A sensitive approach to grasping

Lorenzo Natale[*†]
lorenzo@csail.mit.edu

Eduardo Torres-Jara[*]
etorresj@csail.mit.edu

[*]Massachusetts Institute Technology
Computer Science and
Artificial Intelligence Laboratory
Cambridge, MA 02139 US

[†]University of Genoa
LIRA-Lab, DIST
Viale F. Causa 13
16145 Genoa, Italy

Abstract

Experimental results in psychology have shown the important role of manipulation in guiding infant development. This has inspired work in developmental robotics as well. In this case, however, the benefits of this approach has been limited by the intrinsic difficulties of the task. Controlling the interaction between the robot and the environment in a meaningful and safe way is hard especially when little prior knowledge is available. We push the idea that haptic feedback can enhance the way robots interact with unmodeled environments. We approach grasping and manipulation as tasks driven mainly by tactile and force feedback. We implemented a grasping behavior on a robotic platform with sensitive tactile sensors and compliant actuators; the behavior allows the robot to grasp objects placed on a table. Finally, we demonstrate that the haptic feedback originated by the interaction with the objects carries implicit information about their shape and can be useful for learning.

1. Introduction

Recent work in developmental robotics has emphasized the role of action for perception and learning (Metta and Fitzpatrick, 2003, Natale et al., 2004, Natale et al., 2005). Developmental psychology, on the other hand, recognizes that motor activity is of paramount importance for the correct emergence of cognition and intelligent behavior (Gibson, 1988, Streri, 1993, Bushnell and Boudreau, 1993, von Hofsten, 2004). All embodied agents, either artificial or natural, have numerous ways to exploit the physical interaction with the environment to their advantage.

In robotics actions like pushing, prodding, and tapping have been used for visual and auditory perception respectively (Metta and Fitzpatrick, 2003, Torres-Jara et al., 2005). More articulated explorative actions or grasping might increase these benefits, as they give direct access to physical properties of objects like shape, volume and weight. Unfortunately, all these aspects have not been extensively investigated yet. One of the reasons for this is that controlling the interaction between the robot and the environment is a difficult problem (Volpe, 1990), especially in absence of accurate models of either the robot or the environment (as it is often the case in developmental robotics). The design of the robot can ease these problems. We know for example that having a certain degrees of elasticity in the limbs helps to "smooth" and control the forces that originate upon contact. Another approach is to enhance the perceptual abilities of the robot. Traditional robotic systems in fact have perceptual systems that do not seem adequate for grasping. Haptic feedback in particular is often quite limited or completely absent. This is because, unfortunately, most of the tactile sensors commercially available are inadequate for ro-

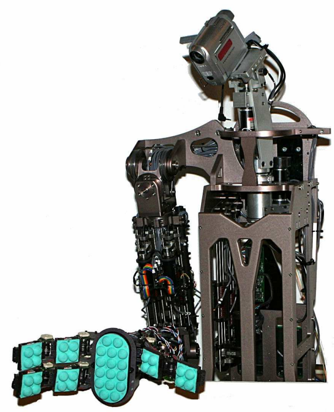

Figure 1: The robot Obrero. The robot has a highly sensitive and force controlled hand, a force controlled arm and a camcorder as a head. Obrero's hand has three fingers, 8 DOF, 5 motors, 8 force sensors, 10 position sensors and 160 tactile sensors.

botics tasks: they are only sensitive to forces coming from a specific angle of incidence, rigid and almost frictionless.

Obrero is an upper body humanoid robot designed to overcome these limitations (Torres-Jara, 2005). It is equipped with series elastic actuators, which provide intrinsic elasticity and force feedback at each joint. The hand is equipped with tactile sensors (Torres-Jara et al., 2006) which provide a deformable and sensitive interface between the fingers and the objects.

We report a series of experiments where Obrero exploits its sensing capabilities to grasp a number of objects individually placed on a table. No prior information about the objects is available to the robot. The use of visual feedback was voluntarily limited. Vision is used at the beginning of the task to direct the attention of the robot and to give a rough estimation of the position of the object. Next, the robot moves its limb towards the object and explores with the hand the area around it. During exploration, the robot exploits tactile feedback to find the actual position of the object and grasp it. The mechanical compliance of the robot and the control facilitate the exploration by allowing a smooth and safe interaction with the object. Results show that the haptic information acquired by the robot during grasping carries information about the shape of the objects.

