

Babybot: an artificial developing robotic agent

Giorgio Metta, Francesco Panerai, Riccardo Manzotti, Giulio Sandini
LiraLab, DIST – University of Genova
Viale Causa, 13 – 16145 Genova, Italy
pasa@lira.dist.unige.it

Abstract

The study of development, either artificial or biological, can highlight the mechanisms underlying learning and adaptive behavior. 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. It is our opinion that the acquisition of the proper behavior might indeed be facilitated because within an ecological context, the agent, its adaptive structure and the environment dynamically interact thus constraining the otherwise difficult “learning problem”.

In very general terms we shall describe the proposed approach and supporting biological related facts. In order to further analyze these aspects 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 kind of advantages the proposed framework might offer.

1. Introduction

Research efforts linking studies on artificial systems to “brain sciences” are certainly not new (Beer, Chiel, Quinn, & Ritzmann, 1998; Voegtlin & Verschure, 1999). Next to the studies on artificial neural networks and genetic algorithms, some authors suggested that a synthetic approach to the study of biological systems might be helpful in deepening our understanding of the otherwise patchy data (Sandini, 1997). In particular some argued that by building “physical models” of segment of biological systems, we might be able either to suggest novel solutions to robotics or processing problems and to advance our understanding of human brain functions (Brooks, 1996; Sandini, 1997). For example, the control of robot heads and visually guided manipulation tasks were studied with reference to psychophysical performance data of humans and animals (Bajcsy, 1985; Ballard & Brown, 1992; Capurro, Panerai, Grosso, & Sandini, 1993). In this respect, many levels of similarity with biological systems can be considered: from emulation to a vague resemblance. The important point is to

grasp, somehow, the relevant aspects of biological systems so that we can both address specific biological questions and propose new methodologies to robotics.

One might argue that a pure simulation could very well serve the scope, however, it is important to note that a proper simulation would be very difficult, if not impossible. Consequently the main advantage of using robots, at least in the study of the motor system, is that the physics of the environment comes “for free”. The concept of embodiment further support this view or, at least, it is thought to be necessary to situate the “brain” into a real physical body (Pfeifer & Scheier, 1997). For instance, an influential model, derived from biology, is the so called “equilibrium point model” (EP), which has direct applicability to robotics (Gomi & Kawato, 1997; Mussa-Ivaldi & Giszter, 1992). Williamson (Williamson, 1996) applied this approach to control the motion of the arms of COG – the humanoid robot being built at MIT. Metta and coworkers used a similar biologically inspired approach to control the movements of the LiraLab Babybot (Metta, Sandini, & Konczak, 1999).

As mentioned above, the study of orienting behavior and ocular movements attracted considerable attention. Many implementations have been proposed, using visual, acoustic, and inertial sensory systems (Berthouze, Bakker, & Kuniyoshi, 1996; Crowley, Bobet, & Mesrabi, 1992; Panerai, Metta, & Sandini, 2000b). For example, Panerai et al. (Panerai & Sandini, 1998) used inertial sensors to simulate the vestibular organs, and employed both visual and inertial information inside the control loop of a binocular robot head, demonstrating superior performance compared to a purely vision based controller.

On the other hand, beside all these efforts, few researchers addressed the problem of adaptive behavior from a developmental point of view (Kuniyoshi & Cheng, 1999; Metta et al., 1999; Sandini, Metta, & Konczak, 1997). In our view in spite of such advances, the systems implemented are still far from achieving human-like performance levels and task flexibility. More importantly, the integration of different behaviors, such as manipulation and gaze control, proved to be more difficult than expected. This difficulty arises, at least in part, from the approach followed to construct complex systems: to make the problem more tractable, sensori-motor coordination is broken down into a set of sub-problems defined by a specific sensory modality

(e.g. vision, audition, touch, etc) or specific motor skills (e.g. manipulation, gaze control, navigation).

A different solution is used in humans and many other vertebrates, where flexible and efficient levels of performance are achieved through the simultaneous development of sensory, motor, and cognitive abilities. This process is not simply caused by the maturation of single components or by learning a progressively more sophisticated set of skills. Instead, it is marked, particularly in the very early stages, by a sequence of changes of the neural circuitry and by a strategic exploitation of the environment with the limited set of motor skills that are present at each developmental stage.

We conjecture that the neural arrangement and, in particular, the temporal evolution of the control structure might constrain learning thus potentially simplifying it. All subparts are harmoniously integrated in adulthood, in spite of the fact that during development they become active at different stages. On the other hand, simple integrative behaviors are already present and functional at birth (e.g. some of the early reflexes), being the foundation of further learning.

The remainder of the paper is organized as follow: section 2 highlights some issues concerning biological development. Sections 3 and 4 describe the experimental setup and the robot's adaptive control structure developed within the proposed framework. Finally, section 5 draws the conclusions of the study.

2. Development, constraints and motor control

This section outlines some relevant developmental issues related to motor control, in particular we focus on the acquisition of eye-head-arm coordination. Though this is not meant to be an exhaustive survey of the biological literature, it deals with the main aspects of development, and eventually we shall make the point that development is not the same as learning. From the agent perspective, the fact that different control pathways develop following distinct time courses might pose constraints on what can be learned and how hard is to learn appropriately.

