Development in artificial systems

Giorgio Metta (pasa@ai.mit.edu)
AI-Lab – Massachusetts Institute of Technology, Cambridge MA, USA

Giulio Sandini (giulio@dist.unige.it)
LIRA-Lab – DIST, University of Genova, Genova, Italy

Lorenzo Natale (nat@lira.dist.unige.it)
LIRA-Lab – DIST, University of Genova, Genova, Italy

Riccardo Manzotti (manzotti@lira.dist.unige.it)
LIRA-Lab – DIST, University of Genova, Genova, Italy

Francesco Panerai (francesco.panerai@college-de-france.fr)
LPPA-Collége de France, Paris, France

Abstract

AI and robotics tackled the problem of building autonomous creatures (human-like sometimes) from many different directions. In spite of this, our systems are still primitive, little adaptive and prone to failure. We present a possible alternative to the construction of embodied artificial systems where the process of *development* is explicitly considered. This is in contrast with the common practice of *modular* design. We describe also an instantiation of this principle in an anthropomorphic robot interacting and *acting* in a real environment. The robot is equipped with diverse sensory modalities such as vision, proprioception, audition, and it interacts with the environment through a sophisticated motor apparatus.

The approach

Sensori-motor development in biological systems is a structured process, where all the sensing mechanisms are gradually integrated with the actuating mechanisms in order to achieve a mature performing system. An extremely interesting aspect of this process is its underlying ability to deal with and solve complex multi-level integration problems. From an engineering point of view, development can be thought of as a hierarchy of integration processes which involve at different times different action systems. The whole process leads each single newborn through intermediate developmental stages characterized by different level of functionalities and capabilities. Eventually the procedure results in the achievement of the final performing and mature individual (Piaget, 1936; Thelen & Smith, 1998).

In robotics, the attempts to build complex robots (e.g. humanoids robots) have always followed principles of modular design: often, to make the problem tractable, engineers applied a "divide and conquer" strategy. Following this approach, each *module* deals with a specific sensory modality and/or a particular motor skill. Although, this has been successful in some cases (Brooks, 1996; Hirai, Hirose, Haikawa, & Takenaka, 1998), the final integration proved to be very difficult, and often failed. Moreover, the system was assumed to be time-invariant and consequently *adaptation and learning* were not explicitly taken into account.

This paper presents an alternative design methodology, which although in its infancy, it is suited to deal with the complexity, adaptation, and integration issues of artificial and biological systems. Taking into account the lesson learned from the modular design approach and trying to overcome the problem of the final integration, we propose the use of *engineered principles of development* to build complex systems.

The aim of the theory should be, firstly, to achieve efficient and adaptive robots and, secondly, to develop a "neurobotic toolbox" in which ideas about neural computation can be tested by engineering working systems. The latter issue should be considered of particular interest for both robotics and neuroscience due to the possibility to foster an interdisciplinary approach to the study of complex systems.

The theory should address two level of analysis: i) system analysis proper, where the system is described, constituents individuated and the functioning explained, and ii) construction theory: that is, how to design a working system; how to take into account time-variance, adaptation and learning. The first issue has been mainly the approach taken by biologists; the latter concerned mostly engineers – we propose a tighter integration of the two fields.

The study of the biology – the modeling of brain functions – can certainly suggest how to build more successful and adaptable "artificial beings". Adaptation raises the issue of learning; in other words, how can the learner acquire useful information in order to accomplish a given task? Which sensors does it need? Is learning always feasible? Until now, robotics and AI have failed to give a definitive answer (assuming it does exist) and indirectly they have also failed to produce truly autonomous and flexible agents. In spite of many successes in building robots of various shape, size, abilities, sensory types, etc. there seems to be something lacking in terms of "cognitive abilities", as well as adaptability of the system to the dynamic of the environment.

In recent times "brain sciences" also face an increasingly and intricate picture, where it is hard to discover the general underlying principles, which eventually will bear the real explanatory power. The panorama consists of a huge number of brain areas and intricate interconnections between them (the so-called "telephone switchboard" model). Here lies the significance of modeling: many researchers have employed computational models to explain functions, to derive general rules, and to integrate data gathered by using different methodologies. Perhaps, not everything is suitable to be simulated, and the world – the "external" environment – is something, which is too complex to be replicated appropriately. In this sense a robot can be a suitable tool where to condense our knowledge of brain functions and to test our models, ultimately, against a real environment.

