Development: Is it the right way towards humanoid robotics?

Giorgio Metta[1], Riccardo Manzotti[1], Francesco Panerai[2], Giulio Sandini[1]

[1]LIRA Lab, DIST – University of Genova – Italy
[2]LPPA, CNRS – Collège de France, Paris – France

Abstract: **The studies presented in this paper stem from an interdisciplinary approach covering aspects of "brain sciences" and robotics, with the goal of answering several questions, namely: is it possible to test hypotheses of the brain functions involved in a particular task, by implementing biologically plausible models on a real physical system such as a robot? How can we design more adaptable and potentially efficient robots? Is it possible to build a truly human-like robot? And finally, is development the right way towards humanoid robots?**

We shall argue whether developmental studies might provide a different and potentially interesting perspective either on how to build an artificial adaptive agent, or on understanding how the brain solves sensory, motor, and cognitive tasks.

From the modeling point of view we shall demonstrate how a twelve degrees of freedom "baby" humanoid robot acquires orienting and reaching behaviors, and what the advantage of the proposed framework over traditional learning paradigms is.

INTRODUCTION

Research activity linking studies on artificial systems to "brain sciences" is certainly not new. Beside the studies on artificial neural networks, substantial efforts are devoted worldwide in building "physical models" of segment of biological systems, with the aim of suggesting novel solutions to robotics or processing problems and to advance our understanding of human brain functions [1, 2]. The main advantage of using robots rather than pure computer simulations, at least in the study of the motor system, is that the physics of the environment comes "for free" – a proper simulation would be difficult, if not impossible.

In our opinion, there might be commonalities, sometimes due to the nature of the tasks, sometimes to the physics itself, which suggest that both artificial and biological agents could consistently employ the same solutions. The study of the "biology" – the modeling of brain functions – could suggest how to build more successful and adaptable "artificial beings".

On the other hand, the quest for 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? Up to now, robotics and AI have not provided a definitive answer, and consequently truly autonomous and flexible agents are still very limited. In spite of many successes on building robots of various shape, size, abilities, sensory types, etc there seems to be something lacking in terms of "cognitive abilities", as well as system adaptability to the dynamic of the environment. Moreover, even for those "successful" robots, the integration of different behaviors and sensory modalities gave rise to a series of unexpected problems. The traditional artificial learning paradigm faced such difficulties, perhaps because of some wrong assumptions about the learning process itself, rather than the lack of proper models and algorithms.

In this context, "brain scientists" have studied, since a long time, the acquisition of behavior and cognitive abilities, and nobody is surprised by the fact that newborns are not simply a sort of "reduced size human beings". What is more surprising is that, even at that age, infants show a series of "innate" behaviors, basic control synergies, and reflexes. On this basis, more sophisticated behaviors develop, and this process undergoes through stages, where the limited abilities already formed are efficiently exploited in order to simplify the learning process itself.

On the contrary, the approach followed in robotics is mainly that of designing the "complete manufact" (i.e. the adult-like robot). One might wonder what about that is wrong. Perhaps, something was underestimated, and from a purely engineering point of view, this "something" was the whole process of design. Dennet [3], for instance, thinks that the overall design process must be included in the specifications of the manufact. This approach shifts the emphasis from the "final product" to the "process" of building it; the goal of the designer becomes that of devising a suitable initial state (at time t_0), and the appropriate developmental rules to get some close approximation of the desired "final product". Given these considerations, we conjecture that developing systems, either artificial or natural, can be or are equipped to tackle the problem of learning consistently.

In the following sections, we shall highlight first what the advantages of a developing agent over a traditional learner might be. Secondly, we will describe briefly those aspects, which are relevant for the design of a developing robot, next to a significant real-world experiment.