For the sake of this argumentation let us consider two different subtasks: gazing and reaching. Gazing requires the coordination of the redundant head/eye degrees of freedom, and the transformation of the visual information into the appropriate motor commands. In the context of eye-head coordination, many authors suggested the hypothesis that newborn's motor acts are controlled only by sub-cortical structures, with cortical control taking over at about two to three months (Atkinson, 1998). Other models consider different visual functions (e.g. "what and where" streams (Goodale, 1989)) as developing with a slightly different timing. The issue of differential timing in development can be seen in the context of learning as reducing the

exploration space¹ by constraining it. In fact, sub-cortical structures, active since birth, could actually guide/constrain learning in cortical areas.

There is further evidence that newborns plan saccades in a retino-centric coordinate frame. This has been taken sometimes as the hallmark of the Superior Colliculus activity (Gilmore & Johnson, 1998). Von Hofsten showed also that the evolution of gaze control shifts from a stage, where head motion is used a little, to successive stages where neck and eye movements are combined effectively (Von Hofsten & Rosander, 1997). The reason for this pattern of development might be twofold: first, neck and eyes form a kinematically redundant system so that, coupling (i.e. reducing the number of independent degrees of freedom) the controllers could be an effective learning strategy. Second, coordinate motion of both head and eyes requires proprioceptive (or efferent copied) information, which is thought to be unreliable at birth.

Concerning reaching, at birth, a human infant can neither reach nor grasp. This should not be surprising if we consider that many of the structures devoted to motor control complete development only postnatally (Leary, 1992; Quartz & Sejnowski, 1997). Furthermore, newborns have a very limited knowledge about the physical makeup of their bodies (i.e. in terms of masses, link lengths, etc), so that their first "goal directed" movements appear ataxic. It is thought that the control structure at birth is mainly reflex driven. That is, infants possess only a limited repertoire of reflexes such as grasping, sucking, swimming, etc. In spite of this simple control arrangement, babies as young as few weeks of age, try to extend their arms towards visual targets.

Reaching behavior appears to become consistent at about 4 months of age (Konczak, Borutta, Topka, & Dichgans, 1995; Thelen et al., 1993). By this age many cortical structure had the chance to develop, thus reaching starts to be voluntarily triggered rather than purely reflex based. In any case, the presence of reflexes in early ontogenesis is perhaps both useful and necessary. First, at the very beginning the baby can start interacting with the environment, and he/she can try to discover those regularities in the sensori-motor data that are the basis of learning. Second, as mentioned above, the exploration space might be reduced. In practice, the presence of the early reflexes can serve as a sort of bootstrap procedure for subsequent learning stages.

Some questions arise: how are these considerations useful for building complex artificial systems?

The first, and perhaps, the major observation relates to various results on the theoretical aspects of learning. It has been recognized that learning from examples is an ill posed problem (Vapnik, 1998). Every learner faces the so-called "theoretical pressures", which require balancing competing needs in order for learning to be feasible. Recently, a few

¹ The part of the state space the agent needs to explore in order to solve a given task. It might or might not correspond to the plant state space depending on the task and learning algorithm considered.

researchers described and formalized these problems (Carpenter & Grossberg, 1986; Sutton & Barto, 1998; Vapnik, 1998).

It is worth stressing that the context we are interested in is that of a real physical agent interacting with the environment (i.e. the training set is collected on-line). This is a major constraint for biological as well as artificial systems. In this situation, 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 “state²” space is worth exploring in order to solve a particular task. As a matter of fact, the size of the state space might consist of hundreds 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 (Schaal & Atkeson, 1998). 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), to unnecessary neural branching (such as in the neuro-muscular junction), and by using random behavior actively (latencies on saccade generation). This role of noise during learning resembles the usage, in system theory, of broadband (e.g. white noise) input signal for system identification purposes. In other words, the system has to be “excited enough” in order for identification to be feasible.

Besides, other researchers provided evidences for the existence of a strong goal-directed behavioral component, even in newborns (Streri, 1993). The fact that the behavior is goal-directed can speed up the acquisition of the appropriate controller, i.e. just imagine a task with a single target state, in this case a goal-directed agent might solve the problem for only a relatively small neighborhood of the target. On the contrary, a random explorative search has to test all possible states, unless some *a priori* knowledge is inserted into the system. Lacking of any constraint, the random explorer needs to visit all possible states prior to any actual control; otherwise, a possibly useful part of the state space might remain unexplored.

Furthermore, the cooperation of many control loops developing with different time spans can help in dealing with the same problem. Roughly speaking, each control loop generates a bias for subsystems that develop later on.

² 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. The discussion presented here effortlessly adapts to all of them.

In the context of computational motor control, one notable example of such a schema is the feedback-error learning model (Kawato, Furukawa, & Suzuki, 1987). 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 be observed in the brain. An example of this process is the already mentioned *cortical take-over*, where cortical areas develop on top of sub-cortical structures. Furthermore, the delay and bandwidth involved in the various structures can be different thus providing the basis for both faster reactions (reflex-like) and accurate control (consider, for instance, the visuo-vestibular integration).