In this context, development can provide a vantage point, for observing the temporal evolution of particular skills, in particular, to determine, what prerequisites, drives, and teaching signals are needed to learn them. This offers the clearest view on learning, not only, in algorithmic terms, but also about its logical consistency.

The following section outlines some of the foundations of a possible theory. Although some of the elements are at the moment speculative, testing and verification are part of the ongoing research.

Further, we describe an instantiation of the approach in a behaving developing agent, which learns sensori-motor coordination during the interaction with the environment. We are aware that this is not yet complete but by adopting a learning by doing philosophy, we hope, in the future, to frame the approach in a more consistent mathematical framework.

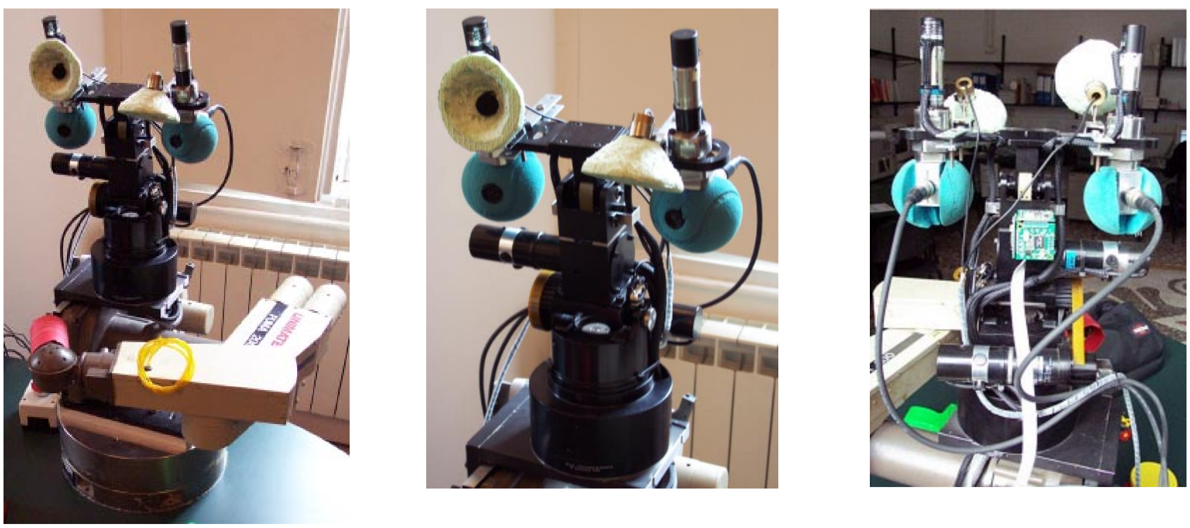

Figure 1: Three shots of the experimental setup. *Left*: complete view of the robot; *middle*: close up view of the head, note the cameras and microphones with earlobes; *right*: rear view of the robot and in particular of the gyros.

Outlines of a theory

The first hints about what to include in a theory of developing systems comes from statistics and learning theory. The main result we refer to is derived from Statistical Learning Theory (Vapnik, 1998). Vapnik et al. suggest that the key of learning in a realistic setting is *complexity control*. That is, not only the parameters of the learning machine need to be optimized, but the structure of the machine as well. We suggest that development is about controlling complexity. This approach justifies also the empirical observation that an increase of complexity is observed during development (Turkewitz & Kenny, 1982). The presence of regulatory processes in the brain, which change the complexity of the learner, have been documented (Johnson, 1997) although rarely interpreted in this way (Quartz & Sejnowski, 1997).

To put development in an "ecological" context, we need further to consider how training data is collected. This is because, data does not come for free: a cost might be associated to the effort of collecting it. Action (and embodiment) needs to be explicitly considered.

The recent analysis of the premotor cortices and the discovery of, for example, mirror neurons (Di Pellegrino, Fadiga, Fogassi, Gallese, & Rizzolatti, 1992; Rizzolatti & Fadiga, 1998) cast a new light in the necessity of action as a prior to cognitive functions. Another element of the theory is thus the view of action as a foundation. Aspects of the development of the ventral vs. dorsal stream also support this notion. Kovacs (Kovacs, 2000), in fact, pointed out that the control of movements (dorsal stream) develops before the ability to categorize (ventral stream). Other examples can be found in (Atkinson, 1998).