The paper is organized as follows. Section 2. briefly reviews the importance of haptic feedback for manipulation in infants and adults. Section 3. describes our robotic platform. Section 4. provides some implementation details and describes the grasping behavior. The latter is evaluated in Section 5. Finally, Section 6. draws the conclusions of this work.

2. Haptic feedback, perception and action

In adults, several studies have revealed the importance of somatosensory input (force and touch). For example Johansson and Westling (Johansson and Westling, 1990) have studied in detail what feedback is provided by the skin during object lifting tasks and how it is used to control the movements of the fingers. The results of these experiments proved the importance of somatosensory feedback: they showed that human subjects had difficulties avoiding object slipping when they had their fingertips anesthetized, even with full vision (Johansson, 1991).

Haptic feedback has an important role for object perception as well. Lederman and Klatzky (Klatzky and Lederman, 1987) identified and described a set of motor strategies *exploratory procedures* used by humans to determine properties of objects such as shape, texture, weight or volume.

Little is known concerning how infants use tactile sensing for manipulation (Streri, 1993). In some circumstances children exploit tactile feedback to learn about objects (Streri and Pêcheux, 1986). Streri and Pêcheux measured the habituation time of newborns (2 months and 5 months old) during tactile exploration of objects placed in their hands. In this experiment children spent more time exploring novel rather than familiar objects, even when they did not make visual contact with the hand.

Motor abilities of children are quite limited during the first months of development. This does not prevent infants from using their hand to engage interaction with the world. The importance of motor activity for perceptual development has been emphasized in developmental psychology (von Hofsten, 2004, Gibson, 1988). Researchers agree on the fact that motor development determines the timing of perceptual development. In other words the ability of infants to explore the environment would determine their capacity to perceive certain properties. Accordingly, perception of object features like temperature, size and hardness is likely to occur relatively early in development, whereas properties requiring more dexterous actions like texture or three dimensional shape would emerge only later on (see (Bushnell and Boudreau, 1993) for a review).

3. The robot Obrero

Obrero (Torres-Jara, 2005) consists of a hand, an arm and a head (Figure 1). Obrero was designed to approach manipulation as a task manly guided by tactile and force feedback. Obrero's limbs are designed to reduce the risk of damages upon contact with objects. The head consists of a commercial camcorder (SONY DCR-HC20) that can move along the pan and tilt directions. The arm has 6 Degrees of Freedom (DOF) distributed in this way: three in the shoulder, one at the elbow and two in the wrist. The arm (Edsinger-Gonzales and Weber, 2004) uses Series Elastic Actuators (Williamson, 1995) which provide low-impedance and force feedback at each joint. Position feedback is provided by potentiometers.

The software controlling Obrero runs on a cluster of computers interconnected through an ethernet network. The connection between the different modules is done using YARP (Metta et al., 2006).

3.1 The hand and the tactile sensors

The hand consists of a palm, a thumb, a middle and an index finger (figure 2). Each one of the fingers has two phalanges that can be opened and closed. The thumb and the middle finger can also rotate. These rotations allow the thumb to oppose to either the index or the middle finger. The to-

(a) **(b)**

Figure 2: Obrero's hand and detail of the tactile sensors. (a) Group of four tactile sensors. The deformation of each of them is measured by a total of four sensors. (b) Tactile sensors mounted on the hand.