It is worth stressing that the exploration-exploitation tradeoff is closely related to the well-known engineering problem called “the curse of the dimensionality” (Bellman, 1956). 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. Moreover, the bigger the space the sparser the data, and indeed it has been shown that there is a limit on the mathematical consistency of topographic mapping based on neighborhood relationships (roughly 20 dimensions) (Scott, 1992), and the latter has been shown as well to be an overall organizing principle of the brain (Hubel & Wiesel, 1977). Hence it might be important to limit the size and dimensionality of the state space, which would allow topographic mapping to be carried on reliably.

3. The experimental setup

Following the proposed framework and, in order to test hypotheses about the development of visuo-motor coordination, we developed the robot architecture shown in Figure 1.




Figure 1 The experimental setup. Dashed lines on the rightmost picture represent the axes of rotation.

The experimental setup 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 “Babybot³” is equipped with two space-variant cameras (Sandini & Tagliasco, 1980), an inertial sensor simulating the vestibular system (Panerai & Sandini, 1998), 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. The software adheres to DCOM, a standard, which allows running objects among the various machines. The reference task, for this discussion, is the coordination of eye-head-arm movements, with the aim of gazing and reaching for visually identified objects in extrapersonal space.

4. Control structure and learning schema

This section outlines Babybot’s adaptive control structure, which has been modeled after the developmental framework introduced in the previous sections. As mentioned before, the idea is to have a system, which at “birth” uses only simple controllers and successively “grows” by employing more sophisticated modules (e.g. an inverse model, maps, etc). Figure 2 below shows this sequence of successive developmental events. The overall system initial state is thus characterized by a very limited number of free parameters, which can be easily estimated on-line.

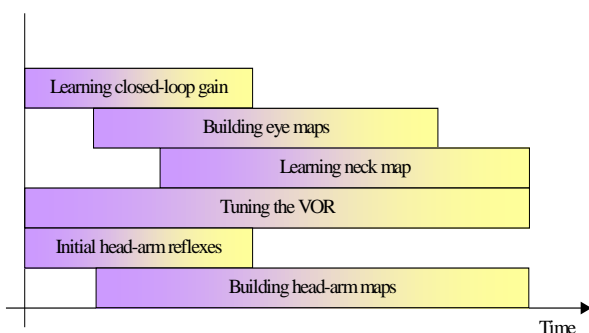


Figure 2 The developmental stages. The diagram above roughly shows the interleaving of the developmental stages; abscissa represents time. The first step is the acquisition of the closed loop gains; reflex-like modules control the arm sub-system. After a while, learning of the saccade control begins. Eventually the eye-head coordination is acquired together with a more effective head-arm coordination map.

The concurrent controllers then learn on the basis of how the simpler loops are behaving. State space exploration is

driven by additive noise, which simulates defective command generation (muscle control). Exploration and exploitation processes are carried out in parallel; in practice, the robot performs system identification and control at the same time. After eye movements reach a reasonable level of performance, the robot starts moving more degrees of freedom (i.e. the neck). Even at this level we stressed the biological parallelism by adopting a schema, which closely resembles the solution found in many species, including humans. We equipped the robot head with an inertial sensor, simulating the vestibular apparatus, which can sense the rate of rotation of the head respect to a vertical axis. Beside the extension of the working bandwidth as shown in (Panerai et al., 2000b), the use of the inertial information also simplifies command generation. Concerning vision, the system employs space-variant images, which resemble the distribution of the photoreceptors in the human retinas. Image resolution is kept at minimum (images are 64×32 pixels), with the general idea of starting the robot with limited sensory capabilities. Without entering into the details of the actual visual processing, we can say that the robot is able to compute the presence of a target and extract its position and velocity relative to the cameras (in retinal coordinates). The sensory abilities are complemented by proprioception, which is provided by optical encoders (one for each mechanical joint). The control variables of the robot are the joint velocities. An appropriate low-level closed loop controller (usually a PID for each joint) generates the motor driving torques. An exception is the arm control schema, which is based on the equilibrium point hypothesis (EP). As in biological systems, the position and the impedance characteristics of the robot arm are the result of the interaction of the stiffness-controlled spring-like simulated actuators.

5. Some experimental results

5.1 Eye movements

As starting point, consider the problem of moving the eyes toward the target. The simplest solution might use positional information to drive a negative feedback loop. The fundamental problem in such a strategy is that of converting the target position, which is expressed in retino-centric coordinates into motor commands. The latter are expressed with respect to a motor (or joint) coordinates system. If this is the case the error is described by:

$$\mathbf{e} = \mathbf{C} \cdot \mathbf{s}(t) \quad (1)$$

where \mathbf{e} is the position error expressed in motor coordinates, $\mathbf{s}(t)$ the retinal error and \mathbf{C} a coordinate conversion matrix. The matrix \mathbf{C} must be designed in order to stabilize the closed-loop system. In this case the generated motor command is:

$$\dot{\mathbf{q}} = -\lambda \cdot \mathbf{e} \quad \lambda > 0 \quad (2)$$

³ Babybot = Baby + Robot

with $\dot{\mathbf{q}}$ the control variable and λ a positive constant gain. \mathbf{C} can be determined imposing an exponential decay rule of the error. A possible choice for \mathbf{C} is:

$$\mathbf{C} = \left(\frac{\partial \mathbf{s}}{\partial \mathbf{q}} \right)^{-1} \quad (3)$$

