



PERGAMON

Neural Networks 12 (1999) 1413–1427

Neural  
Networks

www.elsevier.com/locate/neunet

Contributed article

# A developmental approach to visually-guided reaching in artificial systems

G. Metta<sup>a</sup>, G. Sandini<sup>a,\*</sup>, J. Konczak<sup>b</sup>

<sup>a</sup>*Lira Lab, DIST-University of Genova, Via Opera Pia, 13, 16143 Genova, Italy*

<sup>b</sup>*Department of Psychology, University of Düsseldorf, Düsseldorf, Germany*

Received 5 December 1997; accepted 12 July 1999

## Abstract

The aim of the present paper is to propose that the adoption of a framework of biological development is suitable for the construction of artificial systems. We will argue that a developmental approach does provide unique insights on how to build highly complex and adaptable artificial systems. To illustrate our point, we will use as an example the acquisition of goal-directed reaching. In the initial part of the paper we will outline (a) how mechanisms of biological development can be adapted to the artificial world, and (b) how this artificial development differs from traditional engineering approaches to robotics. An experiment performed on an artificial system initially controlled by motor reflexes is presented, showing the acquisition of visuo-motor maps for ballistic control of reaching without explicit knowledge of the system's kinematic parameters. © 1999 Elsevier Science Ltd. All rights reserved.

*Keywords:* Development; Learning; Motor control; Human infants; Robotics; Artificial systems

## 1. Introduction

Research activity linking studies on artificial systems to 'brain sciences' is certainly not new. Besides the studies on artificial neural networks, substantial efforts are devoted worldwide to building 'physical models' of segments of biological systems with the aim of suggesting novel solutions to robotics or processing problems and to advance our understanding of human brain functions (Brooks, 1996; Sandini, 1997).

On the other hand, the study of sensori-motor coordination in artificial systems has focused mainly on analyzing and implementing skill levels comparable to those of adult humans. For example, the control of robot heads and visually guided manipulation tasks were studied with reference to psychophysical performance data of adult humans and animals (Aloimonos, Weiss & Bandyopadhyay, 1988; Bajcsy, 1985; Bajcsy & Tsikos, 1988; Ballard & Brown, 1992; Crowley, Bobet & Mesrabi, 1992; Capurro, Panerai, Grosso & Sandini, 1993, 1996; Capurro, Panerai & Sandini, 1997; Coombs & Brown, 1990; Gandolfo, Sandini & Tistarelli, 1991; Grosso, Metta, Oddera & Sandini, 1996).

In spite of the recent advances in this area, the systems

implemented are still far from achieving human-like performance levels and task flexibility. More importantly, even for successful implementations the integration of such skills as manipulation and gaze control proved to be more difficult than expected. In our view, 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, by a strategic exploitation of the environment with a limited set of motor skills that are present at each developmental stage, and finally, by the ability of biological systems to calibrate themselves in the presence of ongoing environmental and internal (e.g. anthropometric) changes.

In this view development is not a mere summation of discrete learning events. The degree of learning is determined by the learner's developmental state, which in turn

\* Corresponding author. Tel.: + 39-10-3532779; fax: + 39-10-3532154.

E-mail address: giulio@dist.unige.it (G. Sandini)

Nomenclature	
$\mathbf{q}$	Generalized coordinate vector (all joints are revolute) (rad)
$q_0$	Resting point (in a spring-like equation) (rad)
$\boldsymbol{\tau}$	Torque field generated by an actuator (Nm)
$\kappa$	Stiffness coefficient ( $\text{Kg m}^2/\text{s}^2$ )
$a$	Activation coefficient (for the actuators)
$C$	Activation coefficient (for the controllers)
$\mathbf{I}$	Connection matrix for the controller structure
$\mathbf{T}$	Torque field generated by a controller (Nm)
$\mathbf{x}$	End-point position (m)
$\Lambda$	Direct kinematics
$\mathbf{J}_A$	Manipulator Jacobian
$\mathbf{F}$	Force field for a controller (N)
$f$	Motor–motor map
$\hat{f}$	Approximation of the motor–motor map
$\mathbf{N}$	Zero-mean, uniform distributed noise

is the result of complex interactions of maturational and endogenous processes with the environment.

We have now a reasonable understanding of the developmental progression of reaching in infants. Yet, a mere description of a process of biological development does not shed any light to the issue of how physiological mechanisms of development are helpful for building complex artificial systems. A second step is necessary to bridge this gap. This step is to outline the control problems that have to be solved by human infants when trying to reach for objects in their immediate workspace. As we will see, the control problems for newborn humans and a robot may be expressed in a similar way in spite of the different ‘technologies’ used for the actuators and sensing elements. For the sake of our argument at this point, let us divide the human organism into two distinct elements—the central nervous system (CNS) as the ‘controller’, and the body as the ‘plant’ (we realize that this is an undue simplification).

The first question to ask in the context of motor control is, what physiological or movement parameters does the system actually have to control to achieve its goal of reaching for a target in extrapersonal space? To answer this question, consider that each limb segment of the human arm is moved by a set of actuators with spring-like properties. In order to control the arm in a coordinated manner any controller needs to have at least reasonable approximations of inertia, viscosity and stiffness parameters.

That is, a first step for a control system must be the identification of the plant’s parameters. A second step before goal-directed reaches are possible is the mapping of sensory maps onto available motor maps. That is, the system must be able to localize objects in extrapersonal space, and should have a knowledge where its limbs are positioned relative to the object. In a traditional learning paradigm, these two processes of calibration have to be

completed before the system can begin to work on control (Kalveram, 1991; Kuperstein, 1988).

Today the view of a parallel development of calibration and control processes seems widely accepted by researchers working on neural modeling of adaptive eye–hand coordination (Jordan, 1996; Kawato, Furukawa & Suzuki, 1987). Yet, most researchers model this process as a learning and not as a developmental operation (Jordan & Flash, 1994; Kuperstein, 1988). Implicit to such an approach of artificial eye–hand coordination is the premise that all behaviors of the system have to be learned. However, this assumption is not necessarily true for biological systems. One major difference between a biological and artificial system is that a biological system does not come as a ‘blank slate’. In a wide variety of different species one can observe stereotyped inborn movement sequences that are clearly unlearned. Ethologists have argued for a long time that many behaviors, especially those of lower animals, cannot be explained on the basis of sensori-motor learning alone (Eibl-Eibesfeld, 1970; Gould, 1982). Newborn human infants already possess a repertoire of coordinated movements. For example, they can perform a series of complex multi-joint bilateral movements (i.e. kicking) and have available a set of so-called primitive reflexes. These primitive reflexes resemble a set of complex movement patterns that are triggered by a sensory stimulus.

Yet, these motor primitives may also serve a second function. They help to build up a relationship between vision and proprioception. For example, during pre-reaching the presence of the Asymmetric Tonic Neck Reflex (ATNR) plays a crucial role in allowing babies to see their hand and in increasing visual fixation of the hands (White, Castle & Held, 1964; Bushnell, 1981).<sup>1</sup> Thus, this multi-muscle synergy coupling arm and head movements provides an effective means for linking visual and proprioceptive maps.

## 2. Developmental engineering

Some of the peculiarities of human development outlined in the previous section are particularly relevant from the engineering perspective of building a complex adaptive system.

The first and perhaps major observation relates to the fact that during development the infant is a self-contained, functional system with matched sensorial, motor and cognitive abilities. Motor reflexes and sensory-triggered motions are present at birth, exploiting the still limited sensory and motor abilities and allowing the infant to start some form of interaction with the environment and the acquisition of first sensori-motor experiences. During the initial months,

<sup>1</sup> The stimulus for the ATNR is the turning motion of the head, which subsequently will trigger a complex bilateral synergy. The infant’s arm will be extended to the side the infant is looking, effectively bringing the hand into the field of view. The contralateral arm is flexed as part of crossed extensor reflex spanning both homologous limbs.

some of the abilities are only temporarily present (for example, some of the early reflexes) and they disappear as soon as more mature skills develop. The control structure changes with age starting from an almost purely reflexive system at birth, passing through phases where basal muscular synergies are formed, towards a state where stable kinematic patterns are expressed—a sign that the redundant degrees of freedom of the neuromuscular system are then under voluntary control. From the engineering perspective this means a shift of emphasis from the ‘final product’ to a ‘process-oriented’ approach.

