

Development and Robotics

Giorgio Metta^{1,2}, Giulio Sandini¹, Lorenzo Natale¹, and Francesco Panerai³

¹LIRA-Lab, DIST – University of Genova, Genova – Italy

²AI-Lab, MIT – Cambridge MA – USA

³LPPA – CNRS, College de France, Paris – France

e-mail: pasa@ai.mit.edu

Abstract

We propose that the development of causality can be seen as a *primitive* for understanding and constructing complex systems either biological or artificial. Furthermore, we put forward a view of development in terms of the control of complexity. Although some of these elements are at the moment speculative or barely outlined, the theoretical test and verification are part of the ongoing research. On the artificial side, we will show how developmental principles are used within the architecture of a humanoid robot. The reference problem is the ontogenesis of sensori-motor coordination. Visual, acoustic and inertial cues constitute the sensory repertoire of the robot; computation, in the form of mappings, represents its brain activity. The *continuous and meaningful* adaptation during the natural interaction of the robot with the environment is one of the key aspects of the implementation.

Keywords: computational neuroscience, learning, robotics.

Introduction

We advocated in the past that the principles of biological development are helpful to understand how to design and construct complex artificial systems [1-3]. Although, there might be a consensus on the “general principles” underlying this idea (as others worked pretty much along the same direction; for example [4-9]), we would like to put forward a tentative *manifesto* of what development should mean when applied to, for instance, humanoid robots. Many times developmental principles have been used as a source of inspiration but in a few cases with the

intent of building a truly adaptive system. Machine learning techniques, for example, have been employed in robotics in a few circumstances [10-13]. The usual procedure though was that of collecting the data, training the machine, and eventually controlling the robot. These three steps were performed *off-line* and partially by hand, no further adaptation was included. Our approach is different; it is ecological in the sense that what is significant is the robot and its environment; the robot has to gather its own training set to solve a particular task, and in doing so it shows adaptation. We consider a *time-variant* system, while the previous case can be regarded for any practical purpose as *time-invariant*.

It is fair to say that we do not have any all-encompassing solution already established, but rather we strive to provide principles, which can evolve into a theory of developing systems. This theory should tell us both how to understand and how to construct developing systems – i.e. inherently time-variant (and possibly complicated) artificial systems. Quartz et al. [14] pointed out that time-variance poses difficulties, far too often overlooked.

The “understanding” part of this hypothetical theory should provide means to identify which elements are important to comprehend a biological system. This has been called the *physiology* problem: that is, describe how something works [15]. The more powerful “constructive” part should provide guidelines and design principles on how to build the particular system we investigate on (e.g. a humanoid robot). To address the physiology problem a developmental approach is advantageous. In fact, by studying how the system is constructed, we might be able to explain its functions as a whole, as well as, the relative role of its components. For the second problem, development is essential, because our very goal is to mimic it: i.e. to build a system able to grow and to show adaptation over a long time span (the whole *life-cycle*).

Artificial intelligence, artificial life, as well as artificial vision, just to mention a few, have made clear that non-adaptive systems usually fail in the real

* Giorgio Metta is currently Postdoctoral Associate at MIT, AI-Lab supported by DARPA within the “Natural Tasking of Robots Based on Human Interaction Cues” project under contract number DABT 63-00-C-10102.

world [16], and real adaptation is difficult to obtain and control. What is missing? There are a few proposals [16-18]: we would like to add a new one: that the solution has to be found in *development*. A broad outline of the theoretical aspects addressed in the paper is sketched in figure 1 below. It reflects also the organization of the paper.

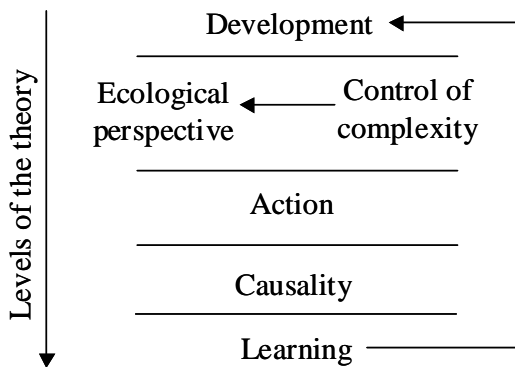


Figure 1: The theoretical aspects addressed in the paper: a general outline.

Development itself

Although a complete account of human sensori-motor ontogenesis is outside the scope of this paper, we would like to introduce a few aspects that can be seen as the first step to theoretically explain development, and to investigate which are the applications to artificial systems.

The main difference in terms of approach can be stated as: *modularity versus integration*.

The critic we put forward to the *modular* approach, especially in engineering, is that to make the problem tractable, very often, complex systems are divided in small parts, which are analyzed in isolation. Complexity is addressed by breaking the system into components. This has been successful so far but it has also hit its own limits. Further it is not guaranteed that these parts are the correct representatives of the characteristics of the system. The separation of one part from another may have strong implications on how each part is built, functions and learns. For these reasons it is not clear whether the tendency of searching the smallest fundamental component would eventually converge.

