Plastic representation of the reachable space for a humanoid robot

Marco Antonelli, Beata J. Grzyb, Vicente Castelló and Angel P. del Pobil

Robotic Intelligence Lab
Universitat Jaume I, Castellón, Spain
{antonell,grzyb,castellv,pobil}@uji.es

Abstract. Reaching a target object requires accurate estimation of the object spatial position and its further transformation into a suitable arm-motor command. In this paper, we propose a framework that provides a robot with a capacity to represent its reachable space in an adaptive way. The location of the target is represented implicitly by both the gaze direction and the angles of arm joints. Two paired neural networks are used to compute the direct and inverse transformations between the arm position and the head position. These networks allow reaching the target either through a ballistic movement or through visually-guided actions. Thanks to the latter skill, the robot can adapt its sensorimotor transformations so as to reflect changes in its body configuration. The proposed framework was implemented on the NAO humanoid robot, and our experimental results provide evidences for its adaptative capabilities.

1 Introduction

Humans live surrounded by objects. Reaching for an object is one of the most common tasks in everyday life. As robots are expected to be active participants in humans' daily life, they also need to have good reaching skills. Moreover, the robots need to be able to constantly learn and autonomously improve their reaching abilities so as to act on unknown objects in new environments or adapt to the changes in their body configurations.

Reaching a target, however, is not an easy task. It requires estimation of the spatial position of the target and its transformation into an apropriate arm motor command. Estimation of the object position is problematic on its own as a three dimensional object is projected into two dimensional surface of camera sensor, which in turn causes the distance to the target to be lost. The common solution is to employ stereopsis to reconstruct the depth of the scene. Human beings, however, are clearly able to perform reaching actions even with a single functioning eye and we are interested in replicating this phenomenon.

Another challenge here is the transformation of the object's spatial location into the arm position that allows reaching the target. These transformations are typically computed analytically by using the known geometric properties of the robotic system provided by the robot's manufacturer or estimated empirically.

This approach to reaching permits to achieve good performance, but only under the assumption that the parameters of the system are time invariant. In practice, it is not always the case, and the system needs to be re-calibrated periodically in order to keep working correctly. Therefore, it is convenient to develop a framework that continuously adapt the sensorimotor mapping to the constantly changing robot parameters.

In previous works, we proposed a framework for the implicit sensorimotor mapping of the peripersonal space that was implemented on a humanoid torso endowed with a pan-tilt vergence stereo head and two multi-joint arms [3, 1]. Instead of using the classical cartesian space, the spatial position of the target was encoded by the gaze direction and by the angular position of the arm joints. Indeed, these variables were implicit because they were directly provided by proprioceptive cues (encoders). This paper presents our reaching framework extended and adapted to work on a monocular robotic setup.

As our previous framework was based only on the depth information provided by stereo cameras, the first objective of this work is to modify the network so as it makes use of distance estimation provided by a monocular camera. Thus, in the proposed adaptation, the target position is represented by the gaze direction that allows bringing the target into the fovea together with the distance to the target. The same position is expressed in terms of the arm posture that allows reaching the target. The direct and inverse transformations between the two frames of reference are learned autonomously by the robot during the exploration of the environment. The results of our computer simulations and robot experiments show that the robot is able to reach correctly the target both by using direct transformation and by visually-guided approach.

Moreover, in this paper we investigate the ability of the system to adapt to the changes in the robot kinematics. Once the robot had been trained to reach the target, its body configuration was changed, that is the position of its elbow joint was rotated about 20 degrees. As this position was assumed to be an invariant configuration of the system, the robot had to autonomously update its sensorimotor maps to reach correctly the target object. The results obtained from our experiments with the robot, showed that the system is able to instantly update its sensorimotor maps to reflect the changes of its body configuration.

The paper is structured as follows. The next section briefly presents the neuroscientific findings that inspired our work and compares our approach to the existing works. The subsequent section describes how the target can be implicitly encoded by the robot sensorimotor maps, which is then followed by the description of the computational model and learning strategies. The next section shows our experimental setup and results obtained from both computer simulations and real robot experiments. We close the paper with the discussion of the results and future work.