Learning of the matrix \mathbf{C} is carried out by acquiring discrete samples of the variation of the retinal error $\Delta \mathbf{s}$ due to a variation of the joint variable $\Delta \mathbf{q}$. A least-square approach is used to compute the components of \mathbf{C} . Besides, it is also important to define how to get the samples. Performing random movements according to the following strategy can easily generate them:

$$\dot{\mathbf{q}} = -\lambda \mathbf{C} \mathbf{s}(t) + \boldsymbol{\eta}(\mathbf{0}, \sigma) \quad (4)$$

The first term is the closed loop formula described above; the second term $\boldsymbol{\eta}$ represents a zero mean uniform noise with standard deviation σ . It is worth noting that, at the beginning noise dominates ($\mathbf{C}=0$, $\sigma \neq 0$), while as learning proceeds the closed loop term takes over the control of movements. An example of the robot behavior in this case is shown in Figure 3. Although the closed loop approach described above is effective, it does not mean it is also efficient. In fact, a position based feedback controller would always lag behind a moving target. Moreover, if the spotted target lies in the periphery of the visual field, the robot would take some control steps to move the cameras towards it. For the same reasons, probably saccade movements evolved in those animals with an efficient visuo-motor apparatus. In our artificial system, the requirement for generating fast movements, emulating saccades, is to know exactly (as precisely as possible) the transformation between retinal error \mathbf{s} and the corresponding motor command $\Delta \mathbf{q}$, that is:

$$\Delta \mathbf{q} = \hat{\mathbf{f}}(\mathbf{s}) \quad (5)$$

Under the hypothesis of a stationary target and a closed

loop control in place as described above, the gathering of training pairs (each of them has the form $(\Delta \mathbf{q}, \mathbf{s})$) is much simplified. The retinal error \mathbf{s} is acquired at the beginning of the movement, while the required motor command can be measured when the retinal error is zeroed. An explicit exploration is not actually needed because the closed loop system is already generating proper commands (directed at reducing the retinal error – and eventually zeroing it).

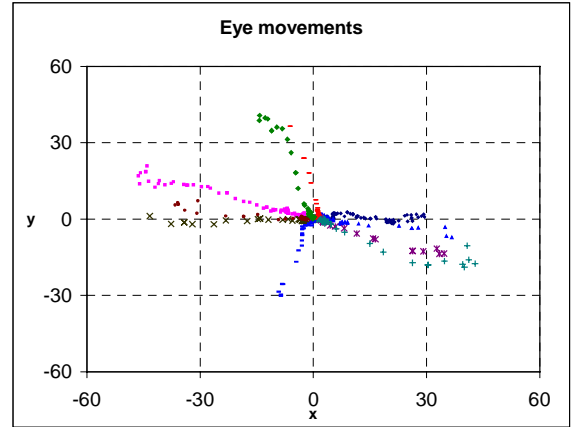


Figure 3 Eye/target trajectories. Abscissa and ordinates represent the image plane, and different graphical signs mark trajectories (the target position at each control step – 40ms period). As expected all the trajectories are converging to the fovea. It is worth noting that in this case the movements are still quite slow, and the number of “points per trajectory” is high.

In order to relax the stationary target hypothesis, it is possible to acquire a new training example, as soon as some control cycles have been performed. In this case motion of the target would influence only a little the measure of the motor command. Furthermore, if we assume that targets generally move with equal probability in each possible direction, the mean of the measure error would be zero. The

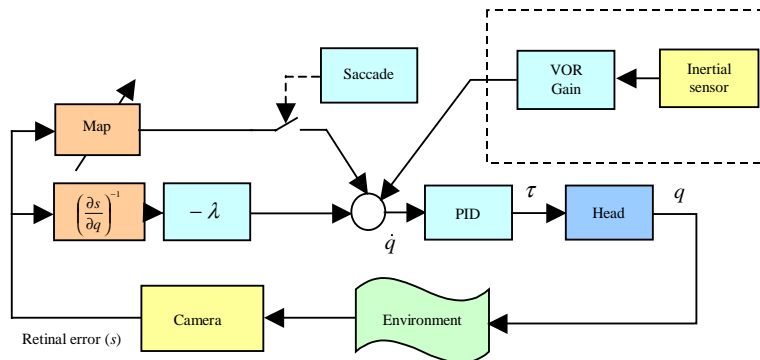


Figure 4 The eye control schema. It consists of a closed loop and a feed-forward secondary loop. The loop using the inverse Jacobian is derived from a classical visual servoing approach. The secondary loop consists of an inverse model (indicated by “Map”). It is activated whenever necessary – retinal error greater than a threshold – and generates a fast motion of the eyes in order to foveate the target. The goal of the network is to learn the inverse model. λ is a positive constant gain. It is tuned in order to obtain stability of the closed loop system. The input to the robot controller is a velocity command. A low level controller (PID) generates the motors’ driving torques τ . The block identified by “Saccade” is the governing logic (i.e. the threshold mechanism issuing the “start” signal for the fast motion).

output of the proposed mapping is then used to generate saccades. A proper velocity command is generated by converting (deriving) its output Δs (i.e. $\dot{\mathbf{q}} \equiv \Delta \mathbf{q} / \Delta T$). Saccade initiation is controlled by another module, which issues a saccade command each time either the retinal error is greater than a fixed threshold (catch-up saccade) or a moving target is detected but a target is not currently being tracked. The overall loop controlling eye motion is shown in Figure 4.

