

	<p style="text-align: center;">ADAPT <i>IST-2001-37173</i> <i>Artificial Development Approach to Presence Technologies</i></p>
---	---

Deliverable Item 4.4
Results of behavioral experiments with the robot


Delivery Date: March 1st, 2005

Classification: Public

Responsible Person: Giorgio Metta, UGDIST

Partners Contributed: ALL

Short Description: This deliverable describes the experiments performed on the robot apart from the study on the unsupervised development of multimodal features which is reported in D5.4 and D5.5.

	<p style="text-align: center;">Project funded by the European Community under the “Information Society Technologies” Programme (1998-2002)</p>
---	---

Contents

Contents	2
Introduction.....	3
Papers.....	3

Introduction

In D5.1 we outlined a possible architecture incorporating various learning modalities (unsupervised, reinforcement learning, etc.). In particular, self-supervised learning was hypothesized as a plausible and efficient method for learning in an autonomous setting when sensorial and motoric data is directly available. In this case, the robot can learn several sensorimotor coordination tasks in a developmental progression. We prepared and implemented a set of experiments allowing the robot to acquire reaching and grasping abilities. It is important to note that the controller is fully integrated merging both online learning and various subtasks simultaneously. Various modules have been designed: in particular, we implemented an attention system, a new segmentation procedure, a simple Bayesian recognition module, reaching, a binocular disparity estimation module, and eventually a grasping system. Two notable aspects: i) the final is fairly complex, distributed and parallelized running on a set of 10 PC, and ii) improvement is still possible, and certainly much work is still needed to make it reliable. The experiments of workpackage 5 have also been implemented either directly on the robot or on data collected while the robot was running.

This deliverable is centered on two papers that describe the system implemented so far. The first manuscript (included in this deliverable) is a journal paper submitted to the International Journal of Humanoid Robotics that describes the system that allows the robot to learn to grasp a visually identified object. The system includes three main modules: a body schema, a reaching/grasping schema and an attention system. In short the robot learns:

- To find its hand by employing a combination of proprioceptive and visual information;
- To reach for a visually identified object;
- A model of the object currently grasped by exploiting the fact it is under active control by the robot;
- To employ this information to bias the attention system to find the object the robot has already manipulated.

In the second paper we present an extension of this approach where binocular disparity is employed to estimate the three-dimensional “shape” of the object. This is plugged onto the existing system and exploits the fact the object is recognized and segmented. Binocular disparity is then estimated on the region of the image belonging to the object and used to calculate the principal axis of the object. This information has been used to choose between two different approaches to the reaching of the object. This is clearly a first step into detecting the object affordances although so far no motor information has been employed. These new experiments are extensions of the haptic exploration we have presented for example in D1.6.

Papers

The following sections consist of these two papers:

1. L. Natale, F. Orabona, G. Metta, G. Sandini. *Exploring the world through manipulation: a developmental approach*. Submitted to International Journal of Humanoid Robotics. 2005.
2. L. Natale, F. Orabona, F. Berton, G. Metta, G. Sandini. *From sensorimotor development to object perception*. Submitted to the International Conference of Humanoids Robotics. 2005.

EXPLORING THE WORLD THROUGH MANIPULATION: A DEVELOPMENTAL APPROACH

LORENZO NATALE, FRANCESCO ORABONA, GIORGIO METTA and GIULIO SANDINI

DIST University of Genova

Viale Causa 13, 16145, Genova, Italy

{nat, bremen, pasa, sandini} @liralab.it

This paper describes a developmental approach to the design of a humanoid robot. The robot, equipped with initial perceptual and motor competencies, explores the “shape” of its own body before devoting its attention to the external environment. The initial form of sensorimotor coordination consists of a set of reflexive and explorative motor behaviors coupled to an ensemble of visual routines providing a bottom-up attention system. The developmental path leads the robot from to construction of a “body schema” to the exploration of the world of objects. The “body schema” allows controlling the arm and hand to reach and touch objects within the robot’s workspace. Eventually, the interaction between the environment and the robot’s body is exploited to acquire a visual model of the objects the robot interacts with. In turn object models can be used to guide a top-down attention system. We discuss the implications of our approach in the study of cognition and our effort to build a cognitive artificial system.

Keywords: humanoid robotics, attention system, body-schema, development, grasping

1. Introduction

In the past few years there has been significant technological advance in computer technology and robotics. Today computers are much more powerful than they used to be and they can be interconnected through fast networks, which allow efficient parallel computation. At the same time digital cameras have higher resolution, better quality and higher frame rate. This notwithstanding, we are still far from achieving the dream of artificial intelligence. Artificial systems (computer programs, expert systems or robots) are not able to face the challenges of the real world. We are still not capable of building devices which are able to cope with the variability of the world where, on the other hand, even the simplest animal can thrive. Likewise there is a growing interest in the scientific community to the study of cognitive systems with the aim of implementing cognitive abilities in artificial systems.¹ However, the study of cognition is still in the pre-paradigmatic stage and, indeed, little agreement has been reached even in defining it (see Clark² for a review). According to *cognitivism*, cognition is “a computational process carried out on a symbolic representation of the world”. Symbols represent the world and can be shared across different entities (artificial or biological); they are a complete characterization of the world in which the entity is located, and as such are independent of the entity itself and its past experience. Somewhat at the other extreme, emergent

approaches define cognition as the result of the interaction and co-development between the agent's body and the environment in which it lives.³⁻⁵

Although the definitive answer is still to be found, the observation of biological systems provides hints to plausible solutions. Two aspects look crucial: i) the existence of a body (embodiment) and ii) the fact that the internal representation of the world is acquired by acting in the environment. The two requirements are obviously intertwined, as the interaction between the agent and the environment is possible only by means of a physical body. As a consequence, the internal representations become function of the particular embodiment and, perhaps more importantly, of the experience of the agent. For example the space-variant visual system of primates forces them to move the eyes to fixate the objects of interest by placing the fovea at those locations. Attention is modulated so that only one stimulus is selected among the wealth of possible choices; the importance that is assigned to different stimuli is influenced by experience.

Significantly, subscribing to the emergent approach implies that internal representations cannot be forced into the system "by design"; instead the cognitive system has to be able to create these representations by directly interacting with the environment or, indirectly, with other agents. Through action, the embodiment and environment co-determine the resulting representations.

Motivated by these considerations this paper proposes a developmental approach to the realization of a number of cognitive abilities in a humanoid robot. Although a fair amount of cognitivism is still present, especially in the realization of the visual system, learning permeates the implementation at various levels. Learning and a certain degree of adaptation is clearly the prerequisite to a fully emergent design, although not yet an end or a definite answer.

We identified the minimum requirements for our robot as having a sophisticated oculomotor system, an arm, and a hand. Although simplified this configuration is enough to allow active manipulation of the world through reaching and grasping. The robot follows a developmental route that goes initially through the exploration of its body and ends into the characterization of external objects (e.g. segmentation) by effect of grasping.

Conceptually this process can be divided in three phases. The first stage is devoted to learning the internal models of the body (*learning a body-map*) which provides basic motor and perceptual skills like gaze control, eye-head coordination and reaching. Based on these abilities the interaction with the external world is investigated in the second phase where the robot discovers properties of objects and ways of handling them (*learning to interact*). The robot tries simple stereotyped actions like pushing/pulling and grasping of objects; this allows it to start acquiring information about the entities that populate its environment and simultaneously discover new more efficient ways of interaction (for example different grasp types). Finally the third stage concerns learning to *understand and interpret events*; the robot has associated its actions with the resulting perceptual consequences. Interpretation is achieved by inverting this association;

perceptions are projected into the corresponding actions which work as a reference frame to give meaning to what happens in the environment.

In our past work we have addressed some of the aspects related to the third phase.^{6,7} In this paper we focus on the two first phases: learning a body-map and learning to interact. We show how the robot can acquire an internal model of its hand which allows the robot to localize it and anticipate its position in the visual scene during action execution. The hand internal model is then used to learn to reach a point in space and to accommodate the position of the hand with respect to the object during grasping. The robot uses these abilities to build a visual model of the objects it grasps. Once an object is grasped, in fact, the robot can move and rotate it to build a statistical model of its appearance.

The remainder of the paper is organized as follow. Section 2 describes the experimental setup used in the experiments. Section 3 details the robot's visual system and, in particular, the visual attention system. Section 4 deals with the problem of learning the *body-schema* and describes the procedure to learn how to reach visually identified objects. Section 5 describes the probabilistic modeling of objects and the object recognition algorithm. Finally Section 6 presents a full experiment where all the modules described previously are integrated in a meaningful behavior. Section 7 draws some conclusions and highlights future directions.

2. Experimental Platform

The experiments reported in this paper were carried out on a robotic platform called Babybot (Figure 1). The Babybot is an upper torso humanoid robot which consists of a head, an arm and a hand. The head has 5 degrees of freedom, two of which control the neck in the pan and tilt direction, whereas the other three actuate two eyes to pan independently and tilt on a common axis. The arm is the well known Unimate PUMA 260, an industrial manipulator with 6 degrees of freedom; it is mounted horizontally to better mimic a human kinematics. The hand has 5 fingers; each finger has three phalanges, the thumb has an additional degree of freedom which allows it to perform a rotation toward the palm. Overall the number of joints is 16 but for reasons of space and weight they are controlled by using only six motors. Two motors are connected to the index fingers: they are linked to the first (proximal) and second phalanges. The distal (small) phalange is mechanically coupled to the preceding one so that the two bend together (see Figure 1). Two motors control the motion of middle, ring and little finger. As in the case of the index finger, the proximal phalanges are actuated by a single motor, while the second and third phalanges are actuated by a second motor. The mechanical coupling between the joints is realized by means of springs to allow a certain degree of adaptation. For example, during a movement of flexion of the fingers toward the palm, if the middle finger were to be blocked by an obstacle the others would continue to bend up to the equilibrium of the motor and the spring torques (Figure 1 b) and c)). Likewise, the

same would happen in case the distal phalanges had hit the obstacle. The thumb is different as one motor controls the rotation around an axis parallel to the palm and a second motor is connected to the three phalanges, whose independent motion is permitted by elastic coupling as for the other fingers.

The sensory system of the Babybot consists of two cameras and two microphones for visual and auditory feedback. Tactile feedback is provided by 17 force sensing resistors mounted on the hand, five of which are placed on the palm and the remaining 12 evenly distributed on the thumb, index, middle and ring fingers. A JR3 force sensor provides torque and force feedback at the wrist. Further proprioceptive information is provided to the robot by optic and magnetic encoders mounted on all motors and by a three-axis gyroscope mounted on the head. More details about the Babybot architecture can be found elsewhere.⁸




Figure 1. a) the experimental setup, the Babybot. Left: details of the hand. b) and c): elastic compliance. d)-f): mechanical coupling between phalanges.

3. Visual System

One of the first steps of any visual system is that of locating suitable interest points in the scene (“salient regions” or events) and eventually direct gaze toward these locations. It was already recognized years ago that the ability to move the cameras helps in solving a range of computer vision problems.^{9, 10} This paradigmatic shift was so important that led to the development of “active vision” as a *modus operandi* for a good part of the computer vision community. Human beings and many animals do not have a uniform resolution view of the visual world but rather only a series of snapshots acquired through a small high-resolution sensor (e.g. our fovea). This leads to two questions: i) how to move the eyes efficiently to important locations in the visual scene, and ii) how to decide what is important and, as a consequence, where to look next.

The literature follows two different approaches in the attempt of accounting for these facts. On the one hand, the space-based attention theory holds that attention is allocated to a region of space, with processing carried out only within a certain spatial window of attention. This theory sees attention as a “spotlight”, an internal eye or a sort of “zoom

lens". The most influential evidences for the spatial selection come from the experiments of Posner, Snyder and Davidson¹¹ and Downing and Pinker.¹² In a pointing experiment, they showed that anticipating the appearance of a target with a cue (for example an arrow) sped the response of the subjects. The opposite occurred, that is the subjects' responses were significantly slowed down, when the cue was in the wrong direction (invalid cue). This means that attention might be directed to a region of space even in absence of a real target. Moreover on invalid cues, the response slowed down monotonically as the distance between the cue and the actual target increased. These results suggest that attention is deployed as a spatial gradient, centered on a particular location.