2 Background

Our approach to the sensorimotor transformation problem is inspired by neuroscientific findings, mainly concerned with human and primates' brain. Two types of visual processing exist in the brain, that is visual processing to obtain information about the features of objects such as color, size, shape ("vision for perception") in the ventral stream of visual cortex, and visual processing needed to guide movements such as reaching and grasping ("vision for action") in the dorsal stream of visual cortex [9, 4]. The main cortical areas related to reaching actions are V6A and MIP [8, 5, 2], both located in the parietal lobe. Findings in V6A neurons showed neurons that encoded the gaze directions and the distance of the target [6, 15]. Moreover, some neurons seemed to be involved in the execution of reaching movements [8]. These findings indicate that V6A is in charge of performing the sensorimotor transformations required for reaching a given target in 3D space.

The radial basis function networks are suitable to model the parietal cortex neurons as they are able to naturally reproduce the gain-field effects often observed in parietal neurons [20]. Moreover, it was suggested that locations of objects in the peripersonal space are coded through the activity of parietal neurons that act as basis functions, and any coordinate frame can be read out from such population coding according to the task requirements [19]. Because of their biological plausibility, and their ability to approximate any kind of non-linear function [17], the direct and inverse transformations in our framework are encoded by two radial basis function networks (RBFNs).

In robotics, even though extensive literature describes the problem of learning eye-hand coordination [10, 7, 16, 14, 21], to the best of our knowledge only few papers describes the use of RBF networks [14, 21]. Our model differes from these works in various points. For example, Marjanovic et al. learned the transformation only on a surface of the space, in such a way that the target distance was not explicitly taken into account [14]. Sun at al. used a stereo system to compute the cartesian position of the target, while our system employes implicit variables [21]. Moreover, our model allows to learn directly both the inverse and direct transformations between the arm position and the gaze direction. Finally, neither of the cited works show how to update on-line the sensorimotor transformations in a goal-directed behavior.

3 Representation of the peripersonal space

In the proposed framework, the spatial position of the target object was maintained by two global frames of reference (f.o.r.). One f.o.r. is head-centered and it consists of a spherical-like coordinate system in which the azimuth and the inclination angles are replaced by the gaze direction, while the radius is the estimated distance of the target.

One important remark should be given about the use of the distance in the RBFN framework. Indeed, the distance is not directly observable by the

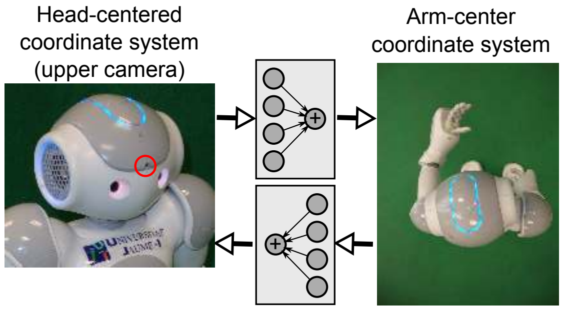

Fig. 1. Computational framework of the sensorimotor integration model. Two transformations allow for converting the head position into arm-motor position and vice-versa.

robot, that is, it is not an implicit variable. However, primates have access to several cues that can be used to estimate the distance, such as stereopsis, familiar size, motion parallax and so on [13]. These cues are implicit and related to the distance, and could be used in our framework in place of the distance. For example, our previous work, used vergence alone [3]; however, when multiple cues are available, it seems more reliable to integrate the cues before calculating the arm position. Such a computation can be performed by a three layer neural network with reward-mediated learning similarly to what is done in [11]. Thus, in our framework, it is possible to replace the distance with the output of another computation as long as it provides neural activation which is correlated with the distance of the target. In this way, the framework becomes more general and can be used independently from the cues available to estimate the distance.

The arm position also provides the spatial position of the target when the robot is touching the object. In this case, the position is described by the joint angle of the arm, provided by the proprioceptive signals. Usually the arm-centered f.o.r. is redundant in the representation of the position, because many joint configurations can bring the hand to the same spatial position. The implication is that the direct mapping (A→H) between the arm-centered f.o.r. and the head-centered f.o.r. is a single-valued function whereas the inverse one (H→A) is not.

As the main focus of this work is learning the sensorimotor transformations for a humanoid robot, the redundancy problem here was bypassed by simplifying our experimental setup. Therefore, only three joints of the arm, two for the shoulder and one for the elbow were used. In this way, also the inverse transformation became a single-valued function.

