International Journal of Humanoid Robotics

ONLINE LEARNING OF THE BODY SCHEMA

Micha Hersch, Eric Sauser and Aude Billard
LASA, School of Engineering, EPFL
Station 9 - 1015 Lausanne - Switzerland
{*micha.hersch, eric.sauser, aude.billard*}*@epfl.ch*

Received Day Month Year
Revised Day Month Year
Accepted Day Month Year

We present an algorithm enabling a humanoid robot to visually learn its body schema, knowing only the number of degrees of freedom in each limb. By body schema we mean the joint positions and orientations and thus the kinematic function. The learning is performed by visually observing its end-effectors when moving them. With simulations involving a body schema of more than 20 degrees of freedom, results show that the system is scalable to a high number of degrees of freedom. Real robot experiments confirm the practicality of our approach. Our results illustrate how subjective space representation can develop as a result of sensory-motor contingencies.

1. Introduction

One of the major components of human agility and dexterity is arguably its ability to merge multiple sensory informations and learn the relationships between information coming from different modalities. In traditional robotics, the importance of learning the mapping between different sensory modalities is often neglected. The robot geometry is generally assumed to be known, and the sensor properties are acquired offline, via a calibration process.

For humanoid robots, which usually comprise a high number of degrees of freedom (DOF) and of sensors, and are intended to operate in human environments, learning the sensorimotor contingencies is likely to bring significant advantages, as it amounts to a continuous self-calibration. For example, by learning its body schema, a humanoid robot can adapt errors in the vision calibration, or to the use of tools. It makes the robotic system less dependent on the human measurements of the system, such as body shape, visual marker positions, and so on.

On the other hand, it has been argued that humanoid robots offer an interesting platform for testing hypotheses on human cognition.[28] Despite their limitations, humanoid robots can be taken as (rough) models of human bodies in that they have mutiple DOFs and mutiple sensory inputs. It is then possible to implement computational processes putatively taking place in the brain and see how they perform for humanoids. This is likely to yield interesting indications on the effectiveness of those hypothesized computational processes.

Bearing in mind this twofolded approach, we propose a mechanism by which a humanoid robot can learn its body schema by combining information from the proprioception (motor encoders), the stereo-vision and possibly tactile sensors. The body schema is modeled by a hierarchy of frames of reference (FoR) transformations which are continuously adapted as sensory information is acquired by the robot. The rest of this paper is organized as follows. Section 2 provides a very succinct overview of previous approaches to learning the body schema. Section 3 recalls some well-known facts about how multi-sensory information is integrated in primates to form a coherent view of the peripersonal space. Section 4 introduces a new model of an adaptive body schema and provides an algorithm for learning it. In section 5, experiments are performed first in simulation on a 24 DOFs humanoids robot and second with a real humanoid robot that adapts to the use of a tool. Those results are discussed in section 6, which also highlights the relevance of the model for the study of human cognition, and a brief conclusion is presented in section 7.

2. Related work

The robot geometry is the key element which determines the forward and inverse kinematic functions of this robot. The forward kinematic function K is defined by the relationship between the vector of joint angles $\mathbf{q}$ defining a given arm configuration and the corresponding position $\mathbf{x}$ of the end-effector in space

$$\mathbf{x} = \mathbf{K}(\mathbf{q}). \tag{1}$$

The inverse kinematic function $\mathbf{K}^{-1}$ is the inverse function, $\mathbf{q} = \mathbf{K}^{-1}(\mathbf{x})$. Finding this inverse function is an ill-posed problem in the case of redundant manipulators. Most work dealing with the learning of a robot geometry directly tackle the issue of learning the inverse kinematics function. This is indeed one of standard issues in robotics, namely knowing that we want the robot to reach a particular location $\mathbf{x}$ in space, what are the arm configurations $\mathbf{q}$ (the vector of arm joint angles) that bring the end-effector to this location. This is a fundamental issue, and many solutions have been suggested. If the forward kinematic function $\mathbf{K}$ is known, it is possible to use local solutions that iteratively bring the robot end-effector to the desired location by computing some (pseudo-) inverse of the Jacobian matrix of $\mathbf{K}$. [34,23,32,29] If $\mathbf{K}$ is unknown, there are global solutions that directly learn the mapping $\mathbf{K}^{-1}$ between the end-effector position and the corresponding joint angles. This is typically done using some function approximators, like multi-layer perceptron [12,20], locally weighted projection regression [7], self-organizing maps [31,10], possibly combined with quantum clustering [17].

Finding the forward kinematic function is much easier. It can be done by measuring the segment lengths and computing the chain of successive rotations and translations (as done in section 4.1). Therefore, it has not attracted much attention in the robotics community. There are, however, some works, mostly coming from the epigenetic robotics community, where an artificial system learns to control its motions. For example Kuperstein[18] learns a visuo-motor coordination, and in Metta *et al.* [22] the motor torques for reaching a point are learned. In those examples, the number of degrees of freedom of the arm are generally quite low (between 2 and 5). Using self-organizing maps Fuke *et al.* [9] learn the correspondence between visual, proprioceptive, and tactile informations in a simulated arm-face system.

The work presented here differs from previous approaches in that the learning is performed entirely online and it can deal with a high number of degrees of freedom. Moreover, the model does not focus solely on determining the position of the end-effector, but also yields the position of each segment and can compute the associated Jacobians. It thus provides additional information, that can be very useful, e.g. for obstacle avoidance or for computing iterative local inverse kinematics. Finally, the system presented here can offer interesting explanations of the way humans represent their peripersonal space, in a multi-modal way.

3. The body schema

In this section we review some evidence based on psychophysical and neurophysiological studies that suggest that, in primates, multi-sensory information is integrated through a hierarchy of frames of reference that reflects the body structure. This hierarchy allows a mapping across the visual, proprioceptive, motor and tactile modalities and is highly adaptive. A more comprehensive review has been written by Holmes and Spence.[14].

There is no doubt that there is a strong interaction between visual, tactile and proprioceptive sensory information. The existence of bimodal visuo-tactile neurons in the monkey and psychophysical experiments involving cross-modal extinctions have put in evidence the existence a visuo-tactile representation,[19] while the discovery of body-part centered visual fields[27] shows that there is a strong interaction between proprioception and vision.

It is believed that each sensory-motor modality receives and provides information represented within different FoRs. For instance, proprioception, touch and motor commands are coded in a FoR centered on the specific body part they represent and control,[1,26,30] i.e., a local representation, whereas visual information is perceived in an eye-centered or retinotopic manner, [16] i.e., visual representation. Since visual, tactile and proprioceptive feedback are tightly coupled in space and time and form together the representation of one's own body, multisensory information has to be integrated across modalities so that a coherent view of the body can emerge. It has been suggested that this transfer is made through a series of transformation across intermediary frames of reference, located in between those provided by the sensory receptors.[6,25] Indeed, neurons coding position in FoR centered on body parts have been reported by Graziano and Gross. [11]

The adaptiveness of those transformations is particularly evident in psychophysical experiments involving prism adaptation.[33] It has long been known, that when subjected to a visual shift or distortion caused by a prism, human subjects first tend to reach, expectedly, to the seen position, rather than to the actual position of the reaching target. After a while, however, they can correct for the visual distortion and accurately reach to the target. For this to occur, visual and proprioceptive feedback of the hand is necessary. When the visual distortion is removed, the subjects show so-called after-effects, i.e., they still reach for the virtual target, as if the visual distortion was still active. This occurs although they are aware that it is not the case. This adaptability has been demonstrated for visual shifts (rotations), reflections, and stretches but could not be shown for more complicated deformations which do not preserve the space topology.[3] Furthermore, people could also adapt to transformations expressed in intrinsic (joint angle) coordinates. [15]
Another kind of experiment emphasizing the adaptiveness of the body schema, involves the use of tools. It has been shown that after some practice with a tool, the monkey integrates this tool into his body schema. [21] The somatosensory receptive field of given neurons where observed to be expanded by the tool, after some practice.
Finally, the "fake limb" experiments also highlight the adaptive and tight connection between different sensory modalities and the feeling of one's own body.[4] In those experiments, a subject sees a fake limb being touched synchronously with his real unseen arm and feels that the fake arm is his.