On the other hand, object-based attention theories argue that attention is directed to an object or a group of objects, and that the attention system processes properties of object(s), rather than regions of space. Growing behavioral and neurophysiological evidence has shown, in fact, that selective attention frequently operates on an object-based representational medium in which the boundaries of segmented objects, and not just their spatial position, determine what is selected and how attention is deployed (see Scholl¹³ for a review). In other words, the visual system seems optimized for segmenting complex three-dimensional scenes into representations of (often partly occluded) objects for recognition and action. Indeed, perceivers must interact with objects in the world and not with disembodied locations. For example, attention to one part of an object confers an attentional advantage to other parts of that object.¹⁴ Similarly, attention to one aspect of an object (e.g. its shape) enhances the cortical response to other aspects of that object (e.g. its color or motion); thus, all the attributes of an attended object seem to be bound together into a single entity. This concept holds even when the attended and ignored objects are spatially superimposed. O'Craven et al.¹⁵ have observed the effects of object-based attention using fMRI. In this study, observers viewed a display containing a sequence of semitransparent grayscale images of spatially superimposed faces and houses. At any given moment, either the house or the face moved with an oscillatory motion. Observers were asked to decide whether the currently visible house (or face) matched the one immediately preceding it; this required them to attend closely to the relevant object type. A spatial "spotlight of attention" could not select one of the two superimposed objects; it would necessarily select both or neither.

Finally, another classification can be made depending on which cues are actually used in modulating attention. One approach uses bottom-up information including basic features such as color, orientation, motion, depth, and conjunctions of features. A feature or a stimulus catches the attention of the system if it differs from its immediate surrounding in some dimensions and the surround is reasonably homogeneous in those dimensions. However higher level mechanisms are involved as well; a bottom-up stimulus, for example, may be ignored if attention is already focused elsewhere. In this case attention is also influenced by top-down information relevant to a particular task.

In the literature a number of attention models that use the first hypothesis have been proposed;¹⁶⁻¹⁸ most of them are derived from Treisman's Feature Integration Theory

(FIT)¹⁹. This model employs a separate set of low-level feature maps which are combined together by a spatial attention window operating in a master saliency map. An important alternative model is given by Sun and Fisher,²⁰ who proposed a combination of object- and feature-based theory.

While it is known that the human visual system extracts basic information from images such as lines, edges, local orientation etc., vision not only represents visual features but also the items that such features characterize. But to segment a scene into items, objects, that is to group parts of the visual field as units, the concept of “object” must be known by the system. In particular, there is an intriguing discussion underway in vision science about reference to entities that have come to be known as “proto-objects” or “pre-attentive objects”.²¹ These are steps up from mere localized features, and they have some but not all of the characteristics of “objecthood”. The visual attention model we propose starts by considering the first stages of the human visual system, using then a concept of salience based on “proto-objects” defined as blob of uniform color in the images. Then, since the robot can act on the world, it can do something more: once an object is grasped, in fact, the robot can move and rotate it to build a statistical model of the color blobs, thus effectively constructing a representation of the object in terms of proto-objects and their spatial relationships. This internal representation feeds then back to the attention system of the robot in a top-down way; as an example we show how the latter can be used to direct attention to spot one particular object among others that are visible on a table in front of the robot.

We propose an object-based approach that integrates bottom-up and top-down cues; in particular bottom-up information suggests/identifies possible regions in the image where attention could be directed, whereas top-down information works as a prime for those regions during the visual search task (i.e. when the robot seeks for a known object in the environment).

3.1. *Log-polar images*

Figure 2 shows the block diagram of the first stage of the visual processing of the robot. The input data is a sequence of color log-polar images.²² The log-polar transformation models the mapping of the primate visual pathways from the retina to the visual cortex. The idea of employing space-variant vision is derived from the observation that the distribution of the cones, i.e. the photoreceptors of the retina involved in diurnal vision, is not uniform: cones have a higher density in the central region called fovea, while they are sparser in the periphery. Consequently the resolution is higher and uniform in the center while it decreases in the periphery proportionally to the distance from the fovea. This particular distribution of the receptors seems to influence the scan-paths of an observer,²³ so it has to be taken into account to better model the overt visual attention.

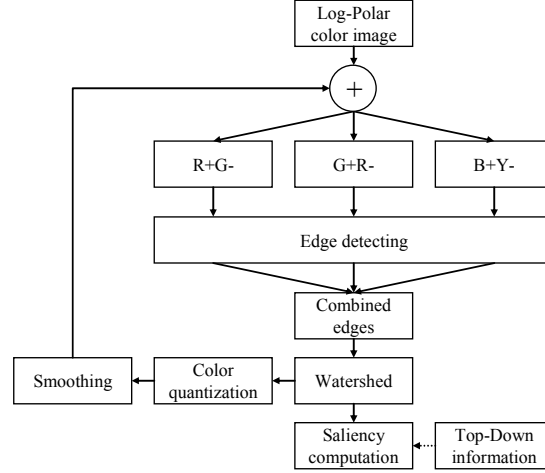


Figure 2. The visual attention system: block diagram (see text for details).

The radial symmetry of the distribution of the cones can be approximated by a polar distribution, whereas their projection to the primary visual cortex is well represented by a logarithmic-polar (log-polar) distribution mapped onto a rectangular-like surface (the cortex). Here the representation of the fovea is expanded several times, i.e. more neurons are devoted to the fovea while the periphery is represented at a coarser resolution.

From the mathematical point of view the log-polar mapping can be expressed as a transformation between the polar plane (ρ, θ) (retinal plane), the log-polar plane (ξ, η) (cortical plane) and the Cartesian plane (x, y) (image plane), as follows:²²

$$\begin{cases} \eta = q \cdot \theta, \\ \xi = \log_a \frac{\rho}{\rho_0}. \end{cases} \quad (1)$$

where ρ_0 is the radius of the innermost circle, $1/q$ is the minimum angular resolution of the log-polar layout and (ρ, θ) are the polar co-ordinates. These are related to the conventional Cartesian reference system by:

$$\begin{cases} x = \rho \cos \theta, \\ y = \rho \sin \theta. \end{cases} \quad (2)$$

Figure 3 illustrates the log-polar layout as derived from equation (1) and (2). Figure 3 shows a Cartesian image and its log-polar counterpart. It is worth noting that the flower's petals, that have a polar structure, are mapped horizontally in the log-polar image. Circles, on the other hand, are mapped vertically. Furthermore, the stamens that lie in the center of the image of the flower, occupy about half of the corresponding log-polar image (the cortical magnification).




Figure 3. Log-polar mapping. The original image (left) and the result of the log-polar mapping in the cortical plane (right).

3.2. Visual attention

As a first step the input image is smoothed, by taking the average between the current frame and the output of the color quantization (see later) on the previous frame. Then the red, green, and blue channels of each image are separated, and the yellow channel is calculated as the mean of the red and green one. These four channels are combined to generate three color opponent channels, similar to those of the retina. Each of these channels, typically indicated as (R+G-, G+R-, B+Y-), has a center-surround receptive field (RF) with spectrally opponent color responses. That is, for example, a red input in the center of a particular RF increases the response of the channel R+G-, while a green one in the surrounding decreases its response. The spatial response profile of the RF is expressed by a Difference-of-Gaussians (DoG) function. Each pixel is considered as the center of a RF, so that the output of the RF filtering is simply obtained by a convolution of the whole image with a DoG kernel, generating an output image of the same size of the input. This computation, considering for example the R+G- channel, is expressed by:

$$R^+G^-(\mathbf{x}) = a \cdot R(\mathbf{x}) \otimes \gamma_c(\mathbf{x}, \sigma_c) - b \cdot G(\mathbf{x}) \otimes \gamma_s(\mathbf{x}, \sigma_s). \quad (3)$$

The two Gaussian functions $\gamma_c(\mathbf{x}, \sigma_c)$ and $\gamma_s(\mathbf{x}, \sigma_s)$ are not balanced and the ratio b/a is 1.5, consistent with the study of Smirnakis et al.²⁴ Similarly to what happens in the human retina²⁵ the unbalanced ratio implicitly code the achromatic information. It is worth noting that filtering the log-polar images with a standard space-invariant filter corresponds to a space-variant filtering in the original Cartesian image.²⁶

Edges are then extracted on the three channels separately by employing a generalization of the Sobel filter due to Li et al.²⁷ The resulting edge maps are combined together to generate a single map as follows:

$$E(\mathbf{x}) = \max \{ \text{abs}(E_{RG}(\mathbf{x})), \text{abs}(E_{GR}(\mathbf{x})), \text{abs}(E_{BY}(\mathbf{x})) \}. \quad (4)$$

It has to be noted that the log-polar transform has the side effect of sharpening the edges near the fovea due to the already mentioned magnification factor. To compensate for this effect the edge map is multiplied by an exponential function, and normalized to a fixed range (0-255).

It has been speculated, that synchronizations of visual cortical neurons may serve as the carrier for the observed perceptual grouping phenomenon.^{28, 29} The differences in oscillator phase between spatially neighboring spiking cells could be used in principle to label different objects in the scene. We have used a watershed transform (rainfalling variant)^{30, 31} on the edge map to simulate the result of this synchronization and to generate the proto-objects. The activation is spread from the center of the image (in the edge map) until all spaces between edges are filled in. As a result the image is segmented into blobs with either uniform color or uniform gradient of color.

Each blob is then tagged with the mean color of the pixels within its internal area (this leads to a sort of quantized image). The result is blurred with a Gaussian filter and stored: it will be averaged with the next frame to obtain a temporal smoothing and reduce the effect of noise. After an initial startup delay of 4-5 frames, the number of blobs and their size stabilizes.

As discussed above, it is known that a feature or stimulus is salient if it differs from its immediate surrounding area. We chose to calculate the bottom-up saliency as the Euclidean distance in the color opponent space between each blob and the average color in a ball surrounding it. The radius of the ball (the spot or focus of attention) is not fixed: it changes with the size of the objects in the scene. In the same way the definition of “immediate surrounding area” should be relative to the size of the focus of attention. For this reason the greater part of the visual attention models in the literature uses a multi-scale approach and filters the saliency map with suitable filters, or “blob” detectors.³² These approaches lack continuity in the choice of the size of the attention focus. We propose instead to vary dynamically the region of interest depending on the size of the blobs. In other words, we compute the saliency of each blob in relation to a neighborhood region whose size is proportional to that of the blob itself. In our implementation we use a rectangular region 3 times the size of the bounding box of the blob. The choice of a rectangular window is not incidental, it was chosen because filters over rectangular regions can be computed efficiently by employing the integral image as in.³³ Blobs that are too small or too big are discarded from the saliency computation and will not be considered as possible candidates to be part of objects (proto-objects).

The bottom-up saliency is computed as:

$$S_{\text{bottom-up}} = \frac{1}{\sqrt{3}} \sqrt{\left(\left\langle R^+ G^- \right\rangle_{\text{blob}} - \left\langle R^+ G^- \right\rangle_{\text{surround}} \right)^2 + \left(\left\langle G^+ R^- \right\rangle_{\text{blob}} - \left\langle G^+ R^- \right\rangle_{\text{surround}} \right)^2 + \left(\left\langle B^+ Y^- \right\rangle_{\text{blob}} - \left\langle B^+ Y^- \right\rangle_{\text{surround}} \right)^2}. \quad (5)$$

where $\langle \rangle$ indicates the average of the pixel values over a certain area (as in the subscripts).

The top-down influence on attention is, at the moment, calculated in relation to the visual search task. When the robot has acquired a model of the object and begins searching for it, it uses the visual information of the object to bias the saliency map. In practice, the top-down saliency map is computed as the distance between the average color of each blob and that of the target:

$$S_{top-down} = 255 - \frac{1}{\sqrt{3}} \sqrt{\left(\left\langle \begin{matrix} R^+ \\ blob \end{matrix} G^- \right\rangle - \left\langle \begin{matrix} R^+ \\ object \end{matrix} G^- \right\rangle \right)^2 + \left(\left\langle \begin{matrix} G^+ \\ blob \end{matrix} R^- \right\rangle - \left\langle \begin{matrix} G^+ \\ object \end{matrix} R^- \right\rangle \right)^2 + \left(\left\langle \begin{matrix} B^+ \\ blob \end{matrix} Y^- \right\rangle - \left\langle \begin{matrix} B^+ \\ object \end{matrix} Y^- \right\rangle \right)^2}. \quad (6)$$

The total saliency is simply estimated as the linear combination of the two terms above:

$$S = \alpha \cdot S_{top-down} + \beta \cdot S_{bottom-up}. \quad (7)$$

The total saliency map S is eventually normalized in the range 0-255, as a consequence the saliency of each blob in the image is relative to the most salient one. The target of the next saccade is the center of mass of the most salient blob.