In robotics the same approach has been applied by designing, for instance, separated sensory processing modules (e.g. vision, audition, touch, etc.) and distinct behaviors (e.g. reaching, obstacle avoidance and so on). Most of the time large-scale system integration either failed or was successful at the expenses of generality and adaptation.

A different approach is taken by biological systems. Newborns, for example, at birth are an already *integrated* system. Many “modules” are still non-functional or they function differently from their “adult” counterpart: neural growth is not completed [19], motor control limited [20], but the sensorial, motor and cognitive abilities are nicely matched. A crude sensory system tailored to a rough motor controller. Further, the hypothetical sub-modules develop simultaneously and harmonically resulting in a system whose components always fit one to another during growth. Adaptation is in the very fabric of the system: we can observe the general tendency of a smooth shift from *simpler to more complicated*; as we will see this might be another key feature for stable adaptation. Examples are the newborn motor reflexes that although clearly not that much adaptive provide the simplest controller yet able to generate useful sensori-motor synergies.

Newborns are maximally efficient in collecting data (making new experiences) and their behaviors are optimized for learning rather than simply for efficiency in absolute terms; even “negative” aspects such as noise contribute instead to the long-term efficiency by supporting the exploratory behaviors.

Thus, developmental studies have the chance to provide both theoretical contributions to the understanding of time-variance in large-scale systems, and a more complete account on what living systems are made of. This alternative approach seeks for rules that govern the dynamics of the system at different levels [21], where not necessarily the system has to be decomposed in any particular way.

At this point, also the word “development” needs some comments especially when compared to “learning”. As it will be clear in the next section, learning is intended as the acquisition of a particular skill, while development encompasses the optimal sequence of functional and structural changes timed to properly carry out adaptation. Learning and development thus run on different time scales [22] – and perhaps on more than two – spanning the continuum from less than a second to years. The reason to point out at the differences is that learning takes place within a context set up by development: i.e. what and how can be learnt is determined by the developmental stage. For example, adult level response to wide field moving stimuli – the Opto-Kinetic Nystagmus (OKN) – is thought to depend on optic flow processing. Atkinson and colleagues [23] have related the development of the response to the neural maturation of various cortical pathways responsible for dealing with motion information, and their interaction with subcortical nuclei. It is clear that the OKN could not develop before motion

processing. The OKN can be tuned only at a particular stage, when all its prerequisites are met. To get closer to the development of causality as a primitive, the next step is to try to reformulate development in a more computational framework; this can be done starting from machine learning theory and statistics.

Complexity control as a model of development

It has been recognized that learning from examples is an ill-posed problem [24, 25]. Recently a probabilistic-theoretical analysis formalized this problem, and we suggest here that it provides hints on the nature of development (as introduced in the previous section). This is true of course if we admit that a good part of learning can be subsumed under the *function approximation* problem¹. The typical problem of learning from examples is, generally speaking, that of tuning the parameters of the approximator in order to get the output as close as possible to the examples. While this is sound, and would eventually work if we were provided with an infinite training set (in the limit), a more accurate analysis for a finite training set reveals a different story. In practice if we use a learner which has too many parameters to tune with respect to the number of training samples, the results is what is called “over-fitting”: that is, a very good approximation but a very poor generalization. Vice versa, being too cautious might lead to an over-smoothing: an inadequacy of the model to grasp the complexity of the problem. The central issue is thus that of balancing the model complexity in order to do what is not too bad. The theory, which formalizes this situation, is known as Statistical Learning Theory (SLT) [24, 26]. The main result of interest to us is the fact that learning in order to be stable must balance the complexity of the learning machine – structural risk minimization in the words of Vapnik [24], or practically speaking cross-validation. To be precise, although, we talk about parameters here, it has been shown that an adequate measure of complexity is not related to the actual number of parameters but rather to the shape of the approximator being used. In some cases the intuition is correct and the complexity is a function of the number of free parameters.

Given SLT, how do we relate it to development? The fundamental idea is that complexity control is what development is all about. Neural processes, as

suggested in [27], come in two flavors: growing and pruning [18]. As two views originated from these different modes of controlling complexity, often they were at odd one to another [14]. SLT instead tells us that the two processes are equally necessary. A slight predominance should be attributed to growth because the quantity of data (experience of the system in the environment) builds up over time, and thus simpler models should be used at the beginning. A real growth is not strictly necessary – in biological systems – a recruitment process would do as well with new functional units connected during development. Starting from a psychological perspective, Turkewitz et al. [28] pointed out that the limitation of newborns’ sensori-motor abilities might be beneficial for learning – we explain this in terms of the control of complexity: it is better to start with a simpler system.