This argues in favor of the existence of a comprehensive framework, which allows to combine information across visual, tactile and proprioceptuo-motor modalities, and to perform the appropriate FoR transformations required for the integration of this information. Those transformations are highly adaptive, and are constantly learned as a result of sensory experience.

4. A model of adaptive body schema

4.1. *Kinematic chains*

Considering a serial manipulator with n rotative joints, it is possible to compute how a position given in the end-effector frame of reference (FoR) can be expressed in the manipulator base FoR. In other words, we can compute the FoR transformation from a FoR centered on the distal segment to the FoR centered on the proximal segment. This is done by considering the rotation and translation corresponding to each joint and segment, as is commonly done for computing kinematic functions, e.g. the Denavit-Hartenberg kinematic chain parametrization. This transformation can be seen as a series of successive rotations and translations, where the rotation angles are given by the manipulator joint angle and the translations are given by the vector difference between the joints. Thus, it is possible to transform a vector

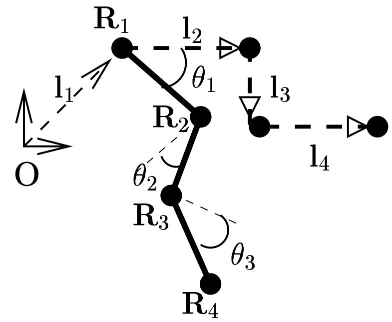

Fig. 1. The parametrization of a kinematic chain. The dashed line represents the kinematic chain in the zero position (when all angles are equal to zero). The solid line represent the same chain with different rotation angles. O refers to the origin.

$\mathbf{v}_n$ from a frame of reference centered on the end-effector to a vector $\mathbf{v}_0$ in a frame of reference centered on the other side of the chain by a transformation $\mathcal{T}$ described by the following equation:

$$\mathbf{v}_0 = \mathcal{T}(\mathbf{v}_n) = \mathbf{T}_1 \circ \mathbf{R}_1 \circ \mathbf{T}_2 \circ \mathbf{R}_2 \circ \cdots \circ \mathbf{T}_n \circ \mathbf{R}_n(\mathbf{v}_n) = \mathbf{l}_1 + \mathbf{R}_1(\mathbf{l}_2 + \mathbf{R}_2(\ldots(\mathbf{l}_n + \mathbf{R}_n(\mathbf{v}_n))\ldots)) \tag{2}$$

where $\mathbf{T}_i$ and $\mathbf{R}_i$ represents respectively the translation and rotation corresponding to segment i and joint i, and $\mathbf{l}_i$ denotes the vector representing the link proximal to joint i at the zero position, and $\mathbf{R}_i$ is the rotation caused by the joint i. Fig. 1 illustrates how the segments are numbered and how this FoR is computed. Note that, similarly to the Denavit-Hartenberg parametrization of kinematic chains, $\mathbf{l}_i$ can be zero if joints $i-1$ and i have the same rotation center.

4.2. *Single segment adaptation*

We consider the following problem regarding a single joint manipulator (see Fig. 2). We assume that we have an initial guess of the unit rotation axis $\mathbf{a}$ and the joint position $\mathbf{l}$. Now, given a vector $\mathbf{v}$ in a FoR centered on the distal segment, its actual transform $\mathbf{v}'$ centered on the proximal segment and the rotation angle θ, how is it possible to adapt $\mathbf{a}$ and $\mathbf{l}$ so that they account better for the actual transformation induced by the manipulator?

In order to do so, we perform a simple gradient descent on the squared distance between the actual and simulated transform vector:

$$\Delta\mathbf{l} = -\epsilon \frac{\partial}{\partial \mathbf{l}} \frac{1}{2} \left|\left| \mathbf{v}' - \left(\mathbf{l} + \mathbf{R}_{\mathbf{a}}^{\theta}(\mathbf{v})\right) \right|\right|^2 \tag{3}$$

$$\Delta\mathbf{a} = -\epsilon \frac{\partial}{\partial \mathbf{a}} \frac{1}{2} \left|\left| \mathbf{v}' - \left(\mathbf{l} + \mathbf{R}_{\mathbf{a}}^{\theta}(\mathbf{v})\right) \right|\right|^2, \tag{4}$$

where $\mathbf{R}_{\mathbf{a}}^{\theta}$ is the rotation of angle θ around axis $\mathbf{a}$ and the learning step ϵ is a small positive scalar. The derivative with respect to $\mathbf{l}$ in Eq. (3) is straightforward to

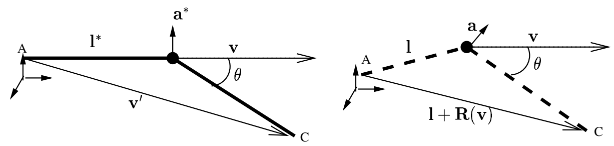

Fig. 2. The learning problem for a single segment. The real rigid transformation is shown on the left and is parametrized by unknown vectors $\mathbf{l}^*$ and $\mathbf{a}^*$ and known angle θ. The current guess of this rigid transformation appears on the right and is parametrized by $\mathbf{l}$, $\mathbf{a}$ and θ. Knowing a vector $\mathbf{v}$ and its real transform $\mathbf{v}'$ and the rotation angle θ, we try to update our guess of $\mathbf{l}$ and $\mathbf{a}$. The letters A and C indicate respectively the origin and the end-effector of the manipulator.

compute. We have

$$\Delta\mathbf{l} = \epsilon\Big(\mathbf{v}' - \big(\mathbf{l} + \mathbf{R}_{\mathbf{a}}^{\theta}(\mathbf{v})\big)\Big). \tag{5}$$

In order to compute the derivative with respect to $\mathbf{a}$ in Eq. (4), we make use of the Rodrigues formula [2]

$$\mathbf{R}_{\mathbf{a}}^{\theta}(\mathbf{v}) = \cos(\theta)\mathbf{v} + \sin(\theta)\mathbf{a}\times\mathbf{v} + (1-\cos(\theta))\mathbf{a}^T\mathbf{v}\mathbf{a}. \tag{6}$$

Hence,

$$\bar{\mathbf{R}}_{\mathbf{a}}^{\theta} \doteq \frac{\partial}{\partial\mathbf{a}}\mathbf{R}_{\mathbf{a}}^{\theta}(\mathbf{v}) = \sin(\theta)\mathbf{v}\!\uparrow + (1-\cos(\theta))(\mathbf{a}\mathbf{v}^T + (\mathbf{a}^T\mathbf{v})\mathbf{I}), \tag{7}$$

where $\mathbf{I}$ is the 3×3 identity matrix and the unary operator$\uparrow$ is defined as

$$\mathbf{v}\!\uparrow \doteq \frac{\partial}{\partial\mathbf{a}}(\mathbf{a}\times\mathbf{v}) = \begin{pmatrix} 0 & v_3 & -v_2 \\ -v_3 & 0 & v_1 \\ v_2 & -v_1 & 0 \end{pmatrix}, \quad \text{with} \quad \mathbf{v} = [v_1\ v_2\ v_3]^T. \tag{8}$$

Thus

$$\Delta\mathbf{a} = \epsilon\Big(\mathbf{v}' - \big(\mathbf{l} + \mathbf{R}_{\mathbf{a}}^{\theta}(\mathbf{v})\big)\Big)^T\Big(\sin(\theta)\mathbf{v}\!\uparrow + \big(1-\cos(\theta)\big)\big(\mathbf{a}\mathbf{v}^T + (\mathbf{a}^T\mathbf{v})\mathbf{I}\big)\Big). \tag{9}$$

Since $\mathbf{a}$ must be of unit norm, it is normalized to 1 after being updated. This solves our problem. Using Eq. (5) and Eq. (9), it is possible to adapt the representation of the joint position and orientation online, as examples of positions in the distal FoR and the corresponding position in the proximal FoR are provided. This algorithm always converges to the correct translation and rotation axis when provided with enough different values of $\mathbf{v}$ and $\mathbf{v}'$, (see section 4.4).

