hall-effect magnet pair. In addition, tactile sensors are under development [43]. From the sensory point of view, the iCub is equipped with digital cameras, gyroscopes and accelerometers, microphones, and force/torque sensors. A distributed sensorized skin is under development using capacitive sensor technology. Each joint is instrumented with positional sensors, in most cases using absolute position encoders.*

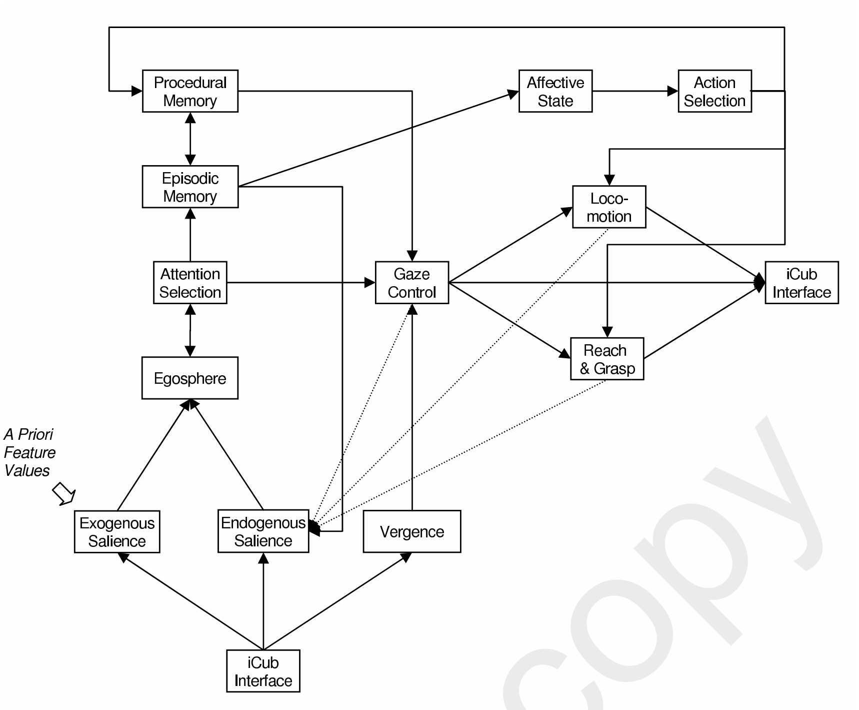

Figure 4. *The iCub cognitive architecture.*

10. Vergence
11. Reach & Grasp
12. Locomotion
13. iCub Interface

Together, the Exogenous Salience, Endogenous Salience, Egosphere, and Attention Selection components comprise the iCub's perception system. Similarly, Gaze Control, Vergence, Reach & Grasp, Locomotion comprise the iCub's actions system. The Episodic Memory and the Procedural Memory together provide the iCub's principle mechanism for anticipation and adaptation. The Affective State component effects the iCub motivations which together with the Action Selection component provide a very simple homeostatic process which regulates the autonomous behaviour of the iCub. The iCub Interface component completes the architecture and reflect the embodiment of the iCub from an architecture point of view.

As we noted above, the iCub cognitive architecture is a work-in-progress and what we have described is only a partial implementation of the roadmap guidelines. Below, for each guideline, we will discuss what aspects have been followed and how they have been implemented, what aspects have not been followed, and how they might be.

Guideline 1 stipulates a rich array of physical sensory and motor interfaces. While the iCub has several extereoceptive sensors, including binocular vision, binaural hearing, a 3 degree-of-freedom vestibular sense, as well as the soon-to-be-deployed skin with cutaneous sensing, we consider only vision and hearing in the initial version of the iCub cognitive architecture. The remaining senses will be integrated at a later date. Concerning proprioceptive sensing, the iCub is specified with absolute position sensors on each joint for accurate servo-control. Force/torque sensors have been designed and are being deployed in the latest version of the robot but they are not yet considered in the cognitive architecture although they have been used in stand-alone mode to demonstrate compliant manipulation.