3.3. IOR

Local inhibition is transiently activated in the saliency map. This prevents the focus of attention to be redirected immediately to a location that was previously attended. Such an “inhibition of return” (IOR) has been demonstrated in human visual psychophysics.

Posner and Cohen,³⁴ for example, favor the hypothesis that the IOR does not function in retinal coordinates but it is instead represented in an allocentric reference frame. Together with Klein,³⁵ they proposed that the IOR is required to allow an efficient visual search by discouraging shifting the attention toward locations that have already been inspected. Static scenes, however, are seldom encountered in real life: objects move and a “tagging system” that merely inhibited environmental locations would be almost useless in any real situation. Tipper et al.³⁶ were among the firsts to demonstrate that the IOR could be attached to moving objects, and this finding has been replicated and extended ever since.³⁷⁻³⁹ These results bring to the conclusion that in humans the inhibition of return works by anchoring tags to objects as they move; in other words this process seems to be coded in an object-based reference frame.

Our system implements a simple object-based IOR. The robot maintains a list of the last five positions it has visited,⁴⁰ coded in a body centered coordinate system. The color information of the relative blobs is also stored in the list which is updated with a First-In First-Out policy. When the robot moves its gaze – for example by moving the eyes or the head in coordination – it keeps memory of the blobs it has visited earlier. Inhibition occurs only if the blob presents the same color that is stored in the list; in case the object moves or its color changes the location becomes available for fixation.

4. Learning about the Self

Internal models are thought to be available to the brain and responsible for formulating predictions about the world or simulating the body.⁴¹ In general the collection of the internal models required to represent the body is called the *body-schema*: it involves, for example, the relative positions of the limbs, and their weight and size. The existence of a body-schema in the brain has been proven, for example, by Graziano and colleagues^{42, 43} who found neurons in the primate's motor cortex (area 5) coding the position of the hand in the visual field. More precisely, these neurons had receptive fields anchored to the body that responded to visual stimuli and to the placement of the limbs at certain positions. Some of them were sensitive to visual cues only, and fired when the hand was visible in a particular region of space. The pure visual nature of these neurons was proven by a control experiment where the neural response was elicited with a mockup similar in appearance to the arm of the monkey. A second class of neurons, however, relied on proprioceptive cues only: that is, they responded even when the arm was invisible, and only if it was placed in a particular position (a barrier was used to occlude the view of the arm). Taken together these neurons represent an example of an internal model of the arm. The visual response corresponds to a model of the appearance of the arm (e.g. its shape, texture, color) whereas the proprioceptive model reflects an internal simulation of the spatial position of the arm. Although it is not clear how the two models are combined, it is clear that together they provide the brain with the ability to localize the arm in space with high accuracy and robustness to the environmental conditions (e.g. in dark vs. light). The fact that area 5 projects directly to the primary motor cortex and the spinal cord suggests also that this representation contributes to the control of limb movement.

In humans and biological systems the internal representation of the body is shaped during development and maintained adapted to the physical modification occurring in life. In artificial agents (where the body does not change with time) adaptation can spare the tedious operation of manually tuning the system's internal models and their calibration. The latter might be required to compensate changes in the visual appearance of the body or drift in the sensors (e.g. the motor encoders). In infants this sense of the body emerges a few months after birth; indeed experiments have shown that, for example, five-month-old infants are already able to recognize the movement of their own legs on a mirror.⁴⁴ However this ability is not present at birth but it is acquired during development.

This is a cause-effect problem because on the one hand the brain uses internal models to recognize the body whereas on the other it has to acquire the body-schema and maintain it up to date. To solve this problem, the brain needs a "bootstrapping" mechanism which allows the identification of the body and, in this way, the acquisition of the internal representation. To distinguish the body from the rest of the world the brain is thought to take advantage of extra information. For example, while a child waves the hand in front of his eyes, his brains "knows" what kind of motion is producing since it has exclusive access to the motor commands it sends to the muscles and the relative

proprioceptive feedback. Pattern similarities between this information and other sensory feedback (mainly vision) may allow the brain to identify the hand (or any other body part) and distinguish it from other entities that move differently. The identification of similarities between different sensory channels, that is the perception of *intermodal forms*, is a possible candidate for this purpose. Other factors could be used as well, like *timing* or time coincidence of events (two events happening at the same time are more likely to have been originated by the same source). However detection of intermodal forms seems to play a dominant role whereas timing has a more flexible contribution during development. In other words, events happening in a relatively long time window are often considered by the brain as if they were originated from the same cause. The reason for this is that, probably, coincidence in time is used to detect causalities at different time scales and link more complicated actions with their relative perceptual consequences (consider for example the action of switching on a neon light).⁴⁴

In robotics there have been attempts to replicate self-recognition mechanisms. Yoshikawa and colleagues⁴⁵ exploit the invariance of the body with respect to the external world to train a neural network to segment the arm of the robot. Their idea is that during learning, when the robot moves in the environment, the background changes, whereas the arms remain stationary with respect to the proprioceptive feedback.

Instead, the active behavior of the robot is used by Metta and Fitzpatrick;⁴⁶ in this case the robot identifies its body because it moves with respect to the background. Since motion alone is not sufficient to segment out external objects that move in the environment, the system seeks similarities between proprioceptive and visual feedback. Among the others, periodic actions may add robustness because offer the possibility to exploit repeatability.⁴⁷

4.1. *Segmentation of the hand*

We propose here an approach similar to Fitzpatrick and Arsenio⁴⁷ and Metta and Fitzpatrick⁴⁶ for visually segmenting the hand of the robot from the background. Repeated, self-generated actions were performed by the robot during the learning phase. In particular the robot was programmed to execute periodic movements of the wrist. The resulting motion of the hand was detected by computing image difference between the current frame and an adaptive model of the background. The period of motion of each pixel in the resulting motion image was then computed with a zero-crossing algorithm; similar information was extracted from the proprioceptive feedback of each motor encoder. As a result, the hand of the robot was segmented by selecting, among the pixels that moved periodically, those whose period matched that of the wrist joints. Conversely non-periodic pixels or pixels moving with different periods were identified as being externally originated and discarded. Figure 4 shows an example of the detection for two




Figure 4. Correlated versus uncorrelated motion, an example. The plots represent the time course of the variables involved in the detection procedure for two exemplar pixels whose motion matched (a) and did not match (b) that of the hand. (a1) and (b1) show the value of the motion for the pixel (normalized between 0 and 1). The result of the zero-crossing algorithm is reported in (a2) and (b2). The same procedure is replicated for the wrist proprioceptive feedback: (a3) and (b3) show the speed of the joint (normalized arbitrary scale), whereas (a4) and (b4) show the result of the zero-crossing algorithm. Compare (a2) to (a4) and (b2) to (b4).

different pixels whose motion was (a) correlated and (b) uncorrelated with that of the robot's hand. Low-pass filtering and a threshold was applied after the detection to obtain a dense segmented image (see Figure 5).

This algorithm forces the robot to stop and wait until the periodic movement of the wrist is performed. For this reason it is not useful during action or to drive a feedback control loop; it is instead ideally suited as a bootstrapping mechanism to acquire an internal model of the hand which can provide faster localization. In practice this was implemented with two neural networks: one trained to compute the position of the hand in the visual field given the current arm and head posture, and another to estimate the hand's shape and orientation (in this case the hand was represented as an ellipse). Indeed, these neural networks can also predict the expected location and the (simplified) appearance of the hand in the visual field given the current posture of the robot (its "felt" position).

4.2. *The hand internal model, expectation and prediction*

To gather the training data the robot moved the arm randomly and then waved the hand for a few seconds; for each spatial location the segmentation of the hand was performed as described in the previous section. For each trial the center of mass of the segmented area was computed along with the best fitting ellipse parameters. The complete algorithm




Figure 5. An example of the detection procedure. From left to right: the original image at the beginning of the procedure, the result of the detection (that is the pixels whose motion was correlated with that of the hand), the result of the low-pass filtering, the segmentation after the ellipse fitting. Notice that the ellipse tends to collapse towards the center, because the log-polar transformation gives more weight to the pixels close to the fovea.

is reported in Figure 6. The resulting (x, y) coordinates were used to train the first neural network whereas the ellipse parameters (orientation, major and minor axis) constituted the training samples for the second neural networks. It is important to take into account that the position of the hand in the visual field depends both on the posture of the arm and hand (this is not true, for example, for the orientation and size of the hand, if we do not consider translational effects). Unfortunately this enlarges the learning space and increases the time required for exploration (to collect the training set) and learning (higher dimensionality). For this reason the position of the hand was projected into an egocentric reference frame before being used to train the neural network. This last operation significantly reduced the dimensionality of the input space of the neural network. When necessary the output of the neural network is projected back to the retinocentric reference frame. Both projections (back and forth from egocentric

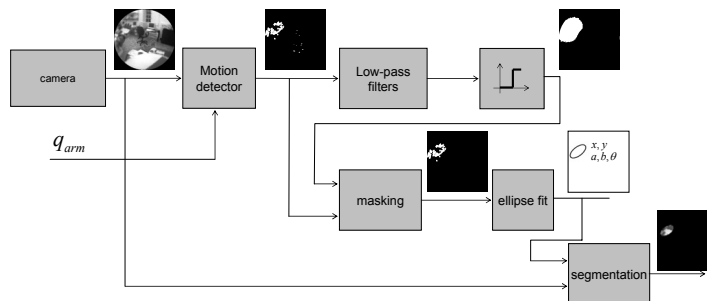


Figure 6. Detection algorithm, block schema. Images are captured from the camera. The “motion detector” block compares the motion in the image with the proprioceptive feedback from the arm (the wrist). A series of low-pass filters identify the blob which contains the hand. The blob is used to mask the result of the “motion detector” to remove possible outliers. An ellipse shape is fitted on the remaining pixels and, eventually, the hand is segmented.




Figure 7. Left: hand position predictor. Right: hand shape predictor.

and retinocentric reference frame) require knowledge of head inverse and direct kinematics. In the experiments reported here they were hardwired in the system, a possible procedure to learn a model of them is suggested by Arsenio⁴⁸. Figure 7 reports the block diagrams of the two models.

Learning was performed on-line by using the Schaal et al. model which is especially suited for incremental learning.⁴⁹ The learning process was validated by testing the ability of the network to predict new samples; when a new sample was obtained, and before being used to perform a learning step, the network was used to predict the output given the input. The resulting output was compared to the current sample and the error computed. The increasing ability of the network to predict new samples proved that learning was effective. Figure 8 (left) reports the plot of the error during an experiment (in this case the error is computed in the image plane to simplify understanding).

At the end of the exploration phase the robot had trained an internal model of the hand by which it could i) localize its center of mass ii) estimate its orientation and approximate size. The output of these models is not based on actual visual feedback, but on the mere projection of the proprioceptive information about the hand: they represent the expectation the robot possesses about its body (in this case, the hand). These

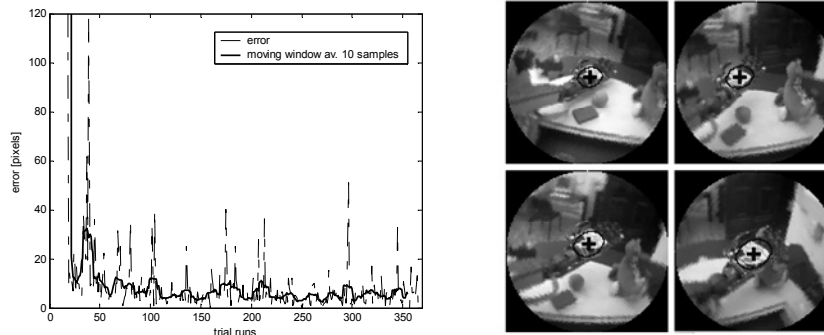


Figure 8. Hand localization error trend (left). As new examples are presented to the network the performance improves. Example of the localization after learning (right). The cross corresponds to the position of the hand, whereas the ellipse represents its approximate shape and orientation.

measures were used in numerous ways. The center of mass was employed to close a visual loop to direct gaze towards the hand (see Figure 8 right). For this task the internal model was addressed with the proprioceptive feedback of the arm. Another possibility was to address the model with the arm motor command (final joint position) to obtain where the hand would be at the end of the movement. In general this model offers a means of computing a prediction of the position, size and orientation of the hand from a given arm configuration or, in other words, of simulating a motor action. In the next section this will be used to learn the reaching map and estimate the visuomotor Jacobian matrix for a reaching task.