tal number of degrees of freedom in the hand is 8. All joints in the hand are equipped with an optimized version of the Series Elastic Actuators (Torres-Jara and Banks, 2004); the fingers have low mechanical compliance to soften the contact with the objects during grasping. The hand is underactuated and driven by only 5 motors: three motors open and close each finger, whereas two motors control the rotation of the thumb and middle finger. The phalanges of each finger are mechanically coupled. However, due to the presence of a Series Elastic Actuator in the joint, independent motion is achieved when the proximal phalange blocks (for example, as a result of contact with an object). This elastic coupling allows the hand to automatically adapt to the object it grasps. Finally, position feedback is obtained through potentiometers mounted in all joints and encoders in the motors. The tactile sensors mounted on the hand were designed to satisfy the needs of robotic tasks. Each unit has a dome-like sensor (see figure 2a) made of silicon rubber. At the tip of the dome we embedded a small magnet, whose position is measured by four hall-effect sensors placed at the dome's base. By sensing the position of the magnet the defomation of the dome is estimated. The sensors are very sensitive and capable of detecting a minimum normal force of 0.098N. The shape of the sensors favors contact with the environment from any direction, as opposed to most of the tactile sensors which are flat. The high deformability and the properties of the silicon rubber allow the sensors to conform to the objects, thus increasing friction and improving contact detection. In this particular implementation, we used the "magnetic" version of these tactile sensors, however, an optical version has also been tested. The description of the design and the analysis of these sensors can be found in (Torres-Jara et al., 2006).

Groups of tactile sensors were placed on the hand. Two groups of four were placed on each finger (a group in each of the two phalanges) and 16 on the palm. A detail of the palm and fingers can be observed in figure 2b. Each one of these tactile units uses four sensors to determine the contact forces. This means that overall the tactile feedback consists of 160 signals. At the base of the palm, where for practical reasons, we were not able to mount these tactile sensors, we placed a smaller infrared proximity sensor. To summarize, the hand has 5 motors, 8 DOF, 8 force sensors, 10 position sensors, 160 tactile sensors and an infrared proximity sensor.

4. Controlling the body

In this section we describe a few perceptual and motor competencies required for the robot to be able to control the body in a meaningful and safe way: this includes a simple attention system to spot the objects to be grasped and the ability to control the body to reach out for them. A the end of the section we describe how the these capabilities are integrated in the grasping behavior.

4.1 Attention System

Motion is a simple yet powerful cue to select points of interest in the visual scene; for an active camera system this is still true assuming we can estimate the motion of the background and account for it. In this paper we use the algorithm proposed by (Kemp, 2005), which uses a 2D affine model to robustly estimate the image motion resulting from the background. In short, the algorithm measures the motion of each pixel with a block matching procedure, and performs a least square fitting of the global affine model. Using the affine model the algorithm predicts the motion of each edge, and marks as foreground those edges that poorly match this prediction. Under the assumption that the majority of the image motion is due to the background, these edges can be used to build a saliency map to direct the attention of the robot.

4.2 Eye-hand coordination

We decided to focus on explorative actions rather than precise, goal directed, actions towards the target objects. This was also motivated by the fact that the monocular visual system of Obrero makes depth estimation very difficult. This situation is actually quite common in robotics, as depth estimation in real time is a challenging problem even with stereo vision. However, we cannot hope to program the robot to perform a blind exploration of the entire workspace. A possible solution is to constrain the exploration to the area of the workspace where the object is detected visually. Since the 3D location of the object is not available, reaching is performed in 2D; the exploration procedure allows the robot to find the actual position of the object. The motor skills required for reaching and exploring can be learned from the visual ability to localize the hand and compute the orientation of the arm.

4.3 Hand Localization

A visual module detects the hand and computes the orientation of the arm in the image. The initial step of the hand detector consists in running a high frequency filter. All points whose frequency is below a certain threshold (fixed a priori) are discarded. A blob detector is run on the resulting image and the biggest blob is selected as the arm. The orientation of the arm is computed as the orientation of the line passing through the top-most and bottom-most pixels of the arm area. Next, specific features (the small circular black and white potentiometers on the fingers) are searched on the arm area. The hand is identified if more than two of these features are found. The detection just described proved reliable enough for our purposes and was used as a short-cut in place of other, more general, methods (Metta and Fitzpatrick, 2003, Natale et al., 2005).