LEARNING AND DEVELOPMENT

It has been recognized that learning from examples is an ill posed problem [4]. Every learner faces the so-called "theoretical pressures", which require balancing competing needs in order for learning to be feasible. Recently, a number of theories on learning formalized these problems [5]. Generally speaking, a learner should be able to learn from incomplete information, using a limited number of samples, and quickly enough to cope with changes in the environment, as well as of its internal physical parameters (e.g. growth, malfunctions, etc). The first step for any learning agent is that of acquiring information through the interaction with the environment. However, without any a-priori information, it is hard to tell which part of the "state[1]" space is worth exploring in order to solve a particular task. In fact, the size of the state space might consist of dozens of dimensions, which precludes whatever sort of enumerative search for a solution. It is not always true that the solution belongs to the whole state space; on the contrary, in many cases the actual problem rests on a lower dimensionality manifold [6]. This suggests that, if the learning process is carried out together with the identification of the relevant sub-manifold, a complete exhaustive search can be avoided.

It turns out that learners have two competing requirements in terms of exploring the control/state space, and in responding as much as possible appropriately to stimuli (i.e. exploit their knowledge). Recent research on human development suggests that such exploration component might be provided naturally by noise. In fact, newborns show several noise sources due to incomplete structures (non-myelinated neurons are an example [7]), to unnecessary neural branching [8], and to the use of random behavior actively [9]. This role of noise during learning resembles the use, in system theory, of broadband (e.g. white noise) input signals for system identification purposes.

Besides, other researchers provided evidences for the existence of a strong "goal-directed" behavioral component, even in newborns [10]. The fact that the behavior is goal-directed can speed up the acquisition of the appropriate controller with respect to a fully random explorative search. In this last case, the learner has to visit all possible states prior to any actual control; otherwise, a possibly useful part of the state space might remain unexplored. It is also worth noting that the cooperation of many control loops developing with different time spans can help in reducing the already mentioned exploration space in the sense that each control loop may generate a bias for subsystems that develop later. In the context of "computational motor control", one notable example of such a schema is the feedback-error learning model [11]. In this case, an inverse modeling is carried out through the interaction of a learner with a much simpler feedback loop. Similar multi-loop structures can also be observed in the brain. An example of this process is the so-called cortical take-over, where cortical areas develop on top of sub-cortical structures [12]. The delay and bandwidth involved in the various structures can be different thus providing the basis for faster reactions (reflex-like) and accurate control at the same time (consider, for instance, the visuo-vestibular integration).

Implicit in the preceding discussion is the assumption that the learning agent is "functional" from the beginning, which means that the training data must be collected "on-line". This is a major constraint for biological as well as artificial systems. Concerning biological systems, it is clear that they cannot be some sorts of "blank slate" at birth; they rather need to have some useful bootstrap functionality. These "initial behaviors" are usually reflex-like and stimulus bound in nature. They can be thought as the initial "bias[2]", and perhaps their role is indeed that of guiding the system through feasible regions of the state space.

It is worth stressing that the exploration-exploitation trade-off is closely related to the well-known engineering problem called "the curse of the dimensionality" [13]. In fact, the need for representational resources grows exponentially with respect to a linear growth of the number of dimensions. For an on-line learner, the time to explore the state space would suffer of this remarkable growth.

Of course, this is far from being a complete argument on what Sejnowski and colleagues [14] called the "theoretical pressures", although we may conjecture on why development is a necessary procedure to simplify learning, and why developing systems can be superior in terms of skill acquisition. For instance, the "bias-variance" dilemma can be handled by controlling the complexity of the learning structure [5]. In biological systems, this is thought to happen during development; it is not only a competitive learning but also a growth process. As already mentioned, the brain is not a monolithic controller, and consequently the sequence of developmental events (i.e. when different subparts get activated) is important. We also argued that noise plays a fundamental role in driving exploration, therefore, the fact that the newborn is perceptually and motorically limited can be seen as a positive factor.

In summary, we argue that we can borrow these mechanisms in the process of designing an artificial autonomous system and, furthermore, we can notice that part of this machinery is only effective in the context of a developing agent interacting with the environment [14-16].