4 Encoding the sensorimotor associations

As introduced in Section 2 , the sensorimotor associations between the A→H and H→A transformations are maintained by two RBFNs (see Fig. 1).

RBFNs are three-layer feed-forward neural networks whose hidden units($\mathbf{h}$) perform a non-linear transformation of the input data($\mathbf{x}$), whereas the output($\mathbf{y}$)

is computed as a linear combination of the hidden units:

$$\mathbf{y} = \mathbf{h}(\mathbf{x}) \cdot \mathbf{W} \tag{1}$$

where $\mathbf{W}$ is the matrix of the weights.

In this work, the hidden units performs a Gaussian activation which is characterized by their centers $\mathbf{c}_i$ and their spread Σ:

$$h_i(\mathbf{x}) = h(||\mathbf{x} - \mathbf{c}_i||) = e^{-(\mathbf{x}-\mathbf{c}_i)^T \Sigma^{-1} (\mathbf{x}-\mathbf{c}_i)} \tag{2}$$

Once the activation of the hidden units is fixed, the learning process can be stated as finding the weights that best approximate the sensorimotor transformation. Given a set of m input-output samples of the target function, the weights of the *j-th* component of the output can be calculated by minimizing the sum of the square error. In this work we use the recursive least square (RLS) algorithm, as proposed in [12].

A new training sample for both maps is obtained when the hand position and the gaze direction are pointing to the same spatial location. The robot autonomously verifies such a condition by checking whether the visual position of the hand is in the center of the visual field (see Fig. 2). If the hand is visible but it is not in the fovea, the robot can gaze the hand to reinforce the head-arm association.

The mapping between the distinct sensorimotor modalities is learned during the interaction with the environment, through gazing and reaching movements. After each performed movement, visual feedback is used to verify the coordination of gaze and arm. At the beginning, the system does not have any previous knowledge of the sensorimotor transformation so random movements are used to begin the exploration of the environment.

Successively, these random movements are suppressed and the system keeps adapting during the goal-directed exploration. In this phase, when the robot fails to reach the target with a ballistic movement (H→A), it starts to use vision to guide the arm movements. This can be done by locally inverting the transformation A→H (for the details please refer to [21]) to calculate the increment of the arm position that is necessary to approximate the target. While the robot is reaching for the target, it tracts its own hand to update its sensorimotor transformations.

5 Experimental Framework

5.1 Robotic Setup

Aldebaran's commercially available humanoid robot NAO was used as platform for testing the proposed framework. The robot is provided with 25 degrees of freedom (d.o.f.s) among which two are placed in the head (pan and tilt) and five in each arm. In this work, we have used three d.o.f.s for the arm and just the upper camera for the vision system.

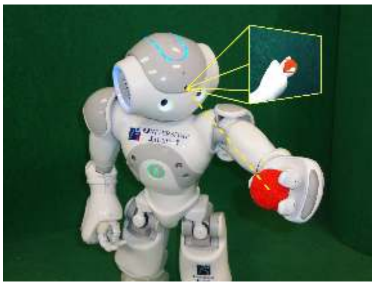

Fig. 2. Association between the oculomotor and arm-motor signals.

5.2 Exploration-based learning

Learning of the sensorimotor transformations is essentially approximating the function through training samples of the form $(d, \theta_{head}, \theta_{arm})i = 1, 2, ..., n$, where d is the distance from the camera to the robot's hand, $\theta_{head}, \theta_{arm}$ are the joint angles of the head and arm, respectively, and n is the size of the training set. Such a training set was generated by random movements of the arm, while the robot was gazing at the hand. Herein, a visual marker (a ball) was used to facilitate the recognition of the hand in a visual field of view. Subsequently, the visual position of the hand was converted into a head movement in order to foveate the hand. The distance to the hand was computed using the familiar size of the ball. That is, knowing the physical size of an object ($S_{physical}$), its absolute depth (d) was calculated by using the following equation: $d = f \times S_{physical}/S_{observed}$ where $S_{observed}$ is the size of the object observed in the image, while f is the focal length. Both $S_{observed}$ and f are expressed in terms of pixels.