4.3. *Reaching*

The solution we propose is based on the use of a direct mapping between the eye-head motor plant and the arm motor plant.⁵⁰ Flanders and colleagues⁵¹ suggested that the information about gaze direction might be employed by the brain to establish a reference point for reaching. They analyzed the error when reaching in the dark and showed how this correlates to the error of the gaze (the gaze drifts away from the target in the dark). Accordingly one premise we make is that the position of the fixation point coincides with the object to be reached. In other words, reaching for an object starts by looking at it. Under this assumption, the fixation point can be considered as the “end-effector” of the eye-head system. The position of the eyes with respect to the head, determines uniquely the position of the fixation point in space relative to the shoulder. The arm motor command can be obtained by a transformation of the eye-head motor/positional variables. We called this approach “motor-motor coordination”, because the coordinated action is obtained by mapping motor variables into motor variables:

$$q_{arm} = f(q_{head}). \quad (8)$$

where q_{head} and q_{arm} are head and arm posture respectively (joint space).

What is interesting in this approach is not equation (8) per se, which, after all, implements the inverse kinematics of the arm, but the mechanisms used to learn it. In fact, this mapping can be easily learnt when the tracking behavior described in the previous section is active. The robot explored the workspace by moving the arm randomly, while simultaneously, it tracked its hand; whenever the eyes fixated the hand a new sample consisting of the arm and head joint angles was acquired and used to train a neural network approximating equation (8). The exploration was conducted in two ways. A first movement of the arm was performed by sampling a random uniform distribution within the part of the arm workspace in front of the robot. Small subsequent movements were performed randomly with Gaussian distribution with zero mean and standard deviation equal to 5 degrees. This last step while not strictly required sped up learning by

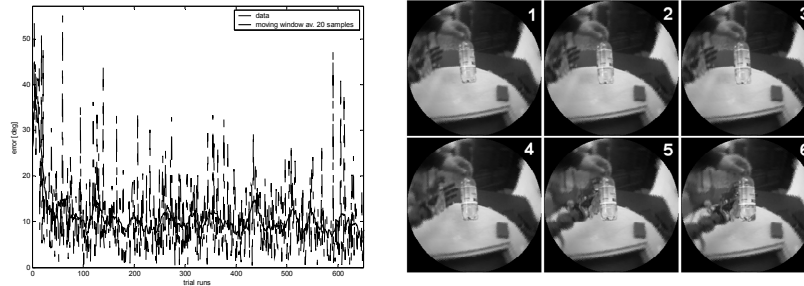


Figure 9. Reaching error (left). As new examples are gathered and presented to the network the performance increases. This improvement is less remarkable; we believe this is due to noise in the training data which affects not only learning, but also the measure of performance. An exemplar sequence of a reaching action after the learning is reported on the right.

sampling quickly large portions of the arm’s workspace: i.e. for small movements of the order of 5 degrees the arm fixation was achieved rapidly and thus a new sample was added to the training set. When a sufficient number of samples were acquired, the robot started using the motor-motor map to actively reach for visually identified objects while learning could continue.

Learning can be further improved by reducing the dimensionality of the input vector q_{head} . In fact, only three variables are needed to code the position of the fixation point; for this purpose we decided to use azimuth, elevation, and distance – in substitution for the five angles of the head joints. This transformation is motivated by practical reasons, but it is also biologically plausible.⁵²

Similarly to the previous section, learning was tested by comparing every new sample to the output of the network (see Section 4.2 for details). The graph of the error during learning is reported in Figure 9 (left) for each sample (dotted line) and the moving window average over 20 samples. From the first plot it is hard to determine a real increment of performance as several samples at the end of the learning session present relatively large errors. This is due to noise in the training data, which affects not only learning, but also the measure of performance. In particular noise is higher in those configurations of the arm where the hand is closer to the head and the system fails to control the angle of vergence between the eyes. In these situations the error is large because the position of the fixation points varies significantly (from very far to very close). The average error, however, has a distinguishable uniform trend. Figure 9 (right) shows a sequence of images taken from the robot left eye during an exemplar reaching action.

It is worth mentioning that there is no need to separate the exploration/training phase and reaching (exploitation). An initial “reflex” can be employed as substitute for the reaching map at the very beginning; this simple behavior could, for example, populate the robot workspace with three positions (left, center and right). Exploration in this case

would still be guaranteed by a random procedure, similar to the one described earlier. This approach was followed by Metta et al. and Metta.^{50, 53}

The reaching problem can also be solved in the image plane. Consider the planar case (i.e. no 3D information is available) and suppose to measure the position of the end point in the image plane (x, y) . We want to control the arm to reach a target point (x^*, y^*) . We can solve the problem by means of a closed loop controller, by following a fairly standard visual servoing approach:

$$\Delta \mathbf{q} = -k \cdot \mathbf{J}(\mathbf{q}) \Delta \mathbf{x}. \quad (9)$$

where:

$$\Delta \mathbf{x} = \begin{bmatrix} x - x^* \\ y - y^* \end{bmatrix}. \quad (10)$$

$k > 0$ is a scalar and $\mathbf{J}(\mathbf{q})$ is the Jacobian of the transformation between the image plane and the arm joint space. For a 2 dimensional arm $\mathbf{J}(\mathbf{q})$ is a 2 by 2 matrix whose elements are a non-linear function of the arm joint angles. Given the Jacobian of the manipulator it is possible to drive the endpoint toward any point in the image plane. At least locally, the Jacobian can be approximated by a constant matrix. In our case:

$$\Delta \mathbf{q} \approx -k \cdot \hat{\mathbf{J}}(\bar{\mathbf{q}}) \Delta \mathbf{x} = -k \cdot \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \Delta \mathbf{x}. \quad (11)$$

Convergence is guaranteed if the following condition is met:

$$J^{-1}(q) \hat{J}(\bar{q}) > 0. \quad (12)$$

Since following the procedure of Section 4.2 the robot has learnt a direct transformation between the arm joint angles and the image plane (see for example Section 4.2, Figure 7), it can now recover the position of the endpoint in the image plane from a given joint configuration:

$$\begin{bmatrix} x \\ y \end{bmatrix} = f(q). \quad (13)$$

Indeed, to compute a local approximation of the Jacobian, a random sampling of the arm joint space around a given point $(\bar{\mathbf{x}}, \bar{\mathbf{q}})$ can be performed:

$$\mathbf{q}_i = \bar{\mathbf{q}} + \Delta \mathbf{q}_i. \quad (14)$$

with

$$\Delta \mathbf{q}_i = \boldsymbol{\eta}(\mathbf{0}, \boldsymbol{\sigma}). \quad (15)$$

and where $\boldsymbol{\eta}(\mathbf{0}, \boldsymbol{\sigma})$ follows a normal distribution of zero mean and standard deviation of 5 degrees.

For each sample, by applying equation (13) we obtain a new value $\mathbf{x}_i = \bar{\mathbf{x}} + \Delta\mathbf{x}_i$ that can be used to estimate the Jacobian around \mathbf{q} with a least squares procedure:

$$\Delta\mathbf{q}_i = \begin{bmatrix} \Delta\mathbf{x}_i^T & \mathbf{0} \\ \mathbf{0} & \Delta\mathbf{x}_i^T \end{bmatrix} \cdot \begin{bmatrix} a_{11} \\ a_{12} \\ a_{21} \\ a_{22} \end{bmatrix}. \quad (16)$$

$\hat{\mathbf{J}}(\bar{\mathbf{q}})$ can then be used in the closed loop controller to drive the arm toward a specific position in the image plane. However, there is no need to close the loop with the actual visual feedback. By using the map in equation (13), in fact, we can substitute the actual visual feedback with the internal simulation provided by this model. From the output of the closed loop controller we can estimate the position of the arm at the next step, by assuming a pure kinematic model of the arm; in this way the procedure can be iterated several times to obtain the joint motor commands required to perform a reaching movement. The flowchart below explains this procedure.

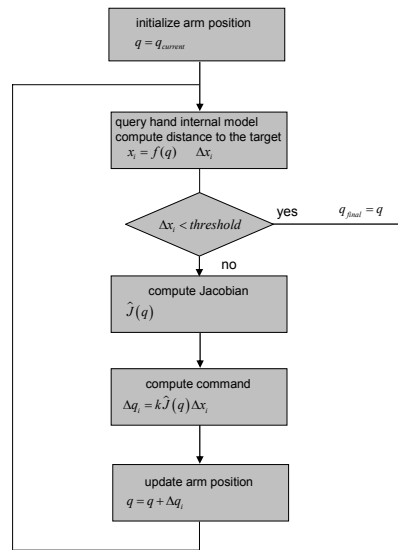


Figure 10. Closed-loop approach to reaching, flowchart. See text for further details.




Figure 11. Arm trajectories for two reaching actions (a) and (b). T0 marks the position of the hand at the beginning of the action. Crosses correspond to the position of the palm; circles show the position of the fingers. The action is divided in three phases. From T0 to T1 arm prepositioning. From T1 to T2, reaching: in this case the motor-motor map is used to move the palm towards the center of the visual field (the target). A small adjustment with the arm Jacobian is performed to position the fingers on the target (T2 to T3).

The main limitation of this approach is that we do not make use of three-dimensional visual information; while this is a clear limitation of this implementation, the same approach can be easily extended to the full 3D case. The implementation is consistent with the hand internal model which provides the position of the hand in the image plane of one of the eyes only (left). Since in the Babybot the hand position is uniquely described by three degrees of freedom (the first three joints of the Puma arm), this technique was used to control only two of them (arm and forearm). Given the kinematics of the Puma arm this allowed to perform movements on the plane defined by the shoulder joint. Another point worth discussing is that the closed loop controller does not use real visual feedback, and, therefore, its accuracy depends on the precision of the hand internal model. To achieve better performances, visual feedback might be required.

Let us summarize what we have described in this section. We have introduced two approaches to solving the inverse kinematics of the manipulator. The first method uses a mapping between the posture of the head (whose fixation point implicitly identifies the target) and the arm motor commands; it allows controlling the arm to reach any point fixated by the robot.^a The second approach uses the hand internal model to compute a piecewise constant approximation of the inverse Jacobian and simulate small movements of the arm in the neighborhood of the desired target. The procedure is iterated several times to compute the motor command required for reaching the target. Reaching in this

^a During the learning of the motor-motor map, the robot tracks the palm of the hand.

case is planned in the image plane; however, since the internal model is two dimensional, the approach is limited to the plane identified by the shoulder. For these reasons, the two methods were mixed in the experiment reported in the next section. The motor-motor mapping is employed to plan a first gross movement to approach the target, whereas the “closed-loop approach” allows a finer positioning of the fingers on the target. This second part of the movement is planned by considering the point of the ellipse at maximum distance from the robot’s body (which corresponds to the fingers) as the arm endpoint (Figure 11). This strategy proved successful because it substantially increased the probability to grasp the objects on the table.

Once the robot has computed the final arm posture it is still required to plan the actual movement. This was done with a simple linear interpolation between the current and final arm configuration. The trajectory was divided in steps which were then effected by the low level controller; to this purpose we employed a low-stiffness PD controller with gravity compensation. The gravity load term for each joint was learnt online as described in Natale.⁸

5. Learning about Objects

In this section we describe a method for building a model of the object the robot grasps. We assume for a moment that the robot has already grasped an object; this can happen because a collaborative human has given the object to the robot (as we describe in the next section) or because the robot has autonomously grasped the object. In this case the robot may spot a region of interest in the visual scene and apply a stereotyped action with the arm and the hand to catch it. Both solutions are valid bootstrapping behaviors for the acquisition of an internal model of the object. When the robot holds the object it can explore it by moving and rotating it.