Another issue worth stressing is the role of the infant’s own body in development. In biological systems, development is very much dependent on the ability of the system to interact with the external world. Many early sensori-motor experiences are stimulated by the newborn’s own body motions which becomes an essential tool to establish a coupling between perception and action. On a neuronal level this organism–environment interaction is necessary to establish new connections, and to prune those that proved to be unfunctional (O’Leary, 1992). Self-generated motor commands elicit sensory feedback (like proprioceptive signals, motion in the visual field, tactile and acoustic stimuli) that not only give the newborn a motivation to repeat or to avoid a specific action, but also serve to adjust and refine the voluntary motor commands. The emergence of the distinction between what is and what is not controllable, is one of the first achievements of a system whose survival depends on its interaction with the outside world.

Based upon these observations we designed an experiment with a robot to demonstrate the utility of a developmental approach for the implementation of an adaptive system engaged in a visually guided reaching task.

### 2.1. Development of visually guided reaching

If we consider an artificial system engaged in a reaching task, the number of motor degrees of freedom (DOF) that have to be controlled in parallel can be as high as 10 or more. Learning how to control all these DOF simultaneously, can be done, but requires most probably a very high number of trials to be performed, in order to reach the required sequence of motor commands.

In a classic control theory approach, some of the system’s parameters required by the controller are estimated a priori while others are identified at run-time by means of ad-hoc calibration procedures (Yoshikawa, 1990). Subsequently, a general purpose or a customized control law is applied. After the calibration procedure the system becomes fully functional and changes of the internal parameters require re-calibration of the system (Papanikolopoulos & Khosla, 1993; Samson, Borgue & Espiau, 1991). Following this approach the problem is simplified by an accurate design but it might require a lot of preliminary effort during the design process.

Other solutions suitable to learn how to control a complex

systems have been proposed. Among them, the *feedback error learning* (Kawato et al., 1987) and the *reinforcement learning* (Sutton & Barto, 1998). Along the same line of research we designed an alternative solution based on direct motor primitives representing multi-joint synergies (i.e. for arm extension). In this case a single command may produce complex multi-joint coordinated movements without the voluntary control of each individual DOF. Examples of such multi-joint synergies are the ATNR (see the previous section) and the grasping reflex which activates a coordinated grasping movement of the hand when the palm touches an object. In order for this approach to be feasible and effective, the crucial points are how to represent the motor primitives, the mechanisms of sensori-motor mapping, and their developmental rules.

#### 2.1.1. Motor primitives

Isolated skeletal muscles act like non-linear (visco-elastic) actuators whose length-tension properties are modulated by neuromuscular activation (Rack & Westbury, 1969). For the purpose of the present work, however, a simplified linear model of the muscles (Kandel, Schwartz & Jessel, 1991) has been used to express the torque exerted by a muscle on each joint:

$$\tau = -a\kappa(q - q_0) \quad (1)$$

where  $q_0$  is the actuator’s resting position,  $a$  is the activation value which modulates the overall stiffness  $\kappa$  (i.e. the spring constant of the muscles).

Assuming this model, a possible procedure for coding motor primitives is the so-called force fields approach proposed by Mussa-Ivaldi and Bizzi ((Gandolfo & Mussa-Ivaldi, 1993; Mussa-Ivaldi & Giszter, 1992; Mussa-Ivaldi, Giszter & Bizzi, 1993) recently extended to nonlinear visco-elastic actuators (Mussa-Ivaldi, 1997)). According to this theory, the neuromuscular control of each joint can be described by means of a torque field expressed by:

$$\boldsymbol{\tau}(\mathbf{q}, a) \quad (2)$$

where  $\mathbf{q}$  is the vector of generalized coordinate,  $a$  is the activation value and  $\boldsymbol{\tau}$  is the generalized torque field.

In case of a multi-joint structure (such as a limb) the overall torque is expressed by:

$$\boldsymbol{\tau} = \sum_i \boldsymbol{\tau}_i(\mathbf{q}, a_i) \quad (3)$$

where  $a_i$  are the control parameters.

From the mechanical point of view, the system controlled by these actuators is passive. Consequently it has a stable Equilibrium Point (EP) in its space state  $(\mathbf{q}, \dot{\mathbf{q}})$ . The EP is (a) the point where the torque field described by Eq. (3) is zero, (b) the intersection of the actuators’ angle–tension curves. If we apply the torques described by Eq. (3) to the multi-joint structure, its state will eventually reach the EP (at equilibrium). Thus, the EP can be thought of as the point toward which the configuration is aiming at each instant of time.

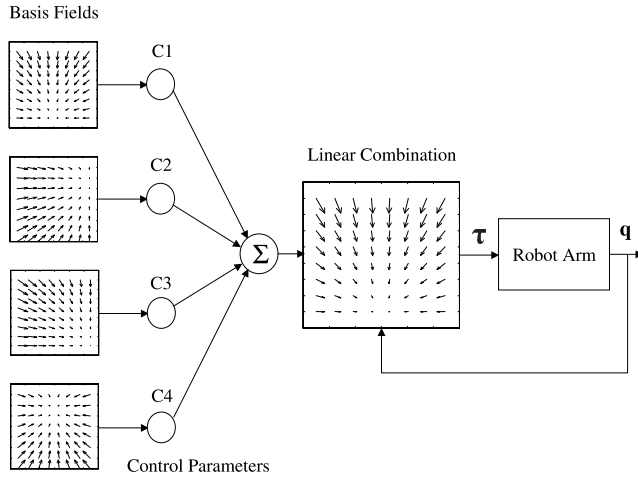


Fig. 1. Controller structure: motor primitives, represented by torque fields are combined (weighted by  $C_1$ ,  $C_2$ ,  $C_3$  and  $C_4$ ). The overall field ‘guides’ the arm end-point toward its EP.

In theory, specification of the EP is enough to drive the system to a given configuration. On the other hand, experimental results in animals and humans (Mussa-Ivaldi et al., 1993) support a rather different view. In fact, it seems that movement is obtained by shifting the EP smoothly from the start to the end, rather than suddenly moving it to the target position. The sequence of EPs defines what is called a *virtual trajectory* (Hogan, 1985). It is worth noting that, due to the dynamic parameters of the arm, the actual arm’s trajectory is different from the virtual trajectory (in other words it is like pulling a toy car with a rubber band: the trajectory in space of the pulling hand is different from the trajectory of the car because of the stretching of the rubber band).

The simplification (and in some sense the feasibility) of this schema comes from the experimental observation that any position of the EP in the arm’s configuration space (and consequently its motion) can be obtained by a linear combination of a small number of motor primitives each represented as a torque field (the so-called




Fig. 2. Two of the four basis fields represented as torque fields in joints coordinates (see also Figs. 1 and 3). Ordinate and abscissa show joint position (joint 1 and 2) in degrees. Arrows point to the common EP of the two joints. Actual resting position of each actuator  $q_0$  was preset by the experimenter.

*basis field* (Mussa-Ivaldi & Giszter, 1992; Mussa-Ivaldi et al., 1993)).

In our model, each motor primitive is a structure which activates a single or a group of actuators (see also Fig. 1). It is actually a synergy which combines (linearly) the effect of a set of actuators by activating them synchronously by means of only one control parameter (i.e.  $a_i = C_j$ ). Primitives can be described by the following torque field:

$$\mathbf{T}_j(\mathbf{q}, C_j) = \sum_i \mathbf{I}_{ji} \boldsymbol{\tau}_i(\mathbf{q}, C_j) \quad (4)$$

where  $\boldsymbol{\tau}_i$  is the  $i$ th actuator field,  $C_j$  the activation value and

$$\mathbf{I}_{ij} = \begin{cases} 1 & \text{if the } j\text{th controller activates the } i\text{th actuator} \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

$\mathbf{T}_j$  are exactly the basis fields as shown in Fig. 2. The total field  $\mathbf{T}$  is expressed by the following:

$$\mathbf{T}(\mathbf{q}) = \sum_j \mathbf{T}_j(\mathbf{q}, C_j) = \sum_j \sum_i \mathbf{I}_{ji} \boldsymbol{\tau}_i(\mathbf{q}, C_j) \quad (6)$$

We designed a priori the connections between actuators and primitives (through  $\mathbf{I}_{ij}$ ). Thus in our case the basis fields are fixed and embedded into the system. The following table represents the  $\mathbf{I}$  matrix:

Primitive ( $j$ )	Actuators ( $i$ )			
	1	2	3	4
1	1	0	1	0
2	0	1	0	1
3	1	0	0	1
4	0	1	1	0

where primitives are labeled with  $j$  and actuators with  $i$  as in Eq. (5) connections.