The time scale issue raised in the previous section in distinguishing between learning and development fits this schema. Development is concerned specifically with the control of the complexity and the structure of the learner. Learning is the mere adaptation of the parameters irrespective of the structure itself. Learning can and must be fast to adapt to impelling exigencies –think about the amazing ability of biological systems to learn from a single example. This procedure alone though can be either unstable or prevent generalization. A slower data-dependent procedure is needed to tailor the structure of the learner to the environment.

A more fundamental difference, and concurrently, a powerful constraint posed on the developing agent is that the training data does not come for free: gathering information has always a cost. To get representative data of the whole state space can take an infinitely long time (the time required is exponential in the number of dimensions). The agent cannot devote all its effort to exploration because otherwise no task would be ever achieved. This issue has been called the *exploration-exploitation dilemma* [29]. Any *real* learning system has to face this problem and adopt strategies to cope with it. It is also worth noting that for a real agent, the cost of failure might be very high and thus errors must be weighed accurately. The general pattern of development, “from simpler to more complex” now makes even more sense. We suggest that the initial sensori-motor coordination schemes, which are mostly reflex-like, might serve exactly this purpose: i.e. bootstrapping safely the system – an example is the generation of movement without exceeding the working ranges of the system (joint limits, energy consumption, etc). The role of the early reflexes is thus that of constraining exploration towards particular directions along the state space.



¹ That this is a reasonable assumption has to be demonstrated. It is sufficient to note here that functions can represent any dynamical entity, including timing related learning or modifications due to diffusible factors.

It is clear that to put development in its “ecological” context shifts the problem towards that of collecting the training data. This is to say that we have to deal with action, and the very capacity to learn is that of finding representative data without incurring into severe penalties in terms of basic drives (e.g. feeding, mating, etc).

Further, the way exploration is performed – the quality of training – depends strongly on how the system acts. Because of this also perception is doomed to be derived from actions. The capacity of categorization emerges out of the sensori-motor coordination patterns. Sensory processing alone does not make that much sense – action must be a prior of perception.

The primacy of action

The reason why this position is tenable comes from one of the most fascinating discoveries of the neurophysiology of the last decade: that is *mirror neurons* [30, 31]. A lot of discussion has originated from this finding, because it is thought to provide the missing link between action and perception. Mirror neurons are a class of neurons found in the monkey’s frontal cortex (area F5). A particular mirror neuron is activated both when the monkey executes an action and when it observes the same action performed by somebody else: hence the name *mirror*. Recently Fadiga et al. [31] provided evidences of the existence of a mirror system in humans.

The importance of the discovery lies in the possibility to relate mirror neurons to gesture recognition (e.g. grasping, tearing, holding, etc), language [32], and imitation learning [33]. Roboticians too were attracted by this possibility [34].

Another older but important discovery is related to another class of neurons in area F5 called *canonical* [35]. They are what might be called the coding of Gibsonian *affordances* in neural terms. These neurons are responsive to action execution (grasping for instance) but also when the monkey sees the “graspable” object. There is an explicit coding of how a given object can be grasped or manipulated: e.g. area F5 distinguishes between a precision grip and a full palm grip.

To frame these discoveries, a few complementary neural pathways need to be described. Area F5 is within what is called the *dorsal* pathway. The link from vision to action and more cognitive functions is believed to split as early as in the primary visual cortex (V1) into two complementary streams [36]. The first one is called the *dorsal* stream and it is devoted to action execution and the visuo to motor transformations needed to perform object-oriented movements; the second one – the *ventral* stream –

takes care of providing the perceptual judgment correlated with the shape of the object. Besides being segregated, of course the two streams do interact. Numerous studies (for a review see [37]) have shown an interesting interaction of perceptual quantities (e.g. size) with action execution (e.g. grasping). That is, perception might be bound to action and vice versa, cognition sometimes helps action. That the dorsal stream goes further than merely transforming coordinate reference frames, comes also from a deficit called *ideomotor apraxia*, which follows parietal lesions. These patients have difficulty in executing gestures, but interestingly also in recognizing the same gestures when performed by others. It is noteworthy that the same principles are not only involved in grasping and reaching actions, but rather, they are thought to be relevant also for highly cognitive abilities such as language (Fadiga 2001, personal communication).

To relate this description back to development we need to analyze what is learnable and which conditions are required for learning to take place. The advantage of this line of reasoning is that we may consider only those models where temporal consistency is preserved. If we assume certain skills and motor control abilities at a given age, we may only employ those to further develop new modules. Automatically, we rule out impossible solutions where a particular feature is used before being learned. For example, without any further assumption mirror neurons are not learnable. They are at the same time needed for imitation, and vice versa, imitation is required to build a mirror system (a solution is proposed in the following section).