5.2 Neck movements

Following this hypothetical line of developmental events, there is a stage when the neck comes into play in the orienting behavior. It is at the same moment when the proprioception becomes reliable and consistent. Gilmore et al. (Gilmore & Johnson, 1998) suggested that a shift could be also observed in the coordinates system governing eye/head movements (from retino-centric to head-centric). Whatever motion strategy, it has to deal with the “degrees of freedom” problem. The head-eye system is kinematically redundant; consequently, a further constraint has to be employed. We required the head system to achieve a symmetrical vergence configuration: that is, the neck should move in order to be roughly heading toward the target. Each degree of freedom was allowed to respect its physical constraints: roughly speaking, the eyes move faster than the neck because of the different inertias. Limiting the accelerations appropriately also enforced this behavior.

Concerning neck motion control, the proposed two-loop system can deal with the situation. A PID controller governs the closed loop module as before. Its goal is that of zeroing the difference between the two eye angles:

$$\dot{q}_1 = PD(q_4 - q_5) \quad (6)$$

Also in this case an “inverse model” map can improve performances. It maps the predicted eye positions to the proper neck motion. The control diagram is sketched in Figure 5. Formally:

$$\Delta q_1 = \hat{f}({}^{pred}q_4, {}^{pred}q_5) \quad (7)$$

where Δq_1 is the neck motion command, ${}^{pred}q_{(4,5)}$ the predicted eye positions.

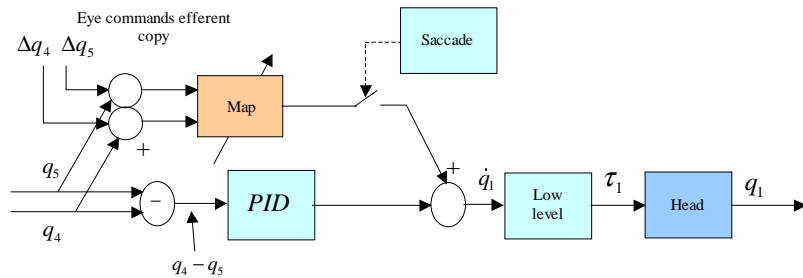


Figure 5 The neck control schema. It employs the same working principle of the eye controller, besides there are a few important differences. First, there is not direct visual feedback, on the contrary, eyes position drive the movement of the head – the PID controller has to move the head in order to maintain as much as possible a symmetric vergence configuration. Second, the saccade-like movement is based on the prediction of the eye positions at the end of the saccade – i.e. efferent copy.

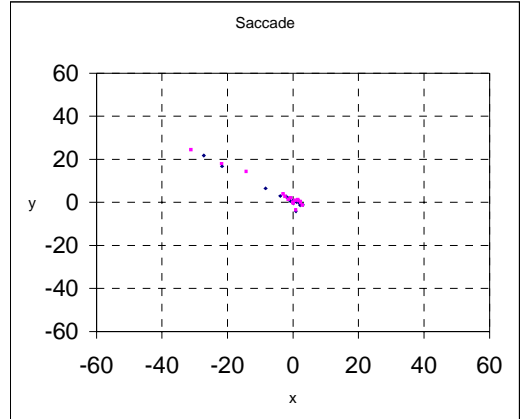


Figure 6 An exemplar trajectory after learning. Note as the first three steps are enough to reduce the retinal error to less than five pixels, afterwards the target remains in the fovea. In this case, it is clear that, the target appeared on the left side of the robot.

What does “predicted eye positions” mean? They are the current eye positions updated by the saccadic eye motion. The eye movement can be recovered using the eye maps, even before any actual motion has been started. In equation form:

$${}^{pred}q_{(4,5)} = q_{(4,5)} + {}^{saccade}\Delta q_{(4,5)} \quad (8)$$

It is worth noting that, although the proposed neck motion strategy can work, there are still some performance related issues to consider; in fact, the motion of the head is likely to disturb the eye movement process (either saccade or tracking). This is especially true if the neck is performing relatively fast movements. In that case, by applying the described control strategy, the robot would likely overshoot the target. The overshoot is eventually compensated by the visual feedback. However, vision is slow compared to the kind of motion we are dealing with (i.e. a saccade might last from 40 to 100ms). Thus the resulting motion, though convergent (i.e. stable), would have poor performance (oscillations).

We can observe that many species developed dedicated sensory systems devoted to measuring the motion of the head/body in space. We equipped our system with a similar

device: an inertial sensor (a solid state gyroscope), as described in (Panerai & Sandini, 1998). Our artificial vestibular sensor can measure neck angular velocity and, in the context of head-eye coordination, it comes into play by counter-rotating the eyes whenever the head moves (VOR). The VOR loop is sketched in Figure 4 together with the saccade control schema. In order to illustrate the principle only the simple constant-gain VOR case is considered. In practice, even this simple schema improves the robot performance considerably. More sophisticated strategies are considered in details in (Panerai et al., 2000b). A learning based framework for the VOR is presented in (Panerai, Metta, & Sandini, 2000a).