4.3. *Multi-segment adaptation*

We can now apply the same principle to multi-segment manipulators. Starting from Eq. (2) it is possible to compute

$$\Delta\mathbf{l}_i = -\epsilon\frac{\partial}{\partial\mathbf{l}_i}\frac{1}{2}||\mathbf{v}_n' - \mathcal{T}(\mathbf{v}_n)||^2 \tag{10}$$

$$\Delta\mathbf{a}_i = -\epsilon\frac{\partial}{\partial\mathbf{a}_i}\frac{1}{2}||\mathbf{v}_n' - \mathcal{T}(\mathbf{v}_n)||^2, \tag{11}$$

where $\mathbf{a}_i$ is the rotation axis of $\mathbf{R}_i$. If $\mathbf{R}_j$ is the rotation matrix corresponding to joint i (i.e., of axis $\mathbf{a}_i$ and angle θ_i), we have

$$\frac{\partial}{\partial \mathbf{l}_i}\mathcal{T}(\mathbf{v}_n) = \prod_{j=1}^{i-1} \mathbf{R}_j \tag{12}$$

$$\frac{\partial}{\partial \mathbf{a}_i}\mathcal{T}(\mathbf{v}_n) = (\prod_{j=1}^{i-1} \mathbf{R}_j)\frac{\partial}{\partial \mathbf{a}_i}\Big(\mathbf{R}_i(\mathbf{T}_{i+1} \circ \mathbf{R}_{i+1} \cdots \circ \mathbf{T}_n \circ \mathbf{R}_n(\mathbf{v}_n))\Big) \tag{13}$$

where the derivative on the right-hand side of the last equation is obtained by applying Eq. (7). All the rotation axes and translation vectors can thus be simultaneously updated using

$$\Delta \mathbf{l}_i = \epsilon\Big(\mathbf{v}'_n - \mathcal{T}(\mathbf{v}_n)\Big)^T \prod_{j=1}^{i-1} \mathbf{R}_j \tag{14}$$

$$\Delta \mathbf{a}_i = \epsilon\Big(\mathbf{v}'_n - \mathcal{T}(\mathbf{v}_n)\Big)^T \Big((\prod_{j=1}^{i-1} \mathbf{R}_j)\frac{\partial}{\partial \mathbf{a}_i}\Big(\mathbf{R}_i(\mathbf{T}_{i+1} \circ \mathbf{R}_{i+1} \cdots \circ \mathbf{T}_n \circ \mathbf{R}_n(\mathbf{v}_n))\Big)\Big) \tag{15}$$

4.4. *Convergence*

4.4.1. *Single joint case*

Theorem 1. *Assuming that we run the algorithm on a set of configurations given by* $\{\mathbf{v}, \mathcal{T}^*(\mathbf{v}), \theta_j\}_{j=1}^J$, *where the* θ_j *follow a symmetric probability density function (pdf) centered on* 0, *such that* $\mathrm{var}(\cos\theta_j) \leq 2\mathrm{var}(sin\theta_j)$, *the algorithm described by iteratively applying (3) and (4) converges to a correct estimate of* $\mathbf{a}$ *and* $\mathbf{l}$.

The proof is given in the appendix.

4.4.2. *Multi-segment case*

The convergence for the multi-segment case cannot be proven. In order to have an idea of the convergence properties, simulations were performed. In a single simulation run, the rotation axes $\mathbf{a}_i^*$ and $\mathbf{a}_i$ of two kinematic chains were randomly generated. The $\mathbf{l}_i$ were initialized with small random values and the algorithm was run in order to see whether the $\mathbf{a}_i$ converge to the $\mathbf{a}_i^*$. Convergence is considered to be attained if the distance between the real limb position and the modeled limb position remains smaller than a threshold (around 1% of chain length) over 500 different configurations. Ten thousands of those runs were performed for 1,3,5 and 7 DOFs kinematic chains. The results can be seen in Figure 3, left. Expectedly, if there is only one joint, the algorithm always converges. When there are more DOfs, the algorithm fails to converge after a million iterations in less than 1% of the cases. The time it takes for convergence is plotted in Fig. 3, right. Figure 4 gives an example of the evolution of the estimate of the rotation axes for a kinematic chain containing 3 DOFs.

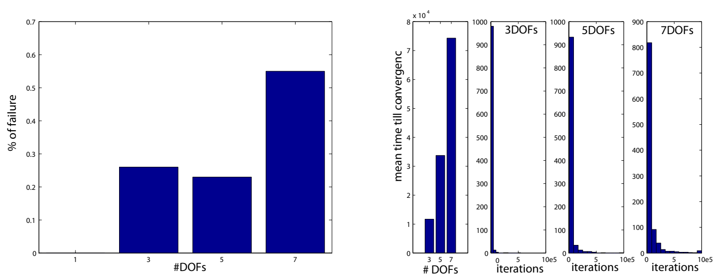

Fig. 3. Left: the percentage of trials that did not converge after a million iterations. Right: the time needed for convergence depending on the number of DOFs. The bars on the left show the mean number of iterations until convergence, and the three histograms on the right show the distribution of convergence time. The distribution have quite a long tail, indicating that in some cases it takes much longer than average to converge.

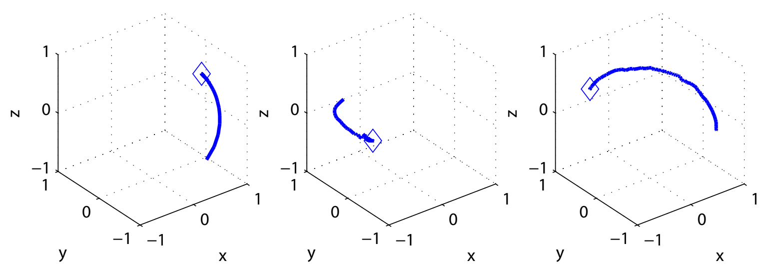

Fig. 4. The evolution of the rotation axes for a 3-DOF kinematic chain. The real axes are indicated by the three diamonds, and each graph show the evolution of one estimated axis on a sphere of radius one

4.5. *Adaptive body schema*

4.5.1. *Kinematic Tree*

The humanoid body schema can be represented as a tree of rigid transformations reflecting the limbs structure, as shown in Fig. 5. We thus have a kinematic tree with adaptive joint positions and orientations. Note that the structure of this tree (i.e., the number of joints and the ordering) is given and remains fixed.
Out of this tree, kinematic chains can easily be extracted as paths in the tree. It is possible to compute the FoR transformation from a FoR attached to any joint of the kinematic tree to a FoR attached to any other joint. This is done by first finding the path joining the two corresponding nodes. To each edge along this path, there corresponds a FoR transformation. Depending on the direction an edge is taken, the transformation or its inverse is considered. An example is provided in Fig. 6, showing how a kinematic chain is extracted from the kinematic tree.

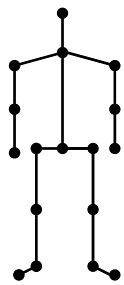

Fig. 5. A kinematic tree representing a humanoid. Nodes represent rotations and edges represent translations.

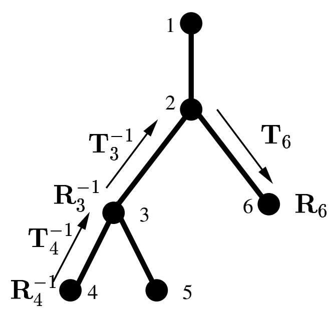

Fig. 6. The generating of the kinematic chain for the FoR transformation relating the FoR attached to two nodes in the tree. When taking an edge up (to the root), one takes the inverse transform, and when taking an edge down (to the leaf) the actual transform is taken. In this example, the transformation from a FoR centered on joint 6 to a FoR centered on joint 4 is given by $\mathbf{R}_4^{-1} \circ \mathbf{T}_4^{-1} \circ \mathbf{R}_3^{-1} \circ \mathbf{T}_3^{-1} \circ \mathbf{T}_6 \circ \mathbf{R}_6$.

4.5.2. *Body schema learning*

We assume that the robot is endowed with a stereo-vision system that can track the position of its end-effectors. This position is provided in a head-centered frame of reference. Within the kinematic tree, the path going from the head to the end-effector corresponds to a kinematic chain that transforms positions and orientations from a frame of reference centered on the end-effector to a visual or head-centered frame of reference. Using (15) and (14), it is possible to update all the rigid transformations along this chain. As input the $\mathbf{v}'_n$ are given by the stereo-vision system and $\mathbf{v}_n$ is the position of the end-effector in its own frame of reference. This is illustred in Fig 7, left. This figure also illustrates how the same algorithm could be used with tactile sensors for tactile body schema learning.