On a more practical basis, what we propose is to put action at the foundation of more cognitive functions, such as categorization. Developmentally, if action has to be a prerequisite of perception, we should observe a different developmental progression of the dorsal versus the ventral stream. This question has been investigated, for example, by Kovacs [38] who provided supporting evidence in this direction. Wexler et al. [39] addressed a similar issue in human behavior and they have recently shown how self-motion can influence perception. From the evolutionary perspective this whole schema makes sense, since it is likely that cognition emerged on top of a preexisting motor control substratum.

In robotics, theories where action had a sort of primacy have been already proposed in the past; for example the active vision or purposive vision paradigms [40, 41]. The critic we move to those proposals (although we generally agree with them) is that action was eventually exploited a little. It was never a fundamental component in the sense we are proposing here. Other approaches where a model of,

for instance, mirror neurons has been attempted [34], have neglected the learning aspects and the learn-act-perceive (in random order) loop that we argue it is fundamental. Yet this is not the most fundamental element we need to “close the loop”.

Looking for primitives

To recap we proposed that: i) development can be modeled as complexity control, and linked to action by the necessity to develop the capacity to collect “good data” while interacting with the environment, and ii) the ability to act in the real world assumes a pivotal role for the acquisition of complex cognitive abilities.

Although not immediately obvious the key aspect of learning and development has to be found in *causality*. Causality has to be intended as the ability to relate action execution – as determined by the efferent copy and sensory afferences relative to the movement (visual, kinesthetic, etc.) – to its sensorial consequences. To put things together either biological or artificial systems have to use an at least very basic “understanding” of causality.

If we examine, simple motor control abilities, where, qualitatively speaking, coordinate transformations regulates the behavior [42], we find that all what is needed for learning is to causally relate the relevant quantities. Behaviorally, Von Hofsten and colleagues [43], for example, have shown that the coordination of eye and head smooth pursuit develops by first synchronizing the movements, and subsequently, by tuning the amplitudes. This can be interpreted in terms of causality as the necessity to learn the correct relationship between the causes and the effects – in engineering terms this is called credit assignment. The first step (i.e. synchronization) is motivated by the fact that the newborn has to firstly determine *when* there is a causal relationship between events. In the eye-head coordination example cited above this means that he/she needs to relate the efferent copy of the control commands to the displacement of the target during pursuit due to the self-generated movement. Amplitude tuning can be carried out later and it is most likely error driven: e.g. the *tracking error* can be used to measure the performance and tune the pursuit gain.

Beside motor control, we can show that by exploiting a basic understanding of causality we can conceptually solve other learning problems such as categorization or build a mirror representation. The first problem can be tackled by observing that objects can be classified pragmatically on the basis of the *affordances* as previously mentioned. Affordances are the characteristics of the object available for exploitation by action. For example, a glass has the

grasping, pouring water into, holding, breaking, and so on affordances. In order to learn affordances, action is necessary. Learning signal are measured directly at the sensory level: for example, grasping is successful if the baby acquires possession of the object. An object, on this basis is defined as the spatio-temporal boundaries of the sensory and motor information due to the particular action being performed (see figure 2).




Figure 2: A cartoon representation of the meaning of *affordances*. The action system defines the object (a cup) in terms of the specific actions that can be applied to it.

Further, this definition avoids the problem of defining what criteria have to be used to separate an object from other parts of the physical world, and as a consequence, it tells us that explicit segmentation comes as a consequence rather than being a prerequisite for categorization. The figure-ground segmentation problem must not be addressed first, but rather it is an emergent property of categorization. Sensory data subsets that participate in the definition of this “action to effects” mapping are the “object”; the remaining is the background. Though an inversion can be an awkward procedure, in theory, the analysis of the effect allows recovering the action to get that particular effect. A decision, recognition, and execution process can be implemented (e.g. if A produces B then to get B just do A).

At this point mirror neurons come “simply” as an additional association of two possible causes generating the same consequence: i.e. either the monkey or the experimenter performs the action but the consequence is the same if expressed in terms of goal (e.g. a “small” object is lifted therefore a pinch grasp was used). In fact, in monkeys, mirror neurons fire only when the goal is explicit. Mimicking the actual action does not make the mirror neurons to respond.

Finally, a new consequence for interpreting the driving forces behind development emerges: the progression of the ability to detect causally related events becomes itself the driving force motivating the learning of more and more complex skills. A link of causality to the basic “drives” of the agent is realized at this point. Rephrasing the grasping example, getting possession of the object is reinforced because it is motivationally pleasant. Most of the difficulty of the task is though in the understanding of causality. The role of the motivations appears to be conceptually simpler.

By rethinking Piaget [44], where he described the ontogenesis of causal understanding, we can focus on the *development of causality* as the most difficult of our mental tasks. If we solve it, we have grasped a powerful principle to guide a stable learning. This would be the building block of a system, which is not limited to any particular task, but open to progressive learning.




Figure 3: The Babybot. See text for details.