BUILDING THE DEVELOPING SYSTEM

The experimental setup (see Figure 1 below) consists of a five degrees of freedom robot head (designed and realized at Lira Lab), and an off-the-shelf six degrees of freedom robot manipulator (an Unimation Puma260), both mounted on a rotating base: i.e. the torso. The kinematics resembles that of the upper part of the human body although with less degrees of freedom. From the sensory point of view, the Babybot is equipped with two space-variant cameras [17], an inertial sensor simulating the vestibular system [18], and proprioceptive information through motor encoders. The robot is controlled by a set of PCs – ranging from Pentium II to Pentium III processors – each running Windows NT and connected by a fast Ethernet link. In order to provide the necessary interface with the hardware (i.e. sensors and motors) some machines are equipped with motion control boards, frame grabbers, AD converters, etc. In particular one machine controls the robot arm and the torso, another one the head, and a third computer carries on the visual processing. Concerning the software, it adheres to DCOM, a standard, which allows running objects among the various machines. The reference task, for this discussion, is the coordination of

[1] The state space can be a proper state space, the parameters manifold, or a combination of the two depending on the kind of learning algorithm considered.

[2] Proper bias selection leads to another impasse usually called the bias-variance dilemma.

eye-head-arm movements, with the aim of gazing and reaching for visually identified objects in extrapersonal space.

Figure 1 – The experimental setup. The 12 degrees of freedom humanoid robot described in the experiment.

In practice, the system is able "at birth" to move the eyes only. Control, at that stage, is a mixture of random and goal-directed movements. Concerning the head-arm coordination, the robot possesses at the beginning only a reflexive behavior simulating basic muscular synergies and spinal reflexes. The initial task of the control process is that of calibrating the closed loop gains. It is worth stressing that even at the very beginning the system is already moving in a "goal-directed" manner, although noise dominates the actual movements. In successive phases, the robot starts learning fast eye movements (saccades), but only the eyes are moving. Indeed, this is necessary because otherwise the neck motion would disturb the estimation of the required eye commands (i.e. part of the required eye movement would be indirectly performed by the head motion).

Once eyes are under "proper" control, the whole head starts moving, at this point, the eye controllers are well formed and can be used to help coordination of the redundant eye-head degrees of freedom. Concurrently, reaching steadily improves by storing more information in a head-arm coordination map; as a result, the initial reflexes become part themselves of the coordinative action. Because reaching depends on gazing, during the initial phases, reaching improves slowly. Later on, as soon as gazing gets to a reasonable performance level, also reaching improves quickly. It is worth stressing that, from the robot's point of view, motor control can be seen as "learning" to combine the initial "skills" – i.e. reflexes – in order to obtain voluntary goal-directed movements.

A sort of vestibulo-ocular reflex (VOR) is always on. The robot learns the appropriate eye compensatory responses by minimizing a performance measure of image stabilization (i.e. the optic flow). When the first multi-joints eye-head movements are practiced, the VOR is already effective in facilitating coordination [19].

Once head and arm controls are in place, the robot can orient appropriately toward moving stimuli, follow them while moving, and eventually, touches the tracked object. Roughly speaking, Babybot starts by looking at objects, which are identified by means of color and motion. It can correctly perform saccadic eye movements, and it possesses a sort of smooth pursuit ability. It is worth mentioning that only the eyes are controlled directly by means of visual information (the neck and arm follow). The redundant DOF are easily "centrally" coordinated. This ability to gaze is the first step toward yet another visually driven behavior: i.e. reaching. By mapping gaze direction into appropriate motor commands, the robot can effectively reach for objects in extrapersonal space. This map, at least initially, does not necessarily, brings the end-effector near the fixation point. However, instead of correcting the error by moving the arm, the direction of gaze is redirected to the end-effector and the arm motor command previously issued is associated to the new eye position. As the learning process proceeds, the initial arm motion gets closer and closer to the visual target, and eventually, the corrective gaze shift will not be necessary unless kinematic changes occur. For a complete description of the learning sequence, see [20]. Moreover, thanks to a low stiffness controller, Babybot can safely interact with humans and the external environment. If our will is to build a truly autonomous system, this robot-environment-humans interaction is of paramount importance.