The structure and parameters of the radial basis function networks were chosen using a heuristic search on a simulated model of the robot. We decided to employ fixed centers, uniformly distributed on a lattice in the input space. We employed Gaussian receptive fields, where the matrix Σ was a diagonal matrix $\sigma\mathbf{I}$. The input space of A→H was the shoulder(1,2)-elbow space normalized between 0 and 1, the lattice was composed of 7x7x7 neurons and σ was set to 0.3. The input space of H→A was the pan-tilt-distance space normalized between 0 and 1, the lattice was composed of 7x7x7 neurons and σ was set to 0.28. In this work weights of each network were learned using the recursive least square algorithm on the training samples.

After the exploration process, the networks were tested on the acquired sample points using K-Fold cross validation with K set to 5. The error was calculated as Euclidean distance in the cartesian space between sampled and computed values. This was done using the kinematic model of the robot, which was built using the parameters provided by the manufacturer. The performances of the networks are reported in Table 1.

The transformation of the head-centered to the arm-centered f.o.r. performed worse than the other transformation and, in general, seemed more difficult to

Table 1. Performances of the RBFNs using the K-Fold cross validation (K=5). Mean error and standard deviation ($\mu \pm \sigma$) are expressed in mm in the cartesian space.

Transf.	N. points	K=1 $\mu \pm \sigma$	K=2 $\mu \pm \sigma$	K=3 $\mu \pm \sigma$	K=4 $\mu \pm \sigma$	K=5 $\mu \pm \sigma$
A→H	1458	3.8 ± 2.4	3.9 ± 2.6	4.2 ± 3.2	3.9 ± 2.8	4.2 ± 3.2
H→A	1458	5.0 ± 3.5	5.0 ± 3.2	5.3 ± 3.9	5.0 ± 4.3	5.4 ± 4.3

approximate. Nevertheless, in each case, the magnitude of the error was small enough to allow the robot to grasp the target in most cases (see next session).

5.3 Grasping task

The performance of the system was tested on a grasping task. The robot had to localize and to grasp a red ball. The ball was placed on two lattices of 3 by 3 points that covered a region of 5 cm by 8 cm (x,y) on the left side of the robot (see Fig. 3). The two lattices were placed at different altitudes. Each movement of the arm was initiated from a safe position that allowed reaching the ball with a ballistic movement without any collision. During training of the H↔A transformations, the robot was gazing at the hand, so we expected that a correct arm movement would bring the center of the hand near the target. For each location of the ball, the robot had to gaze at it and to calculate the arm position by means of the H→A transformation. After the training, the robot grasped correctly the ball for every position on both lattices.

5.4 Goal-directed learning

Until now we have demonstrated the capability of the system to encode the sensorimotor transformations. The next step is to demonstrate the plasticity of the system for updating its internal representation to the changes of the body parameters. For this purpose, we changed the body configuration by modifying the position of the elbow roll motor some 20°. The position of the motor is not accesible for the RBFNs, thus the networks require to be adapted to the new configuration of the joint. Indeed, with this new configuration the robot failed to grasp the ball in every position, with a mean error of about $3.2cm$ (see Fig.3).

However, when the robot try to grasp the ball, it can recognize the failure through its vision, by checking if the hand position and the ball position are the same. If it is not the case, the system can thus multiply the distance between the hand and the ball (expressed in the head centered f.o.r.) by the pseudo-inverse of the jacobian of A→H to obtain an increment of the arm position. In this way, the robot produces a sequence of visually-guided arm movements until the target is grasped. At each arm movement, the robot looks at its hand (using visual feed-back) and updates both the A→H and H→A transformations. After, three visually-guided executions of the grasping task, the robot was able to correct its sensorimotor transformations and to perform correctly the ballistic grasping (see Fig. 3).