Guideline 2 states that system should exhibit structural determination: that is, the system should have a range of autonomy-preserving processes of homeostasis that maintain the system's operational identity and thereby determine the meaning of the system's interactions. As noted above, at present, there is just one very simple homeostatic process in the action selection component.

Guideline 3 stipulates a humanoid morphology. The iCub follows this guideline faithfully, especially as the dimensions of the iCub are modelled on those of a human child and as it has such a high number of degrees of freedom, particularly in the hands and the head.

Guideline 4 stipulates that the system must support developmental processes that modify the system's structure so that its dynamics of interaction are altered to effect an increase in the space of viable actions and an extension of the time horizon of the system's anticipatory capability. The internal simulation system comprising the episodic and procedural memories in the iCub cognitive architecture accomplishes this to a limited extent. As the iCub explores its environment, as it looks around, guided by attentive processes that are triggered by both internal and external stimuli, as it moves, reaches, grasps, manipulates, it learns to associate perceptions with actions and actions with perceptions and thereby develops an understanding of its environment which as we have seen can then be used to predict and act. However, it is a weak form of development: it learns from experience how things are, rather than how things might be. That is, the current iCub cognitive architecture has no capacity for generalization. While recurrent action-perception associations are indeed strengthened by exploration and experience, there is no generative mechanism which constructs models of these action-perception associations that go beyond these particular instances to capture a more encompassing lawfulness in the iCub's interactions. Put another way, the iCub's cognitive architecture currently has no way of building a model by extrapolating from experience and then validating, refining, or discarding that extrapolated model.

Guideline 5 builds on Guideline 4 by requiring that the system should operate autonomously so that developmental changes are not a deterministic reaction to an external stimulus but result from an internal process of generative model construction. Notwithstanding the fact that the present iCub cognitive architecture implements Guideline 4 in a weak manner, the mechanism which governs the construction of these procedural models are nontheless autonomous: they depend only on the affective state of the system which depend in turn on how well the outcome of its explorative actions match its expectations. The system's goals are driven entirely by the internal affective processes.

Guideline 6 stipulates that development must be driven by internally-generated social and exploratory motives which enable the discovery of novelty and regularities in the world and the potential of the system's own actions. This guideline has been partially followed in that explorative motives have been implemented in the affective state but as yet social motives have not. Two forms of explorative motive have been implemented: curiousity and experimentation, focussing on exogenous and endogenous events respectively. It is envisaged that social motives will balance the two, the main idea being that in social interaction, a cognitive agent is trying to establish a common epistemology with the social partners and this requires equal attention to interactions generated by the partner (which have to be assimilated into the model the agent is constructing) and interactions generated by the agent (which are attempts to ground that model by interacting with the partner to see if the agent's expression of that model in the interaction is understood by the partner).

Guideline 7 requires the ability to learn object affordances [44]. Affordances can be modelled as associations between objects, action, and effects [41, 45–47]. This has been implemented on the iCub on a stand-alone basis [46–48] and it remains to integrate it into the iCub cognitive architecture in the internal simulation subsystem comprising the episodic memory and the procedural memory.

Guideline 8 — which states that the system should incorporate processes of internal simulation to scaffold knowledge and to facilitate prediction of future events, explanation of observed events, and the imagination of new events — is implemented directly in the cognitive architecture through procedural memory with prediction being effected by following a sequence of perception-action-perception associations $P, A, P, A, \ldots$ forward in time along the path with the strongest associative connections. Explanation (or reconstruction) follows the path backward in time and imagination follows it forward along a path with weak associative connections. However, the lack of a capacity for generalization limits the power of this internal simulation at present.