Although not immediately obvious, the theory requires, at an even more fundamental level, goal-directness and *causal understanding*. In fact, the very basic element of every learning procedure is that of detecting regularities in the sensory data together with the ability to link the effects with the appropriate causes. Goal-directness has to be seen as a way to understand when the result of an action is within a particular class irrespective of the viewpoint or sensory modality – classes are, for example, grasping, holding, tearing, etc. Any low-level sensory-motor coordination task (e.g. reaching, eye movements, etc.) can be learnt in principle by repetitive trials. The efferent copy and the sensory reafferences generated during action need to be appropriately linked and synchronized if the cause and effect ought to be correctly interpreted. Appropriate mapping can fill in the gaps (i.e. the transformations) and their inverses to recover the correct movement to apply in a given situation. This procedure can only take into account the learning of approximately direct mappings between the sensory and the motor space. The cause-effect principle has conversely a much broader applicability.

At the moment, we regard these considerations, mostly as speculative, because a real theory should be backed up by appropriate data. To recap, the goals of defining a theoretical developmental framework are:

i) Engineered principled: define what are the necessary elements of the theory of learning. Examples we put forward are complexity control, the primacy of action, and understanding of causality. The latter has to be intended in practical terms rather than in general ones.

ii) Definition of a neurobotic toolbox: provide an environment where to test biologically plausible theories of development. It is not required to be humanoid-shaped, although we believe it has to be an embodied device with the ability to act in the real world. We already discussed, why we need action as a constituent of the theory.

iii) Ongoing investigation is devoted mostly to the identification of the constituents of the model. From the robotic side, a model incorporating some of the discussed learning and development aspects has been realized (see next section).

iv) It is clear that more experiments need to be carried out in order to answer specific questions developmentally. This might include, in the future, experiments with animals or human subjects. A measure of the success of the method is certainly the possibility of suggesting particular aspects of the theory that need to be worked out, or conversely, where the biological explanations are missing. The farther we will get along this line, the more successful the *integrated approach* we are proposing.

Delivering the promises (or part of them)

The artificial system, named Babybot (Metta, Sandini, & Konczak, 1999; Sandini, Metta, & Konczak, 1997), consists of a twelve degree-of-freedom eye-head-arm platform equipped with vision, audition, proprioception and vestibular senses. The robot acquires and processes images in a space variant format also known as log-polar (Sandini & Tagliasco, 1980). The robot's eyes observe the world through a high-resolution *fovea* and a lower resolution periphery. Acoustic sensation is provided by a couple of microphones and plastic earlobes. The vestibular sensor consists of three gyros arranged along mutually orthogonal axes and it provides a measure of the angular velocity of the head in space (Panerai, Metta, & Sandini, 2000). A global view of the setup is shown in figure 1, together with a close-up view of the head and the vestibular sensor.

The process we focused on is how Babybot learns to coordinate the eyes, head and arm from an initial stage where control is mostly reflex-based. Figure 2 shows the temporal dependences of the different maps controlling the Babybot and where learning is applied.

The specific learning tasks the robot is involved in (see (Metta, 2000)) steer the system from the initial reflex-driven stage, with the help of some noise (Metta, Carlevarino, Martinotti, & Sandini, 2000), toward more stable patterns of sensori-motor coordination. For example, initially the robot learns the relationship between eye movements and image displacements, which is needed to correctly foveate a spotted location, but also to move the system into a second stage. In fact, from the initial short eye movements the robot learns how to generate correct saccades. As the eyes get under control, more complicated coordination patterns might emerge, the neck can be moved, and thus head-eyes coordination can be acquired. This latter stage is accomplished by learning both the proper eye to head coordination map, but also a vestibulo-ocular reflex (VOR). VOR itself is fine tuned automatically on the basis of a visuo-motor performance measure based on optical flow. Intuitively if the whole coordination schema is not performing well, the lack of stability is translated into a residual optical flow, which drives a corresponding teaching signal tailored at reducing the instability. In parallel, the system practice also with reaching by trying roughly to touch whatever object is located in its field of view. Of course, at least initially, most of these tentative movements are wrong. Babybot, though, improves by measuring its errors (only visual in this case), and after some hundreds of trials learns how to reach appropriately within the arm's workspace. It is worth noting that vision is not the only available sensory modality; sound for example is used to redirect the attention of the robot toward possibly interesting events. Also in this case, maps linking the spatial perception of sound with action systems are learnt. Further, audition and vision are integrated in a coherent percept, which allows the robot to consistently redirect gaze when both sound and vision are present or when only one of them is available.