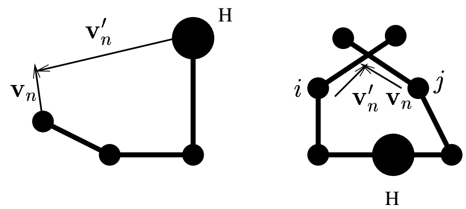

Fig. 7. Humanoid body schema adaptation. The big circle marked H is the head of the humanoid. Left: visual learning, side-view of the robot. The position of the limb in its own frame of reference $\mathbf{v}_n$ is transformed into a visual head-centered frame of reference $\mathbf{v}'_n$. Right: tactile learning, top-view. The position of touch sensors in the frame of reference of their limb $\mathbf{v}_n$ and $\mathbf{v}'_n$ transform into one another.

4.5.3. *Subjective body schema*

Traditionally, the body schema has been considered as an objective account of the body characteristics, such as the arrangement of the limbs, their lengths or the positioning and effect of the joints. To this "objective" body schema, it is possible to oppose a "subjective" view of the body schema, which would be dependent on the perceptual abilities of the robot. In this view, which is adopted in this paper, the body schema only indirectly deals with physical properties of the body. It primarily deals with the frame of reference transformations associated to the sensory signals. For example, given a proprioceptive input corresponding to a particular posture, the body schema can predict the corresponding visual perception. This depends not only on the physical properties of the body, but also on the properties of the sensory system. Moreover, it can be that a precise account of the physical properties of the body is not necessay to the "subjective body schema", depending on the sensory system. For example, in our case, if the robot can track only end-effector positions, many different body geometries will yield the same "subjective" body schema. In the simple example depicted in Fig. 8, such a robot will not be able to differentiate between the two "objective" body schemes. They will both correspond to the same "subjective" body schema. As the corresponding kinematic function and Jacobian remain the same in both cases, using one or the other body schema for controlling its movements will produce the same end-effector trajectories.

5. Humanoid experiments

5.1. *Setup*

In order to validate the algorithm described above, we tested it in simulation on a 24 DOFs humanoid robot. Testing the algorithm on a real robot with that many DOFs would be quite unpractical due to the following reasons.

- It requires the availability of such a sophisticated robot, with that many DOFs, stereo-vision and a tactile skin.
- The high number of positions that the robot has to visit, would require a

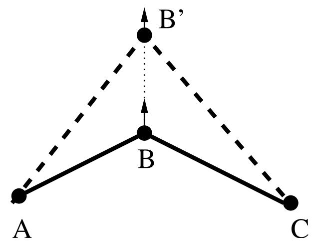

Fig. 8. A simple example of two different geometries (solid and dashed) yielding the same FoR transformation from a A-centered to a C-centered FoR.

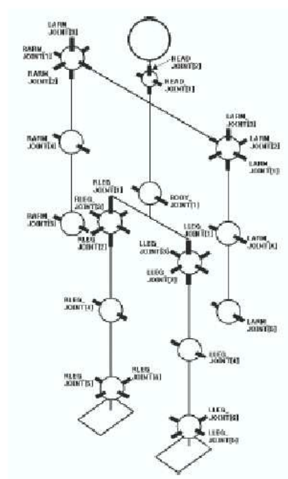

Fig. 9. The structure of the Hoap3 robot. Spheres show the joints, with the rotation axis shown as dark lines going through them. The hand rotations were not used, and an additional head roll joint was modeled. This makes 24 DOFs. This picture is taken from the Hoap documentation provided by Fujitsu.

real robot many days of continuous exploration for the algorithm to converge.

- A real robot may not be able to see all its end-effector (in particular its feet) due to a limited visual angle and joint range and possible occlusions.

The simulated humanoid (or avatar) has the shape of the Fujitsu Hoap3 robot and comprises 24 DOFs. A schematics of the robot is drawn in Fig 9. When learning the body schema, the avatar configuration space was randomly sampled with a uniform distribution. The joint angles and corresponding visual or tactile position were fed into the algorithm. The body schema was initialized with random joint orientations and small random body segments.

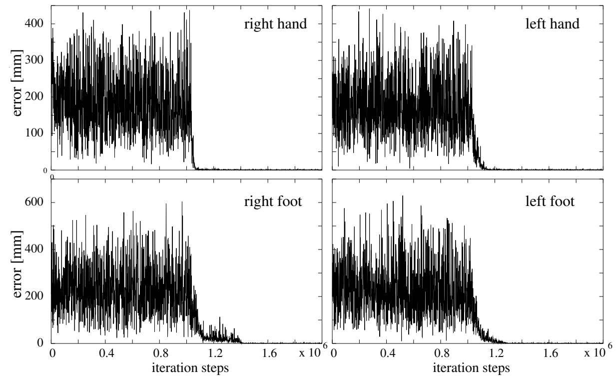

Fig. 10. The convergence of the learning algorithm. On the y-axis the error on the computation of limb position, in a head-centered frame of reference.

5.2. *Results*

In this section, we present a set of experiments intended to first validate the framework described above. In a first validation step, we used the algorithm described above to learn the body schema of the Hoap3 robot described in Section 5.1. In this first experiment, only the two hands and feet were tracked. This means that four kinematic chains where concurrently used: head - right hand, head - left hand, head - right foot, head - left foot. At each time step a joint angle configuration was randomly chosen, and the corresponding position of the hands and feet in a head-centered FoR where computed for the Hoap3 robot. Along with the joint angle values, those four positions where fed as $\mathbf{v}'_n$ (see Eq. 14, 15) for adapting the corresponding kinematic chains. The result can be seen in Fig. 10 which plots the error of the kinematic function, i.e. $\|\mathbf{v}'_n - \mathcal{T}(\mathbf{v}_n)\|$ at each iteration. The after many iterations, the system converges to the appropriate subjective body schema as the error converges to zero. Note however, that the objective body schemes differ, as can be seen in Fig. 11.

In this experiment, a minimal amount of information is provided by the vision system as it always only tracks end-effectors, like hands and feet. But it is also possible for the vision system to track non-terminal body parts, like elbows and knees, as well. This, of course, is expected to make the system converge to the right geometry and also to considerably speed-up the learning process, as much more information is available to the system. Indeed, tracking non-terminal body parts amounts to having shorter kinematic chains which significantly reduces the dimension of the problem.

So, in a second experiment, the vision system alternatively tracks terminal (hand and feet) and non-terminal (knees,elbows, shoulders and waist) body segments. In this case, convergence is faster and the robot geometry is correctly retrieved as there is no ambiguity on the joint locations (see Fig 11).

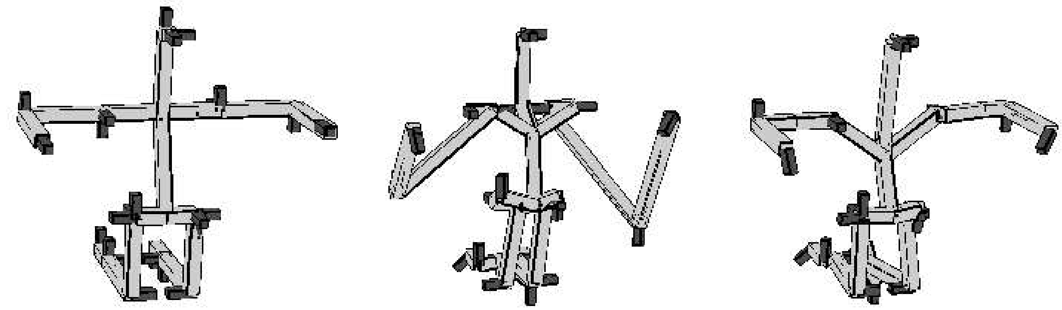

Fig. 11. Left: the Hoap3 "real" body schema, when standing on its knees. Middle: The objective body schema that was learned by the system when looking only at its hands and feet. It is not exactly the same, but the sensory relationships of both schemes (i.e., the subjective body schemes) are the same, when considering only vision of the hands and feet. Right: The body schema learned when also looking at elbows, knees, shoulders and the waist. The darker sticks indicate the rotation axis of each joint. There are three DOFs at the shoulders and the hips.