5.3 Reaching

Concerning reaching the approach we shall follow here is directly based on motor primitives, representing multi-joint synergies (e.g. arm extension). In this case, a single command may produce complex multi-joint coordinated movements without the voluntary control of each individual degree of freedom (DOF). In order for this approach to be feasible and effective, the crucial points are how to represent the motor primitives and the mechanisms of sensori-motor mapping. As far as the coding of motor primitives is concerned, one possible procedure is the so-called force fields approach originally proposed by Mussa-Ivaldi and Bizzi (Mussa-Ivaldi & Giszter, 1992), (Mussa-Ivaldi, Giszter, & Bizzi, 1993). It is not our intention to present the whole mathematical details here. Anyway, for the sake of the argumentation, according to the force field theory, the action of reaching a point in space can be described by a force vector field converging to an Equilibrium Point (EP). The EP can be thought of as the point towards which the end-point of the limb is moving at each instant of time, and a limb trajectory can be represented by a sequence of EPs.

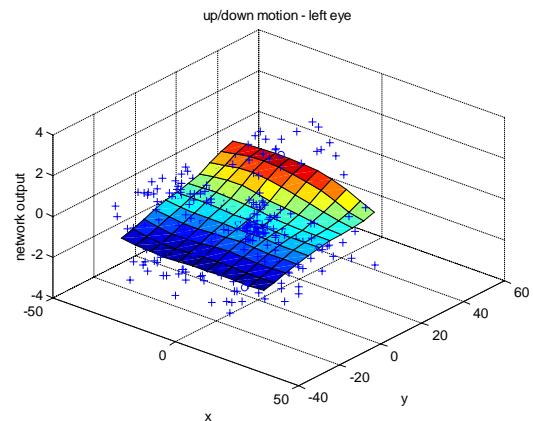


Figure 7 A saccade map implemented by a growing neural gas network. The “+” sign represents the most recent 300 samples of the training set. The input space (x,y) is the image plane (in pixels), the output (the height of the surface plot) is the angle required to foveate a target appearing in the corresponding (x,y) image position.

The trajectory in space of the EP does not correspond to the actual trajectory of the arm and is, therefore, called “virtual trajectory”. Of course, there is nothing magic about the origin of this “force field” formalism. In practice, each force field is simply the result of the action of muscles and the EP is simply the intersection of muscles’ torque-length characteristics.

The mechanism is well suited to implement the kind of motor reflexes present at birth. Each reflex can be represented by a force field; this is in turn obtained by activating a synergy of simulated muscles. Of course, the robot’s actuators are not muscles, but they allow being torque controlled – i.e. by programming the current that flows into the motors, it is possible to simulate whatever characteristic in software, although with some limitations (Metta et al., 1999).

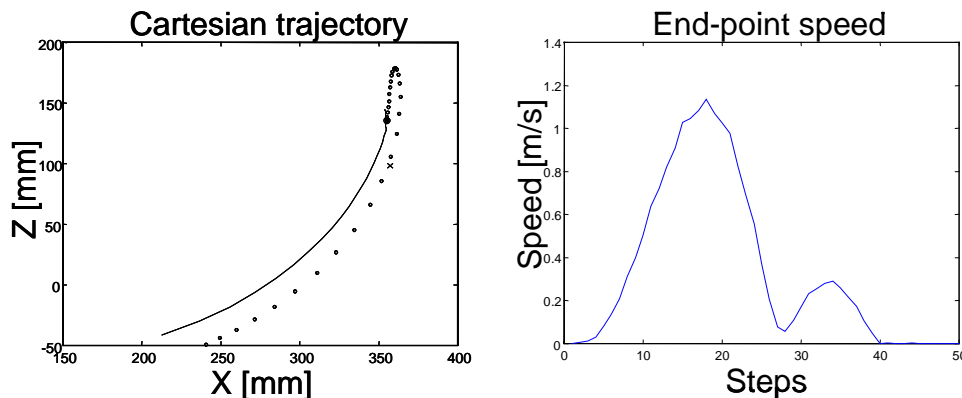


Figure 8 Trajectory showed a bell-shaped velocity profile. On the left, trajectory of the arm end-point: abscissa and ordinate represent the plane where the arm motion was constrained. The dashed line is the actual trajectory sampled at 40Hz, the solid line is the “virtual trajectory”. In spite of the virtual trajectory that moves directly to the target, the actual motion showed an overshoot. On the right, the hand speed has a bell-shaped profile. Note the two bumps corresponding to the first large “transport phase” and a second corrective movement. Time is expressed in control cycles (25ms), and speed in meters per second.

With regard to the mapping, the solution we propose here is based on the use of a direct transform between the eye-head motor plant and the arm motor plant. One premise we make is that the position of the fixation point coincides – at least at some stage of the control process – with the object to be reached. In other words, the reaching for an object starts by looking at it. Under this assumption, the fixation point can be seen as the “end-effector” of the eye-head system. The positions of the head with respect to the torso, and that of the eyes with respect to the head, uniquely determine its position in space relative to the shoulder. That is, the position of the fixation point in space is uniquely coded by motor state variables – i.e. the position of the head-eye system joints. Consequently, at least in principle, the arm force fields can be obtained by a “transformation” of these “plant” variables. We will call this approach “motor-motor coordination”, because the coordinated action is obtained by mapping motor plant variables into motor plant variables. A complete description of the control structure in the 2D case can be found in (Metta et al., 1999). For the purpose of this paper it is important to bear in mind that the learning component has to map the gaze direction (expressed in particular by the vergence, version, and tilt angles) into the arm control values (the position of the EP in space).