Given this assumption, the task of the controller is to combine the basis fields by providing, for each point of the configuration space, a set of control parameters  $C_j$ . A schematic diagram of the controller is shown in Fig. 1 in the case of four basis fields and two joints.

A further simplification allowed by the force field approach comes from the fact that control parameters are not dependent on any particular frame of reference (Mussa-Ivaldi & Giszter, 1992). This is easily shown converting Eq. (6) into extrinsic coordinates. Let  $\mathbf{x} = \Lambda(\mathbf{q})$  be the direct kinematics mapping of the arm and  $\mathbf{J}_\Lambda = \partial\Lambda/\partial\mathbf{q}$  its Jacobian. For any configuration where  $\mathbf{J}_\Lambda$  is not singular we can write:

$$\mathbf{J}_\Lambda^{-T} \boldsymbol{\tau} = \mathbf{F} \quad (7)$$

where  $\mathbf{J}_\Lambda^{-T}$  is the transposed inverted Jacobian,  $\boldsymbol{\tau}$  the torque vector and  $\mathbf{F}$  the corresponding force vector in extrinsic coordinate. Substituting Eq. (7) into Eq. (6) and considering

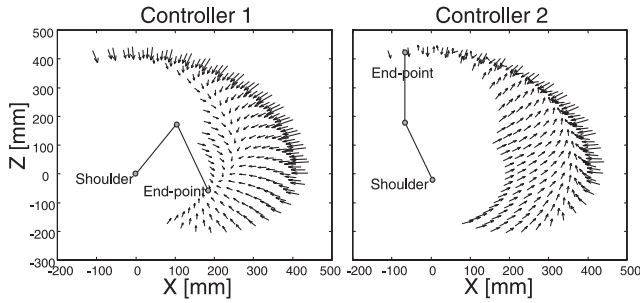


Fig. 3. Two of the four basis fields represented in Cartesian coordinates (see also Fig. 2). Ordinate and abscissa represents the plane where the arm motion has been constrained.

linear actuators (see Eq. (1)), yields:

$$\mathbf{J}_A^{-T} \mathbf{T} = \sum_j C_j \mathbf{J}_A^{-T} \mathbf{T}_j \quad (8)$$

where  $\mathbf{J}_A^{-T} \mathbf{T} = \mathbf{F}$  is the total force field and  $\mathbf{J}_A^{-T} \mathbf{T}_j = \mathbf{F}_j$  are the basis fields in extrinsic coordinate. Substituting yields

$$\mathbf{F} = \sum_j C_j \mathbf{F}_j \quad (9)$$

Eq. (9) shows that the control coefficients  $C_j$  are invariant under coordinate transformation (for a discussion of the underlying conditions see Mussa-Ivaldi and Giszter, 1992).

A similar result applies for the redundant case (where  $\mathbf{J}_A$  is not invertible) depending on the motor primitives considered (Gandolfo & Mussa-Ivaldi, 1993). Given these results it is correct to freely exchange torque fields generated by actuators with force fields applied to the arm end-point because the two representations are indeed equivalent.

From the developmental point of view, this approach is advantageous for two reasons: (i) the kinematic parameters are embedded in the resulting force field (Hogan, 1985); and (ii) each force field corresponds to the activation of a synergy of muscles and does not require the coordinated control of each degree of freedom. Furthermore the ‘innate’

motor synergies can be easily represented through basis fields (or a combination of them). Figs. 2 and 3 show two exemplar basis fields as used in our experiments.

Fig. 2 shows two torque fields in joint coordinates while Fig. 3 plots the corresponding two fields converted in Cartesian coordinates. The picture in Fig. 4 shows the corresponding positions of our robot at equilibrium.

### 2.1.2. Motor–motor coordination

Let us now address the issue of how to drive the motor plant with positional information obtained by vision. In other words, we want to define a way of transforming the visually specified, spatial position of the target into the control parameters  $C_j$ . If this task were implemented on the basis of the Cartesian position of the target in space, the kinematics of the eye–head system as well as of the arm, would have to be explicitly considered in order to select (or combine) the appropriate force fields. The solution we propose is based on the use of a direct mapping 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 of 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 (Fig. 5) and its position in space with respect to the shoulder is uniquely determined by the motor commands controlling the position of the head with respect to the torso and that of the eyes with respect to the head. That is, the position in space of the fixation point can be coded directly using motor commands from any parameter set used to control gaze direction. Mapping these motor parameters into the arm’s force field is what is required to coordinate visual information and motor control.

The system learns the transformation by collecting vector pairs of the form ‘head control vector’–‘arm control vector’ while interacting with the environment. We call this approach motor–motor coordination (Gandolfo, Sandini &




Fig. 4. Equilibrium robot positions corresponding to the exemplar basis fields depicted in Figs. 2 and 3.




Fig. 5. The experimental setup. It consists of a 10 degrees of freedom robot, links are indicated by continuous lines. The two links originating from the cameras are 'virtual' and their intersection represent the fixation point (the end-effector of the head system). Joints are depicted as small quadrilaterals, their axis of rotation are along the corresponding links. The four joints utilized in our experiments are marked with a superimposed circle.

Bizzi, 1996), because the coordinated action is obtained by mapping motor commands onto motor commands. It is also worth noting that the resulting map is, indeed, a representation of the end-effector command in egocentric coordinates and it is consequently in agreement with recent biological findings (Laquaniti & Caminiti, 1998).

### 3. The experiment

Based upon the choices presented in the previous sections an experiment was designed to show how reaching behavior could be acquired by building a map from the head activation values to arm activation values.

In this experiment a four DOF set-up was used: two DOF to control the gaze direction of the camera and two, controlling the position of the end-effector on a plane. The position of the head and arm resembles an anthropomorphic structure and their relative position is fixed but unknown to the system (see Fig. 5).

The visual part is based on a color camera with a space-variant distribution of sensing element generating images with about 2000 pixels in a log-polar format (Capurro et al., 1997; Panerai & Sandini, 1998; Sandini & Tagliasco, 1980; Sandini, Gandolfo, Grosso & Tistarelli, 1993). The

processing is distributed between a Pentium 200 computer and a Sun workstation controlling the arm using the RCCL software package (Lloyd, 1992). The two systems are linked through a TCP/IP ethernet connection.

The following constraints were imposed to obtain the real-time performance required. First, the visual localization of target and end-effector is based on a simple color segmentation algorithm. The target is identified by a green region and the end-effector by a red one (in both cases the position of the segmented region in the image plane is identified by the center of gravity). The second constraint is that only the arm motion is learned while the mapping between the position of the target in the image plane and the eye's motor command required to fixate the target is tuned beforehand. Lastly, the representative control parameters we choose in order to describe the head plant (gaze direction) were the joint angular positions. It is fair to say that there are at least two possible choices: (1) joint positions; (2) head activation values. Both solutions provide their own advantages and drawbacks. We choose to use joint positions in order to keep implementation as simple as possible. Integration of a more interesting gaze control strategy such as those described in (Berthouze & Kuniyoshi, 1998; Capurro et al., 1997; Panerai & Sandini, 1998; Panerai, Metta & Sandini, 1999) is currently under implementation.