As mentioned before, every sensory modality available to the robot is activated from the very beginning. Conversely the sequence of possible action systems is limited by various factors (i.e. only the eyes are appropriately coordinated at the very beginning, the neck moves subsequently, and the torso later in time (Bertenthal & Von Hofsten, 1998)). This dependence shapes in a particular way both learning and the kind of data gathered while learning.

The different "cortical" maps of Babybot are tuned in order to work together and integration is already embedded into the system. The sequence of learning events is constrained in part by the physics of the system, in part by how the modules are designed. The resulting maps can develop in many different ways depending on the particular history of learning itself. For example, a lack of vestibular stimulation, or a failure of the sensor itself, would lead to a reduced coordination of the eye-head system. As a consequence, the gain of the movement would be automatically retuned because of the diminished stability and increased residual optic flow. This is a very diverse adaptation with respect to a system where the different parts are designed in isolation. Further, because the maps are embodied and mutually interacting the state space experienced by the learner might rest in a sub-manifold of the actual parameter space: i.e. not all combinations of parameters are valid or physically generable. The benefit is that the exploration space might be reduced as well as the time required to acquire particular skills.

Conclusion and future work

In conclusion, development might be a necessary characteristic of adaptive systems: in other words, the particular sequence of developmental events simplifies learning by constraining the possible "roads" towards skilled behavior. For example, newborns do not control all available degrees of freedom but rather start by practicing only a few. External/physical (like gravity) and internal/neural limitations (like reflexes or initial synergies) effectively shape how learning has to be carried out.

From the theoretical point of view, statistical learning theory (Vapnik, 1998) actually endorses a procedure where the complexity of the system is controlled or increased as a function of the "experience" of the system (i.e. the number of training samples). A developing system might take advantage of this situation by increasing the complexity of the controller over time. This insight can be also employed at a different level of analysis to justify a growth of complexity of the controlling neural networks (Quartz & Sejnowski, 1997). This is not in contrast with prevailing theories on neural selection (Edelman, 1988). The key aspect in fact is the control of complexity, thus, two processes (growth and selection) might balance and influenced by the environment optimally tune the "size" of the networks involved in sensori-motor as well as cognitive processes.

Future work is directed toward a more formal definition of the many facets involved in the developmental paradigm. It is worth mentioning that the sensori-motor layer represents only a very first stage of development, which is at the moment artificially carried out without cognition. This is clearly unrealistic (in human development); a value system, drives and other higher-level processes are certainly needed to complement the model and give the robot the abilities to bootstrap autonomously to new levels of development.

References

Atkinson, J. (1998). The 'Where and What' or 'Who and How' of Visual Development. In F. Simion & G. Butterworth (Eds.), *The Development of Sensory, Motor and Cognitive Capacities in Early Infancy* (pp. 3-24). Hove, East Sussex: Psychology Press Ltd.

Bertenthal, B., & Von Hofsten, C. (1998). Eye, Head and Trunk Control: the Foundation for Manual Development. *Neuroscience and Behavioral Reviews, 22*(4), 515-520.

Brooks, R. (1996, November 4-8). *Behavior-Based Humanoid Robotics.* Paper presented at the IEEE/RSJ IROS'96, Osaka, Japan.

Di Pellegrino, G., Fadiga, L., Fogassi, L., Gallese, V., & Rizzolatti, G. (1992). Understanding motor events: a neurophysiological study. *Experimental Brain Research*(91), 176-180.

Edelman, G. M. (1988). *Neural Darwinism: The Theory of Neuronal Group Selection*. Oxford: Oxford University Press.

Hirai, K., Hirose, M., Haikawa, Y., & Takenaka, T. (1998, May, 16-20). *The development of Honda humanoid robot.* Paper presented at the IEEE International Conference on Robotics and Automation, Leuven, Belgium.

Johnson, M. H. (1997). *Developmental Cognitive Neuroscience* (3 ed. Vol. 1). Malden, MA and Oxford UK: Balckwell Publisher Inc.