A practical implementation

As a first step for our theories of development we designed and constructed a test bed shaped as a humanoid robotic system: Babybot [45]. Babybot consists of a robot head, arm, and torso for a total of twelve degrees of freedom (see figure 3). The sensory system is composed of cameras, microphones, an

inertial device and motor encoders. Babybot’s vision is space variant: the robot observes the world through a high-resolution *fovea* and a progressively lower resolution periphery [46, 47]. Sound is acquired by means of two microphones, and plastic earlobes provide directionality. The inertial sensors mimic the corresponding biological vestibular system and the motor encoders give a sort of kinesthetic sensation to the robot. Actuators are torque controlled whenever relevant to the robot behavior to simulate the low-stiffness characteristics of muscles and to provide a natural mechanical compliance. Learning is carried out by growing neural networks as described in [48]. The reference problem we investigated on Babybot has been that of sensori-motor coordination, and in particular orienting and reaching towards an object identified because of vision and/or audition.

Learning to act

Babybot is capable of learning coordinate movements starting from an initial stage where the control is mostly reflex driven. The initial reflexes here provide a way to initially direct learning along a particular route. Examples are an initial tendency to perform small saccades although embedded in a strong noise, and a complex synergy mimicking an early human reflex called Asymmetric Tonic Neck Reflex [45]. The exploration of the environment, besides being initiated from the very beginning of the robot’s “life”, is driven by a combination of the reflexes with an endogenous noise generation process. The latter is meant to mimic all the limitations and the defective control present at birth because of uncompleted neuronal growth (e.g. myelination) or excessive connectivity [19].

The specific sequence of developmental events is roughly similar to that observed by Von Hofsten and colleagues [49] in humans. During the first stages only eye movements are attempted and an appropriate map is learned which causes the robot to improve its gazing performance. Concurrently, self or externally generated movements contribute in stimulating the vestibular system. Another neural network, together with basic visual processing abilities learn how to tune the robot’s vestibulo-ocular reflex (VOR) [50]. In a successive stage, head movements are initiated, and a new map connects them with the VOR and the already formed eye movements. This developmental trend is beneficial to the robot for at least two reasons: i) learning can address one problem at the time, and consequently the correct explorative strategy can be applied without disturbing the functioning of other modules; ii) the remaining degrees of freedom are coordinated simply by stereotyped reflexes that although non-adaptive provide a way to keep the system in a status

of ongoing exploration. During a successive stage Babybot learns another map to convert gaze direction into reaching commands [45]. The control schema of the arm is biologically inspired from the theory of Bizzi and colleagues [51, 52]. More maps, with the help of vision, can be learned: i) to control the head movements in order to attend an auditory stimulus; ii) to align auditory and visual map of space in order to integrate, when feasible, the sensory cues.

It is worth stressing that learning here is continuously carried out *on-line* – this is not commonly the case in robotics. We did not distinguish between the training and testing stage. The robot explores and acts (exploits) at the same time.

Figure 4 shows the relationships and time sequence of learning events of the Babybot. As an example of the behavior of the algorithms employed within the Babybot's architecture, figure 5 shows the acquisition of the movement of the head to attend an acoustic stimulus. In particular, the mean and standard deviation of the error are shown (both computed over a moving window of 150 samples). After the activation of the learning procedure (vertical solid line) a sharp increase of the motor performance can be noticed. The topmost panel shows the map as obtained at the end of the learning phase: the output is the required saccadic command, the input the initial error in terms of the two sensory cues used to localize a sound source in space. The first is the interaural time difference (ITD) and can be associated to the position of the target along the horizontal direction. The second is the interaural level difference (ILD) and under certain conditions measures the location of the sound source in the vertical direction. Note, as in figure 5, vectors point toward the origin (0,0) of the map. This is expected since the movement has to zero the error between the gaze and the target. Finally figure 6 presents the results of the tuning of the maps used to generate saccades by applying a variable resolution schema. In this case, we tested the effect of visual resolution in learning to perform saccadic eye movements. In particular we compared the learning performance in two conditions. In the first the resolution of the retina is maximum and does not vary with time. In the second we simulated the maturation of the retina and the corresponding increase of resolution with time. Although eventually the two maps should converge to the same final result, it is easy to see that the time-varying resolution case converges much faster to a usable map (learning is faster). The relative error between the two is more than 50% of the total learning for the period tested.