The important characteristics of the whole procedure are: i) the initialization of the map allows embedding a sort of “infant” reflexes into the system. They are chosen in order to facilitate learning (by initializing the head-arm map so that the arm is extended roughly in the direction the head is turned – this has the role of maintaining the arm within the field of view). ii) Trajectory generation is not explicitly addressed: a simple linear interpolation between control values has been employed.

The motor-motor mapping, at least initially, does not necessarily bring the end-effector near the fixation point (it will bring the arm as close as possible to the target on the basis of what has been learned so far). 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. In other words, the role of the visual target appearing in the environment has the only function of initiating the arm motion, while the learning process is based on the act of looking at the end-effector. 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. Furthermore, it is important to note that if the procedure were noise free, the motion of the arm towards the target would always bring the end-effector in the same final position and the system would not be able to learn.

An exemplar trajectory after about 130 trials is shown in Figure 8 together with the end-point speed profile. On the other hand, the trajectory itself is not straight, and shows remarkable overshoots, perhaps because of the lack of dynamic compensation. This is to say, as pointed out by some authors (Gomi & Kawato, 1997), that the CNS might need to take into account dynamics when moving a limb, because under such low-stiffness control self-generated forces might contribute substantially to the total torque at joint level. In our model, dynamics is not explicitly considered, thus it is not surprising that trajectories are neither straight nor precise on the target. Anyway, the “force field” approach allows considering even more basis fields, and we may imagine that some of them are built by a learning procedure. In other words, the robot can acquire more basis fields, which compensate for the self-generated forces, or other external disturbances – such as gravity.

As for the 2D case we recorded the trajectory of the arm end-point and the gaze direction. From the latter it is possible to recover the position of the fixation point in space. The two quantities have been plotted during a reaching movement in order to illustrate qualitatively the behavior of the robot (see Figure 9). The simple wire-frame model represents the robot. Small circles indicate joints; solid lines are the links. Concerning the fixation point, two different marks can be distinguished: the “+” marks




Figure 9 Reaching trajectories: two representative trajectories of the robot behavior. Two different marks can be distinguished: the crosses mark smooth pursuit control; the small squares are related to saccadic control. Note as in the leftmost picture the end-point and fixation position almost coincide.

represent the time instants when tracking was of smooth pursuit type, the small boxes are related to saccadic control.

6. Conclusions

In conclusion, we presented a proposal for a general approach aimed at the design and comprehension of complex adaptive 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 the 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”. Other aspects have been discussed, for instance, the presence of “innate” behaviors, which later disappear. This is an open question: do they really disappear? It might very well be that those initial modules get embedded into more complex control structures. In this sense, voluntary control can be seen as learning to combine the initial reflexes, in order to solve a particular task.

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 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 period” adaptation can effectively take place without disturbing too 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 very 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”.

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).

Of course, most of these conjectures need to be verified, from either the theoretical side (e.g. learning theory) or the biological point of view (e.g. by designing new experiments, for example to determine how gazing correlates with reaching). In this light, the most sensible prosecution of this work would be that of investigating all these open questions, in order to formalize a theory about developing systems. On the other hand, many hypotheses arose, which might be worth testing on “real brains”.

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.
- Bajcsy, R. K. (1985). *Active Perception vs. Passive Perception*. Paper presented at the Third IEEE Workshop on Computer Vision: Representation and Control, Bellaire (MI).
- Ballard, D. H., & Brown, C. M. (1992). Principles of Animate Vision. *Computer Vision Graphics and Image Processing*, 56(1), 3-21.
- Beer, R. D., Chiel, H. J., Quinn, R. D., & Ritzmann, R. E. (1998). Biorobotic approaches to the study of motor systems. *Current Opinion in Neurobiology*, 8(6), 777-782.
- Bellman, R. E. (1956). *Dynamic Programming*. Princeton: Princeton University Press.
- Berthouze, L., Bakker, P., & Kuniyoshi, Y. (1996). *Learning of Oculo-Motor Control: a Prelude to Robotic Imitation*. Paper presented at the IEEE-RSJ IROS 1996.
- Brooks, R. (1996). *Behavior-Based Humanoid Robotics*. Paper presented at the IEEE/R SJ IROS'96.
- Capurro, C., Panerai, F., Grosso, E., & Sandini, G. (1993, July). *A Binocular Active Vision System Using Space Variant Sensors: Exploiting Autonomous Behaviors for Space Application*. Paper presented at the International Conference on Digital Signal Processing, Nicosia, Cyprus.
- Carpenter, G., & Grossberg, S. (1986). *Adaptive Resonance Theory: stable self-organization of neural recognition codes in response to arbitrary list of input patterns*. Paper presented at the Eighth Annual Conference on the Cognitive Science.
- Crowley, J. L., Bobet, P., & Mesrabi, M. (1992, May). *Gaze Control with a Binocular Camera Head*. Paper presented at the European Conference of Computer Vision, Santa Margherita Ligure, Italy.
- Gilmore, R. O., & Johnson, M. H. (1998). Learning What is Where: Oculomotor Contributions to the Development of Spatial Cognition. In F. Simion &