These constraints do not affect, in our view, the main




Fig. 6. Initial configuration of the head–arm mapping. With reference to Fig. 1 the four plots show (coded in gray levels) the weights  $C_1$ ,  $C_2$ ,  $C_3$  and  $C_4$  used to linearly combine the basis fields as described in Section 2.1. Ordinate and abscissa are the head joint angles (in degrees). In this figure only three points are defined, corresponding to the three initial reflexes. The resulting three force fields and arm's equilibrium positions are presented in Fig. 7. For example, the point corresponding to pan and tilt equal to zero, which represents the look-ahead condition, holds the controller's activation values  $C_1$ ,  $C_2$ ,  $C_3$  and  $C_4$  causing the arm to be extended in front of the body. See Fig. 10 for the corresponding plots after training.

points of the approach proposed. However it is fair to say that removing some of these constraints (e.g. introducing a redundant manipulator) may likely introduce new problems, not accounted for at this point.

Given these considerations the map, in this particular experiment, can be represented by:

$$\mathbf{C} = f(\mathbf{q}) \quad (10)$$

where  $f$  is the unknown true function which must be approximated by learning,  $\mathbf{q} \in \mathcal{R}^2$  is the head joint angle's vector and  $\mathbf{C} \in \mathcal{R}^4$  is the arm activation vector.

The  $(\mathbf{q}, \mathbf{C})$  pairs required to estimate the function  $f$  are measured whenever the system is fixating its own hand (and not when the gaze is fixating the target). The values of the activation vectors  $\mathbf{C}$  are stored in a look-up table (the motor–motor coordination map) whose input space  $q$  is sampled with uniform resolution (a vector  $\mathbf{C}$  can be stored in each location or 'cell'). If a cell has never been visited, but the function value for that input position is needed, the value stored in the nearest visited cell is used instead. Learning proceeds by updating the values of  $\mathbf{C}$  each time the corresponding head joint vector is used to fixate the end-effector.

### 3.1. Initialization of the motor–motor map

The first problem to be solved is how to initialize the map in a meaningful way (or, in other words, what type of motor primitives should be used as the basis of the learning procedure). In natural systems this is obtained by reflexive mechanisms like the ATNR which has the role of maintaining the arm within the field of view.

In our experiment, the robot utilizes a discrete approximation of the ATNR by initializing the head–arm map so that the arm is extended roughly in the direction the head is turned. As shown in Fig. 6, each map stores three initial values of each of the four elements of the activation vector  $\mathbf{C}$  corresponding to the three head positions. The coordinate axes represent the head's tilt (abscissa) and pan (ordinate) angles.

Each map is virtually empty apart from the three 'dots' representing the values  $\mathbf{C}$  corresponding to three head positions (see Fig. 6). The three activation vectors span uniformly the arm workspace and were computed so that whenever they are used the arm end-point would move into the camera field of view. Consequently, even if the choice of just three positions is arbitrary, this initialization of the



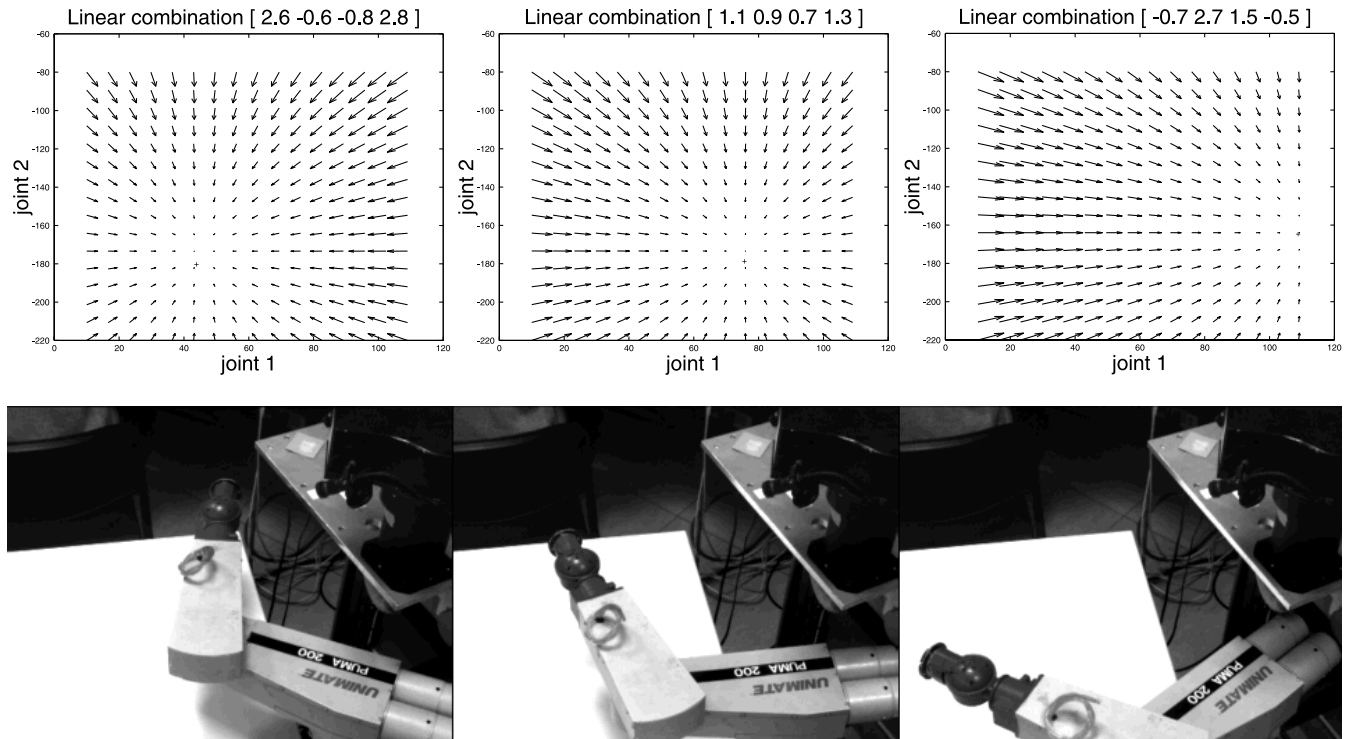


Fig. 7. Force fields (upper row) and corresponding equilibrium arm position (lower row) of the three initial reflexes. Force fields are represented in joint's coordinates and arrows point to the common equilibrium position of the two arm's joints. See also Fig. 6.

head–arm mapping is advantageous with respect to a random sampling of the workspace for two reasons. First the system is put in the conditions to be able to learn from visually measured errors (the arm is kept in the field of view) and second the initial values implicitly limit the exploitation space to accessible and safe regions of the workspace.

The arm's postures corresponding to the initial values of  $\mathbf{C}$  are shown in Fig. 7. It is worth noting that initially the head can explore its entire workspace while only three positions of the arm are possible. The goal of the learning procedure is to fill the empty space of the maps.

### 3.2. Trajectory generation

Both head and arm motions are controlled by torque values representing head and arm force fields. For the head, the instantaneous torque  $\tau$  is obtained from the activation values  $C_i$  derived from the gaze error (see Fig. 8). The gaze error is measured as the target position with respect to the image center expressed in image coordinates.

For example, if the target is located to the right of the fovea, the activation of the right muscle is increased by an amount proportional to the error, while the antagonist is inhibited (activation is decreased by the same amount). The same applies for the tilt axis controlling the up–down motion.

Because the arm has higher inertia and friction of the reduction gears, the extracted activation vectors  $\mathbf{C}$  cannot

be applied instantaneously. An instantaneous application of a torque step would bring the force outside the operational range of the motor. To avoid this situation, a mechanism transforming the activation values obtained by the map into smooth sequences is required. Such gradual rise in force is also observed in biological motion (Kandel et al., 1991). A possible biological mechanism for incremental rise in force levels is motor unit size, with smaller units discharging first during the contraction (Hennemann, Mendell & Brooks, 1981).

To achieve a smooth rise in torque, we applied a linear interpolation for a fixed number of steps between the initial and the final activation values:

$$\mathbf{C}_{i+1} = \mathbf{C}_i + \frac{\mathbf{C}_{\text{final}} - \mathbf{C}_{\text{initial}}}{n_{\text{steps}}} \quad (11)$$

where  $\mathbf{C}_i$  is the activation vector at the  $i$ th time step,  $\mathbf{C}_{\text{final}}$  the target activation vector,  $\mathbf{C}_{\text{initial}}$  its value when the command was issued, and  $n_{\text{steps}}$  the number of steps.