The visual feedback of the hand could be used for closed-loop control. However closed-loop control is not always suitable. This happens for example in presence of occlusions or when the hand is not within the visual field. Open-loop control is an alternative solution. A possible open-loop control consists of a mapping between the fixation point of the head and the arm end-point (Metta, 2000). The advantage of this approach is that the mapping can be easily learned if the robot is able to look at the hand. Another approach uses the output of the hand detector to learn a direct mapping between the arm proprioception (encoder feedback) and the position of the hand in the image (Natale et al., 2005). The direct (forward) mapping can be inverted locally to control the arm to reach for a visually identified target. The solution we adopt here is similar: in a discovery phase the robot tracks the hand as the arm moves to randomly explore the workspace. This behavior allows the robot to acquire samples in the form:

$$\left(\begin{array}{ccccc} x & y & \alpha & q_{head} & q_{arm} \end{array} \right)_{0,1 \ldots, k}$$

where x, y and α are the coordinates of the hand and the orientation of the arm in the image, q_{head} and q_{arm} are the position of the head and arm respectively. Given q_{head} it is possible to convert x and y into an egocentric reference frame:

$$\left[\begin{array}{cc} \theta_h & \phi_h \end{array} \right]^T = f_{head}^{-1}\left(\left[\begin{array}{ccc} x & y & q_{head} \end{array} \right]^T \right) \quad (1)$$

θ_h and ϕ_h represents the polar coordinates of the hand in the reference frame centered at the base of the head (azimuth and elevation). Basically f_{head}^{-1} includes knowledge of the inverse kinematics of the head and the parameters of the camera. The opposite transformation maps polar coordinates into the image plane:

$$\left[\begin{array}{cc} x & y \end{array} \right]^T = f_{head}\left(\left[\begin{array}{ccc} \theta_h & \phi_h & q_{head} \end{array} \right]^T \right) \quad (2)$$

Given these two transformations a neural network can be trained to learn the following mapping:

$$\left[\begin{array}{ccc} \theta_h & \phi_h & \alpha \end{array} \right]^T = f\left(q_{arm}\right) \quad (3)$$

which links the arm posture q_{arm} to the polar coordinates of the hand $\left[\theta_h, \phi_h\right]^T$ and the orientation of the arm α. This mapping was learnt online by using the neural network proposed by (Schaal and Atkenson, 1998).

The mapping of equation 3 allows computing the polar coordinates of the hand with respect of the robot from the encoders of the arm. Whenever required equation 2 maps the polar coordinates back onto the image plane.

4.4 Reaching

Suppose we want to move the arm towards a location of the workspace identified visually. Let $\left[\begin{array}{cc} x_t & y_t \end{array} \right]^T$ be such position. Knowing q_{head} from equation 1 we can convert the target position into the body centered reference frame $\left[\begin{array}{cc} \theta_t & \phi_t \end{array} \right]^T$. The reaching problem can now be stated as a the minimization of the following cost function:

$$\min_{q_{arm}}(C) = \min_{q_{arm}} \left\| \left[\begin{array}{cc} \theta_t & \phi_t \end{array} \right]^T - \left[\begin{array}{cc} \theta_h & \phi_h \end{array} \right]^T \right\|^2 \quad (4)$$

where θ_h and ϕ_h are computed from equation 3.

Assuming a stationary target the minimum of equation 4 can be found by gradient descent. The gradient of C is proportional to the Jacobian transposed of the manipulator, that is:

$$\nabla C = -2\nabla f\left(q_{arm}\right) = -2J^T\left(q_{arm}\right) \quad (5)$$

$\nabla f\left(q_{arm}\right)$ was approximated by partial differentiation of equation 3. Because the basis functions used by the neural network are guassians this was easily done analytically (another approach is to perform numerical differentiation).

To summarize we have described a method to compute the arm commands required to reach for a visual target. The method employs the forwards kinematics of the arm. The direct kinematics is learned by the robot as described in the previous section. The reaching problem is solved iteratively by using an approximation of the arm Jacobian. The latter is obtained by differentiating the basis functions of the neural network approximating the direct kinematics. This procedure is carried out online without using the real visual feedback of the hand.