Learning of the various maps is carried out by a growing neural gas type network, which is able to tune its growth rate on the basis of the approximation error. By tailoring the model complexity to the actual approximation requests, the network can avoid over-fitting and over-smoothing.

Figure 2 shows the trajectories of the fixation point and the arm end-point. They have been acquired during an unrestrained experiment: i.e. an experimenter handled a target, in such a way to cause the robot to react. The whole experiment endured for about half an hour during which joint positions were recorded at 25*Hz* rate. Figure 3 plots saccades during the initial phases of learning. Abscissa and ordinates, in this case, represents the image plane. Note that all trajectories are actually converging to the fovea.

CONCLUSION

This paper presented a proposal for a novel approach aimed at the design and comprehension of complex systems. This approach arose by observing how biological systems solve the problem of learning and adaptation during the early stages of their lives. We tried to isolate those aspects, which may be relevant both for the construction of artificial systems and for advancing our understanding of the corresponding brain functions. An important point worth stressing is that the brain cannot be seen as a monolithic structure, but rather we need to look at it as a developing system, where many subparts optimally interact. This internal organization might indeed facilitate learning and in this sense it is worth copying when one goes through the design of an "artificial adaptive agent". We are aware that we did not provide any formal justification, but at least we provided hints on what aspects might be relevant. Finally, by using a "learning by doing" philosophy, we built a humanoid robot, and "programmed" it following some of the biological aspects we denoted as "relevant" for artificial development. The robot indeed faced problems, such as moving many degrees of freedom by employing many different cooperating controllers. This is exactly the point, how should we connect

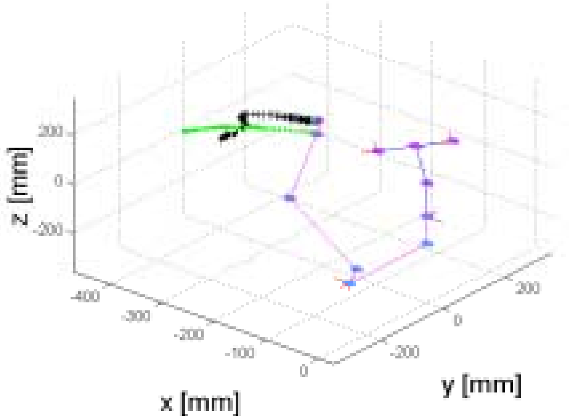

Figure 2 – Gazing and reaching. Two trajectories are shown, the fixation point and the arm end-point (100 samples are displayed). The simple wire-frame model represents the robot. Small circles indicate joints; solid lines are the links.

all these modules together? Consider that they are not separated because all of them act on the same non-linear physical plant. Consequently, interactions must be explicitly taken into account. We devised a solution, where the timing of adaptation is carefully (but not too much) programmed. That is, the solution goes by creating a proper time slot for each subpart (slots do not need to be temporally separated one from another). Inside this "critical periods", adaptation can effectively take place without disturbing much the other modules. This is important, especially in the early phases, when plasticity must be high (i.e. exploration) in order to quickly acquire a consistent behavior. Yet another type of interaction occurs: modules that develop first influence modules that develop later. Consequently, the "explored state space" depends much on how these early controllers behave. Each module can function as a "bootstrap" procedure for other subsystems. This is exactly "constructive learning" on a coarse scale, where entire streams, areas, controllers can be considered as "basis modules". Constructive learning is thought to be superior to other learning techniques (pruning based). So, the spotlight moved from learning itself to the process of learning: i.e. development. What and how could be learned is determined by the learner's developmental stage, that is, by what the state of the whole system is in terms of the other subparts (e.g. the robot could not move the neck without controlling the eyes first).