Guideline 9 says that the system should also incorporate processes for grounding internal simulations in actions to establish by observation their validity. This is accomplished in the iCub cognitive architecture by the Affective State, Action Selection, Endogenous Salience, and Episodic Memory. Specifically, when the affective state is in an explorative state, the endogenous salience is primed by an episodic memory representing the expected outcome of an action which is about to be performed. If the subsequently acquired percept matches the expectation, then the perception-action association is strengthened. If it isn't then the affective state changes from exploration to curiousity and is driven by exogenous factors, not internally-generated endogenous ones.

6. Challenges

We can identify two types of challenge posed by the adoption of an enactive approach to cognitive systems: general challenges inherent to the paradigm itself and more specific technical challenges associated with the nine guidelines set out above. As we have already commented on the challenges associated with the nine guidelines in the context of their application to the iCub cognitive architecture, we focus here on the more general challenges.

The first general challenge is the identification of the phylogenetic configuration and the ontogenetic processes. Phylogeny — the evolution of the system configuration from generation to generation — determines the sensory-motor capabilities that a system is configured with at the outset and that facilitate the system's innate behaviours. Ontogenetic development — the adaptation and learning of the system during its lifetime — gives rise to the cognitive capabilities that we seek. To enable development, we must somehow identify a minimal phylogenetic state of the system. In practice, this means that we must identify and effect perceptuo-motor capabilities for the minimal behaviours that ontogenetic development will subsequently build on to achieve cognitive behaviour. The nine guidelines set out above go some way towards addressing this problem. However, they are necessary but not sufficient and must be augmented with others derived from psychology, neurophysiology, and computational modelling [40].

The requirements of real-time synchronous system-environment coupling and historical, situated, and embodied development pose a second general challenge. Specifically, the maximum rate of ontogenetic development is constrained by the speed of coupling (*i.e.* the interaction) and not by the speed at which internal processing can occur [14]. Natural cognitive systems have a learning cycle measured in weeks, months, and years and, while it might be possible to condense these into minutes and hours for an artificial system because of increases in the rate of internal adaptation and change, it cannot be reduced below the time-scale of the interaction. You cannot short-circuit ontogenetic development because it is the agent's own experience that defines its cognitive understanding of the world in which it is embedded. This places a natural limit on the rate at which cognitive development can proceed. It is unlikely that this rate can be exceeded without discarding the principles of enaction.

Development implies the progressive acquisition of anticipatory capabilities by a system over its lifetime through experiential learning. Development depends crucially on the motives which underpin the goals of actions. The two most important motives that drive actions and development are social and exploratory. There are at least two exploratory motives, one focussing on the discovery of novelty and regularities in the world, and one focussing on the potential of one's own actions. A third challenge that faces all developmental embodied robotic cognitive systems is to model these motivations and their interplay, to identify how they influence action, and thereby build on the system's phylogeny through ontogenesis to develop every-richer cognitive capabilities. For enactive systems, this challenge can be addressed by tackling the twin problems of homeostasis and self-modification (cf. Guidelines 2 & 4). Since this homeostasis — autonomy-preserving self-regulation — entails structural determination, the homeostatic processes need to regulate the system's actions to ensure that the conditions required for the maintenance of autonomy are preserved in the environment.

This, in turn, depends both on the system's internal structures and its physical realization, and both must figure in whatever homeostatic processes are embedded in the system's cognitive architecture. It also means that the conditions required for the maintenance of autonomy must be explicitly identified.[14] Since homeostasis is concerned with the maintenance of autonomy by structural coupling through the system's phylogenetic repertoire of actions and anticipatory capability, the space of environmental perturbations it can withstand is consequently limited. The purpose of self-modification is to develop the system so that it has a larger repertoire of actions and a greater degree of anticipation to enable it to withstand a larger space of perturbations by the environment. As Bickhard puts it when discussing recursive self-maintenant systems — systems that contribute actively to the conditions for persistence — these systems can deploy different processes of self-maintenance depending on environmental conditions: "they shift their self-maintenant processes so as to maintain self-maintenance as the environment shifts" [22]. Viewed in this way, development and self-modification are intrinsically linked to the processes of homeostasis, giving them more degrees of freedom in the manner in which autonomy is perserved and endowing the systems with a greater ability to 'make sense' of its world through enaction.