Conclusions

This paper dealt with the problem of defining the foundations of a theory that should encompass both the design and understanding of complex systems. The pillars of the theory are to be found in development. We showed that learning theory could describe (by means of the concept of complexity control) one of the goals of a developing agent – i.e. stable and effective adaptation. We argued also that learning in the real world could only be obtained if a further optimization of resource expenditure is carried out in order to gather “good” training data. It is clear that the training data has to be collected by means of actions. As soon as action becomes the concern, we are forced to consider the perceptual component as bounded to action itself. We believe that this provides a new vantage point to interpret difficult problems both in neuroscience and robotics. By framing categorization and action under the same explanatory principle, we expect to get new “algorithms” to solve old problems such as those found in artificial vision. At least, this is now completely formulated in an ecological context: the robot within its environment. Yet to solve the learning problems, we have to resort to an even more basic principle that is causality. Goal directness and causality are shown to conceptually solve learning in a general sense (not in the sense of providing a new algorithm). Object recognition, mirror neurons, reaching and motor control learning problems become all subsumed under the general problem of understanding cause-effect relationships. The development of the “understanding” of causality can now be seen as the driving force of stable learning. Finally, we presented the initial implementation of the theory on a humanoid robotic platform. We have shown that i) the robot successfully employs some developmental principles to learn sensori-motor coordination; ii) learning is completely carried out on-line; iii) the system is already integrated. Clearly motor control is only the very first problem to be solved.

Acknowledgements

The research described in this paper has been supported by the EU Project COGVIS (IST-2000-29375) and MIRROR (IST-2000-28159), and by the Italian Space Agency (ASI).

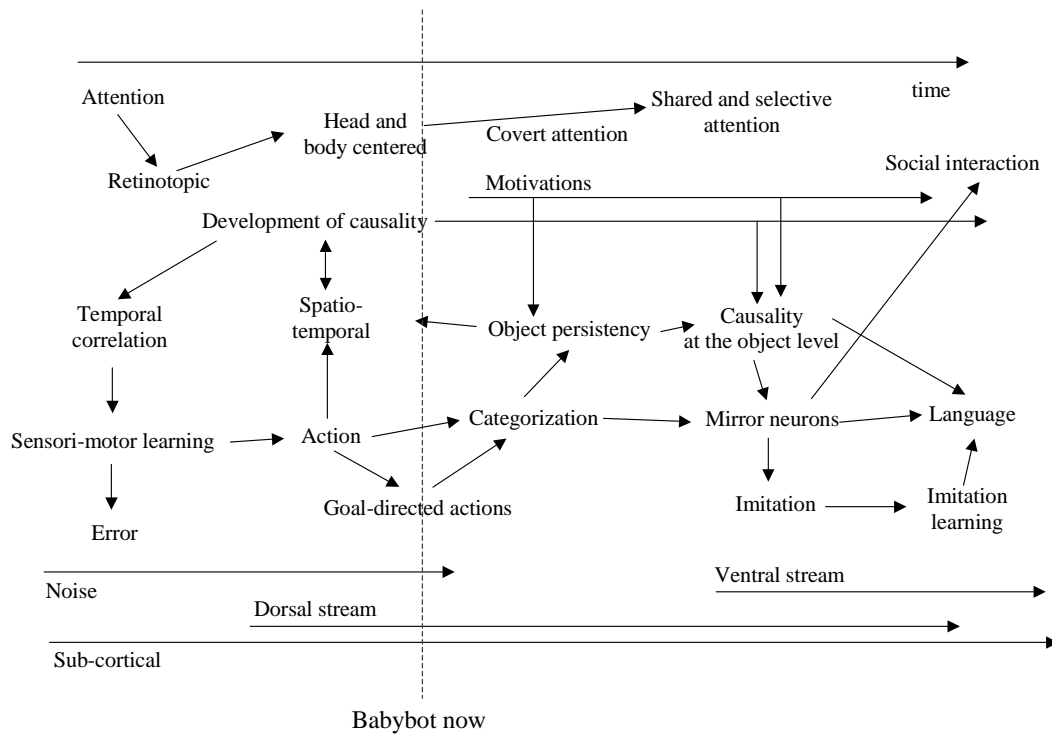


Figure 4: The developmental progression of the Babybot. A representation of how the various phases of the robot's development are interrelated. See text for further description.

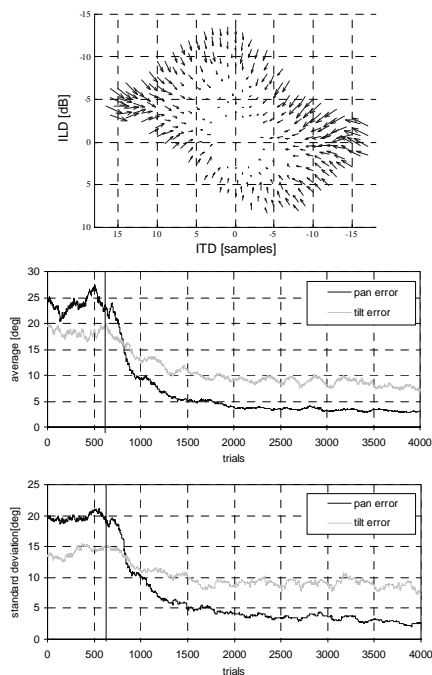


Figure 5: Example of learning curves relative to the acquisition of the controller of the head.




Figure 6: Variable vs. constant resolution map. From top to bottom: i) the motor error measured at the end of the saccade (in radians); ii) the resolution of the map over time (expressed in pixels); the maximum resolution is 1 pixel; iii) The relative error between the variable and constant resolution map.