Kovacs, I. (2000). Human development of perceptual organization. *Vision Research, 40*, 1301-1310.

Metta, G. (2000). *Babybot: a Study on Sensori-motor Development.* Unpublished Ph.D. Thesis, University of Genova, Genova.

Metta, G., Carlevarino, A., Martinotti, R., & Sandini, G. (2000). *An Incremental Growing Neural Network and its Application to Robot Control.* Paper presented at the International Joint Conference on Neural Networks, Como, Italy.

Metta, G., Sandini, G., & Konczak, J. (1999). A Developmental Approach to Visually-Guided Reaching in Artificial Systems. *Neural Networks, 12*(10), 1413-1427.

Panerai, F., Metta, G., & Sandini, G. (2000). Visuo-inertial Stabilization in Space-variant Binocular Systems. *Robotics and Autonomous Systems, 30*(1-2), 195-214.

Piaget, J. (1936). *The origin of intelligence in children* (M. Cook, Trans.). London: Interntational University Press, Inc. and Routledge & Kegan Paul Ltd.

Quartz, S. R., & Sejnowski, T. J. (1997). The neural basis of cognitive development: A constructivist manifesto. *Behavioral and Brain Sciences*(20), 537-596.

Rizzolatti, G., & Fadiga, L. (1998). Grasping objects and grasping action meanings: the dual role of monkey rostroventral premotor cortex (area F5). In G. R. Bock & J. A. Goode (Eds.), *Sensory guidance of movement, Novartis Foundation Symposium* (pp. 81-103). Chichester: Wiley and Sons.

Sandini, G., Metta, G., & Konczak, J. (1997, November). *Human Sensori-motor Development and Artificial Systems.* Paper presented at the AIR&IHAS '97, Japan.

Sandini, G., & Tagliasco, V. (1980). An Anthropomorphic Retina-like Structure for Scene Analysis. *Computer Vision, Graphics and Image Processing, 14*(3), 365-372.

Thelen, E., & Smith, L. B. (1998). *A Dynamic System Approach to the Development of Cognition and Action* (3 ed.). Cambridge, MA: MIT Press.

Turkewitz, G., & Kenny, P. A. (1982). Limitation on Input as a Basis for Neural Organization and Perceptual Development: a Preliminary Theoretical Statement. *Developmental Psychology, 15*, 357-368.

Vapnik, V. N. (1998). *Statistical Learning Theory*. New York: Wiley.

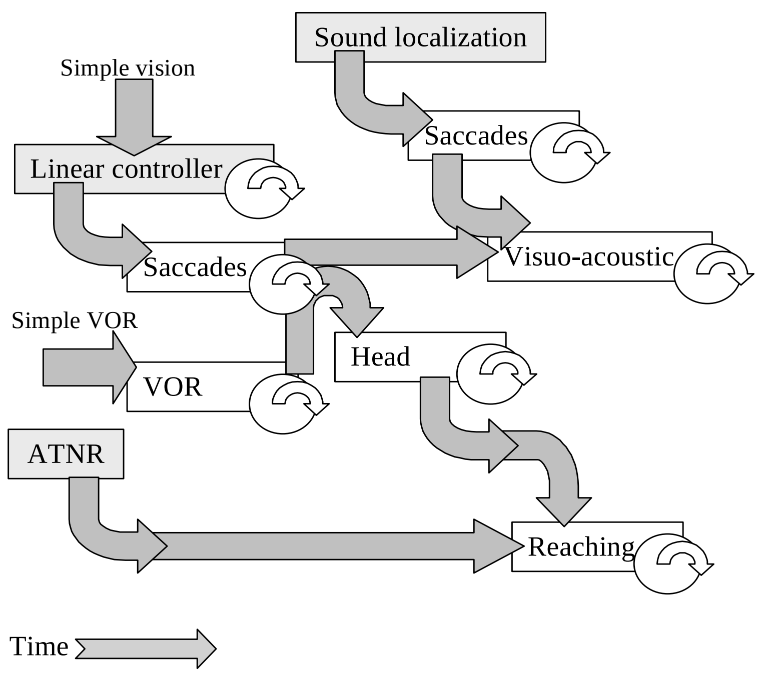

Figure 2 The developmental progression of the Babybot. A representation of how the various phases of the robot's development are interrelated. The light gray blocks are the initial configuration. See text for description of the learning blocks.