Acknowledgments

The work described in this article was carried out as a part of the RobotCub research project to develop the iCub, an open and widely-adopted humanoid robot for cognitive systems research. This project was supported by the European Commission, Project IST-004370, under Strategic Objective 2.3.2.4: Cognitive Systems and I gratefully acknowledge the funding that made this research possible. Sincere thanks go to reviewers of an earlier version of this paper for their helpful and constructive comments.

References

[1] D. Vernon, G. Metta, and G. Sandini. A survey of artificial cognitive systems: Implications for the autonomous development of mental capabilities in computational agents. IEEE Transaction on Evolutionary Computation, 11(2):151–180, 2007.

[2] W. J. Freeman and R. Núñez. Restoring to cognition the forgotten primacy of action, intention and emotion. Journal of Consciousness Studies, 6(11-12):ix–xix, 1999.

[3] J. R. Anderson, D. Bothell, M. D. Byrne, S. Douglass, C. Lebiere, and Y. Qin. An integrated theory of the mind. Psychological Review, 111(4):1036–1060, 2004.

[4] P. Langley. An adaptive architecture for physical agents. In IEEE/WIC/ACM International Conference on Intelligent Agent Technology, pages 18–25, Compiegne, France, 2005. IEEE Computer Society Press.

[5] H. Maturana and F. Varela. The Tree of Knowledge – The Biological Roots of Human Understanding. New Science Library, Boston & London, 1987.

[6] D. Vernon. Cognitive vision: The case for embodied perception. Image and Vision Computing, 26(1):127–141, 2008.

[7] D. Vernon and D. Furlong. Philosophical foundations of enactive AI. In M. Lungarella, F. Iida, J. C. Bongard, and R. Pfeifer, editors, 50 Years of AI, volume LNAI 4850, pages 53–62. Springer-Verlag, Heidelberg, 2007.

[8] D. Vernon, G. Metta, and G. Sandini. Embodiment in cognitive systems: on the mutual dependence of cognition and robotics. In J. Gray and S. Nefti-Meziani, editors, Embodied Cognitive Systems. Institution of Engineering and Technology (IET), UK, 2010.

[9] H. Maturana. Biology of cognition. Research Report BCL 9.0, University of Illinois, Urbana, Illinois, 1970.

[10] H. Maturana. The organization of the living: a theory of the living organization. Int. Journal of Man-Machine Studies, 7(3):313–332, 1975.

[11] H. R. Maturana and F. J. Varela. Autopoiesis and Cognition — The Realization of the Living. Boston Studies on the Philosophy of Science. D. Reidel Publishing Company, Dordrecht, Holland, 1980.

[12] F. Varela. Principles of Biological Autonomy. Elsevier North Holland, New York, 1979.

[13] F. J. Varela. Whence perceptual meaning? A cartography of current ideas. In F. J. Varela and J.-P. Dupuy, editors, Understanding Origins – Contemporary Views on the Origin of Life, Mind and Society, Boston Studies in the Philosophy of Science, pages 235–263, Dordrecht, 1992. Kluwer Academic Publishers.

[14] T. Winograd and F. Flores. Understanding Computers and Cognition – A New Foundation for Design. Addison-Wesley Publishing Company, Inc., Reading, Massachusetts, 1986.

[15] E. Di Paolo, M. Rohde, and H. De Jaegher. Horizons for the enactive mind: Values, social interaction, and play. In J. Stewart, O. Gapenne, and E. Di Paolo, editors, Enaction: Towards a New Paradigm for Cognitive Science, Cambridge, MA, 2008. MIT Press.

[16] E. Thompson. Mind in Life: Biology, Phenomenology, and the Sciences of Mind. Harvard University Press, Boston, 2007.