5.3. *Real robot experiment*

In order to evaluate the practicality of our approach, we conducted an experiment on a real robotic setting. In this experiment, the robot carries an unknown tool. By looking at the tip of the tool the robot can integrate this tool into his body schema, thus enabling it to manipulate it adequately. This setting is shown on Fig. 12. We use the Hoap3 robot, which is endowed with a stereo-vision system.

In the results presented here, the robot is initialized with its "real" body schema. The robot holds a 340 mm long stick in the hand with a color blob at its tip. We then make the robot passively move its arm, with the tip of the stick remaining within the field of view of the cameras. The stereo-vision system tracks the tip of the stick and the robot joint angles are read from the motor encoders. Those two set of values (position of the tip of stick and joint angles value) are then continuously fed into the learning algorithm, like in the simulation experiments. The vector $\mathbf{v}$ in Eq. (14) and (15) is the position of the end-effector in its own FoR, in other words zero. Two cases (using the same data) are tested. In the first case only the terminal limb is adaptive, i.e. we do not change the position and direction of the non-terminal joints. In the second case the whole arm is adaptive, as in the simulations presented above. The results are displayed in Fig. 13. This figure shows the distance between the position of the end-effector as seen by the stereo-vision on one hand, and as computed by the body schema on the other hand. In other words, the y-axis plots $\|\mathbf{v}'_n - \mathcal{T}(\mathbf{v}_n)\|$, where $\mathbf{v}'_n$ is given by the stereo-vision and $\mathbf{v}_n$ is zero. One can see that in both cases the system starts with an error of about

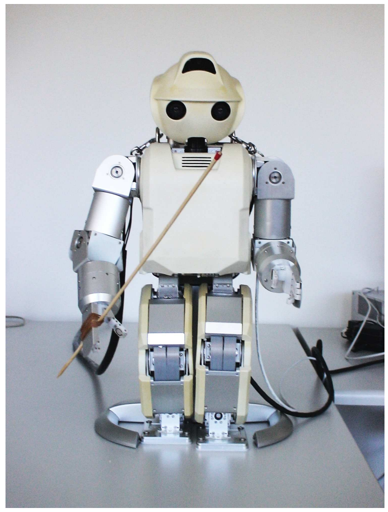

Fig. 12. The setup of the real robot experiment. The robot holds a stick and visually tracks its tip and records its arm joint angle values, while its arm is passively moved.

40cm, which corresponds approximately to the length of the stick, and reduces this error to about 5cm. This means that the stick has been incorporated in the body schema, although imprecision in the stereo-vision system keeps this error of about 5cm. In the plot, one can notice a few peaks of large errors. Those are outliers of the stereo-vision system, and correspond to a bad tracking of the color blob. Of course, the resulting error values are inevitable, but it is interesting to notice that the system is quite robust to such outliers, as the next error values are again in a reasonable range. It takes about 2000 steps to reduce this error if only the terminal limb transformation is adaptive and 1000 steps if all the limb transformation are adaptive. This takes two to three minutes as updates are performed at a rate of approximately 10 Herz.

6. Discussion

Our results show that the suggested framework for body schema learning is effective. Indeed, simulation results show that it can learn the structure of a 24 DOFs robot by tracking only the end-effectors. Furthermore, real robot experiments show that the method is applicable in a real setting. Given the fast developments of actuator and sensor technologies, it may possible in a near future to use this algorithim with all DOFs of real humanoid robot. For practical pruposes, as it may be cumbersome to gather millions of points in the training set, it may be more efficient to apply this algorithm to a smaller set of data, well distributed in the joint ranges and perform the iterations on this more restricted set.

The model presented here differs from earlier work on this topic mainly in that the

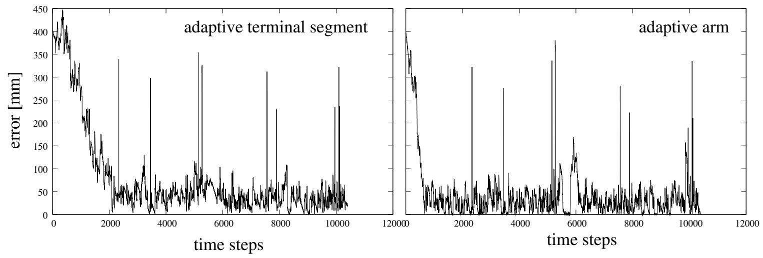

Fig. 13. Incorporating the tool in the body schema. Those graphs show the evolution of the distance between the end-effector position seen by the hand on one hand, and computed using proprioception and the body schema on the other hand. On the left, only the terminal limb is adaptive, on the right all the limbs are adaptive. In both cases the body schema adaptation reduces the error from approximately 40 cm to 5 cm.

knowledge of the kinematic structure is given in advance. In other words, the robot not only knows the number of DOFs, but also how they are arranged (serially or in parallel). Moreover, the model takes explicitly advantage of the fact that those are rotative joints, which was usually not done in other works. Thus the effectiveness of the learning algorithm relies on this a priori knowledge. However, we believe that it is reasonable to assume such an a priori knowledge, as the kinematic structure of the humanoids usually do not change over time. Similarly the kinematic structure of humans is fixed and does not evolve. Limbs grow, but new joints do not appear in a lifetime.

This last point illustrate that the results presented in this paper put into light several interesting questions related to the learning of the body schema in robot and humans. In its present state our model cannot learn a new body structure, and to our knowledge, it is unknown whether humans can adapt to different body structures. Another point raised by our results is that some body structures seem to be easier to learn than others, irrespective of the number of joints. It would be very informative to investigate whether this is an artifact of the learning procedure of if there is some intrinsic complexity dependent on the succession of rotation axes. Such various degrees of complexity have been described within the context of manipulator inverse kinematics,[5] where manipulators with the same number of DOFs can have different number of self-motion manifolds. If this complexity is intrinsic, it would be interesting to look at the complexity of the human body schema.

Because this model of the body schema consists in a hierarchy of coordinate system transformations, it can also be seen as a model of peripersonal space representation. Although the concepts of body schema and peripersonal space representation have traditionally been considered separately, the recent theories of motor perception[8] and of sensorimotor account of consciousness[24] suggest that those two concepts should be unified in a single framework. Indeed perception and action are

densely intertwined processes,[13] that cannot be easily isolated one from another. However, within this work, we did not yet integrate the motor modality in our framework. This can be done by deriving the Jacobian from the body schema, by differentiating (2) with respect to the angles. This is one of the next steps intended for a further improvement of this model.
Another intended development is to extend the implementation this model on the humanoid robot, and try learn the whole body schema in a real setting. This is certainly not an easy task, as explained in section 5.1. It is interesting to note that most of the difficulties mentioned there are alleviated when considering babies learning their body schema. Indeed, they have stereo-vision and a tactile skin, it takes them several months to learn their body schema, they are very flexible and have thus no problem to see their feet, hands and legs, and their compliant limbs allow them to explore their workspace much less cautiously (and thus faster) than the robot is allowed to. So in a way, it may be much harder for a robot to learn its body schema than for a human.

7. Conclusion

In this paper, we have provided a general model of the body schema for humanoid robots, based on a hierarchy of frames of reference transformations reflecting the body structure. This model integrates the visual and proprioceptive modalities and allows a robot to form a coherent image of itself. Moreover, we provide a way to adapt this image as a result of sensory experience. This is done in a completely online and self-supervised manner so that the system can continuously adapt itself, as is the case for humans, who can adapt to visual shifts, tool use and so on. The effectivity of the model has been shown in simulations with a 24-DOF humanoid robot and in a real setting for tool adaptation.

The results obtained suggest that the kinematic function (and for that matter its inverse) of the robot should not be dissociated from its sensory abilities, as it precisely provides the relationship between multiple sensory information. Learning this relationship allows a permanent recalibration of the sensors and is likely to contribute to the precision, autonomy and adaptability of humanoid robots.

Acknowledgments

This work was funded by the Swiss National Science Foundation and EU Project IST-2004-004370 RobotCub.