This interpolation procedure is particularly effective in




Fig. 8. Head control scheme. The position error is extracted from the image, converted and used directly to change the activation values of head actuators. The system in this case is using a closed loop approach as in the visual servoing paradigm.






Fig. 9. The overall arm's control scheme. The position of the head ( $q_{head}$ ) addresses the map which outputs the activation vector for the arm. This stage is followed by a trajectory generation which interpolates linearly between activation vectors. The resulting force field is then computed and used to generate the torques which drive the arm motors.

our case, because even if new activation values are issued at approximately 1 Hz (due to the visual and learning processes), the arm controller (running at 50 Hz) can easily generate the interpolated values.

At each time instant  $t$ , it is possible to determine an EP which is a function of  $\mathbf{C}_t$  by imposing:

$$\mathbf{T} = \sum_j \mathbf{T}_j = \sum_j C_j = \sum_i \mathbf{I}_{ji} \tau_i = 0 \quad (12)$$

and solving for  $q$ .

The sequence of EPs defines the arm's virtual trajectory (see also Section 3.4). However, the sequence of  $\mathbf{C}$  through time also determines the shape of the trajectory.  $\mathbf{C}_t$  can be considered as a set of parameters which are in principle learnable. In fact, they could be tuned in order to straighten the trajectories or to reduce overshoots.

Consider the usual Lagrange equation for a planar manipulator:

$$\mathbf{T} = A(\mathbf{q})\ddot{\mathbf{q}} + B(\mathbf{q}, \dot{\mathbf{q}}) \quad (13)$$

where  $\mathbf{T}$  is the generalized torque applied to the arm. Substituting the expression for  $\mathbf{T}$ , generated by the set of elastic actuators and controllers as previously defined, yields:

$$\sum_j C_j \sum_j \mathbf{I}_{ji} \tau_i = A(\mathbf{q})\ddot{\mathbf{q}} + B(\mathbf{q}, \dot{\mathbf{q}}) \quad (14)$$

Two considerations stem from the previous equation: (i) the real trajectory of the arm is determined by the shape and evolution in time of the torque field (left hand side of equation); (ii) as already pointed out, the shape of the torque field is controlled by  $\mathbf{C}_t$ . If the system were able to tune  $\mathbf{C}_t$  beside the simple linear interpolation, it could also change the resulting arm trajectory precisely. Although this could be a sensible strategy (for example to learn how to get a straight trajectory instead of a curved one) it was beyond the available computational power of our system.

The overall control scheme is shown in Fig. 9. The first stage of the processing is implemented in the map containing the arm activation vector. These values are interpolated and the output from the trajectory generator is sent to the actuators simulator (identified by the block 'Field generation') which generates the torque commands.

### 3.3. The learning procedure

The learning algorithm can be formally described as follows:

Repeat forever.

1. A proper stimulus appears in the field of view.
2. By fixating the visual target the robot also initiates arm motion by computing the arm activation vector  $\mathbf{C}$  in the following way:

$$\hat{f}_i(\mathbf{q}) + \mathbf{n} \quad (15)$$

The term  $\mathbf{n}$  describes a zero-mean uniform noise component introduced to simulate errors in the arm control.  $\hat{f}_i$  is the estimate of  $f$  at the  $i$ th iteration.

3. The vector  $\mathbf{C}$  is used by the arm controller which computes the actual torques to drive the motors and consequently the arm moves toward the new EP.
4. At this point the arm is as close as possible to the target (initially it is not very close but certainly it is in the field of view) so that the system can re-direct the gaze to its own hand.
5. As a result of the previous step a new pair  $(\mathbf{q}, \mathbf{C})$  is available which is used to update the map by computing the value  $\hat{f}_{i+1}(\mathbf{q})$  in the following way:

$$\hat{f}_{i+1}(\mathbf{q}) = \hat{f}_i(\mathbf{q}) \frac{n_v - 1}{n_v} + \frac{\mathbf{C}}{n_v} \quad (16)$$

where  $n_v$  is the number of visits of the cell corresponding to  $\mathbf{q}$ .

6. The arm then returns to a fixed resting position (near the chest).

It is important to note that if the procedure were noise-free the motion of the arm toward the target (end of step 3) would always bring the end-effector in the same final position and the system would not be able to learn (in fact it would always update the same cell of the map with the same vector  $\mathbf{C}$ ). In human development, the system is not noise free at birth. For example, nerve growth is not completed giving rise to slowed nerve conduction that is prone to interference, which ultimately may lead to poor sensory data (i.e. low visual acuity) and noisy motor output (Kinney, Brody, Kloman & Gilles, 1988; Konczak & Dichgans, 1997; Konczak, Borutta & Dichgans, 1995, 1997). In our experimental condition errors were introduced 'naturally' by the




Fig. 10. Head–arm mapping after training (refer also to Fig. 6 to compare the maps shown here with those before training). Ordinate and abscissa are the head joint angles (in degrees). Grey levels represent activation value  $C_1$ ,  $C_2$ ,  $C_3$  and  $C_4$  of each controller (see Section 2.1).

friction of the arm reduction gears and ‘artificially’ by introducing a noise term in the motor–motor transformation (see step 2). During learning and subsequent retention trials friction of the arm could not be altered. However, the artificial noise was removed, whenever we tested the accuracy of reaching movements.

A map obtained after a training period of about 5 min and consisting of approximately 100 trials is shown in Fig. 10. It is worth noting that, in spite of the sparse initialization of the map (see Fig. 6), the distribution of the activation values after 100 trials is rather continuous considerably extending the reaching skills.

A key point of the overall process is that the body plays the role of both the actor (by exploring the environment) and the role of the environment (by using the eye’s fixation point as the target of the reaching process). This allows the overall system and the learning process to be self-contained and adaptive to kinematic and dynamic changes of the internal parameters (such as body segments length and weight). Moreover, the process is intrinsically egocentric. The motor–motor mapping, at least initially, does not necessarily bring the end-effector near to the fixation point (in fact 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.

### 3.4. Experimental results

Two different experiments were performed to illustrate the performance of the proposed approach. The first describes the learning of ballistic reaching movements toward static visual targets, the second presents the results of smooth coordinated eye–hand movements toward moving targets emerging from the learned ballistic behavior.

In order to test the performance of the system at different learning stages, the position in the arm’s workspace of three targets was calibrated beforehand by manually positioning the end-effector at target center and storing the corresponding joint angle values measured by the encoders. Each target consisted of a piece of cardboard about  $5 \times 5$  cm in size.

Table 1

Endpoint positioning before learning and after 51 and 134 trials. The error expresses the Euclidean distance in millimeters between the end-effector and the target at the end of reach

Number of trials	Before	51	134
Error (mm)	77.8 ± 15.0	39.5 ± 12.0	28.8 ± 8.9

During the training phase the target of the reaching task was manually moved by the experimenter over the arm's workspace while the reaching behavior was continuously activated. From time to time training was suspended and the performance evaluated.

During the evaluation phase the three targets in the calibrated positions were activated one at a time and the trajectory of the arm stored. The reaching error was measured by computing the Euclidean distance between the pre-calibrated target positions and the position of end-effector at the end of the reaching movement. At least 30 trials (10 for each target) were executed and the average error and standard deviation were computed. During this evaluation phase the map update was stopped and the noise term removed (see Eq. (15)).

The reaching error before learning, after 51 trials and after 134 trials is reported in Table 1.

It is important to note that trajectories are not learned by the system. They are just a consequence of the applied control strategy as described in Section 2.

A typical arm trajectory after training in joint and Cartesian coordinates is shown in Fig. 11. In both graphs, the presence of overshooting of the real trajectories is observed. This is the effect of not knowing the dynamic parameters of the arm and particularly of the arm's inertia. As a consequence the torque applied in the initial part of the movement brings the end-effector beyond the target. The 'force field approach', however, corrects this overshoot by applying a force in the opposite direction and partially compensates

this lack of dynamic information. In our current schema there is no possibility to 'learn' how to avoid this overshoot because this would require the tuning of other parameters such as the stiffness or the presence of compensating modules which explicitly take into account dynamics (Ghez, Gordon, Ghilardi & Sainburg, 1996).