References

1. Sandini, G. *Artificial Systems and Neuroscience*. in *Proc. of the Otto and Martha Fischbeck Seminar on Active Vision*. 1997.
2. Sandini, G., G. Metta, and J. Konczak. *Human Sensori-motor Development and Artificial Systems*. in *AIR&IHAS '97*. 1997.
3. Metta, G., *Babybot: a Study on Sensori-motor Development*, in *DIST*. 2000, University of Genova: Genova. p. 176.
4. Berthouze, L. and Y. Kuniyoshi, *Emergence and Categorization of Coordinated Visual Behavior Through Embodied Interaction*. *Machine Learning*, 1998(31): p. 187-200.
5. Pfeifer, R. and C. Scheier. *Representation in Natural and Artificial Agents: an Embodied Cognitive Science Perspective*. in *Natural Organisms, Artificial Organisms, and Their Brains*. 1998. Bielefeld, Germany: Verlag.
6. Brooks, R. *Behavior-Based Humanoid Robotics*. in *IEEE/RSJ IROS'96*. 1996. Osaka, Japan.
7. Brezeal, C.L., *Sociable Machines: Expressive Social Exchange Between Humans and Robots*, in *Electrical Engineering and Computer Science*. 2000, MIT: Cambridge, MA. p. 264.
8. Kozima, H. and J. Zlatev. *An Epigenetic Approach to Human-Robot Communication*. in *IEEE International Workshop on Robot and Human Communication*. 2000. Osaka, Japan.
9. Takanobu, H., et al. *Waseda-SSSA joint research for human and humanoid robot interaction*. in *First IEEE-RAS Conference on Humanoids - Humanoids 2000*. 2000. Cambridge, MA.
10. Miyamoto, H., et al., *A kendama learning robot based on bi-directional theory*. *Neural Networks*, 1996. **9**(8): p. 1281-1302.
11. Schaal, S. and C.G. Atkeson, *Constructive Incremental Learning from Only Local Information*. *Neural Computation*, 1998(10): p. 2047-2084.
12. Atkeson, C.G., et al., *Using Humanoid Robots to Study Human Behavior*, in *IEEE Intelligent Systems*. 2000. p. 46-56.
13. Schaal, S., C. Atkeson, and S. Vijayakumar. *Real-time robot learning with locally weighted statistical learning*. in *International Conference on Robotics and Automation*. 2000. San Francisco.
14. Quartz, S.R. and T.J. Sejnowski, *The neural basis of cognitive development: A constructivist manifesto*. *Behavioral and Brain Sciences*, 1997(20): p. 537-596.
15. Rosen, R., *Life Itself*, in *Life Itself*, C.U. Press, Editor. 1991. p. 1-38.
16. Brooks, R., *The relationship between matter and life*. *Nature*, 2001. **409**(1): p. 409-411.
17. Penrose, R., *The Emperor's New Mind*. 1989, Oxford: Oxford University Press. 466.
18. Edelman, G.M., *Neural Darwinism: The Theory of Neuronal Group Selection*. 1988, Oxford: Oxford University Press. 371.
19. Leary, D.D.M.O., *Development of connectional diversity and specificity in the mammalian brain by the pruning of collateral projections*. *Current Opinion in Neurobiology*, 1992(2): p. 70-77.
20. Konczak, J., et al., *Development of goal-directed reaching in infants: Hand trajectory formation and joint force control*. *Experimental Brain Research*, 1995. **106**: p. 156-168.
21. Kelso, J.A., *Dynamic Patterns: the self organization of brain and behavior*. 1999, Cambridge, MA: MIT Press.
22. Thelen, E. and L.B. Smith, *A Dynamic System Approach to the Development of Cognition and Action*. 3 ed. 1998, Cambridge, MA: MIT Press. 376.
23. Atkinson, J., *The 'Where and What' or 'Who and How' of Visual Development*, in *The Development of Sensory, Motor and Cognitive Capacities in Early Infancy*, F. Simion and G. Butterworth, Editors. 1998, Psychology Press Ltd: Hove, East Sussex. p. 3-24.
24. Vapnik, V.N., *Statistical Learning Theory*, ed. Wiley. 1998, New York: Wiley. > 700.
25. Poggio, T. and V. Torre, *Ill-Posed Problems and Regularization Analysis in Early Vision*. 1990, MIT A.I. Laboratory,.
26. Evgeniou, T., M. Pontil, and T. Poggio, *A unified framework for Regularization Networks and Support Vector Machines*. 1999, MIT: Boston, MA. p. 1-40.
27. Johnson, M.H., *Developmental Cognitive Neuroscience*. 3 ed. *Fundamentals of Cognitive Neuroscience*, ed. M.J. Farah and M.H. Johnson. Vol. 1. 1997, Malden, MA and Oxford UK: Balckwell Publisher Inc. 234.
28. Turkewitz, G. and P.A. Kenny, *Limitation on Input as a Basis for Neural Organization and Perceptual Development: a Preliminary*