Appendix A. Proof of Theorem 1

We consider equations (3) and (4) as a continuous time dynamical system in the parameter space defined by $\mathbf{l}$ and $\mathbf{a}$. Indeed, the learning step ϵ can be interpreted as

the integration constant of the dynamical system, in which case $\epsilon^{-1}\Delta\mathbf{l}$ and $\epsilon^{-1}\Delta\mathbf{a}$ tend respectively to $\frac{\partial}{\partial t}\mathbf{l}$ and $\frac{\partial}{\partial t}\mathbf{a}$, where t denotes the time.
In order to simplfy the notation, we put the two parameter vectors $\mathbf{l}$ and $\mathbf{a}$ into a single parameter vector $\mathbf{p} = [\mathbf{l}^T \mathbf{a}^T]^T$. We then consider the function $\mathbf{E}(\mathbf{p})$ defined by:

$$\mathbf{E}(\mathbf{p}) = \langle \mathbf{E}^\theta(\mathbf{p})\rangle, \quad \text{with} \quad \mathbf{E}^\theta(p) = \frac{1}{2}\left|\left|\mathcal{T}^*(\mathbf{v}_n) - \mathcal{T}_\mathbf{p}(\mathbf{v}_n)\right|\right|^2, \tag{A.1}$$

where $\mathcal{T}_\mathbf{p}$ is the estimated rigid-body transformation parametrized by $\mathbf{p}$ and dependent on the joint configuration θ ($\mathcal{T}_\mathbf{p} = \mathbf{l} + \mathbf{R}^\theta_\mathbf{a}(\mathbf{v}_n)$),$\mathcal{T}^*(\mathbf{v}_n) = \mathbf{l}^* + \mathbf{R}^\theta_{\mathbf{a}^*}(\mathbf{v}_n)$ is the real underlying transformation, and $\langle\rangle$ denotes the expectation operator assuming θ follows symmetric pdf centered on 0. Since grad is a linear operator, the $\langle[\frac{\partial}{\partial t}\mathbf{l}\ \frac{\partial}{\partial t}\mathbf{a}]\rangle = -\text{grad}\mathbf{E}(\mathbf{p})$. So (3) and (4) correspond to a gradient descent on $\mathbf{E}$.
It remains to be shown that if $\frac{\partial}{\partial t}\mathbf{E}(\mathbf{p}) = 0$, then $\mathcal{T}_p = \mathcal{T}^*$, which amounts to saying that there is no local minima to $\mathbf{E}$. To this end we compute the gradient of $\mathbf{E}$.

$$\text{grad}_\mathbf{p}\mathbf{E} = \langle\big(\mathcal{T}_\mathbf{p}(\mathbf{v}) - \mathcal{T}^*(\mathbf{v})\big)^T \frac{\partial}{\partial \mathbf{p}}\mathcal{T}_\mathbf{p}\rangle \tag{A.2}$$

$$= \langle\big(\mathcal{T}_\mathbf{p}(\mathbf{v}) - \mathcal{T}^*(\mathbf{v})\big)^T [\mathbf{I}\ \bar{\mathbf{R}}^\theta_\mathbf{a}(\mathbf{v})]\rangle \tag{A.3}$$

We can divide the gradient vector onto its two components $\text{grad}_\mathbf{p}\mathbf{E} = [\text{grad}_\mathbf{l}\mathbf{E}\ \text{grad}_\mathbf{a}\mathbf{E}]$. We then have

$$\text{grad}_\mathbf{l}\mathbf{E} = \big(\langle\sin\theta\rangle(\mathbf{a}-\mathbf{a}^*)\times\mathbf{v} + \langle 1-\cos\theta\rangle(\mathbf{v}^T\mathbf{a}\mathbf{a} - \mathbf{v}^T\mathbf{a}^*\mathbf{a}^*) + \mathbf{l} - \mathbf{l}^*\big)^T \tag{A.4}$$

$$\begin{aligned}\text{grad}_\mathbf{a}\mathbf{E} = \langle\big(&\sin\theta(\mathbf{a}-\mathbf{a}^*)\times\mathbf{v} + (1-\cos\theta)\big(\mathbf{v}^T\mathbf{a}\mathbf{a} - \mathbf{v}^T\mathbf{a}^*\mathbf{a}^*\big) + \mathbf{l} - \mathbf{l}^*\big)^T \\ &\big(\sin\theta\mathbf{v}\!\uparrow + (1-\cos\theta)(\mathbf{a}\mathbf{v}^T + (\mathbf{a}^T\mathbf{v})\mathbf{I})\big)\rangle\end{aligned} \tag{A.5}$$

Developping this product leads to the following sum:

$$\begin{aligned}\text{grad}_\mathbf{a}\mathbf{E} = &\langle\sin^2\theta\rangle\mathbf{v}\times\big((\mathbf{a}-\mathbf{a}^*)\times\mathbf{v}\big) + \langle(1-\cos\theta)\sin\theta\rangle((\mathbf{v}\times(\mathbf{a}-\mathbf{a}^*))^T(\mathbf{a}\mathbf{v}^T + \mathbf{a}^T\mathbf{v}\mathbf{I}) \\ &+\langle\sin\theta(1-\cos\theta)\rangle\big(\mathbf{v}\times(\mathbf{v}^T\mathbf{a}\mathbf{a} - \mathbf{v}^T\mathbf{a}^*\mathbf{a}^*)\big) \\ &+\langle(1-\cos\theta)^2\rangle(\mathbf{v}^T\mathbf{a}\mathbf{a} - \mathbf{v}^T\mathbf{a}^*\mathbf{a}^*)^T(\mathbf{a}\mathbf{v}^T + \mathbf{a}^T\mathbf{v}\mathbf{I}) \\ &+(\mathbf{l}-\mathbf{l}^*)^T\big(\langle\sin\theta\rangle\mathbf{v}\!\uparrow + \langle 1-\cos\theta\rangle(\mathbf{a}\mathbf{v}^T + \mathbf{a}^T\mathbf{v}\mathbf{I})\big)\end{aligned} \tag{A.6}$$

As the θ are zero-centered and symmetric, we have $\langle\sin\theta\rangle = \langle\sin(\theta)(1-\cos\theta)\rangle = 0$. Thus, we have

$$\text{grad}_\mathbf{l}\mathbf{E} = \big(\langle 1-\cos\theta\rangle(\mathbf{v}^T\mathbf{a}\mathbf{a} - \mathbf{v}^T\mathbf{a}^*\mathbf{a}^*) + \mathbf{l} - \mathbf{l}^*\big)^T \tag{A.7}$$

$$\begin{aligned}\text{grad}_\mathbf{a}\mathbf{E} = &\langle\sin^2\theta\rangle\big(||\mathbf{v}||^2(\mathbf{a}-\mathbf{a}^*) - \mathbf{v}^T(\mathbf{a}-\mathbf{a}^*)\mathbf{v}\big) \\ &+\langle(1-\cos\theta)^2\rangle(\mathbf{v}^T\mathbf{a}\mathbf{a} - \mathbf{v}^T\mathbf{a}^*\mathbf{a}^*)^T(\mathbf{a}\mathbf{v}^T + \mathbf{a}^T\mathbf{v}\mathbf{I}) \\ &+\langle 1-\cos\theta\rangle(\mathbf{l}-\mathbf{l}^*)^T(\mathbf{a}\mathbf{v}^T + \mathbf{a}^T\mathbf{v}\mathbf{I})\end{aligned} \tag{A.8}$$

Now, $\mathbf{p}$ is a fixed point of the dynamical system if and only $\text{grad}_\mathbf{l}\mathbf{E} = \mathbf{0}$ and $\hat{\mathbf{a}}$ is colinear to $\mathbf{a}$ (thus taking the normalization of $\mathbf{a}$ into account). Setting $\text{grad}_\mathbf{l}\mathbf{E}$ to zero, we obtain

$$
\begin{aligned}
\mathbf{l}-\mathbf{l}^* &= -\langle 1-\cos\theta\rangle(\mathbf{v}^T\mathbf{a}\mathbf{a}-\mathbf{v}^T\mathbf{a}^*\mathbf{a}^*) && \text{(A.9)}\\
\text{grad}_{\mathbf{a}}\mathbf{E} &= \langle\sin^2\theta\rangle\big(\|\mathbf{v}\|^2(\mathbf{a}-\mathbf{a}^*)-\mathbf{v}^T(\mathbf{a}-\mathbf{a}^*)\mathbf{v}\big)^T + \\
&\quad \big(\langle(1-\cos\theta)^2\rangle-\langle(1-\cos\theta)\rangle^2\big)(\mathbf{v}^T\mathbf{a}\mathbf{a}-\mathbf{v}^T\mathbf{a}^*\mathbf{a}^*)^T(\mathbf{a}\mathbf{v}^T+\mathbf{a}^T\mathbf{v}\mathbf{I}) && \text{(A.10)}\\
&= \langle\sin^2\theta\rangle\big(\|\mathbf{v}\|^2(\mathbf{a}-\mathbf{a}^*)-\mathbf{v}^T(\mathbf{a}-\mathbf{a}^*)\mathbf{v}\big)^T\\
&\quad +\text{var}(\cos\theta)\big(\mathbf{v}^T\mathbf{a}\mathbf{v}+(\mathbf{v}^T\mathbf{a})^2\mathbf{a}-\mathbf{v}^T\mathbf{a}^*\mathbf{a}^T\mathbf{a}^*\mathbf{v}-\mathbf{v}^T\mathbf{a}^*\mathbf{v}^T\mathbf{a}\mathbf{a}^*\big)^T && \text{(A.11)}
\end{aligned}
$$