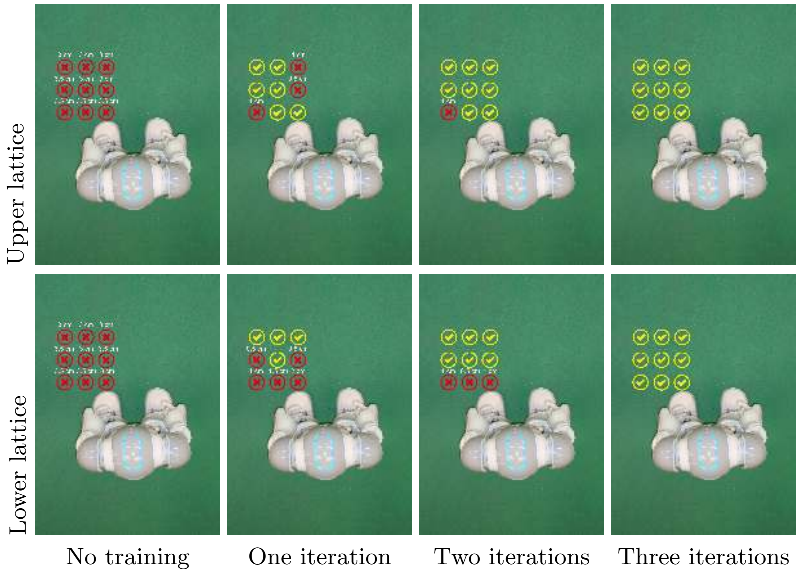

Fig. 3. Experimental setup. After changing the elbow roll position of 20 degrees the robot was tested in a ballistic grasping task (A→H). The target ball was put on two lattices at different hight. The figure shows the grasping error without the on-line adaptation and after one, two and three goal-directed training sessions.

6 Discussion and Future Work

This work was focused on the encoding of the visuomotor transformations for a reaching behavior. The RBFNs were trained with the real data collected while the robot was gazing its hand. The real data, however, are usually quite noisy, which has an impact on the learning process of the neural networks and its performance afterwards. The overall error of the direct transformation (A→H), that is a transformation from arm to eye position was much smaller that the error of the inverse transformation (H→A), that was the transformation from eye to arm position. This can result from the uniform distribution of the centers that, for the H→A transformations, is not so efficient as for the A→H ones. Thus, regularization algorithms [18] that adjust the centers and the spread of the neural activation can improve the performances of the algorithm.

Experimetal results showed that the robot is able to update its performance in the goal-directed behavior by exploiting visual feedback to correct the trajectory of the arm. It is done by inverting the forward model that converts the arm position into head position [21].

In the currently implemented framework, the distance was calculated using the familiar size of the object. Such a distance, however, can be estimated by other cues, e.g. motion parallax, kinetic depth effect and so on, which can be combined in the spirit of the Bayesian theorem in order to obtain a reliable distance estimation. Moreover, more implicit distance observations can be used directly as input to our RBFNs. Thus, our future work will focus on the integration of the proposed sensorimotor framework with another model that implicitly encodes the perceived distance.

This work is part of a larger framework that is inspired by infant development. The final goal is to provide the robot with a coherent near and far space representation. The visuomotor knowledge of the peripersonal and extrapersonal space should be built in a dynamical way, through the active interaction with the environment, in a similar fashion as infants do. Following this approach, the robot has to be able to keep learning during its normal behavior, by interacting with the world and continually update its representation of the world itself. Moreover, the learning process should be self-supervised in order to avoid the need of an external teacher. That is, the robot should be able to improve its capabilities by observing the outcome of its actions.

7 Conclusions

This paper presented a framework for sensorimotor transformations that is inspired by neuroscientific findings. The plausibility of our framework was tested with the NAO humanoid robot. The proposed representations of the space are plastic, indeed the robot was able to update and to improve its performance during interaction with the environment. Moreover, the adaptation of our framework on the NAO robot provides further support for the extendability and generality of our approach.

8 Acknowledgements

This research was partly supported by WCU (World Class University) program through the National Research Foundation of Korea funded by the Ministry of Education, Science and Technology (Grant No. R31-2008-000-10062-0), by Ministerio de Ciencia e Innovación (FPU grant AP2007-02565, FPI grant BES-2009-027151, DPI2011-27846), by Generalitat Valenciana (PROMETEO/2009/052) and by Fundaciò Caixa Castello-Bancaixa (P1-1B2011-54).