In short, the idea is to represent objects as collections of blobs generated by the visual attention system and their relative positions (neighboring relations). The model is created statistically by looking at the same object several times from different points of view. At the same time the system estimates the probability that each blob belongs to the object by counting the number of times each blob appears during the exploration.

In the following, we use the probabilistic framework proposed by Schiele and Crowley.^{54, 55} We want to calculate the probability of the object O given a certain local measurement M . This probability $P(O|M)$ can be calculated using Bayes’ formula:

$$P(O|M) = \frac{P(M|O)P(O)}{P(M)}, \quad (17)$$

$$O_{MAP} = \arg \max_{O, \sim O} \{P(O|M), P(\sim O|M)\}.$$

where $P(O)$ is the *a priori* probability of the object O , $P(M)$ the *a priori* probability of the local measurement M , and $P(M|O)$ is the probability of the local measurement M when

the object O is fixated. In the following experiments we carried out only a single detection experiment, there are consequently only two classes, one representing the object and another representing the background. $P(O)$ and $P(\sim O)$ are simply set to 0.5 because this choice does not affect the maximization.

Since a single blob is not discriminative enough, we considered the probabilities of observing pairs of blobs; the local measurement M becomes the event of observing both a central (i.e. fixated) and surrounding blobs:

$$P(M | O) = P(B_i | B_c \text{ and } (B_i \text{ adjacent } B_c)). \quad (18)$$

where B_i is the i^{th} blob surrounding the central blob B_c which belongs to the object O . That is, we exploit the fact the robot is fixating the object and assume B_c to be constant across fixations of the same object – this is guaranteed by the fact the object is being held by the hand. In practice this corresponds to estimating the probability that all blobs B_i adjacent to B_c (which we take as a reference) belong to the object. Moreover the color of the central blob B_c will be stored to be used during visual search to bias the salience map. This procedure, although requiring the “active participation” of the robot (through gazing) is less computationally expensive compared to the estimation of all probabilities for all possible pairs of blobs of the fixated object. Estimation of the full joint probabilities would require a larger training set than the one we used in our experiments. The probabilities $P(M|\sim O)$ are estimated during the exploration phase with the blobs not adjacent to the central blob. The local measurements were considered independent, because they refer to different blobs, so the total probability $P(M_1, \dots, M_N | O)$ can be factorized in the product of the probabilities $P(M_i | O)$. An object is detected if the probability $P(O | M_1, \dots, M_N)$ is greater than a fixed threshold.

Our requirement was that of building the object model with the shortest possible exploration procedure. Unfortunately, the small training set might give histograms $P(M|*)$ with many empty bins zero counts bins. To overcome this problem a probability smoothing method was used. A popular method of zero smoothing is Lidstone’s law of succession:⁵⁶

$$P(M | O) = \frac{\text{count}(M \wedge O) + \lambda}{\text{count}(O) + v\lambda}. \quad (19)$$

for a v valued problem. With $\lambda=1$ and a two valued problem ($v=2$), we obtain the well-known Laplace’s law of succession. Following the results of Kohavi et al.,⁵⁷ we choose $\lambda=1/n$ where n is equal to the number of frames utilized during the training. Then our probability estimator becomes:

$$P(M | O) = \frac{\text{count}(M \wedge O) + 1/n}{\text{count}(O) + v/n}. \quad (20)$$

When an object is detected after visual search, a possible figure-ground segmentation is attempted, using the information gathered during the exploration phase. Each blob is

segmented from the background if it is adjacent to the central blob and if its probability to belong to the object is greater than 0.5. This probability is approximated using the estimated probability as follows:

$$P(B_i \in O | B_c \text{ and } (B_i \text{ adjacent } B_c)) \approx P(B_i | B_c \text{ and } (B_i \text{ adjacent } B_c)). \quad (21)$$

As an example Figure 13 shows the result of the segmentation procedure.

In table 1, results are shown of using a toy car and a toy airplane as target objects; 50 training sessions were performed for each object. The first column shows the recognition rate, the second the average number of saccades (mean \pm standard deviations) it takes the robot to locate the target in case of successful recognition.

Table 1. Performance of the recognition system measured from a set of 50 trials.

Object	Recognition rate	Number of saccades when recognized
Toy car	94%	3.19 \pm 2.17
Toy airplane	88%	3.02 \pm 2.84




Figure 12. Object exploration and corresponding segmentation 1-3 and 4-6 respectively. The segmentation consists in the object central blob together with the relative adjacent ones. Notice that fixation is maintained on the object by using the hand localization module as explained in Section 4.




Figure 13. Visual search. The robot has acquired a model of the airplane toy during an exploration phase (not shown); this information primes the attention system. The blue blob at the center of the airplane is selected and a saccade performed. (a) and (b) show the visual scene before and after the saccade. (d) and (e) show the output of the visual attention system synchronized with (a) and (b) respectively. The result of the segmentation after the saccade is in (c).

6. A Grasping Behavior: where Everything Comes to a Sense

The modules described in the previous sections can be integrated to achieve an autonomous grasping behavior. Figure 14 can be used as a reference for the following discussion. The action starts when an object is placed in the robot's hand and the robot detects pressure in the palm (frame 1). This elicits a clutching action of the fingers; the hand follows a preprogrammed trajectory, the fingers bend around the object toward the palm. If the object is of some appropriate size, the intrinsic elasticity of the hand (as described in Section 2) facilitates the action and the grasping of the object. The robot moves the arm to bring the object close to the cameras and begins its exploration. The object is placed in four positions with different orientations and background (frames between 2 and 6). During the exploration, the robot tracks the hand/object; when the object is stationary and fixation is achieved, a few frames are acquired and the model of the object trained as explained in Section 5. At the end of the exploration the object is released (frame 4). At this point the robot has acquired the visual model of the object and starts searching for it in the visual scene. To do this, it selects the blob whose features better match those of the object's main blob and perform a saccade. After the saccade the model of the object is matched against the blob that is being fixated and its surrounding. If the match is not positive search continues with another blob, otherwise grasping starts (frames 7-8-9). At the end of the grasp the robot uses haptic information to detect whether it is holding the object or the action failed. In this process the weight of the object and its consistence in the hand is checked (the shape of the fingers holding the

object). If the action is successful the robot waits for another object, otherwise it performs another trial (search and reach).

It is fair to say that part of the controller was preprogrammed. The hand was controlled with stereotyped motor commands. Three primitives were used: one to close the hand after pressure was detected, and two during the grasping to pre-shape the hand and actually clasp the object. The robot relied on the elasticity of the hand to achieve the correct grasping. To facilitate grasping, the trajectory of the arm was also programmed beforehand; waypoints relative to the final position of the arm were included in the joint space to approach the object from the top.




Figure 14. A sequence of the robot grasping an object. The action starts when an object is placed on the palm (1). The robot grasps the object and moves the eyes to fixate the hand (2). The exploration starts in (3) when the robot brings the object close to the camera. The object is moved in four different positions while maintaining fixation; at the same time the object model is trained (3-6). The robot drops the object and starts searching for it (7). The object is identified and a saccade performed (7-9). The robot eventually grasps the toy (10-12).

7. Discussion and Conclusions

In this paper we have presented a developmental approach to the realization of cognitive abilities in a humanoid robot which starts from the exploration of the body and unfolds by eventually exploring the external world. The robot starts from a limited set of initial motor and perceptual competencies and autonomously develops more sophisticated ways to interact with the environment. This knowledge is used to begin the exploration of the environment and build a visual model of the objects that are grasped.

We have presented an implementation of a visual attention system properly taking into account top-down and bottom-up information. The top-down system divides the visual scene into color blobs; each blob is assigned a saliency depending on the ratio between its color and the color of the area surrounding it. The robot actively explores the visual appearance of the objects it grasps: every time an object is placed on the palm a statistical model of the blobs that are part of it is constructed. This information is subsequently fed to the attention system as a bottom-up primer to control the visual search of the same object. Thus the robot experience allows it to build a representation of the object with which it interacts while, at the same time, modulates the visual attention system. The robot's ability to act is used together with the body internal model to drive the exploration of the environment. This facilitates learning in different ways. Firstly it helps the robot to focus attention both in space and in time. During the acquisition of the object visual model, in fact, the robot can track the object because it knows the position of the hand from its proprioceptive feedback. The latter is also useful to detect when the acquisition of the model can be initiated because the object does not move and the eyes have acquired a stable fixation on it. Finally, the fact that the object is being held by the hand guarantees the link between different sensory modalities (for example the sight of the object and the kinesthetic information from the hand). The object model makes use of visual information; in Natale et al.⁵⁸ we show how it is possible to build a model of the objects based only on haptic information. In the future we would like to investigate the integration of the two approaches.

We support the enactive view of cognition in showing how much the body and the ability to build the representation of the external world through the interaction between the body and the environment can be useful for an autonomous agent. Even a simple set of behaviors (such as the one initially provided to the robot) is sufficient to begin the exploration of the environment and acquire an internal representation of it. On the other hand it is fair to say that much of the system presented in this paper is still "cognitivist" and more or less carefully handcrafted into the robot. For practical reasons, our implementation lays in between a full emergent and a cognitivist approach although biologically informed choices were made when possible.

We have also shown how this initial body-environment interaction is sufficient to start linking actions with their resulting consequences to form prediction about the behavior of the robot. Very often prospective control is required to plan a successful action. During grasping, for example, the correct timing of preshaping and closure of the fingers is required; the lags in the sensory streams (visual and tactile) typical of artificial and natural systems make feedback control ineffective. To be able to anticipate the impact of the hand with the object, the robot is required to control the timing between preshaping and actual grasping; clearly this cannot be based only on visual and tactile feedback. Prospective control, however, is not only important for action. It gives an agent the possibility to create expectations on which to base the interpretation of the world and the actions performed by others. By means of the interaction with the world the agent builds a model of the behavior of external entities (objects, people, etc.) and the

associated sensory feedback. This link can be used afterward to anticipate the consequences of a similar action and, eventually, to compare them with the real feedback. In the same way new situations can be interpreted by matching them against the robot's past experience. For example, the event of a ball that falls on the floor (and the resulting visual and auditory sensations) can be associated to the action of dropping it. Anticipation and predictions enhance the agents' ability to understand and interact with the environment and, for this reason, are important aspects of cognition. The results of this paper represent the first steps into the implementation of cognitive abilities in an artificial system. It is difficult to think, at least from an emergent perspective, of a shortcut that prescinds from sensorimotor coordination in achieving cognitive skills to be used in the real world.

To conclude, we would like to comment on the effort required to build a complete robotic platform on the one hand, and the software architecture on the other. Presently the Babybot is an integrated robotic platform where it is extremely easy for software modules controlling different subparts (arm, head or hand to mention just a few) to exchange information and coordinate with each other.⁵⁹ This is not very common, as usually in the literature papers report single experiments where the robotic platform is specifically programmed to perform the desired task, but care is not taken to realize a system which can grow in complexity as new modules are added. The experiment reported in Section 6 does not only show the integration between the visual attention system and the motor system but also the complexity of the system as a whole. We believe that this is a necessary prerequisite to carry out research in humanoid robotics as the complexity and number of skills increase.

Acknowledgments

This work is funded by the European Commission's Cognition Unit, Directorate-General Information Society, as part of project no. IST-2004-004370: RobotCub — ROBotic Open-architecture Technology for Cognition, Understanding, and Behaviour and by the European Commission Future and Emerging Technologies project ADAPT project no. IST-2001-37173.

References

1. *Unit E5: Cognition*, http://www.cordis.lu/ist/directorate_e/cognition/index.htm.
2. A. Clark, *Mindware: an introduction to the philosophy of cognitive science*, (Oxford University Press, Oxford, UK, 2001).
3. R. Beer, D., Dynamical approaches to cognitive science, *TRENDS in Cognitive Sciences*, 4(3) (2000) 91-99.
4. H. Maturana, R. and F. Varela, J., *The tree of knowledge, the biological roots of human understanding*, (Shambhala Publications, Inc., Boston & London, 1998).