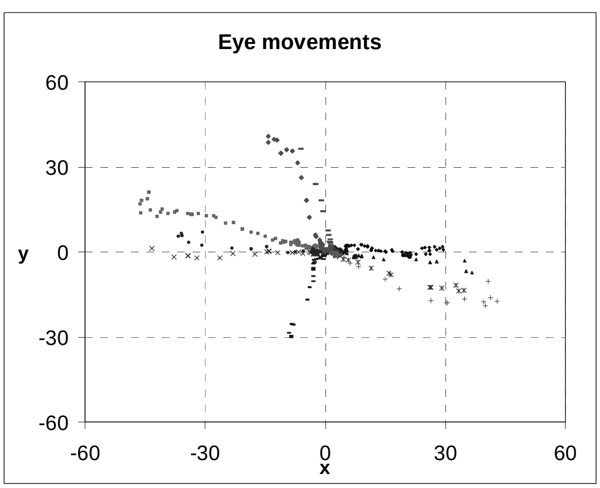

Figure 3 Eyes trajectories in the image plane. Coordinates are expressed in pixels.

REFERENCES

[1] R. Brooks, "Behavior-Based Humanoid Robotics," presented at IEEE/RSJ IROS'96, 1996.

[2] G. Sandini, "Artificial Systems and Neuroscience," presented at Proc. of the Otto and Martha Fischbeck Seminar on Active Vision, 1997.

[3] D. C. Dennet, *L'idea pericolosa di Darwin - Orig. Darwin's dangerous idea. Evolution and the meaning of life.*, First edition ed. Torino: Bollati Boringhieri, 1997.

[4] T. Poggio, V. Torre, and C. Koch, "Computational Vision and Regularization Theory," *Nature*, vol. 317, pp. 314-319, 1985.

[5] V. N. Vapnik, *Statistical Learning Theory*. New York: Wiley, 1998.

[6] S. Schaal and C. G. Atkeson, "Constructive Incremental Learning from Only Local Information," *Neural Computation*, pp. 2047-2084, 1998.

[7] E. R. Kandel, J. H. Schwartz, and T. M. Jessel, *Principles of Neuroscience*: Elsevier, 1991.

[8] D. D. M. O. Leary, "Development of connectional diversity and specificity in the mammalian brain by the pruning of collateral projections," *Current Opinion in Neurobiology*, pp. 70-77, 1992.

[9] R. H. S. Carpenter, "A Neural Mechanism that Randomises Behaviour," *Journal of Consciousness Studies*, vol. 6, pp. 13-22, 1999.

[10] A. Streri, *Seeing, Reaching, Touching - The Relations between Vision and Touch in Infancy*: MIT Press, 1993.

[11] M. Kawato, K. Furukawa, and R. Suzuki, "A hierarchical neural network model for control and learning of voluntary movement," *Biological Cybernetics*, pp. 169-185, 1987.

[12] J. Atkinson, "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, Eds. Hove, East Sussex: Psychology Press Ltd, 1998, pp. 3-24.

[13] R. E. Bellman, *Dynamic Programming*. Princeton: Princeton University Press, 1956.

[14] S. R. Quartz and T. J. Sejnowski, "The neural basis of cognitive development: A constructivist manifesto," *Behavioral and Brain Sciences*, pp. 537-596, 1997.

[15] R. Pfeifer and C. Scheier, "Sensory-motor coordination: The metaphor and beyond," *Robotics and Autonomous Systems*, vol. 20, pp. 157-178, 1997.

[16] R. S. Sutton and A. Barto, *Reinforcement Learning: an Introduction*. Cambridge: MIT Press, 1998.

[17] G. Sandini and V. Tagliasco, "An Anthropomorphic Retina-like Structure for Scene Analysis," *CVGIP*, vol. 14, pp. 365-372, 1980.

[18] F. Panerai and G. Sandini, "Oculo-Motor Stabilization Reflexes: Integration of Inertial and Visual Information," *Neural Networks*, vol. 11, pp. 1191-1204, 1998.

[19] F. Panerai, G. Metta, and G. Sandini, "Visuo-inertial Stabilization in Space-variant Binocular Systems," *Robotics and Autonomous Systems*, vol. 30, pp. 195-214, 1999.

[20] G. Metta, G. Sandini, and J. Konczak, "A Developmental Approach to visually-guided reaching in artificial systems," *Neural Networks*, vol. 12, pp. 1413-1427, 1999.