Similar considerations can be drawn by observing the plots in Fig. 12. In this case reaching movements toward three different targets at the end of the training phase are shown. Trajectory toward target 1 shows the same overshoot described for Fig. 11. The opposite happens when the most distant target 3 is reached. In this case the end-effector undershoots the target. Also in this case the remaining error can be attributed, in part, to intrinsic errors of the learning process, but also to the accumulation of errors deriving from friction (which is not only unknown but also partly unpredictable). Trajectory toward target 2 shows a back-and-forth motion with the final position reached after a couple of adjustments. This behavior is caused by the fact that the system is continuously operating and, consequently, whenever the end-effector partially covers the target, the head shifts the fixation point over the center of gravity of the remaining visible part. This change of fixation generates a new 'force field' and, consequently a new trajectory. Eventually the visible part of the target does not change (the center of gravity of the target does not move any more) and the arm reaches its final position.

At the end of the training phase the system was also tested in the task of reaching toward a smoothly moving target. In this case both the head and the arm are continuously moving to track the target. It is worth noting, that this condition was totally new to the system which, during the training phase, was programmed to move to a flexed fixed position of the arm after each reaching motion.

In both panels of Fig. 13 the continuous line shows the trajectory of the target moving at constant speed. Before

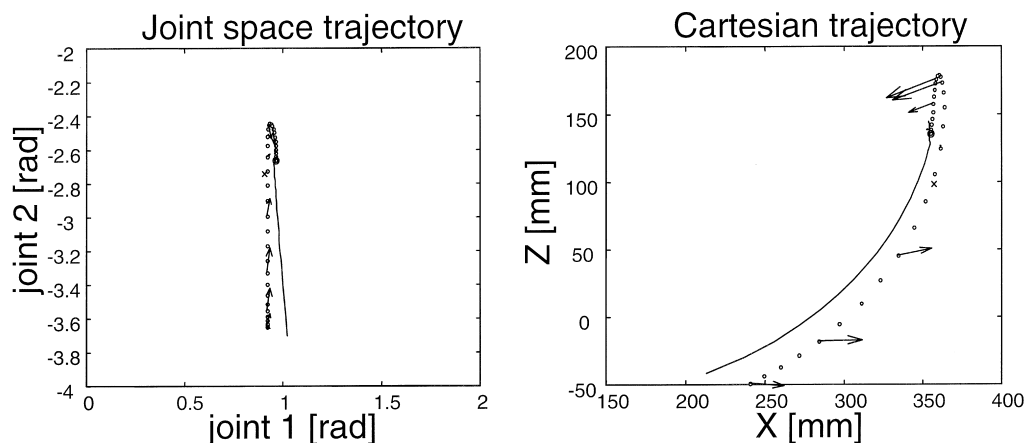


Fig. 11. Typical arm trajectory in joint (left) and Cartesian (right) coordinates. Dotted lines plot the real arm trajectories, continuous lines plot virtual trajectories.  $\times$  is the final end-point position and is the actual position of the target. The vectors plotted every five data points are torques in joint space (as generated by the motors) and forces in the Cartesian space. Left: joint 1 is the shoulder position while joint 2 refers to the elbow. Right: the XZ plane is the horizontal plane representing the arm workspace. This particular motion is a reach to the left involving mainly the motion of the elbow.

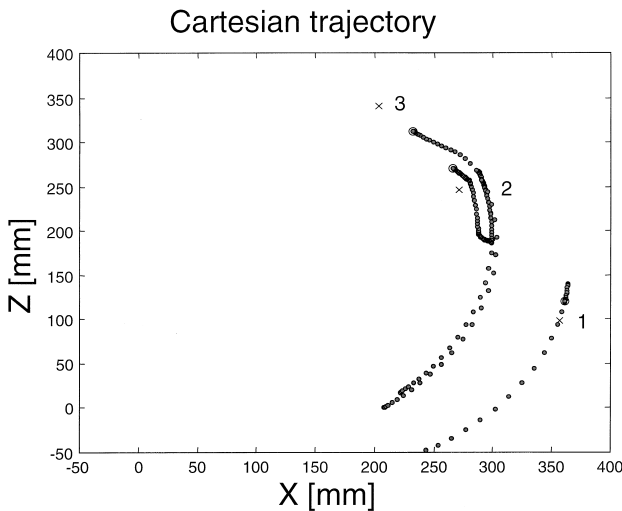


Fig. 12. End-point trajectories in Cartesian coordinates generated during reaching toward three different targets.  $\circ$  represents the final end-point positions and  $\times$  the target positions. In all cases movements started with the arm in a resting position flexed close to the body, i.e. roughly in a (200;0) position.

training the arm motion is composed of two motion units each corresponding to one of the three initial force fields. The movement was dysmetric. The endpoint missed the target by nearly 10 cm. After training the motion of the endpoint is much smoother and the overall tracking performance is certainly improved as the trajectory of the endpoint stays closer to the target's trajectory. Also the final error is drastically reduced.

The remaining small oscillations are mainly due to the discrete nature of the look-up table approximating the map and to some extent to the scheme applied to control the head movements. Beside this defect, the enhancement of the positioning abilities is very noticeable.

#### 4. Conclusion

This paper presents a framework for the implementation

of adaptable sensori-motor strategies for visually guided reaching. The implemented framework is inspired by studies on human development, and we attempt to achieve visuo-motor coordination by adopting biologically plausible control structures.

We demonstrate the development of visually-guided reaching from an initial state characterized by a set of reflexive behaviors (motor primitives). Subsequently, visuo-motor skills are acquired by refining the mapping between sensory information and motor commands. During this developmental progression there is no distinction between plant's calibration and control, and the kinematic and dynamic parameters are not explicitly identified as in classical control theory approaches. This developmental process may be viewed as adaptive change towards competence (Keogh & Sugden, 1985). Here, adaptive change is not the same as learning. In distinguishing between learning and development, we regard learning as a function of development rather than development being the overall summation of a series of learnings (Piaget & Inhelder, 1969). This view implies that learning is shaped by the learner's developmental state. In practice, this means that the system incrementally adapts its learning goals to the evolving developmental state. Engineering such a process means being able to define a sequence of events that cause the system to become incrementally more skilled. One way of looking at it is to model a developmental stage as a set of control variables (in our case motor but, in general, also sensory and cognitive) and to model the process of development as a progressive, adaptive selection of the learning parameters.

Consider the developmental state of a human infant at birth. That state is characterized by an incomplete visuo-motor map, by imprecise knowledge of the plant, by the availability of basal intra- and interlimb synergies and a set of primitive reflexes. This setup allows the infant to explore and to exploit the environment even at that early age. The exact mechanisms of this ontogenic process in the first weeks of life are still not clear. However, visuo-motor coordination is likely based on at least two motor maps: a spinal and a cortical one. We




Fig. 13. Tracking a moving target. The continuous line shows the trajectory of a target moving at constant speed in the arm's workspace. The target was moving at about 8 cm/s from (350;100) to (200;330). Dotted lines represent the endpoint trajectory. The left plot has been obtained using the map before training, the right one refers to the map after 134 training trials.

know today that most cortico-spinal projections are not differentiated at birth (O'Leary, 1992) and that spinal reflex circuits provide the system with the reflexes and basal synergies necessary to initiate simple interactions with the environment.

As maturation progresses, cortical control loops become operational, building a map based on existing reflexes and basal synergies. In this sense, there seems to be no need to assume a cortical suppression mechanism inhibiting spinal reflexes (Gesell, 1946; McGraw, 1946). Instead, the development of visuo-motor coordination looks more like a process during which the system learns how to 'drive' the spinal motor primitives. The autonomy of the spinal motor system becomes evident after the disruption of cortical input. Spinal synergies are preserved and still functional (for a review see Grillner, 1981).