References

1. Antonelli, M., Chinellato, E., del Pobil, A.: Implicit mapping of the peripersonal space of a humanoid robot. IEEE Symposium on Computational Intelligence, Cognitive Algorithms, Mind, and Brain pp. 1–8 (Feb 2011)
2. Caminiti, R., Ferraina, S., Mayer, A.B.: Visuomotor transformations: early cortical mechanisms of reaching. Current opinion in neurobiology 8(6), 753—761 (Jan 1998)
3. Chinellato, E., Antonelli, M., Grzyb, B., del Pobil, A.: Implicit sensorimotor mapping of the peripersonal space by gazing and reaching. Autonomous Mental Development, IEEE Transactions on 3, 45–53 (Jan 2011)
4. Chinellato, E., Del Pobil, A.P.: The neuroscience of vision-based grasping: a functional review for computational modeling and bio-inspired robotics. Journal of Integrative Neuroscience 8(2), 223–254 (2009)

5. Dechent, P., Frahm, J.: Characterization of the human visual V6 complex by functional magnetic resonance imaging. European Journal of Neuroscience 17(10), 2201–2211 (2003)
6. Fattori, P., Kutz, D., Breveglieri, R., Marzocchi, N., Galletti, C.: Spatial tuning of reaching activity in the medial parieto-occipital cortex (area V6A) of macaque monkey. European Journal of Neuroscience 22(4), 956–972 (2005)
7. Fuke, S., Ogino, M., Asada, M.: Acquisition of the head-centered peri-personal spatial representation found in vip neuron. Autonomous Mental Development, IEEE Transactions on 1(2), 131–140 (2009)
8. Galletti, C., Kutz, D., Gamberini, M., Breveglieri, R., Fattori, P.: Role of the medial parieto-occipital cortex in the control of reaching and grasping movements. Experimental Brain Research 153(2), 158–170 (2003)
9. Goodale, M., Westwood, D.: An evolving view of duplex vision: separate but interacting cortical pathways for perception and action. Current opinion in neurobiology 14(2), 203–211 (2004)
10. Jones, M., Vernon, D.: Using neural networks to learn hand-eye co-ordination. Neural Computing and Applications 2(1), 2–12 (1994)
11. Karaoguz, C., Weisswange, T.H., Rodemann, T., Wrede, B., Rothkopf, C.A.: Reward-based learning of optimal cue integration in audio and visual depth estimation. In: The 15th International Conference on Advanced Robotics, Tallinn, Estonia (2011)
12. Karayiannis, N.B., Venetsanopoulos, A.N.: Fast learning algorithms for neural networks. Circuits And Systems-II: analog and digital signal processing, IEEE Transactions on 39(7), 1–22 (Jul 1992)
13. Landy, M.S., Maloney, L.T., Johnston, E.B., Young, M.: Measurement and modeling of depth cue combination: in defense of weak fusion. Vision Research 35(3), 389 – 412 (1995)
14. Marjanovic, M., Scassellati, B., Williamson, M.: Self-taught visually guided pointing for a humanoid robot. From Animals to Animats 4: Proc. Fourth Int l Conf. Simulation of Adaptive Behavior pp. 35—44 (1996)
15. Marzocchi, N., Breveglieri, R., Galletti, C., Fattori, P.: Reaching activity in parietal area V6A of macaque: eye influence on arm activity or retinocentric coding of reaching movements? European Journal of Neuroscience 27(3), 775–789 (2008)
16. Nori, F., Natale, L., Sandini, G., Metta, G.: Autonomous learning of 3d reaching in a humanoid robot. IEEE/RSJ IROS, International Conference on Intelligent Robots and Systems pp. 1142–1147 (2007)
17. Park, J., Sandberg, I.W.: Universal approximation using radial-basis-function networks. Neural Comput. 3(2), 246–257 (Jun 1991)
18. Poggio, T., Girosi, F.: Regularization algorithms for learning that are equivalent to multilayer networks. Science 247(4945), 978–982 (1990)
19. Pouget, A., Sejnowski, T.J.: A new view of hemineglect based on the response properties of parietal neurones. Philosophical Transactions of the Royal Society B: Biological Sciences 352(1360), 1449–1459 (Oct 1997)
20. Salinas, E., Thier, P.: Gain modulation: a major computational principle of the central nervous system. Neuron 27(1), 15–21 (Jul 2000)
21. Sun, G., Scassellati, B.: A fast and efficient model for learning to reach. International Journal of Humanoid Robotics 2(4), 391–414 (2005)