5. G. Sandini, G. Metta, and D. Vernon, RobotCub: An Open Framework for Research in Embodied Cognition, in: *IEEE-RAS/RJS International Conference on Humanoid Robotics*, (IEEE, Santa Monica, CA, 2004).
6. L. Natale, Rao S., and G. Sandini, Learning to act on objects, in: *Second International Workshop, BMCV 2002*, (Springer, Tubingen, Germany, 2002), p. 567-575.
7. P. Fitzpatrick, G. Metta, L. Natale, S. Rao, and G. Sandini, Learning About Objects Through Action: Initial Steps Towards Artificial Cognition, in: *IEEE International Conference on Robotics and Automation (ICRA 2003)*, (Taipei, Taiwan, 2003).
8. L. Natale, *Linking action to perception in a humanoid robot: a developmental approach to grasping*, PhD Thesis, (DIST, University Of Genoa, Genova, 2004)
9. J. Aloimonos, I. Weiss, and A. Bandyopadhyay, Active Vision, *International Journal of Computer Vision*, 1(4) (1988) 333-356.
10. D. H. Ballard and C. M. Brown, Principles of Animate Vision, *Computer Vision Graphics and Image Processing*, 56(1) (1992) 3-21.
11. M. I. Posner, C. R. R. Snyder, and B. J. Davidson, Attention and the detection of signals, *Journal of Experimental Psychology: General*, 109 (1980) 160-174.
12. C. Downing and S. Pinker, *The spatial structure of visual attention*, in: *Attention and Performance*, Vol. XI, (London, 1985), pp. 171-187.
13. B. J. Scholl, Objects and attention: the state of the art, *Cognition*, 80 (2001) 1-46.
14. R. Egly, J. Driver, and R. Rafal, Shifting visual attention between objects and locations: evidence for normal and parietal subjects, *Journal of Experimental Psychology: General*, 123 (1994) 161-177.
15. K. O'Craven, P. Downing, and N. Kanwisher, fMRI evidence for objects as units of attentional selection, *Nature*, 401 (1999) 584-587.
16. R. Milanese, *Detecting Salient Regions in an Image: From Biological Evidence to Computer Implementation*, PhD Thesis, (Department of Computer Science, University of Geneva, Geneva, 1993)
17. L. Itti, C. Koch, and E. Niebur, A Model of Saliency-Based Visual Attention for Rapid Scene Analysis, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(11) (1998) 1254-1259.
18. G. J. Giefing, H. Janssen, and H. A. Mallot, Saccadic object recognition with an active vision system, in: *International Conference on Pattern Recognition*, (1992), p. 664-667.
19. A. M. Treisman and G. Gelade, A feature integration theory of attention, *Cognitive Psychology*, 12 (1980) 97-136.
20. Y. Sun and R. Fisher, Object-based visual attention for computer vision, *Artificial Intelligence*, 146 (2003) 77-123.
21. Z. Pylyshyn, Visual indexes, preconceptual object, and situated vision, *Cognition*, 80 (2001) 127-158.
22. G. Sandini and V. Tagliasco, An Anthropomorphic Retina-like Structure for Scene Analysis, *Computer Vision, Graphics and Image Processing*, 14(3) (1980) 365-372.
23. J. M. Wolfe and G. Gancarz, *Guided Search 3.0 Basic and Clinical Applications of Vision Science*, in: *Basic and Clinical Applications of Vision Science*, (Kluwer Academic, Dordrecht, Netherlands, 1996), pp. 189-192.
24. S. M. Smirnakis, M. J. Berry, D. K. Warland, W. Bialek, and M. Meister, Adaptation of retinal processing to image contrast and spatial scale, *Nature*, 386 (1997) 69-73.

25. V. A. Billock, Cortical Simple Cells Can Extract Achromatic Information from the Multiplexed Chromatic and Achromatic Signals in the Parvocellular Pathway, *Vision Research*, 35(16) (1995) 2359-2369.
26. W. von Seelen and H. A. Mallot, Neural Mapping and Space-Variant Image Processing, in: *International Joint Conference on Neural Networks (IJCNN)*, (San Diego, CA, 1990), p. 417 - 422.
27. X. Li, T. Yuan, N. Yu, and Y. Yuan, Adaptive color quantization based on perceptive edge protection, *Pattern Recognition Letters*, 24 (2003) 3165-3176.
28. R. Eckhorn, R. Bauer, W. Jordan, M. Brosch, M. Kruse, W. Munk, and H. J. Reitboeck., Coherent oscillations: A mechanism of feature linking in the visual cortex?, *Biological Cybernetics*, 60 (1988) 121-130.
29. C. M. Gray, P. König, A. K. Engel, and W. Singer, Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties, *Nature*, 338 (1989) 334-336.
30. P. D. Smet and R. Pires, Implementation and analysis of an optimized rainfalling watershed algorithm, in: *IS&T/SPIE's 12th Annual Symposium Electronic Imaging 2000: Science and Technology*, (San Jose, California, USA, 2000).
31. L. Vincent and P. Soille, Watersheds in digital spaces: an efficient algorithm based on immersion simulations, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(6) (1991) 583-598.
32. L. Itti and C. Koch, Computational modelling of visual attention, *Nature Reviews Neuroscience*, 2(3) (2001) 194-203.
33. P. Viola and M. J. Jones, Robust Real-Time Face Detection, *International Journal of Computer Vision*, 57(2) (2004) 137-154.
34. M. I. Posner and Y. Cohen, *Components of visual orienting*, in: Attention and Performance, Vol. X, (Erlbaum, Hillsdale, NJ, 1984), pp. 531-556.
35. R. M. Klein, Inhibitory tagging system facilitates visual search, *Nature*, 334 (1988) 430-431.
36. S. P. Tipper, Object-centred inhibition of return of visual attention, *Quarterly Journal of Experimental Psychology*, 43A (1991) 289-298.
37. R. A. Abrams and R. S. Dobkin, Inhibition of return: effects of attentional cuing on eye movement latencies, *Journal of Experimental Psychology: Human Perception and Performance*, 20(3) (1994) 467-477.
38. B. S. Gibson and H. Egeth, Inhibition of return to object-based and environment-based locations, *Perception & Psychophysics*, 55(3) (1994) 323-339.
39. S. P. Tipper, Object-based and environment-based inhibition of return of visual attention, *Journal of Experimental Psychology: Human Perception and Performance*, 20(3) (1994) 478-499.
40. J. M. Wolfe, Moving towards solutions to some enduring controversies in visual search, *TRENDS in Cognitive Sciences*, 7(2) (2003) 70-76.
41. D. M. Wolpert and R. C. Miall, Forward models for physiological motor control, *Neural Networks*, 9(8) (1996) 1265-1279.
42. M. S. A. Graziano, Where is my arm? The relative role of vision and proprioception in the neuronal representation of limb position, *Proceedings of the National Academy of Science*, 96 (1999) 10418-10421.
43. M. S. A. Graziano, D. F. Cooke, and C. S. R. Taylor, Coding the location of the arm by sight, *Science*, 290 (2000) 1782-1786.

44. P. Rochat and T. Striano, Perceived self in infancy, *Infant Behavior & Development*, 23 (2000) 513-530.
45. Y. Yoshikawa, K. Hosoda, and M. Asada, Does the invariance in multi-modalities represent the body scheme ? - a case study with vision and proprioception -, in: *2nd Intelligent Symposium on Adaptive Motion of Animals and Machines*, (Kyoto, Japan, 2003).
46. G. Metta and P. Fitzpatrick, Early Integration of Vision and Manipulation, *Adaptive Behavior*, 11(2) (2003) 109-128.
47. P. Fitzpatrick and A. Arsenio, Feel the beat: using cross-modal rhythm to integrate perception of objects, others and self, in: *Fourth International Workshop on Epigenetic Robotics*, (Lund University Cognitive Studies, Genoa, 2004).
48. A. Arsenio, *Cognitive-Developmental Learning for a Humanoid Robot: A Caregiver's Gift*, PhD Thesis, (CSAIL, MIT, Boston, 2004)
49. S. Schaal and C. G. Atkeson, Constructive Incremental Learning from Only Local Information, *Neural Computation*, (10) (1998) 2047-2084.
50. G. Metta, G. Sandini, and J. Konczak, A Developmental Approach to Visually-Guided Reaching in Artificial Systems, *Neural Networks*, 12(10) (1999) 1413-1427.
51. M. Flanders, L. Daghestani, and A. Berthoz, Reaching beyond reach, *Experimental Brain Research*, 126(1) (1999) 19-30.
52. F. Lacquaniti and R. Caminiti, Visuo-motor transformations for arm reaching, *European Journal of Neuroscience*, 10 (1998) 195-203.
53. G. Metta, *Babybot: a Study on Sensori-motor Development*, Ph.D. Thesis Thesis, (DIST, University of Genova, Genova, 2000)
54. B. Schiele and J. L. Crowley, Probabilistic object recognition using multidimensional receptive field histograms, in: *13th International Conference on Pattern Recognition*, (Vienna, Austria, 1996), p. 50-54.
55. B. Schiele and J. L. Crowley, Where to look next and what to look for, in: *IEEE/RSJ International Conference on Intelligent Robots and Systems*, (Osaka, 1996).
56. G. Lidstone, Note on the general case of the Bayes-Laplace formula for inductive or a posteriori probabilities, *Transactions of the Faculty of Actuaries*, 8 (1920) 182-192.
57. R. Kohavi, B. Becker, and D. Sommerfield, Improving simple Bayes, in: *European Conference on Machine Learning*, (1997).
58. L. Natale, G. Metta, and G. Sandini, Learning haptic representation of objects, in: *International Conference on Intelligent Manipulation and Grasping*, (Genoa, Italy, 2004).
59. P. Fitzpatrick, G. Metta, and L. Natale, *YARP: Yet Another Robotic Platform website*, <http://yarp0.sf.net>.

From sensorimotor development to object perception

Lorenzo Natale, Francesco Orabona, Fabio Berton, Giorgio Metta, Giulio Sandini

LIRA-Lab, DIST

University of Genoa

Viale Causa 13, 16145, Genoa, Italy

Email: {nat, bremen, fberton, pasa, sandini}@liralab.it

Abstract—This paper describes a developmental sequence that allows a humanoid robot to learn about the shape of its body and successively about certain parts of the environment. We equipped the humanoid robot with an initial set of motor and perceptual competencies ranging from simple stereotyped actions to more sophisticated visual routines providing a bottom-up attention system. This initial form of sensorimotor coordination is sufficient to initiate the interaction with the environment and allows the robot to improve its motor and perceptual skills by first constructing a “body-schema” and later by learning about objects. The body-schema allows controlling movements to fixate, reach and touch objects in the environment. The interaction is further used to form a visual model of the objects grasped by the robot which eventually modulate the attention system in a top-down way. In another experiment we show an initial effort to study the acquisition of object affordances. We discuss the importance of sensorimotor coordination as a required step not only for the control of action but also, and more importantly, for perceptual development.

I. INTRODUCTION

Manipulation is a unique opportunity to study the interaction between an artificial system and the environment. We focus on manipulation not only as a means to perform useful practical tasks but, also and in particular, because it offers the possibility to investigate active learning. Active learning refers, for example, to the ability of an agent to autonomously perform and guide the exploration of the environment. At a deeper level, action can contribute to change the environment in a direction that is best suited to the agent’s goal, for example, to facilitate perception. In this context manipulation allows the agent to collect information about objects by performing specific actions on them [1]. Even very simple forms of action like poking or pushing an object can be sufficient to this purpose [2].