In the robot visual information (and hence the mapping of equation 3) is two-dimensional and does not carry any information about distance. The solution found by descending the gradient of the direct kinematics takes care of minimizing the distance between the target and the hand *on the image plane*,

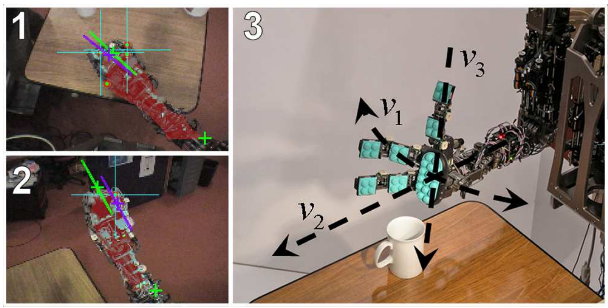

Figure 3: Left, frames 1 and 2: hand localization and arm orientation. Right, frame 3: exploration primitives. Primitives v_1 and v_2 are perpendicular and parallel to the arm orientation. v_3 is along the null space of the arm Jacobian. For simpler understanding these primitives are here sketched in the cartesian plane, but they are actually computed in the joint space (see Section 4. for more details).

and as such, is not concerned with the third dimension R (the distance between the hand and the head, along the optical axis of the camera). In practice however the components of the gradient along R are small compared to the others. The value of R at the end of the reaching movement depends on the initial position of the arm; we chose this value so to keep the hand above the table.

4.5 Exploration

Starting from the direct mapping of the hand position and arm orientation we can identify a set of explorative primitives, that is a set of vectors in joint space that allows the robot to explore the arm workspace. We chose three vectors v_1, v_2 and v_3, as follows (see also Figure 3):

v_1: moves the hand along the direction perpendicular to the arm. It is computed by planning a reaching movement towards a point a few pixels away from the hand along the line perpendicular to the orientation of the arm.

v_2: moves the hand along the direction of the arm. It is computed by planning a reaching movement towards a point a few pixels away from the hand along the arm.

$v_3 \in \ker\left(J\left(q_{arm}\right)\right)$: v_3 lays in the null space of the arm Jacobian; in our case the null space of the Jacobian consists of those vector that do not affect either the projection of the hand onto the visual plane or the orientation of the arm. These vectors produce a movement of the hand along the optical axis of the camera, or, in other word, along R.

4.6 A grasping behavior

In this section we describe the grasping behavior of the robot. The sequence begins when the experimenter waves an object in front of the robot. The head tracks the object until it remains stationary within the workspace of the arm. The robot reaches for the object; motion is planned visually as described in Section 4.4. Reaching is not accurate enough to guarantee a correct grasp. Since no three dimensional information is available the arm reaches a region above the object (see Section 4.4). At this point the exploration starts; the robot computes the explorative primitives v_1, v_2 and v_3. The exploration uses three behaviors:

- *depth behavior*, moves the hand "downwards" along v_3;
- *hovering behavior*, moves the hand back and forth along v_1;
- *pushing behavior*, moves the hand along v_2;

The *depth behavior* moves the hand along the direction of the optical axis of the camera and adjusts the height of the hand with respect to the object/table. To avoid crashing the hand into the table this behavior is inhibited when the infrared proximity sensor detects an obstacle (usually this happens close to the table). The *hovering behavior* and the *depth behavior* are activated at the beginning of the exploration. The goal of this initial phase is to adjust the position of the hand until the index finger touches the object. This allows adjusting the position of the hand along the directions v_1 and v_3. During the exploration the arm stops when the hand detects the object, to avoid pushing it away or knocking it over; if no contact is detected, on the other hand, the amplitude of the exploration is extended (this increases the probability to touch the object in case the reaching error is large). The exploration terminates when the contact with the object is detected by any of the tactile sensors placed on the index finger. At this point the *hovering behavior* is suspended and the *pushing behavior* activated. The "pushing" movement along v_2 brings the palm in contact with the object while the *depth behavior* takes care of maintaining the correct distance with the table. When the robot detects contact on the palm the exploration stops and the *grasping behavior* is activated. The *grasping behavior* simply closes the fingers to a specific position. The low impedance of the joints allows the fingers to adapt to the different objects being grasped.