With no loss of generality and to simplify the notation, we can drop $\|\mathbf{v}\|^2$ and consider $\mathbf{v}$ to be of unit norm. To further lighten the notation, we define $c = \text{var}(\cos\theta)$ and $s = \langle\sin^2\theta\rangle = \text{var}(\sin\theta)$. We have a fixed point if and only if $\text{grad}_{\mathbf{a}}\mathbf{E}$ is colinear to $\mathbf{a}$ or equivalently if it is perpendicular to the projection of any vector $\mathbf{r}$ on the plane orthogonal to $\mathbf{a}$, $(\mathbf{I}-\mathbf{a}\mathbf{a}^T)\mathbf{r}$. Hence we have

$$
\begin{aligned}
\mathbf{0} &= \text{grad}_{\mathbf{a}}\mathbf{E}\cdot(\mathbf{I}-\mathbf{a}\mathbf{a}^T) && \text{(A.12)}\\
&= s(-\mathbf{a}^*+\mathbf{a}^T\mathbf{a}^*\mathbf{a}-\mathbf{v}^T\mathbf{a}\mathbf{v}+\mathbf{v}^T\mathbf{a}^*\mathbf{v}+(\mathbf{v}^T\mathbf{a})^2\mathbf{a}-\mathbf{v}^T\mathbf{a}\mathbf{v}^T\mathbf{a}^*\mathbf{a})+\\
&\quad c(\mathbf{v}^T\mathbf{a}\mathbf{v}-(\mathbf{v}^T\mathbf{a})^2\mathbf{a}-\mathbf{v}^T\mathbf{a}^*\mathbf{a}^T\mathbf{a}^*\mathbf{v}+\mathbf{v}^T\mathbf{a}\mathbf{v}^T\mathbf{a}^*\mathbf{a}^T\mathbf{a}^*\mathbf{a}-\mathbf{v}^T\mathbf{a}^*\mathbf{v}^T\mathbf{a}\mathbf{a}^*+\mathbf{v}^T\mathbf{a}^*\mathbf{v}^T\mathbf{a}\mathbf{a}^T\mathbf{a}^*\mathbf{a}) && \text{(A.13)}\\
&= s(\mathbf{v}^T\mathbf{a}^*-\mathbf{v}^T\mathbf{a})(\mathbf{v}-\mathbf{a}^T\mathbf{v}\mathbf{a})-\mathbf{a}^*+\mathbf{a}^T\mathbf{a}^*\mathbf{a}\\
&\quad +c(\mathbf{v}^T\mathbf{a}-\mathbf{v}^T\mathbf{a}^*\mathbf{a}^T\mathbf{a}^*)(\mathbf{v}-\mathbf{a}^T\mathbf{v}\mathbf{a})-\mathbf{v}^T\mathbf{a}\mathbf{v}^T\mathbf{a}^*(\mathbf{a}^*-\mathbf{a}^T\mathbf{a}^*\mathbf{a}) && \text{(A.14)}\\
&= \big(s(\mathbf{v}^T\mathbf{a}^*-\mathbf{v}^T\mathbf{a})+c(\mathbf{v}^T\mathbf{a}-\mathbf{v}^T\mathbf{a}^*\mathbf{a}^T\mathbf{a}^*)\big)(\mathbf{v}-\mathbf{a}^T\mathbf{v}\mathbf{a})-\big(c\mathbf{v}^T\mathbf{a}\mathbf{v}^T\mathbf{a}^*+s\big)(\mathbf{a}^*-\mathbf{a}^T\mathbf{a}^*\mathbf{a}) && \text{(A.15)}
\end{aligned}
$$

We now notice that this last expression is the weighted sum of two vectors $\mathbf{v}-\mathbf{a}^T\mathbf{v}\mathbf{a}$ and $\mathbf{a}^*-\mathbf{a}^T\mathbf{a}^*\mathbf{a}$. Those are respectively the projections of $\mathbf{v}$ and $\mathbf{a}^*$ on the plane orthognal to $\mathbf{a}$. This implies that $\mathbf{a},\mathbf{v},\mathbf{a}^*$ are coplanar and we can rewrite (A.15) as

$$
0 = \big(s(\cos\alpha^*-\cos\alpha)+c(\cos\alpha-\cos\alpha^*\cos\gamma)\big)\sin\alpha-\big(s+c\cos\alpha\cos\alpha^*\big)\sin\gamma, \qquad \text{(A.16)}
$$

where α is the angle between $\mathbf{v}$ and $\mathbf{a}$, α^* is the angle between $\mathbf{a}^*$ and $\mathbf{v}$, and γ is the angle between $\mathbf{a}$ and $\mathbf{a}^*$, as depicted in Fig. 14. Since $\gamma = \alpha^*-\alpha$, we have

$$
\sin\gamma = \sin\alpha^*\cos\alpha-\cos\alpha^*\sin\alpha \qquad \cos\gamma = \cos\alpha\cos\alpha^*+\sin\alpha\sin\alpha^* \qquad \text{(A.17)}
$$

Inserting (A.17) into (A.16) yields

$$
\begin{aligned}
0 &= \big(s(\cos\alpha^*-\cos\alpha)+c(\cos\alpha-\cos\alpha^*(\cos\alpha\cos\alpha^*+\sin\alpha\sin\alpha^*))\big)\sin\alpha\\
&\quad -\big(s+c\cos\alpha\cos\alpha^*\big)(\sin\alpha^*\cos\alpha-\cos\alpha^*\sin\alpha) && \text{(A.18)}\\
&= \big(c-s\big)\cos\alpha\sin\alpha+2s\cos\alpha^*\sin\alpha-s\sin\alpha^*\cos\alpha+c(-\cos^2\alpha^*\cos\alpha\sin\alpha\\
&\quad -\cos\alpha^*\sin\alpha^*\sin^2\alpha-\cos\alpha^*\sin\alpha^*\cos^2\alpha+\cos^2\alpha^*\cos\alpha\sin\alpha) && \text{(A.19)}\\
&= (c-s)\cos\alpha\sin\alpha+2s\cos\alpha^*\sin\alpha-s\sin\alpha^*\cos\alpha-c\cos\alpha^*\sin\alpha^* && \text{(A.20)}\\
&= \frac{1}{2}(c-s)\sin(2\alpha)+s\sin(\alpha-\alpha^*)+s\cos\alpha^*\sin\alpha-\frac{c}{2}\sin(2\alpha^*) && \text{(A.21)}\\
&= \frac{1}{2}(c-s)\sin(2\alpha)+\frac{3s}{2}\sin(\alpha-\alpha^*)+\frac{s}{2}\sin(\alpha+\alpha^*)-\frac{c}{2}\sin(2\alpha^*) && \text{(A.22)}
\end{aligned}
$$

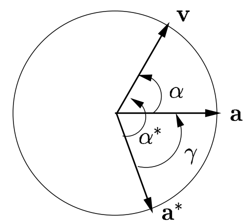

Fig. 14. Left: Going from (A.15) to (A.16). Right: All three fixed points $\mathbf{a}^*$, $\mathbf{a}_1$ and $\mathbf{a}_2$ are located on the plane defined by $\mathbf{a}^*$ and $\mathbf{v}$.