[17] W. D. Christensen and C. A. Hooker. An interactivist-constructivist approach to intelligence: self-directed anticipative learning. Philosophical Psychology, 13(1):5–45, 2000.

[18] T. Froese and T. Ziemke. Enactive artificial intelligence: Investigating the systemic organization of life and mind. Artificial Intelligence, 173:466–500, 2009.

[19] T. Ziemke and R. Lowe. On the role of emotion in embodied cognitive architectures: From organisms to robots. Cognition and Computation, 1:104–117, 2009.

[20] F. Varela, E. Thompson, and E. Rosch. The Embodied Mind. MIT Press, Cambridge, MA, 1991.

[21] J. A. S. Kelso. Dynamic Patterns – The Self-Organization of Brain and Behaviour. MIT Press, Cambridge, MA, 3rd edition, 1995.

[22] M. H. Bickhard. Autonomy, function, and representation. Artificial Intelligence, Special Issue on Communication and Cognition, 17(3-4):111–131, 2000.

[23] R. Grush. The emulation theory of representation: motor control,

---

[14] *At present, robot autonomy is trivialized by two significant simplifications: the assumption that the robot has access to an unrestricted supply of energy and the assumption that the physical system suffers no wear and does not degenerate (or develop). With these assumptions in place, Bickard's view of autonomy, for example, as the property of a system to contribute to its own persistence [22] renders autonomy almost meaningless: if the system can't run out of energy and it can't be damaged, it is trivial to achieve persistence. The implications of this are quite significant in that with without the requirement for non-trivial conditions for the maintenance of autonomy, the associated processes of homeostasis and self-modification also become somewhat trivial. However, this does not invalidate the relevance of enaction and autonomy to robots. On the contrary, it makes enaction even more important when we eventually give up or relax these simplifying assumptions, as we will inevitably do at some point in the future.*

imagery, and perception. Behavioral and Brain Sciences, 27:377–442, 2004.

[24] G. Hesslow. Conscious thought as simulation of behaviour and perception. Trends in Cognitive Sciences, 6(6):242–247, 2002.

[25] M. P. Shanahan. A cognitive architecture that combines internal simulation with a global workspace. Consciousness and Cognition, 15:433–449, 2006.

[26] B. Ogden, K. Dautenhahn, and P. Stribling. Interactional structure applied to the identification and generation of visual interactive behaviour: Robots that (usually) follow the rules. In I. Wachsmuth and T. Sowa, editors, Gesture and Sign Languages in Human-Computer Interaction, volume LNAI 2298 of Lecture Notes LNAI, pages 254–268. Springer, 2002.

[27] H. H. Clark. Managing problems in speaking. Speech Communication, 15:243–250, 1994.

[28] J. Nadel, C. Guerini, A. Peze, and C. Rivet. The evolving nature of imitation as a format for communication. In J. Nadel and G. Butterworth, editors, Imitation in Infancy, pages 209–234. Cambridge University Press, Cambridge, 1999.

[29] G. S. Speidel. Imitation: a bootstrap for learning to speak. In G. E. Speidel and K. E. Nelson, editors, The many faces of imitation in language learning, pages 151–180. Springer Verlag, 1989.

[30] C. Trevarthen, T. Kokkinaki, and G. A. Fiamenghi Jr. What infants' imitations communicate: with mothers, with fathers and with peers. In J. Nadel and G. Butterworth, editors, Imitation in Infancy, pages 61–124. Cambridge University Press, Cambridge, 1999.

[31] E. Thelen and L. B. Smith. A Dynamic Systems Approach to the Development of Cognition and Action. MIT Press / Bradford Books Series in Cognitive Psychology. MIT Press, Cambridge, Massachusetts, 1994.

[32] D. Vernon. The space of cognitive vision. In H. I. Christensen and H.-H. Nagel, editors, Cognitive Vision Systems: Sampling the Spectrum of Approaches, LNCS, pages 7–26, Heidelberg, 2006. Springer-Verlag.