Figure 4 reports an example of the robot grasping a porcelain cup. The grasping behavior proved to be quite reliable, as repetitive tests show in Section 5.

5. Results

The grasping behavior described in Section 4. was evaluated by presenting different objects to the robot and by counting the number of successful grasps.

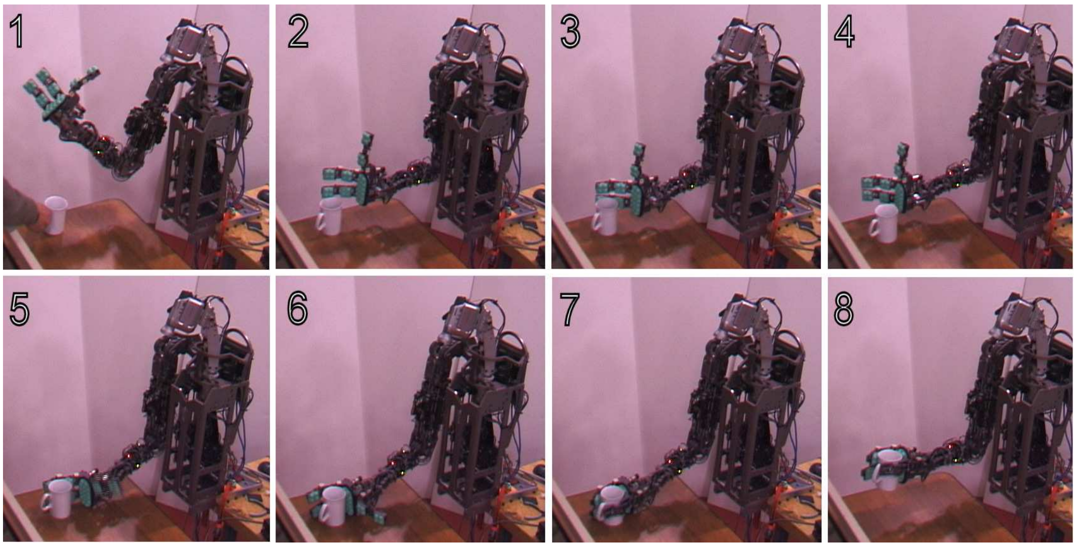

Figure 4: Grasping behavior: an example. Sequence of the robot grasping a porcelain cup. Frame 1: the cup is presented to the robot. Frame 2: the robot reaches for the cup. Frames 3 to 6: the robot explores the space and uses tactile feedback to find the object and adjust the position of the hand. Frames 7 and 8: the robot grasps and lifts the cup.

Table 1: Objects.

	Description	Weight(Kg)	No.Trials	No.Failures	Contains
1	Plastic bottle	0.265	22	0	Vitamins
2	Porcelain cup	0.255	24	1	Nothing
3	Plastic cup (Starbucks)	0.220	24	4	Bolts
4	Rectangular box (Nesquick)	0.240	24	2	Nesquick powder

We chose objects of different size and shape: a plastic bottle, a plastic rectangular box, a porcelain cup and a plastic cup (see figure 5). Some of the objects were partially filled, so that the weight was roughly uniform among all objects (about 220-250 grams, see Table 1). The robot had no prior knowledge about these objects.

Each object was presented to the robot more than 20 times and randomly placed on the table. Overall the number of grasping trials was 94, of which only 7 were not successful. In some of these trials, the robot managed to grasp the object, but was not able to hold it because the grip did not produce enough friction. In a few cases the tactile sensors failed to detect the object and the exploration was aborted before the object was actually grasped (more details are reported in Table 1).

As a further validation, we clustered the haptic information originated from the grasping. We collected the hand feedback at the moment the robot lifted the object; the idea is that given the intrinsic compliance of the hand, its configuration and the force exerted by each joint depend on the shape of the object being grasped. The hand feedback was clustered by means of a Self Organizing Map (SOM). The results show that the bottle, the rectangular box and the cups form three clusters. Unfortunately the clusters formed by the two cups are not clearly distinguishable. This is probably due to the fact that the hand grasped the objects from the top, and that in that part the two objects are quite alike (both are circular with similar diameter). In these cases the limited number of fingers (three) made it hard to distinguish between the cylindrical and conic shape of the cups. Together the results prove that the grasping behavior of the robot is reliable. The high number of successful trials shows that the haptic feedback manages to drive the robot during the exploration until it finds the object and grasps it. This is further demonstrated by the clustering, which show that the behavior allows extracting meaningful information about the physical properties of the objects (i.e. their shape).