By performing the following change of variables: $\beta = \alpha - \alpha^*$, $\phi = \alpha + \alpha^*$, we can rewrite this last equation as

$$
\begin{aligned}
0 &= \frac{1}{2}(c-s)\sin(\beta+\phi) + \frac{3s}{2}\sin\beta + \frac{s}{2}\sin\phi + \frac{c}{2}\sin(\beta-\phi) && \text{(A.23)}\\
&= c\sin\beta\cos\phi + \frac{3s}{2}\sin\beta + \frac{s}{2}\big(\sin\phi - \sin(\beta+\phi)\big) && \text{(A.24)}\\
&= (c\cos\phi + \frac{3s}{2})\sin\beta - s\sin\frac{\beta}{2}\cos(\frac{\beta}{2}+\phi) && \text{(A.25)}\\
&= (2c\cos\phi + 3s)\sin\frac{\beta}{2}\cos\frac{\beta}{2} - s\sin\frac{\beta}{2}\cos(\frac{\beta}{2}+\phi) && \text{(A.26)}\\
&= \sin\frac{\beta}{2}\big((2c\cos\phi+3s)\cos\frac{\beta}{2} - s\cos(\frac{\beta}{2}+\phi)\big) && \text{(A.27)}\\
&= \sin(\tfrac{\alpha-\alpha^*}{2})\big(2c\cos(\alpha+\alpha^*)\cos(\tfrac{\alpha-\alpha^*}{2}) + 3s\cos(\tfrac{\alpha-\alpha^*}{2}) - s\cos(\tfrac{3\alpha+\alpha^*}{2})\big) && \text{(A.28)}\\
&= \sin(\frac{\alpha-\alpha^*}{2})\big((c-s)\cos(\frac{3\alpha+\alpha^*}{2}) + 3s\cos(\frac{\alpha-\alpha^*}{2}) + c\cos(\frac{\alpha+3\alpha^*}{2})\big) && \text{(A.29)}
\end{aligned}
$$

The first factor indicates that, as expected, $\alpha = \alpha^*$ is a solution. The other solutions are given by the zeros of the second factor. If $| c - s | \leq | s |$, i.e $c < 2s$, this factor has at most two zeros, so there are at most three solutions to (A.29). Since they all lie on a circle, there is one minimum, one maximum and possibly a plateau. So there is only one minimum in $\mathbf{E}$, which completes the proof.

References

1. T. N. Aflalo and M. S. A. Graziano. Relationship between unconstrained arm movements and single-neuron firing in the macaque motor cortex. *Journal of Neuroscience*, 27:2760–2780, 2007.
2. O.A. Bauchau and L. Trainelli. The vectorial parameterization of rotation. *Nonlinear Dynamics*, 32:71–92, 2003.
3. F.L. Bedford. Keeping perception accurate. *Trends in Cognitive Science*, 1999.
4. M. Botvinick. Probing the neural basis of body ownership. *Science*, 305(5685):782 – 783, 2004.
5. J.W. Burdick. On the inverse kinematics of redundant manipulators: Characterization of the self-motion manifolds. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 264–270, 1989.

6. Y. Burnod, P. Baraduc, A. Battaglia-Mayer, E. Guigon, E. Koechlin, S. Ferraina, F. Laquaniti, and R. Caminiti. Parieto-frontal coding of reaching: an integrated framework. *Experimental Brain Research*, 129:325–346, 1999.
7. A. D'Souza, S. Vijayakumar, and S. Schaal. Learning inverse kinematics. In *Proceedings of the 2001 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 298–303, 2001.
8. L. Fadiga, L. Fogassi, V. Gallese, and G. Rizzolatti. Visuomotor neurons: ambiguity of the discharge or 'motor' perception? *International Journal of Psychophysiology*, 35:165–177, 2000.
9. S. Fuke, M. Ogino, and M. Asada. Body image constructed from motor and tactile images with visual information. *International Journal of Humanoid Robotics*, 4(2):347–364, 2007.
10. Chris Gaskett and Gordon Cheng. Online learning of a motor map for humanoid robot reaching. In *Proceedings of the 2nd International Conference on Computational Intelligence, Robotics and Autonomous Systems*, 2003.
11. M.S.A. Graziano and C.G. Gross. A bimodal map of space: somatosensory receptive fields in the macaque putamen with corresponding visual receptive fields. *Experimental Brain Research*, 97:96–109, 1993.
12. A. Guez and Z. Ahmad. Accelerated convergence in the inverse kinematics via multilayer feedforward networks. In *Proceedings of the International Joint Conference on Neural Networks*, pages 341–344, 1989.
13. T. Higuchi, K. Imanaka, and A.E. Patla. Action-oriented representation of peripersonal and extrapersonal space: Insights from manual and locomotor action. *Japanese Psychological Research*, 2006.
14. N.P. Holmes and C. Spence. The body schema and multisensory representation(s) of peripersonal space. *Cognitive Process*, 2004.
15. H. Imamizu, Y. Uno, and M. Kawato. Adaptive internal model of intrinsic kinematics involved in learning an aiming task. *Journal of experimental pschology. Human perception and performance*, 24:812–829, 1998.
16. T. Jellema, G. Maassen, and D. I. Perrett. Single cell integration and aminate form, motion and location in the superior temporal cortex of the macaque monkey. *Cerebral Cortex*, 14:781–790, 2004.
17. Nimit Kumar and Laxmidhar Behera. Visual-motor coordination using a quantum clustering based neural control scheme. *Neural Processing Letters*, 20(1):11–22, 2004.
18. M Kuperstein. Neural model for adaptive hand-eye coordination for single postures. *Science*, 1988.
19. E. Làdavas. Functional and dynamical properties of visual peripersonal space. *Trends in Cognitive Sciences*, 2002.
20. M. Lopes and J. Santos-Victor. Visual learning by imitation with motor representations. *IEEE Transactions on Systems, Man and Cybernetics, Part B*, 2005.
21. A. Maravita and A. Iriki. Tools for the body (schema). *Trends in Cognitive Sciences*, 8(2), 2004.
22. G. Metta, G. Sandini, and J. Konczak. A developmental approach to visually-guided reaching in artificial systems. *Neural Networks*, 12(10):1413–1427, 1999.
23. Y. Nakamura and H. Hanafusa. Inverse kinematics solutions with singularity robustness for robot manipulator control. *ASME Journal of Dynamic Systems, Measurement, and Control*, 108:163–171, 1986.
24. J.K. O'Regan and A. Noë. A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, 24(5):939–1011, 2001.
25. A. Pouget, S Deneve, and J. R. Duhamel. A computational perspective on the neural

basis of multisensory spatial representations. *Nature Review Neuroscience*, 3:741–747, 2002.
26. M. J. Prud'homme and J. F. Kalaska. Proprioceptive activity in primate primary somatosensory cortex during active arm reaching movements. *Journal of Neurophysiology*, 72(5):2280–2301, 1994.
27. G. Rizzolatti, L. Fadiga, L. Fogassi, and V. Gallese. The space around us. *Science*, 1997.
28. G. Sandini, G. Metta, and D. Vernon. Robotcub: an open framework for research in embodied cognition. In *Proceedings of the 2004 IEEE/RAS International Conference on Humanoid Robots*, pages 13–32, 2004.
29. L. Sciavicco and B. Siciliano. A solution algorithm to the inverse kinematic problem for redundant manipulators. *IEEE Journal of Robotics and Automation*, 1988.
30. S. I. Tillery, J. F. Soechting, and T. J. Ebner. Somatosensory cortical activity in relation to arm posture: nonuniform spatial tuning. *Journal of Neurophysiology*, 76:2423–2438, 1996.
31. Jorg A. Walter and Klaus J. Schulten. Implementation of self-organizing neural networks for visuo-motor control of an industrial robot. *IEEE Transactions on Neural Networks*, 4(1):86–95, January 1993.
32. C.W. Wampler. Manipulator inverse kinematic solutions based on vector formulations and damped least-squares methods. *IEEE Transactions on Systems, Man and Cybernetics, Part C*, 16(1):93–101, 1986.
33. R.B. Welch. Adaptation of space perception. In K.R. Boff, L. Kaufman, and J.P. Thomas, editors, *Handbook of perception and human performance*, pages 24:1–24:45, 1986.
34. D.E. Whitney. Resolved motion rate control of manipulators and human prosteses. *IEEE Transactions on Man-Machine Systems*, 10(2):47–52, 1969.