[33] J. L. Krichmar and G. M. Edelman. Principles underlying the construction of brain-based devices. In T. Kovacs and J. A. R. Marshall, editors, Proceedings of AISB '06 - Adaptation in Artificial and Biological Systems, volume 2 of Symposium on Grand Challenge 5: Architecture of Brain and Mind, pages 37–42, Bristol, 2006. University of Bristol.

[34] H. Gardner. Multiple Intelligences: The Theory in Practice. Basic Books, New York, 1993.

[35] E. Thelen. Time-scale dynamics and the development of embodied cognition. In R. F. Port and T. van Gelder, editors, Mind as Motion – Explorations in the Dynamics of Cognition, pages 69–100, Cambridge, Massachusetts, 1995. Bradford Books, MIT Press.

[36] T. Ziemke. Are robots embodied? In C. Balkenius, J. Zlatev, K. Dautenhahn, H. Kozima, and C. Breazeal, editors, Proceedings of the First International Workshop on Epigenetic Robotics — Modeling Cognitive Development in Robotic Systems, volume 85 of Lund University Cognitive Studies, pages 75–83, Lund, Sweden, 2001.

[37] T. Ziemke. What's that thing called embodiment? In R. Alterman and D. Kirsh, editors, Proceedings of the 25th Annual Conference of the Cognitive Science Society, Lund University Cognitive Studies, pages 1134–1139, Mahwah, NJ, 2003. Lawrence Erlbaum.

[38] R. A. Brooks. Flesh and Machines: How Robots Will Change Us. Pantheon Books, New York, 2002.

[39] L. Camaioni, M. C. Caselli, E. Longbardi, and V. Volterra. A parent report instrument for early language assessment. First Language, 11:345–360, 1991.

[40] D. Vernon, C. von Hofsten, and L. Fadiga. A Roadmap for Cognitive Development in Humanoid Robots. Cognitive Systems Monographs (COSMOS). Springer, in press, 2010.

[41] G. Metta, L. Natale, F. Nori, G. Sandini, D. Vernon, L. Fadiga, C. von Hofsten, J. Santos-Victor, A. Bernardino, and L. Montesano. The iCub Humanoid Robot: An Open-Systems Platform for Research in Cognitive Development. submitted to Neural Networks, 2010.

[42] G. Sandini, G. Metta, and D. Vernon. The icub cognitive humanoid robot: An open-system research platform for enactive cognition. In M. Lungarella, F. Iida, J. C. Bongard, and R. Pfeifer, editors, 50 Years of AI, volume LNAI 4850, pages 359–370. Springer-Verlag, Heidelberg, 2007.

[43] M. Maggiali, G. Cannata, P. Maiolino, G. Metta, M. Randazzo, and G. Sandini. Embedded distributed capacitive tactile sensor. In The 11th Mechatronics Forum Biennial International Conference, University of Limerick, Ireland, 2008.

[44] J. J. Gibson. The theory of affordances. In R. Shaw and J. Bransford, editors, Perceiving, acting and knowing: toward an ecological psychology, pages 67–82. Lawrence Erlbaum, 1977.

[45] L. Montesano and M. Lopes. Learning grasping affordances from local visual descriptors. In IEEE International Conference on Development and Learning, Shanghai, China, 2009.

[46] L. Montesano, M. Lopes, A. Bernardino, and J. Santos-Victor. Modeling affordances using bayesian networks. In IEEE/RSJ International Conference on Intelligent Robots and Systems, San Diego, USA, 2007.

[47] L. Montesano, M. Lopes, A. Bernardino, and J. Santos-Victor. Learning object affordances: From sensory motor maps to imitation. IEEE Transactions on Robotics, 24(1), 2008.

[48] L. Montesano, M. Lopes, F. Melo, A. Bernardino, and J. Santos-Victor. A Computational Model of Object Affordances. IET, 2009.