With our robot setup we attempted to follow a similar line of developmental events. We equipped our robot with a set of basal synergies represented by basis fields. These synergies were embedded into the system as a set of initial visuo-motor reflexes. Their biological analogs were the Asymmetric Tonic Neck Reflex, where head rotation triggers arm extension, thus bringing the reaching hand in the field of view. In addition, the biological parallelism was stressed at the level of the actuators by simulating the elastic properties of the muscles and exploiting those properties by means of motor primitives based on force fields. Visual processing, although rather simple, was performed on images acquired by an anthropomorphic sensor simulating the space-variant distribution of photoreceptors in the human retina.

The current implementation is still constrained by practical reasons (see Section 3). One such constraint is related to the use of color to drive visual segmentation and the identification of the point to be reached. This is not a significant constraint, at least for the scope of the present article, and can be substituted by any visual primitive driving the fixation point toward the point in space to be reached. The second constraint is the fact that learning is limited to the control of the arm while the mapping between position in the image plane and eye-head motor command is pre-calibrated. This point is certainly more relevant for the present paper. However it is worth noting that our visual system is based on a retina-like sensor which 'only' acquires 2000 pixel images. Consequently, the pre-calibrated values are consistent with the general idea that the initial situation is characterized by reduced resolution. In other words, we do not rely on accurate pre-calibrated measures—an initial situation not far from that of a 'newborn'. The third constraint has been the use of angular positions instead of 'force fields' as control parameters to describe the head-eye plant.

This constraint roughly resembles the difference that exists between the use of an efference copy and a proprioceptive input variable. If fast motion were involved, this might be crucial (i.e. the efference copy information could

be used to produce a sort of state prediction) but in our case the use of either type of signals made no difference. The proprioceptive input though simplified the mapping.

Finally, even if our experiments were based on a non-redundant manipulator, the extension to redundant systems is feasible. A first approach would be to build synergies in order to reduce the number of controllable degrees of freedom. In this case an additional constraint would be added at the actuators level. With reference to our current set-up we can imagine to have a planar 3 DOF robot arm controlled by six elastic actuators arranged in push-pull pairs. At the controller level we can always choose the connections ( $I_{ij}$ ) so that we only have two independent degrees of freedom. In this case the redundant arm would actually behave as a non-redundant device. The learning procedure would, in this case, be the same as described before.

Another possibility is to use the solution proposed by (Gandolfo & Mussa-Ivaldi, 1993) who demonstrated that even in the redundant case (excess of degrees of freedom) the 'force fields' approach might behave adequately. In particular, they showed that a Jacobian pseudo-inverse is able to map basis fields from joint to Cartesian coordinates. Their solution (see Mussa-Ivaldi & Hogan, 1991) is integrable and provides the required additional condition for the Jacobian inversion in terms of a corrective stiffness matrix. In our case we can express the position of the point to be reached (i.e. the fixation point) in Cartesian coordinates and plan the control of the redundant robot arm as proposed by (Mussa-Ivaldi & Hogan, 1991). However this solution requires some form of calibration of the vision system.

As a general comparison with other biologically inspired models one of the alternatives is the direct modeling approach (see Jordan, 1996 for a review). In theory, the direct modeling approach might be able to solve the same problem. However it needs to be trained off-line and the training set has to be generated prior to any actual control. Suppose a direct modeling approach were applied to the visuo-motor coordination task. As usual, the system has to be provided with a training set. The question is how we can generate it. The easier and most used strategy has been that of uniformly sampling the robot control space, execute the relative commands, observe the results and eventually adjust the internal model. This is feasible though it is not efficient.

First, in that case we need to know a priori the size and the limits of the control space. Second, we have also to decide which sampling strategy to adopt (uniform, space-variant). Thus, the resolution is established prior to any actual use of the system. Finally, adaptation is precluded to such a system. Any modification during the control stage (such as a change on a physical system parameter) would require to switch the system into a new learning stage. Unfortunately, this would also require to re-explore the whole control space even if the change might affect only a part of it. This is a major difference. Our schema is intrinsically goal directed while the direct approach even for a single

goal case would explore the entire control space. Perhaps, this is not an issue for the simple 2D case considered here, but it might be a problem when the number of dimensions is increased. It is worth stressing that in our model, as well as in others (for example see Kawato, Furukawa & Suzuki, 1987) control and training are two parallel processes.

Compared with the so-called bi-directional theory (Miyamoto et al., 1998) the method proposed here is simpler and in some parts it does not involve complete forward and inverse models. In Kawato's model each level namely control, planning, coordinate transformation is connected through a forward and an inverse block to another one in a hierarchical fashion. In our case, some of the feedback signals are not present and some loop simplified. For instance the trajectory generation stage in our model is completely feed-forward but also a real visual feedback is absent (i.e. we do not evaluate directly the error to correct the behavior). This is not to say they would not be needed in later stages of development. At the same time it is fair to say that due to hardware limitations we were not able to test the system in fast dynamic tasks, thus a proper comparison was not possible in quantitative terms.

In conclusion we would like to stress again that the implemented artificial system only simulates biological development in a rather primitive fashion. Yet, it possesses a set of features that we think are promising for designing autonomous robots that can act adaptively in a visually specified environment. These features are as follows:

- The system is complete in the sense that specific sensory and motor components develop simultaneously.
- The sight of the hand drives the learning of the motor–motor map. In this respect, robot's body generates enough information to learn a map which guides the reaching for an external object (target). We are not saying that eye-hand coordination in humans is implemented through a direct motor to motor mapping. This remains a researchable question.
- The basal motor repertoire becomes part of the growing motor–motor map without requiring the explicit inhibition of 'innate' motor reflexes. This is not to say that at some later developmental state, inhibitory mechanisms are not to be employed.

In light of the above features, a final note on the distinction between learning and development is warranted, in order to stress that our approach characterizes a developmental rather than of a learning process: during learning a defined relationship between various sensory and motor subsystems constrains the kind of tasks that can be learned. In contrast, development takes place on a larger scale. It is not merely a series of discrete learning experiences, but is characterized as a process where the interrelationships between subsystems are not yet rigidly defined. It is this lack of rigidly defining initial states, that enable developing systems to adapt to a potentially larger range of tasks.

## Acknowledgements

The research described here is supported in part by a National Project of the Italian Ministry of Scientific and Technological Research (MURST), by Italian Space Agency (ASI) and in part by the EU-TMR project Virgo.