- G. Butterworth (Eds.), *The Development of Sensory, Motor and Cognitive Capacities in Early Infancy* (pp. 25-47). Hove, East Sussex: Psychology Press Ltd.
- Gomi, H., & Kawato, M. (1997). Human arm stiffness and equilibrium-point trajectory during multi-joint movement. *Biological Cybernetics*, 76, 163-171.
- Goodale, M. A. (1989). Modularity in Visuomotor Control: from Input to Output. In Pylyshyn (Ed.), *Computational Processes in human vision: an interdisciplinary perspective* (pp. 262-285). Norwood, NJ: Ablex Publishing.
- Hubel, D. H., & Wiesel, T. N. (1977). Functional architecture of macaque monkey cortex. *Proceedings of the Royal Society of London*(198), 1-59.
- Kawato, M., Furukawa, K., & Suzuki, R. (1987). A hierarchical neural network model for control and learning of voluntary movement. *Biological Cybernetics*(57), 169-185.
- Konczak, J., Borutta, M., Topka, H., & Dichgans, J. (1995). Development of goal-directed reaching in infants: Hand trajectory formation and joint force control. *Experimental Brain Research*, 106, 156-168.
- Kuniyoshi, Y., & Cheng, G. (1999). Complex Continuous Meaningful Humanoid Interaction: A Multi Sensory-Cue Based Approach : Personal Communication.
- Leary, D. D. M. O. (1992). Development of connectional diversity and specificity in the mammalian brain by the pruning of collateral projections. *Current Opinion in Neurobiology*(2), 70-77.
- Metta, G., Sandini, G., & Konczak, J. (1999). A Developmental Approach to visually-guided reaching in artificial systems. *Neural Networks*, 12(10), 1413-1427.
- Mussa-Ivaldi, F. A., & Giszter, S. F. (1992). Vector field approximation: a computational paradigm for motor control and learning. *Biological Cybernetics*, 67, 491-500.
- Mussa-Ivaldi, F. A., Giszter, S. F., & Bizzi, E. (1993). Convergent Force Fields Organized in the Frog's Spinal Cord. *The Journal of Neuroscience*, 13(2), 467-491.
- Panerai, F., Metta, G., & Sandini, G. (2000a, 26-28 April 2000). *Learning VOR-like Stabilization Reflexes in Robots*. Paper presented at the 8th European Symposium on Artificial Neural Network, Bruges, Belgium.
- Panerai, F., Metta, G., & Sandini, G. (2000b). Visuo-inertial Stabilization in Space-variant Binocular Systems. *Robotics and Autonomous Systems*, 30(1-2), 195-214.
- Panerai, F., & Sandini, G. (1998). Oculo-Motor Stabilization Reflexes: Integration of Inertial and Visual Information. *Neural Networks*, 11, 1191-1204.
- Pfeifer, R., & Scheier, C. (1997). Sensory-motor coordination: The metaphor and beyond. *Robotics and Autonomous Systems*, 20, 157-178.
- Quartz, S. R., & Sejnowski, T. J. (1997). The neural basis of cognitive development: A constructivist manifesto. *Behavioral and Brain Sciences*(20), 537-596.
- Rucci, M., Wray, J., Tononi, G., & Edelman, G. M. (1997, April, 20-25). *A robotic system emulating the adaptive orienting behavior of the barn owl*. Paper presented at the IEEE International Conference on Robotics and Automation, Albuquerque, USA.
- Sandini, G. (1997, April.). *Artificial Systems and Neuroscience*. Paper presented at the Proc. of the Otto and Martha Fischbeck Seminar on Active Vision.
- Sandini, G., Metta, G., & Konczak, J. (1997, November). *Human Sensori-motor Development and Artificial Systems*. Paper presented at the AIR&IHAS 97.
- Sandini, G., & Tagliasco, V. (1980). An Anthropomorphic Retina-like Structure for Scene Analysis. *CVGIP*, 14(3), 365-372.
- Schaal, S., & Atkeson, C. G. (1998). Constructive Incremental Learning from Only Local Information. *Neural Computation*(10), 2047-2084.
- Scott, D. W. (1992). *Multivariate density estimation*. New York: Wiley.
- Streri, A. (1993). *Seeing, Reaching, Touching - The Relations between Vision and Touch in Infancy*: MIT Press.
- Sutton, R. S., & Barto, A. (1998). *Reinforcement Learning: an Introduction*. Cambridge: MIT Press.
- Thelen, E., Corbetta, D., Kamm, K., JP, J. P. S., Schneider, K., & Zernicke, R. F. (1993). The transition to reaching: mapping intention and intrinsic dynamics. *Child Development*, 64(1058-1098).
- Vapnik, V. N. (1998). *Statistical Learning Theory*. New York: Wiley.
- Voegtlin, T., & Verschure, F. M. J. (1999). What Can Robots Tell Us About Brains? A Synthetic Approach Towards the Study of Learning and Problem Solving. *Reviews in the Neurosciences*, 10, 291-310.
- Von Hofsten, C., & Rosander, K. (1997). Development of Smooth Pursuit Tracking in Young Infants. *Vision Research*, 37(13), 1799-1810.
- Williamson, M. M. (1996). *Postural primitives: interactive behavior for a humanoid robot arm*. Paper presented at the From Animals to Animats: 4th International Conference on Simulation of Adaptive Behavior, Cambridge, Massachusetts USA.