The sensorial experience of a humanoid robot can be quite rich including, for example, the possibility of multi-sensory perception: some features of objects are more naturally perceived through senses other than vision. For example the smoothness of a surface, the weight of an object and its three dimensional structure are naturally determined through tactile experience. This information is extracted by appropriate exploratory actions. The importance of motor activity for perceptual development has been emphasized in developmental psychology [3], [4]. Many researchers agree on the fact that motor development during infancy determines the timing of perceptual development (for a review see [5]). For example perception of object features such as volume, hardness, texture

and weight are unlikely to emerge before 6/9 months of age. Haptic sensitivity of three dimensional shapes appears even later at around 12/15 months. It is perhaps illuminating that this timetable fits surprisingly well with the development of actions in infants: the ability to move the hand is required for infants to begin manipulating objects and consequently perceiving certain properties. Lederman and Klatzky [6] have shown that adults make use of stereotyped hand movement (*exploratory procedures*) to determine certain properties of objects; different procedures were employed by subjects to assess different properties. The ability of infants to correctly execute these procedures can determine their ability to perceive the associated object characteristic. Infants ability to interact with objects is indeed quite limited at birth; early reaching in newborns is pretty inaccurate and only rarely result in actual contact with the object [7]. At the age of three months infants are more reliable in grasping objects, although grasp is usually with the full hand open (as in the power grasp). Only later on at 6/9 months of age infants become skilled in handling objects and grasping them with differentiated grasp types [8]. Accordingly, perception of properties like temperature, size and hardness can occur relatively early in development, whereas properties requiring more dexterous actions like texture or three dimensional shape would emerge only later on.

Similarly, we are pursuing a developmental approach for the design of a humanoid robot. Development of the robot unfolds along three phases: learning about the body, learning to interact with the environment, and learning to interpret events. In the first phase the robot learns properties of its body, which allow recognizing and controlling movement. For example, the robot controls reaching movements by first learning the weight of its arm and to recognize the hand. These abilities are used in the next phase to initiate the interaction with the environment and to learn about it. The robot begins this exploration by reaching for objects and learning their physical properties when it manages to grasp them [9]–[11]. The robot’s experience acquired during this interaction is used in the third phase to interpret events around the robot by matching expectations and perceptions (as for example in [1]). We focus here on the two first phases: learning about the body and learning to interact with the environment. We describe how the robot builds an internal model of its hand which allows localizing it in the visual scene. The hand internal model is used to direct gaze towards the hand and to learn an inverse




Fig. 1. The robotic platform: The Babybot.

model which can be used to control how to reach a point in space. The robot uses these abilities to build a visual model of the objects it happens to grasp. The robot's ability to interact with the environment influences the visual attention system; the visual model of the object grasped by the robot, in fact, is then used as a top-down primer during the search of graspable objects.

The rest of the paper is organized as follow. Section II describes the robotic platform. Section III describes the attention system of the robot, the used model of objects and the method to extract three-dimensional information about objects. Section IV describes how the robot acquires its motor skills. Section V presents an experiment where we show how the modules described in the paper are integrated in a complete behavioral system. Finally in Section VI we draw the conclusions.

II. THE ROBOTIC PLATFORM

The experiments reported in this paper were performed by using an upper torso humanoid robot called Babybot (Figure 1). The Babybot consists of a head, an arm and a hand. The head has five degrees of freedom, and it is equipped with two cameras, two microphones and a set of gyroscopes. The cameras can pan independently and tilt around a common axis; the remaining degrees of freedom allows the head to pan and tilt at the level of the neck. The arm is an industrial PUMA 260 manipulator. The hand is mounted on the arm end point. It consists of a total of 16 degrees of freedom actuated by only 6 motors. Its five fingers are thus largely underactuated: the thumb and index are controlled independently by two motors each, whereas the remaining two motors are connected to the middle, ring and small finger which form a single virtual joint. The coupling between each joint and the motors is achieved by means of springs which give the hand a certain degree




Fig. 2. Details of the hand of the Babybot.

of compliance and elasticity. Magnetic potentiometers provide position and force feedback at each joint whereas force sensing resistors on the palm and fingers provide tactile feedback (see figure 1 and figure 2). A more detailed description of the hand can be found in [12].

III. VISUAL SYSTEM

The robot visual system employs log-polar images as in [13]. The log-polar transformation, applied in this case to the traditional rectangular images coming from the cameras, mimics the distribution of the photoreceptors in the retina and the topological mapping from the retina to the primary visual cortex. Log-polar images have a small central area with maximum resolution (fovea) and a continuously decreasing number of pixels moving toward the periphery. In humans and primates there is the need to move the sensor to take high resolution snapshots of important points in the environment. Likewise, in order to acquire information from the environment, the robot has to move the cameras and place the fovea at interesting locations in the visual scene possibly according to the task at hand. In other words there is the need to have a module that selects information for further visual processing.

In addition, another important requirement for a visual system apt to guide manipulation is that of segmenting objects from a possibly cluttered background. That is, we need both the localization of the object and its segmentation. The problem of segmentation is directly related to the problem of defining what an object is, that is to define which properties distinguish an object from the background.

Our definition of "objecthood" is created in two steps:

- the selection of a set of visual features that combined appropriately can characterize any object of a certain set and allow to segment it from the background;
- the selection of a criterion for grouping features that segment and identify the objects uniquely;

which means, in practice, that we selected certain features and a method for deciding when a specific feature belongs to an object. The features we chose for our implementation are *colored blobs* while the criterion is a consequence of the action of the robot onto objects. In fact, by grasping an

object the robot has the possibility of observing it at will from many points of view and likely with different backgrounds. It is consequently easy to imagine a procedure that selects the constant features across consecutive views of the same object.

By selecting blobs as features we do not propose to define what is an object directly, but rather to consider a sort of “proto-objects”. They are a step above the mere features (e.g. edges), possessing some but not all the characteristics of an object; “proto-objects” in this view are clusters of points on the image “naturally” grouped together. The idea of proto-objects has its roots in psychological [14] and neurobiological literature. In fact it has been proposed that the synchronization of visual cortical neurons can be the carrier of the perceptual grouping phenomenon [15], [16].

In our implementation, we decided that the grouping acted on color and intensity information and thus the grouping of elementary feature leads to the extraction of colored blobs as we mentioned earlier. To simulate the results of the process of grouping, we employed the watershed transform (raining variant) [17] on the edge map resulting from a preliminary feature extraction stage. As a consequence the image is segmented in regions of constant color or a constant gradient of color. A segmentation of this type has been demonstrated to happen in humans before the attention is deployed to the scene [18].

Further, following our definitions, the identity of an object cannot be known without active manipulation, unless some other prior knowledge is inserted into the system. One possible route for learning something about an object autonomously is by means of action. The robot can go beyond the concept of proto-objects, learning a model of any object that is manipulated. In particular, objects are seen here as a collection of proto-objects and their spatial relations. Our solution goes by allowing the robot to manipulate objects and acquire different views of them, and, using the probabilities of occurrence, calculate the probability the collection of blobs being currently fixated is one of the objects the robot is searching for. Then, using these same probabilities, the figure-ground segmentation can be attempted. The complete description of the visual attention model and segmentation procedure can be found in [19].

The segmentation mask is used together with a binocular disparity estimation algorithm which can be used to extract three-dimensional information about the object. The mask defines a region of interest around the object, where a depth map is estimated, and eventually the orientation in space of the object can be extracted and used to guide the behavior of the robot.

In order to achieve a good detection of the object orientation, we developed a fast and robust binocular disparity estimation algorithm, which can work in real-world conditions. The algorithm is based on the work by Van Meerbergen et al. [20], where, given a scanline, all the possible matches between pixels are analyzed by exploring a graph (using dynamic programming) built by assigning a cost to each position pair and each occlusion. The algorithm works at nearly frame rate




Fig. 3. A section of the graph: the columns have constant disparity ($m - n$).

especially when running on a small portion of the image (a pair of images).

With respect to the original formulation of the algorithm we relaxed the hypothesis of similar extension (over the scanline) of the pixel patterns: something that makes sense when the cameras are quasi-parallel, but on the other hand fails dramatically when the cameras are converging and the objects are very close compared to the interocular distance. The consequences in this last conditions are that a surface can be very different in shape between the two images of the stereo pair. As a result of our changes a short sequence of pixels in one image can match a long one in the other image, which is, as we said, reasonable in the case of our robot.

Each node of the graph represents a pixel pair, one on the left scanline (L_x) and one on the right one (R_y), starting from the nodes containing either (L_0) or (R_0) or both (the first pixels of the scanline) and ending with the nodes containing L_N , R_N or both, where N is the length in pixels of the scanline. Each node is connected with all the following nodes according to these rules:

- 1) *Disparity Range*: $\delta_{min} \leq x - y \leq \delta_{max}$
- 2) *Ordering*: Given another pair ($L_{x'}, R_{y'}$), if $x' \geq x$ then $y' \geq y$, assuming that there cannot be duplicate nodes
- 3) *Continuity*: the nodes following (L_x, R_y) are all the nodes containing (L_{x+1}) or (R_{y+1}) that respect the previous two constraints

Each arc has an associated cost:

$$c = |lum(L_x) - lum(R_y)| + k\beta + \alpha$$

where lum represents the luminance of a pixel, β is a cost (linear) associated to disparity jumps (if a node has disparity δ_i and the following one has disparity δ_j , then $k = |\delta_i - \delta_j|$) and α takes into account the differences along the vertical dimension (to penalize large discontinuities in the final disparity map). Once the graph is built, the minimum cost to traverse it from beginning to end is found using dynamic programming. The disparity map is then constructed by considering the pairing of pixels along the minimum cost path. An example of the construction of the graph when considering two successive levels is shown in figure 3.

Since the complexity of dynamic programming, for each graph (i.e. each scanline), is $O(m)$ – where m is the total




Fig. 4. The Disparity Algorithm Left: The Stereo Pair, Right: The Final Disparity Map, masked. The mask image comes from the attention algorithm.

number of arcs in the graph and it is proportional to the length of the scanlines – we considered only the portion of the image around the object of interest, segmenting the object itself by using the information coming from the saliency algorithm. This reduces both m and the total number of lines to be processed. To increase the robustness, once computed the disparity map D_{l-r} (displacement of the pixels in the left image compared to the ones in the right one), we used it to detect the object position on the right image, according to the formula:

$$D_{l-r}(x) = D_{r-l}(x - D_{l-r}(x)) \quad (1)$$

This result is then used to segment the object on the right image. The images are swapped (left with right) and the disparity is computed again. This new estimate is finally used to validate and correct the previous one. An example of a masked depth map is shown in figure 4.

IV. THE BODY

Humans become skillful at controlling their own body after a long period of postnatal development and probably thousands of trials. As we discussed in the introduction, however, motor development is extremely important and enables the correct perceptual development of the child. For this reason the robot spends the first phase of its artificial development learning how to correctly control the head and the arm to perform various tasks such as visual tracking and reaching for a visually identified target.

Control of the body requires implicit knowledge of its structure (relative position of the limbs, their size, etc.) as well as its dynamical characteristics (e.g. the weight of the body segments). The ensemble of this knowledge is called *body-schema*; experiments in neuroscience have given support to the existence of a body-schema in the primate brain [21], [22]. Graziano and co-workers have found neurons in the motor cortex of the monkey, which code the position of the hand in the visual field.




Fig. 5. Hand segmentation: an example.

On the other hand, developmental psychologist have been trying to understand the mechanisms which allow the brain to acquire such a representation. As roboticists we are interested in the same mechanisms as they allow the system (biological or artificial) to autonomously acquire and maintain all parameters required to the control of action and avoid manual estimation and calibration. For this reason the problem has been studied by many authors [1], [23], [24].

We follow here an approach similar to the one of [1], [24] where repeated self-generated actions are exploited for learning. We programmed the robot to perform a periodic movement of the wrist. This motion is observed by the robot visually and *motorically*. In the former case the visual motion is computed by image differencing with an adaptive model of the background. In the latter case the robot computes the first derivative of the encoder feedback. The period of motion of each pixel in the motion image is compared to that of the encoder. Pixels whose motion is periodic and whose period matches that of the joints are selected and grouped together to form the segmentation of the hand. Figure 5 shows an example of the result of this procedure.

We used this procedure to train a neural network to estimate the position of the hand in the visual field given the current robot posture (arm and head joint configuration). Another neural network learns the approximate shape and orientation of the hand given the same segmentation information. Figure 6 show the result of this procedure.

The role of vision during reaching is still debated [25], although experimental results suggest that the sight of the hand is not required for children to start reaching for an object [26], [27] and it is used only relatively late in development to actually adjust the trajectory of the hand during action [28].