## References

- Aloimonos, J., Weiss, I., & Bandyopadhyay, A. (1988). Active vision. *International Journal of Computer Vision*, 1 (4), 333–356.
- Bajcsy, R. K. (1985). Active perception vs. passive perception. *Proceedings of the Third IEEE Workshop on Computer Vision: Representation and Control* (pp. 13–16). Bellaire, MI: USA.
- Bajcsy, R., & Tsikos, C. (1988). Perception via manipulation. *Proceedings of the International Symposium and Exposition on Robots* (pp. 237–244). Sydney: Australia.
- Ballard, D., & Brown, C. (1992). Principles of animate vision. *Computer Vision Graphics and Image Processing*, 56 (1), 3–21.
- Berthouze, L., & Kuniyoshi, Y. (1998). Emergence and categorization of coordinated visual behavior through embodied interaction. *Machine Learning*, 31, 187–200.
- Brooks, R. (1996). Behavior-based humanoid robotics. *Proceedings of the IEEE/RJS Intelligent Robots and Systems IROS'96* (Vol. 1, pp. 1–8).
- Bushnell, E. (1981). The ontogeny of intermodal relations: vision and touch in infancy. In R. Walk & H. Pick, *Intersensory perception and sensory integration*, (pp. 5–37). New York: Plenum Press.
- Capurro, C., Panerai, F., Grosso, E., & Sandini, G. (1993). A binocular active vision system using space variant sensors: exploiting autonomous behaviors for space application. *International Conference on Digital Signal Processing*. Nicosia: Cyprus.
- Capurro, C., Panerai, F., Grosso, E., & Sandini, G. (1996). Visuo-motor coordination for advanced robotics. *Sixth ISIR International Symposium on Industrial Robotics*. Milan: Italy.
- Capurro, C., Panerai, F., & Sandini, G. (1997). Dynamic vergence using log-polar images. *International Journal of Computer Vision*, 24 (1), 79–94.
- Coombs, D., & Brown, C. (1990). Intelligent gaze control in binocular vision. *Proceedings of the Fifth IEEE International Symposium on Intelligent Control*. Philadelphia, PA.
- Crowley, J., Bobet, P., & Mesrabi, M., (1992). Gaze control with a binocular camera head. *Proceedings ECCV92-European Conference on Computer Vision* (Vol. LNCS-588). Santa Margherita Ligure, Italy: Springer-Verlag.
- Eibl-Eibesfeld, I. (1970). *Ethology: the biology of behavior (e.klinghammer trans.)*, New York: Holt, Rinehart & Winston.
- Gandolfo, F., & Mussa-Ivaldi, F. (1993). Vector summation of end-point impedance in kinematically redundant manipulators. *Proceedings of the IEEE/RJS International Conference of Intelligent Robots and Systems* (pp. 1627–1634). Yokohama: Japan.
- Gandolfo, F., Sandini, G., & Bizzi, E. (1996). A field-based approach to visuo-motor coordination. *Proceedings of the Workshop on Sensorimotor Coordination: Amphibians, Models and Comparative Studies*. Sedona, AZ: USA.
- Gandolfo, F., Sandini, G., & Tistarelli, M. (1991). Towards vision guided manipulation. *Proceedings of the International Conference on Advanced Robotics ICAR'91*. Pisa: Italy.
- Gesell, A. (1946). The ontogenesis of infant behavior. In L. Carmichael, *Manual of child psychology*, (pp. 295–331). New York: Wiley.
- Ghez, C., Gordon, J., Ghilardi, M., & Sainburg, R. (1996). Contribution of vision and proprioception to accuracy in limb movements. In M. S. Gazzaniga, *The cognitive neurosciences*, Cambridge: MTP-Press.

- Gould, J. (1982). *Ethology: the mechanism and evolution of behavior*, New York: Norton.
- Grillner, S. (1981). Control of locomotion in bipeds, tetrapods and fish. In V. Brooks, *Handbook of physiology*, (pp. 1179–1236). Bethesda: American Physiological Society (Section 1: The Nervous System).
- Grosso, E., Metta, G., Oddera, A., & Sandini, G. (1996). Robust visual servoing in third reaching tasks. *IEEE Transactions on Robotics and Automation*, 12 (8), 732–742.
- Hennemann, E., & Mendell, L. (1981). In V. Brooks, *Handbook of physiology*, (pp. 423–507). Washington, DC: American Physiological Society (Section 1 Vol. 2 Part 1).
- Hogan, N. (1985). The mechanics of multi-joint posture and movement control. *Biological Cybernetics*, 52, 313–331.
- Jordan, M. (1996). Computational motor control. In M. S. Gazzaniga, *The cognitive neurosciences*, Cambridge: MIT-Press.
- Jordan, M., & Flash, T. (1994). A model of the learning of arm trajectories from spatial deviations. *Journal of Cognitive Neuroscience*, 4 (6), 359–376.
- Kalveram, K. (1991). Controlling the dynamics of a two-joined arm by central patterning and reflex-like processing: a two stage hybrid model. *Biological Cybernetics*, 65, 65–71.
- Kandel, E., Schwartz, J., & Jessel, T. (1991). *Principles of neuroscience*, Amsterdam: Elsevier.
- Kawato, M., Furukawa, K., & Suzuki, R. (1987). A hierarchical neural network model for control and learning of voluntary movement. *Biological Cybernetics*, 57, 169–185.
- Keogh, J., & Sugden, D. (1985). *Movement skill development*, New York: Macmillan.
- Kinney, H., Brody, B., Kloman, A., & Gilles, F. (1988). Sequence of central nervous system myelination in human infancy. Part II Patterns of myelination in autopsied infants. *Journal of Neuropathology and Experimental Neurology*, 47 (3), 217–234.
- Konczak, J., Borutta, M., & Dichgans, J. (1995). Development of goal-directed reaching in infants: hand trajectory formation and joint force control. *Experimental Brain Research*, 106, 156–168.
- Konczak, J., Borutta, M., & Dichgans, J. (1997). Development of goal-directed reaching in infants. Part II Learning to produce task-adequate patterns of joint torque. *Experimental Brain Research*, 113, 465–474.
- Konczak, J., & Dichgans, J. (1997). Goal-directed reaching: development toward stereotypic arm kinematics in the first three years of life. *Experimental Brain Research*, 117, 346–354.
- Kuperstein, M. (1988). Neural model of adaptive hand–eye coordination for single postures. *Science*, 239, 1308–1311.
- Laquaniti, F., & Caminiti, R. (1998). Visuo-motor transformations for arm reaching. *European Journal of Neuroscience*, 10, 195–203.
- Lloyd, J. (1992). *Multi-rccl Reference Manual* (release 4.2). McGill University: Montreal.
- McGraw, M. (1946). Maturation of behavior. In L. Carmichael, *Manual of child psychology*, (pp. 332–369). New York: Wiley.
- Miyamoto, H., Schaal, S., Gandolfo, F., Gomi, H., Koire, Y., Osu, R., Nakano, E., Wada, Y., & Kawato, M. (1996). A kendama learning robot based on bi-directional theory. *Natural Networks*, 9 (8), 1281–1302.
- Mussa-Ivaldi, F. (1997). Nonlinear force fields: a distributed system of control primitives for representing and learning movements. *Proceedings of CIRA'97*. Monterey, CA, USA.
- Mussa-Ivaldi, F., & Giszter, S. (1992). Vector field approximation: a computational paradigm for motor control and learning. *Biological Cybernetics*, 67, 491–500.
- Mussa-Ivaldi, F., & Hogan, N. (1991). Integrable solutions of kinematic redundancy via impedance control. *The International Journal of Robotics Research*, 10, 481–491.
- Mussa-Ivaldi, F., Giszter, S., & Bizzi, E. (1993). Convergent force fields organized in the frog's spinal cord. *The Journal of Neuroscience*, 13 (2), 467–491.
- O'Leary, D. (1992). Development of Connectional diversity and specificity in the mammalian brain by the pruning of collateral projections. *Current Opinion Neurobiology*, 2, 70–77.
- Panerai, F., & Sandini, G. (1998). Oculo-motor stabilization reflexes: integration of inertial and visual information. *Neural Networks*, 11 (7), 1191–1204.
- Panerai, F., Metta, G., & Sandini, G. (1999). Visuo-inertial stabilization in space-variant binocular system. *Robots and Autonomous System*, in press.
- Papanikolopoulos, N. P., & Khosla, P. K. (1993). Adaptive robotic visual tracking: theory and experiments. *IEEE Transactions on Automatic Control*, 38 (3), 429–445.
- Piaget, J., & Inhelder, B. (1969). *The psychology of the child*, New York: Basic Books.
- Rack, P., & Westbury, D. (1969). The effects of length and stimulus rate on tension in the isometric cat soleus muscle. *Journal of Physiology*, 204, 443–460.
- Samson, C., Borgue, M. L., & Espiau, B. (1991). *Robot control: the task function approach*, Oxford: Clarendon Press.
- Sandini, G. (1997). Artificial Systems and Neuroscience. In: M. Srinivasan, *Proceedings of the Otto and Martha Fischbeck Seminar on Active Vision* (abstract).
- Sandini, G., Gandolfo, F., Grosso, E., & Tistarelli, M. (1993). Vision during action. In Y. Aloimonos, *Active perception*, (pp. 151–190). London: Lawrence Erlbaum Associates.
- Sandini, G., & Tagliasco, V. (1980). An anthropomorphic retina-like structure for scene analysis. *Computer Graphics and Image Processing*, 14 (3), 365–372.
- Sutton, R., & Barto, A. (1998). *Reinforcement learning: an introduction*, Cambridge: MIT Press.
- White, B., Castle, P., & Held, R. (1964). Observations on the development of visually guided reaching. *Child Development*, 35, 349–364.
- Yoshikawa, T. (1990). *Foundations of robotics, analysis and control*, Cambridge: MIT Press.