6. Conclusions

We have described the design of a behavior that allows a humanoid robot to grasp objects without prior knowledge about their shape and location. We summarize here our approach and the lessons we learned:

- Give up precision, explore instead. Sometime in robotics we struggle to have robots as precise as

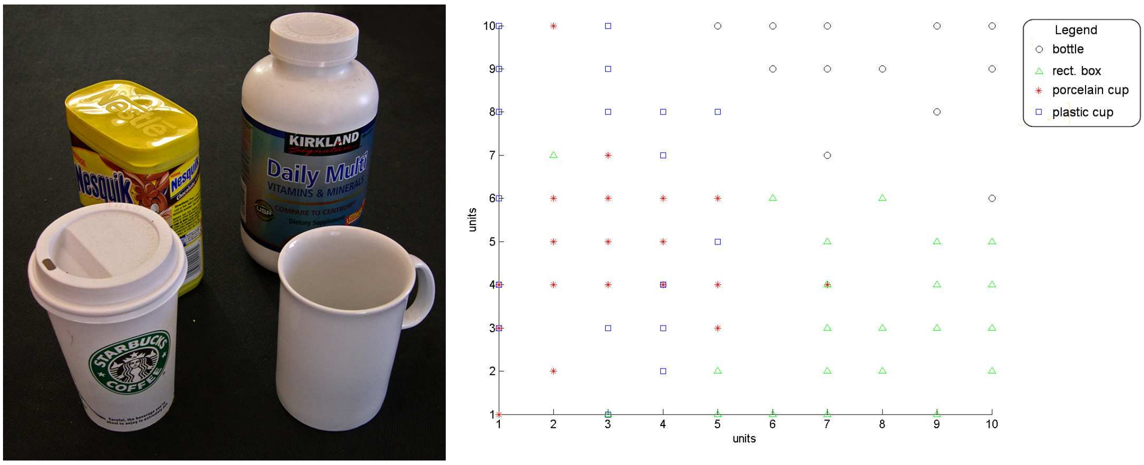

Figure 5: Left: the set of objects used in the experiments: a plastic bottle, a porcelain cup, a plastic cup and a rectangular plastic box. Some objects were partially filled to increase the weight (all objects weighed about 220-250g). Right: result of the clustering. Black circles, green triangles, red stars and blue squares represent respectively the bottle, the rectangular box and the porcelain and the plastic cups. The two cups are not clearly separated because have similar shape in the area where they were grasped.

possible in performing the tasks for which we program them. We found that exploration can be more effective in dealing with uncertainties.

- Be soft. Exploration must be gentle if we want to avoid catastrophic effects on either the robot or the objects/environment. The mechanical design of the robot proved helpful in this respect.

- Sense and exploit the environment. If inquired the world can provide useful feedback; however the robot must be able to ask the right questions (interact) and interpret the answers (have appropriate sensors).

We endowed the robot with the minimum capabilities required to explore the environment. These include a simple ability to detect visual motion, a way to control the arm to roughly reach for objects and a set of explorative primitives. Haptic feedback drives the exploration and allows the robot to successfully grasp objects on a table. We show that the information generated in this ways can be potentially used to learn physical properties of objects like shape.

In the context of epigenetic robotics we are interested in studying methods to improve the perceptual abilities of robots by exploiting the physical interaction with the environment. In this paper we have shown how haptic feedback can significantly improve this interaction thereby enhancing the robot's ability to learn about the environment.

Finally, it is worth saying that, to better illustrate our point, we deliberately took a somewhat extreme approach. We certainly believe that future robots will have to take advantage of the integration of all sensory modalities.