- Theoretical Statement*. Developmental Psychology, 1982. **15**: p. 357-368.
29. Sutton, R.S. and A. Barto, *Reinforcement Learning: an Introduction*. 1998, Cambridge: MIT Press.
 30. Rizzolatti, G., et al., *Premotor cortex and the recognition of motor actions*. Cognitive Brain Research, 1996(3): p. 131-141.
 31. Fadiga, L., et al., *Visuomotor neurons: ambiguity of the discharge of 'motor' perception?* International Journal of Psychophysiology, 2000. **35**(2-3): p. 165-177.
 32. Fadiga, L. and V. Gallese, *Action representation and language in the brain*. Theoretical Linguistics, 1997(23): p. 267-280.
 33. Rizzolatti, G. and L. Fadiga, *Grasping objects and grasping action meanings: the dual role of monkey rostroventral premotor cortex (area F5)*, in *Sensory guidance of movement, Novartis Foundation Symposium*, G.R. Bock and J.A. Goode, Editors. 1998, Wiley and Sons: Chichester. p. 81-103.
 34. Mataric, M.J., *Getting Humanoids to Move and Imitate*, in *IEEE Intelligent Systems*. 2000. p. 18-24.
 35. Rizzolatti, G. and M. Gentilucci, *Motor and visual-motor functions of the premotor cortex*, in *Neurobiology of Neocortex*, P. Rakic and W. Singer, Editors. 1988, Wiley: Chichester. p. 269-284.
 36. Milner, D.A. and M.A. Goodale, *The Visual Brain in Action*. Oxford Psychology. Vol. 27. 1995, Oxford: Oxford University Press.
 37. Jeannerod, M., *The Cognitive Neuroscience of Action*. Fundamentals of Cognitive Neuroscience, ed. M.J. Farah and M.H. Johnson. 1997, Cambridge, MA and Oxford UK: Blackwell Publishers Inc. 236.
 38. Kovacs, I., *Human development of perceptual organization*. Vision Research, 2000. **40**: p. 1301-1310.
 39. Wexler, M., et al., *Self-motion and the perception of stationary objects*. Nature, 2001. **409**: p. 85-88.
 40. Aloimonos, J. *Purposive and Qualitative Active Vision*. in *Proc. 10th Internat. Conf. on Pattern Recognition*,. 1990. Atlantic City, U.S.A.,.
 41. Bajcsy, R.K. *Active Perception vs. Passive Perception*. in *Third IEEE Workshop on Computer Vision: Representation and Control*. 1985. Bellaire (MI).
 42. Laquaniti, F. and R. Caminiti, *Visuo-motor transformations for arm reaching*. European Journal of Neuroscience, 1998. **10**: p. 195-203.
 43. Von Hofsten, C. and K. Rosander, *Development of Smooth Pursuit Tracking in Young Infants*. Vision Research, 1997. **37**(13): p. 1799-1810.
 44. Piaget, J., *The origin of intelligence in children*. 1936, London: International University Press, Inc. and Routledge & Kegan Paul Ltd.
 45. Metta, G., G. Sandini, and J. Konczak, *A Developmental Approach to Visually-Guided Reaching in Artificial Systems*. Neural Networks, 1999. **12**(10): p. 1413-1427.
 46. Sandini, G. and V. Tagliasco, *An Anthropomorphic Retina-like Structure for Scene Analysis*. Computer Vision, Graphics and Image Processing, 1980. **14**(3): p. 365-372.
 47. Schwartz, E.L., *A Quantitative Model of the Functional Architecture of Human Striate Cortex with Application to Visual Illusion and Cortical Texture Analysis*. Biological Cybernetics, 1980. **37**: p. 63-76.
 48. Metta, G., et al. *An Incremental Growing Neural Network and its Application to Robot Control*. in *International Joint Conference on Neural Networks*. 2000. Como, Italy.
 49. Bertenthal, B. and C. Von Hofsten, *Eye, Head and Trunk Control: the Foundation for Manual Development*. Neuroscience and Behavioral Reviews, 1998. **22**(4): p. 515-520.
 50. Panerai, F., G. Metta, and G. Sandini. *Learning VOR-like stabilization reflexes in robots*. in *8th European Symposium on Artificial Neural Networks*. 2000. Bruges, Belgium.
 51. Mussa-Ivaldi, F.A. and S.F. Giszter, *Vector field approximation: a computational paradigm for motor control and learning*. Biological Cybernetics, 1992. **67**: p. 491-500.
 52. Mussa-Ivaldi, F.A. and E. Bizzi, *Motor Learning through the Combination of Primitives*. Philosophical Transaction of the Royal Society: Biological Sciences, 2000. **355**(1404): p. 1755-1769.