Sight of the hand, however, might be used to acquire eye-hand coordination. By tracking the hand the robot builds a mapping between the position of the arm and the corresponding head configuration when fixation is achieved. The hypotheses is that reaching starts by first fixating the object; in this condition the fixation point coincides with the target and uniquely identifies its position with respect to the body. The arm motor command can be obtained by a transformation between the head and arm joint angles, that is by mapping motor variables into motor variables:

$$q_{arm} = f(q_{head}) \quad (2)$$

where q_{arm} and q_{head} are head and arm posture respectively.

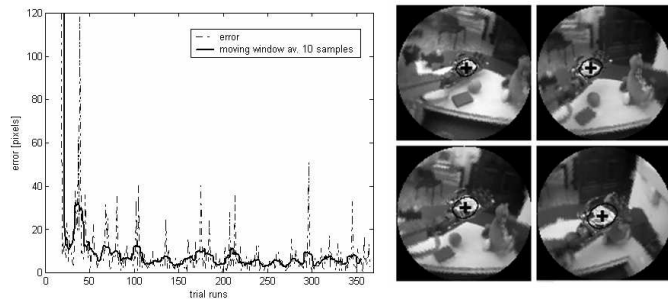


Fig. 6. Learning to localize the hand in the visual field. The convergence of the approximation error as result of training of the neural network (left) and a few exemplar frames of the subsequent prediction (right).

This mapping implicitly implements the inverse kinematics of the arm and it can be learnt when the robot looks at its hand: that is, the robot can maintain the fixation of the hand while moving the arm randomly in the workspace. Every time the arm stops and the eyes have achieved a stable fixation of the hand, a new pair of arm-head positions is acquired and used as a training sample to the neural network approximating the mapping of equation 2. The robot uses the mapping to reach for visually identified objects as soon as a few samples are acquired.

The actual trajectory is computed by linearly interpolating the motor command and the current joint position. The trajectory results in a set of small changes that are effected by a PD controller with gravity compensation. A procedure to learn the gravity component is explained in [12].

V. INTERACTION

We present here a grasping behavior based on the modules described in the previous sections. The interaction with the environment starts when an object is placed in the robot’s hand; the robot detects the object by using the tactile sensors on the palm (see figure 7 frame 1). When pressure on the palm is sensed the fingers close in a stereotyped grasping action. The intrinsic elasticity of the hand (see section II) facilitates grasping, because the fingers automatically adapt to the shape of the object. The robot starts the exploration of the object by bringing it close to the cameras in four different positions and orientations (frames 2-3). During the exploration the robot keeps fixation on the object by tracking the hand. At each position a few frames are acquired and processed as explained in section III to train the model of the object. As the exploration is completed the object is dropped on the table. The robot exploits now the visual model of the object to search for it again (meanwhile the object might have been moved elsewhere by the experimenter) in the visual scene.

The search procedure is driven by a top-down attention module whose contribution exploits the knowledge the robot just acquired about the object. In practice, this happens by selecting the blob whose features better match those of the object’s main blob and performing a saccade motion towards it. After the saccade the object is in the fovea (frame 4 and




Fig. 7. The execution of a grasping experiment.

7) and its model is matched against the blobs that are now fixated. If the match is positive grasping starts otherwise the search continues. The disparity map of the segmented object is computed to determine the orientation of the object (frames 5 and 8); two different actions are then attempted to maximize the possibility to successfully grasp the object. If the principal axis is oriented horizontally the robot moves the hand above the object, otherwise the hand approaches the object from the side (frames 6 and 9). To determine if the grasping is successful, the robot checks the weight of the object and its “consistence” in the hand (the shape of the fingers around the object). In case of failure another grasping trial is attempted, otherwise the robot waits for a new object to be placed in the palm.

VI. DISCUSSION

We have shown results on two phases of the acquisition of sensorimotor coordination in a upper body humanoid robot. The robot includes a visual attention system employing top-down and bottom-up information. The former is introduced in the system beforehand, whereas the latter is modulated by the robot’s interaction with the environment.

We have shown the importance of the interaction between the environment and the robot for learning. This was demonstrated indirectly, when the robot exploited self-produced actions to explore its own body, and directly when the robot actively explored the visual properties of the objects it grasped.

In the experiment discussed here, we start to link different actions to different objects to investigate the possibility for the robot to autonomously learn which actions are more suitable for different contexts (different objects or environment). Although far from completed, this is meant to enlighten us about the possibility of autonomously enriching the robot’s knowledge of the world. This is not only relevant for action, but

also from a perceptual point of view. Indeed, actions establish a link between events and the causes that have generated them. In other words by acting in the world an “active” agent can link the actions it performs with their consequences. This link can be used in two ways i) for planning, to select the particular action required to bring about a desired consequence and ii) for interpretation, to understand the meaning of an attended event. In the first case, the advantage to use such a representation is that, sometimes, it might be convenient to express goals in perceptual terms. For example pushing an object in a particular direction can be represented by means of the resulting visual motion in the image plane [2]. In the second situation, the only available information is the sensorial experience associated with the event. In this case the robot can search its own experience for an event that closely matches what is observed and can select the action(s) that generated those perceptual experience. For example the sound of an object that hits the floor can be associated with the action of dropping it. Either problems, planning and interpretation, are interesting and challenging, and luckily the solution to both appears to be tightly intertwined with sensorimotor development.

It is fair to say that the system developed so far, although complex, still manifests a certain degree of brittleness perhaps associated to the amount of “handcrafted” components we were nonetheless forced to use to reach this level of functioning in a reasonable amount of time. For instance, the choice of color blobs as features clearly limits the visual system in a way that sometimes prevents the robot from perceiving certain object characteristics. Also, in some other circumstances the residual error in reaching goes unnoticed to the robot that eventually fails to grasp the object reliably. On the object recognition side, objects composed of only a few blobs are easily mistaken for other blobs in the background since a single color (or certain blob combinations) is clearly not discriminative enough (the background might present similar combinations).

We are aware of many of these limitations and, in fact, our ongoing work is exactly aimed to improving the overall performance of the robot both motorically and perceptually with a particular emphasis on the manipulation abilities we deem fundamental for autonomous development.

ACKNOWLEDGMENTS

This work was supported by European Union grants RobotCub (IST-2004-004370) and ADAPT (IST-2001-371173).

REFERENCES

- [1] G. Metta and P. Fitzpatrick, “Early integration of vision and manipulation,” *Adaptive Behavior*, vol. 11, no. 2, pp. 109–128, 2003.
- [2] P. Fitzpatrick, G. Metta, L. Natale, S. Rao, and G. Sandini, “Learning about objects through action - initial steps towards artificial cognition,” in *Proc. of the IEEE Internat’l Conf. on Robotics and Automation*, Taipei, Taiwan, May 2003.
- [3] C. von Hofsten, “An action perspective on motor development,” *Trends in cognitive sciences*, vol. 8, no. 6, pp. 266–272, 2004.
- [4] E. J. Gibson, “Exploratory behavior in the development of perceiving, acting, and the acquiring of knowledge,” *Annual Review of Psychology*, vol. 39, pp. 1–41, 1988.

- [5] E. Bushnell and J. Boudreau, “Motor development and the mind: the potential role of motor abilities as a determinant of aspects of perceptual development,” *Child Dev.*, vol. 64, no. 4, pp. 1005–21, 1993.
- [6] S. J. Lederman and R. L. Klatzky, “Hand movements: A window into haptic object recognition,” *Cognitive Psychology*, vol. 19, no. 3, pp. 342–368, 1987.
- [7] C. von Hofsten, “Eye-hand coordination in the newborn,” *Dev. Psychology*, vol. 18, no. 3, pp. 450–61, 1982.
- [8] C. von Hofsten and L. Ronnqvist, “The structuring of neonatal arm movements,” *Child Dev.*, vol. 64, no. 4, pp. 1046–57, 1993.
- [9] L. Natale, G. Metta, and G. Sandini, “Learning haptic representation of objects,” in *International Conference on Intelligent Manipulation and Grasping*, Genoa, Italy, July 2004.
- [10] L. Natale, F. Orabona, G. Metta, and G. Sandini, “Exploring the world through grasping: a developmental approach,” in *Proceedings of the 6th CIRA Symposium*, Espoo, Finland, June 2005.
- [11] E. Torres-Jara, L. Natale, and P. Fitzpatrick, “Tapping into touch,” in *Fifth International Workshop on Epigenetic Robotics (forthcoming)*. Nara, Japan: Lund University Cognitive Studies, July 22–24 2005.
- [12] L. Natale, “Linking action to perception in a humanoid robot: a developmental approach to grasping,” Ph.D. dissertation, University of Genoa, Genoa, Italy, 2004.
- [13] G. Sandini and V. Tagliasco, “An anthropomorphic retina-like structure for scene analysis,” *Computer Vision, Graphics and Image Processing*, vol. 14, no. 3, pp. 365–372, 1980.
- [14] Z. Pylyshyn, “Visual indexes, preconceptual object, and situated vision,” *Cognition*, vol. 80, pp. 127–158, 2001.
- [15] R. Eckhorn, R. Bauer, W. Jordan, M. Brosch, M. Kruse, W. Munk, and H. J. Reitboeck., “Coherent oscillations: A mechanism of feature linking in the visual cortex?” *Biological Cybernetics*, vol. 60, pp. 121–130, 1988.
- [16] C. M. Gray, P. Knig, A. K. Engel, and W. Singer, “Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties,” *Nature*, vol. 338, pp. 334–336, 1989.
- [17] P. D. Smet and R. Pires, “Implementation and analysis of an optimized rainfalling watershed algorithm,” in *IS&T/SPIE’s 12th Annual Symposium Electronic Imaging 2000*, San Jose, California, USA, 2000, pp. 759–766.
- [18] J. Driver, G. Davis, C. Russell, M. Turatto, and E. Freeman, “Segmentation, attention and phenomenal visual objects,” *Cognition*, vol. 80, pp. 61–95, 2001.
- [19] F. Orabona, G. Metta, and G. Sandini, “Object-based visual attention: a model for a behaving robot,” in *3rd International Workshop on Attention and Performance in Computational Vision at the CVPR 2005*, San Diego, CA, USA, 2005.
- [20] G. V. Meerbergen, M. Vergauwen, and M. Pollefeys, “A hierarchical symmetric stereo algorithm using dynamic programming,” *International Journal on Computer Vision Graphics and Image Processing*, vol. 47, no. 1/2/3, pp. 275–285, 2002.
- [21] M. Graziano, “Where is my arm? the relative role of vision and proprioception in the neuronal representation of limb position,” *Proceedings of the National Academy of Science*, vol. 96, pp. 10 418–10 421, 1999.
- [22] M. Graziano, D. Cooke, and C. Taylor, “Coding the location of the arm by sight,” *Science*, vol. 290, pp. 1782–1786, 2000.
- [23] Y. Yoshikawa, K. Hosoda, and M. Asada, “Does the invariance in multi-modalities represent the body scheme ? - a case study with vision and proprioception -,” in *2nd Intelligent Symposium on Adaptive Motion of Animals and Machines*, Kyoto, Japan, 2003.
- [24] P. Fitzpatrick and A. Arsenio, “Feel the beat: using cross-modal rhythm to integrate perception of objects, others and self,” in *Fourth International Workshop on Epigenetic Robotics*, vol. 117. Genoa: Lund University Cognitive Studies, 2004.
- [25] J. Saunders and D. Knill, “Humans use continuous visual feedback from the hand to control fast reaching movements,” *Experimental Brain Research*, vol. 152, pp. 341–352, 2003.
- [26] R. Clifton and M. C. D.W. Muir, D.H. Ashmead, “Is visually guided reaching in early infancy a myth?” *Child Dev.*, vol. 64, no. 4, pp. 1099–110, 1993.
- [27] R. Clifton, P. Rochat, D. Robin, and N. E. Berthier, “Multimodal perception in the control of infant reaching,” *J Exp Psychol Hum Perform.*, vol. 20, no. 4, pp. 876–86, 1994.
- [28] D. Ashmead, M. McCarty, L. Lucas, and M. Belvedere, “Visual guidance in infants’ reaching toward suddenly displaced targets,” *Child Dev.*, vol. 64, no. 4, pp. 1111–27, 1993.