Acknowledgements

This research benefited from discussion with Giorgio Metta. The authors would also like to thank Paul Fitzpatrick, Charles Kemp and Lijin Aryananda for providing some useful code. Funds for this project were provided by ABB and the NASA Systems Mission Directorate, Technical Development program under contract 012461-001. Lorenzo Natale was supported in part by the European Union grant RobotCub (IST-2004-004370).

References

Bushnell, E. and Boudreau, J. (1993). Motor development and the mind: The potential role of motor abilities as a determinant of aspects of perceptual development. *Child Dev.*, 64(4):1005–1021.

Edsinger-Gonzales, A. and Weber, J. (2004). Domo: A force sensing humanoid robot for manipulation research. In *Proc. of the IEEE International Conf. on Humanoid Robotics*, Los Angeles.

Gibson, E. J. (1988). Exploratory behavior in the development of perceiving, acting, and the acquiring of knowledge. *Annual Review of Psychology*, 39:1–41.

Johansson, R. S. (1991). How is grasping modified by somatosensory input? In Humphrey, D. R. and Freund, H. J., (Eds.), *Motor Control: Concepts and Issues*, pages 331–355. John Wiley and Sons Ltd, Chichester.

Johansson, R. S. and Westling, G. (1990). *Attention and performance XIII*, chapter Tactile afferent signals in control of precision grip. Lawrence Erlbaum Assoc., Hillsdale, N.J.

Kemp, C. (2005). *A Wearable System that Learns a Kinematic Model and Finds Structure in Everyday Manipulation by using Absolute Orientation Sensors and a Camera*. PhD thesis, MIT-CSAIL.

Klatzky, R. and Lederman, S. (1987). Hand movements: A window into haptic object recognition. *Cognitive Psychology*, 19:342–368.

Metta, G. (2000). *Babybot: a study into sensorimotor development*. PhD thesis, LIRA-Lab, DIST.

Metta, G. and Fitzpatrick, P. (2003). Early integration of vision and manipulation. *Adaptive Behavior*, 11(2):109–128.

Metta, G., Fitzpatrick, P., and Natale, L. (2006). Yarp: Yet another robot platform. *International Journal of Advanced Robotics Systems, Special Issue on Software Development and Integration in Robotics*, 3(1).

Natale, L., Metta, G., and Sandini, G. (2004). Learning haptic representation of objects. In *International Conference on Intelligent Manipulation and Grasping*, Genoa, Italy.

Natale, L., Orabona, F., Berton, F., Metta, G., and Sandini, G. (2005). From sensorimotor development to object perception. In *IEEE-RAS International Conference on Humanoid Robots*, Tsukuba, Japan.

Schaal, S. and Atkenson, C. (1998). Constructive incremental learning from only local information. *Neural Computation*, (10):2047–2084.

Streri, A. (1993). *Seeing, reaching, touching: the relations between vision and touching in infancy*. MIT Press.

Streri, A. and Pêcheux, M.-G. (1986). Tactual habituation and discrimination of form in infancy: a comparison with vision. *Child Development*, (57):100–104.

Torres-Jara, E. (2005). Obrero: A platform for sensitive manipulation. In *Proceedings of the IEEE-RAS International Conference on Humanoid Robots*.

Torres-Jara, E. and Banks, J. (2004). A simple and scalable force actuator. *35th International Symposioum on Robotics*.

Torres-Jara, E., Natale, L., and Fitzpatrick, P. (2005). Tapping into touch. In *Fith International Workshop on Epigenetic Robotics (forthcoming)*, Nara, Japan. Lund University Cognitive Studies.

Torres-Jara, E., Vasilescu, I., and Coral, R. (2006). A soft touch: Compliant tactile sensors for sensitive manipulation. Technical Report MIT-CSAIL-TR-2006-014, MIT-CSAIL, 32 Vassar St. Cambridge, MA 02319, USA.

Volpe, R. (1990). *Real and Artificial Forces in the Control of Manipulators: Theory and Experiments*. PhD thesis, Carnegie Mellon University, Department of Physics.

von Hofsten, C. (2004). An action perspective on motor development. *Trends in cognitive sciences*, 8(6):266–272.

Williamson, M. (1995). Series elastic actuators. Master's thesis, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA.