Robot body model self calibration through contact events

Rodrigo Zenha[1], Pedro Vicente[1], Lorenzo Jamone[2,1] and Alexandre Bernardino[1]

***Abstract*— Both humans and robots need a model of their body, *i.e.*, either a neural or numerical representation of it, to successfully interact with the environment; notably, such model needs to be continuously calibrated, *i.e.*, though either neural plasticity or some form of parameter estimation, to cope with changes over time. In this work, we present an online strategy that allows a robot to self-calibrate its body model by touching known planar surfaces (*e.g.*, walls). This is achieved through an adaptive parameter estimation (Extended Kalman Filter) which makes use of planar constraints obtained at each contact detection. We compare different update methods using a realistic simulation of the iCub humanoid robot, showing that the model inaccuracies can be reduced by more than 80%.**

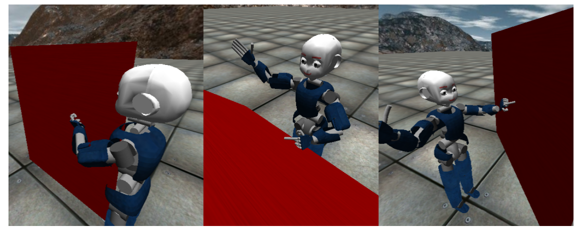

Fig. 1. Simulation of an iCub robot reaching for three different walls in its workspace.

I. INTRODUCTION

Humans develop a neural representation of their body (*i.e.*, a body schema [1]) through an incremental learning process that starts in early infancy [2], and likely even prenatally [3], and goes through continuous adaptations over time, based on multimodal sensorimotor information acquired during motor experience [4]: visual, tactile, proprioceptive. This (physical) body schema is a crucial part of human self-awareness and supports the precise control of body movements, coping with the morphological changes that occur in the body over time, *e.g.*: body growth, tool assimilation. Clearly, endowing artificial agents with similar learning and adaptation capabilities is a major challenge for cognitive robotics and it paves the way for the next generation of robots able to act in complex environments. Indeed, an accurate model of the robot structure (*i.e.*, kinematics) is required for any robotic task. A variety of factors, such as friction, worn joints and bended rigid bodies, induce changes to the robot kinematic model over time. As a consequence, such model needs to be calibrated, either offline or online: clearly, online techniques are desirable as they can be performed by the robot during its normal operations, without requiring to collect data in a separate procedure.

In this work, we develop an incremental calibration strategy that is performed automatically by the robot during the execution of any arm movement that involves contacts on known planar surfaces, using the Extended Kalman Filter for adaptive parameter estimation. To do so, we make use of contact (pressure sensitive fingertips) and proprioceptive (joint encoders) sensors, commonly present in many robots.

[1] Institute for Systems and Robotics, Instituto Superior Técnico, Universidade de Lisboa, Lisbon, Portugal. rodrigo.n.zenha@tecnico.ulisboa.pt , {pvicente,alex}@isr.tecnico.ulisboa.pt

[2] ARQ (Advanced Robotics at Queen Mary), School of Electronic Engineering and Computer Science, Queen Mary University of London, UK l.jamone@qmul.ac.uk

This work was partially supported by FCT (UID/EEA/50009/2013 and PD/BD/135115/2017), and by EPSRC UK (project NCNR, National Centre for Nuclear Robotics, EP/R02572X/1). We acknowledge the support of NVIDIA with the donation of a GPU.

II. RELATED WORK

Body schema learning and adaptation in robots has been widely studied (see [5] and [6] for a survey on this topic).

Online solutions using on-board robot cameras have been studied [7], [8], [9], in which markers are used to easily detect the end-effector position; the inclusion of additional parts into the kinematic chain (*i.e.*, tools) has been considered as well [10]. Interestingly, goal-directed strategies ([11] and [12]) show to enhance body schema learning, for example, by reducing the time necessary for calibration convergence. Marker-free solutions have been explored [13], [14] in which the robot visual and proprioceptive information are compared by using a realistic 3D computer graphics model of the robot, to estimate simultaneously the robot hand pose and kinematic model, with a *Particle Filter*.

Contact information was used in [15] to develop an offline automatic kinematic chain calibration resorting to self-touch events, which was proven to be highly effective to optimize the robot model Denavit-Hartenberg parameters; however, joint angle measurements inaccuracies are not assumed, and a sensitive skin covering the robot body is required. Online calibration based on contact information has also been studied [16], using an implicit *Manifold Particle Filter*.

III. OUR APPROACH

Here we propose a novel online approach to adapt the kinematic model of the robot body using proprioception and contact sensors. Information about contacts on known planar surfaces feeds a low computational cost estimation method (Extended Kalman Filter) enabling real-time model adaptation.

A. Body Model

We identify the robot body model as the kinematic chain from the root reference frame to the end-effector, *i.e.*, the robot forward kinematics: $\mathbf{T}^{\mathrm{e}} = \mathcal{K}(\boldsymbol{\theta})$; with the associated inverse model: $\boldsymbol{\theta} = \mathcal{K}^{-1}(\mathbf{T}^{\mathrm{e}})$. Let us define $\mathbf{T}^{\mathrm{e}}$ as a 4x4 roto-translation matrix which encapsulates the pose of the end-effector on the root reference frame. Due to modulation errors, we only get an estimation of the robot kinematics function ($\widehat{\mathcal{K}}(\cdot)$) based on the joint angles ($\boldsymbol{\theta}$) retrieved: $\mathbf{T}^{\mathrm{e}} = \widehat{\mathcal{K}}(\boldsymbol{\theta})$, where $\mathcal{K}(\boldsymbol{\theta})$ is the true robot kinematics. Moreover, due to the existence of calibration bias, the real joints angles are different from the ones read from proprioception (joints encoders): $\boldsymbol{\theta} = \boldsymbol{\theta}^{\mathrm{p}} + \boldsymbol{\beta}$, where $\boldsymbol{\theta}$ are the real angles values, $\boldsymbol{\theta}^{\mathrm{p}}$ are the encoders readings and $\boldsymbol{\beta}$ are the angular offsets. To better estimate the robot kinematics we take into account the estimate $\widehat{\boldsymbol{\beta}}$:

$$\mathbf{T}^{\mathrm{e}} = \widehat{\mathcal{K}}(\boldsymbol{\theta}^{\mathrm{p}} + \widehat{\boldsymbol{\beta}}). \tag{1}$$

The parameter vector and our state (to be estimated recursively) is defined as follows:

$$\boldsymbol{\beta} = [\beta^1 \beta^2 \ldots \beta^N]^T, \tag{2}$$

where N is the number of degrees of freedom of the robot's manipulator. Assuming the joint offsets to be slowly varying in time, we define the state-transition model as:

$$\boldsymbol{\beta}_t = \boldsymbol{\beta}_{t-1} + \boldsymbol{\varepsilon}_t, \tag{3}$$

where $\boldsymbol{\varepsilon}_t$ is a multivariate zero-mean Gaussian noise.

B. Observation Model

The observation model relates the system state $\boldsymbol{\beta}$ with a single measurement (z_k) from the contact sensors. We assume that there is a known planar surface described by:

$$\mathbf{x} \cdot \mathbf{n} - d = 0, \tag{4}$$

where $\mathbf{n} = [n_x, n_y, n_z]^T$ define the plane's normal vector ($||\mathbf{n}|| = 1$), and d is the plane's minimum distance to the robot root frame. In a simulation environment, both $\mathbf{n}$ and d are known *a priori*. In the real world, vision sensing could be used to estimate the pose of the surface, *e.g.*, using the Aruco marker [17], or computing a planar fit on depth point clouds from stereo vision. When a contact occurs, we are ensuring that the arm's end-effector 3D position ($\mathbf{x}^e$) respects Eq. (4). However, due to errors in the kinematic model, each set of coordinates $\widehat{\mathbf{x}}^e_k$ at an instant k, follows the equation: $\widehat{\mathbf{x}}^e_k \cdot \mathbf{n} - d = \alpha_k$, where α_k is the error produced by the model inaccuracies. The observation model is then defined as:

$$z_k(\boldsymbol{\theta}^p_k + \widehat{\boldsymbol{\beta}}_t) = \alpha_k + \delta_k \tag{5}$$

where $\widehat{\mathbf{x}}^e_k$ is retrieved using the forward kinematics (Eq. (1)), $\widehat{\boldsymbol{\beta}}_t$ are the offsets estimation at time instate t and δ_k is random Gaussian noise associated to an observation.

C. Parameter Estimation - Extended Kalman Filter (EKF)

We estimate the angular offsets $\boldsymbol{\beta}$ by exploiting contact constraints obtained at each end-effector contact with a surface. The strategy devised can be divided into two steps: i) a movement towards the target planar surface, stopping when a contact in the index finger occurs; and ii) a calibration phase in which an Extended Kalman Filter is fed with multisensory input (*i.e.*, proprioception, surface characteristics and contact feedback) adapting the state $\boldsymbol{\beta}$. We incorporate the dynamics (Eq. (3)) and observation model (Eq. (5)) into the EKF algorithm, which outputs the current offsets estimation, $\widehat{\boldsymbol{\beta}}_t$, and respective covariance matrix, Σ_t. The algorithm receives as input $\mathbf{z}_t$ and $\mathbf{H}_t$ which encapsulate a set of observations. When a contact is detected and the joint encoders readings are retrieved, we acquire an observation (z_k) as well as $\mathbf{H}_k = \nabla z_k(\boldsymbol{\theta}^p_k + \widehat{\boldsymbol{\beta}}_t)$.

We evaluate 3 strategies for new data incorporation:

1) Aggregation of Multiple Observations: Coupling together a varying number of contact constraints (k) before a filter update step (t):

$$\mathbf{H}_t = \begin{bmatrix} \mathbf{H}_{k-n} & \cdots & \mathbf{H}_{k-1} & \mathbf{H}_k \end{bmatrix}^T, \tag{6a}$$

$$\mathbf{z}_t = \begin{bmatrix} z_{k-n} & \cdots & z_{k-1} & z_k \end{bmatrix}^T, \tag{6b}$$

where t is the instant when we perform an estimation step. Using only one observation we have $\mathbf{H}_t \equiv \mathbf{H}_k$ and $\mathbf{z}_t \equiv z_k$.

2) Estimation Differential Entropy Evaluation: Upon contact, we compute the predicted next step estimation covariance matrix $\boldsymbol{\Sigma}_t$. Following the approach in [18], we decide to incorporate the new data if the current estimation differential entropy decreases compared to the previous estimation:

$$\frac{1}{2} \log_e \frac{|\boldsymbol{\Sigma}_{t-1}|}{|\boldsymbol{\Sigma}_t|} > 0, \tag{7}$$

and discard new data that does not bring innovative information. Here $|.|$ denotes a matrix determinant and $\boldsymbol{\Sigma}_{t-1}$ is the current estimation covariance matrix.

3) Anti-Windup Control (A-W): In [19] a technique is described in order to avoid windup associated to recursive estimation methods, such as the Recursive Least Squares (which can be recast into a EKF). They propose controlling the parameter random walk covariance matrix, $\mathbf{Q}(t)$ (associated to the EKF algorithm), so as to get $\boldsymbol{\Sigma}_t$ to achieve a constant pre-defined covariance matrix, $\mathbf{P}_d$, thus avoiding it to get unacceptable large eigenvalues. We use the same technique adapted to the EKF framework:

$$\mathbf{Q}(t) = \frac{\mathbf{P}_d \mathbf{H}_t \mathbf{H}_t^T \mathbf{P}_d}{\mathbf{R}(t) + \mathbf{H}_t^T \mathbf{P}_d \mathbf{H}_t}. \tag{8}$$

IV. Experimental Setup

Our system is evaluated on the iCub simulator [20], with a setup composed of three reachable surfaces (see Fig. 1) with *a priori* known parameters, $\mathbf{n}$ and d. For the first set of experiments, the robot reaches for a single surface; then we perform a second set of experiments where the contact events alternate between all three surfaces. We control the robot left arm and define its left index fingertip to be the end effector. The angular offsets (*i.e.*, model errors) on the seven DoFs of the arm are artificially set as: $\boldsymbol{\beta} = \begin{bmatrix} -11, 11, -7, -17, -7, -17, 7 \end{bmatrix}^T$ deg (values consistent with the calibration errors typically encountered on the real robot). We compare the results obtained with all new data incorporation methods (see subsection IV-A), running ten simulations for each method. For each experiment, we register 45 contact events (49 for the 7-contact setting). After each filter update, we compute the global estimation root mean squared error (RMSE) relative to the real offsets:

$$\textbf{RMSE} = \sqrt{\frac{1}{7}\sum_{i=1}^{7}(\widehat{\beta}^i - \beta^i)^2}. \tag{9}$$

A. New Data Incorporation Methods

We evaluate six new data incorporation methods:

a) 7-Contact (7C): Upon each contact detection, z_k and $\mathbf{H}_k$ are stored in $\mathbf{z}_t$ and $\mathbf{H}_t$, respectively. The system performs an estimation step after it collects 7 contact constraints (equal to the number of the iCub's arm DoFs).

b) Single Contact (SC): Upon each contact event, z_k and $\mathbf{H}_k$ are fed to the filter and an update step is performed.

c) Single Contact with Entropy (SC-E): Similar to the previous technique, but now for each new data obtained the system computes the predicted next step estimation covariance matrix and evaluates whether or not the new observation actively contributes to the estimation differential entropy reduction using Eq. (7). If the condition is not satisfied, the new observation is discarded.

d) Varying-Contact with Entropy (VC-E): Equivalent to the previous method, but each time a new observation fails to reduce the global estimation differential entropy, instead of being discarded, z_k and $\mathbf{H}_k$ are added to the matrices $\mathbf{z}_t$ and $\mathbf{H}_t$, respectively. Every time new data is obtained, the system evaluates if $\mathbf{z}_t$ and $\mathbf{H}_t$ are able to reduce the entropy of the next step estimation. If so, an update step is performed regarding all previously stored observations.

e) Single Contact with A-W (SC-AW): Equivalent to the Single Contact Estimation method, but $\mathbf{Q}_t$ is controlled with the anti-windup technique described in Eq. (8), rather than being a predefined matrix.

f) Single Contact with Entropy and A-W (SC-EAW): The system discards every new observation data which fails to reduce the next step estimation differential entropy, and controls $\mathbf{Q}_t$ matrix so as to avoid estimation windup from uncorrelated measures.

V. Results

We evaluate the results of $\boldsymbol{\beta}$ estimation for contacts on a single (Fig. 2(a)), and on three different surfaces[1] (Fig. 2(b)).

A. Contacts on a single surface

Results can be seen in Fig. 2(a). For method **7C**, the estimation error decreases slowly over each estimation step, and the model inaccuracies are reduced by 50% (**RMSE** ≈ 5.87 deg) after 35 contacts. Method **SC** shows a worse performance: the system keeps a slow steady error descend during the whole estimation, reducing the estimation absolute error by 50% after 37 contacts. Moreover, σ remains relatively high, due to the slow reaction of the filter against estimation steps taken in the wrong direction during periods of poor excitation. We attempted to improve method **SC** by evaluating the estimation differential entropy or by using the A-W control technique. For the first solution (method **SC-E**), we conclude that it is able to stabilize the filter performance, since the global estimation σ is reduced compared to the simple single contact setting; however, this happens at the cost of the filter converging sooner to a minimum. The latter solution (method **SC-AW**) achieves a better performance as well, being able to reduce both the time required for convergence, and the absolute error achieved after 45 contacts. Then, with method **VC-E** the estimation keeps a steady error descend up until the 16th contact (reducing the model inaccuracies by 50%), slowing the pace for the next contacts. Finally, the **SC-EAW** method, that combines the estimation differential entropy and the A-W control technique, shows three key features: i) the system is able to converge to a lower overall estimation minimum (reducing the estimation error by 15% compared to the single contact setting, and 65% overall, after 45 contact events), ii) the overall experiments σ is 30% lower compared to method **SC**, and iii) the system converges faster to a minimum (requires 10 steps to reduce the model inaccuracies by 50%).

B. Contacts on three different surfaces

By broadening the robot spatial exploration and performing contacts on 3 surfaces, we expect an overall better $\boldsymbol{\beta}$ estimation performance, since contact constraints obtained in this manner provide richer information to the filter. Fig. 2(b) shows the convergence of the different methods. Evidently, methods **7C** and **SC-E** are the ones which benefit less from the richer information acquired from contacts on 3 different surfaces, since they both converge to the highest estimation errors. The **7C** setting has a steady slow error reduction slope due to not being able to quickly compensate for estimation steps given on wrong directions. The **SC-E** estimation converges early (5th contact) to a local minimum, not being able to easily find relevant observations from there. The best results are obtained with the **SC-EAW** method, reducing both the estimation error by 45% relative to the single surface scenario (and 80% overall), and presenting the lowest overall σ value.

Then, in Fig. 2(c) we see the results of using the method **SC-EAW** for $\boldsymbol{\beta}$ estimation of 3 different sets of artificial offsets (10 experiments for each set). Up until the 60th contact event, all experiments reach an estimation minimum; moreover, independently of the true readings offsets, the filter is always able to reduce the estimation error to approximately 2.5 deg, illustrating the reliability of the devised strategy.

Finally, to evaluate the filter robustness to estimating slowly varying offsets values, we simulate a scenario where we start by introducing the artificial $\boldsymbol{\beta}$ values set described in section IV, inducing random small changes to the offsets values after every 15 contacts events, up until 60 contacts. The offsets values variation is achieved by the addition of random values with Gaussian distribution, with zero mean and a standard deviation of 1.7 deg. By using the **SC-EAW** method, we see the system stabilizing around $\mu[^\circ]$

[1]a video can be found in `https://youtu.be/EFx0OmRKTQg`

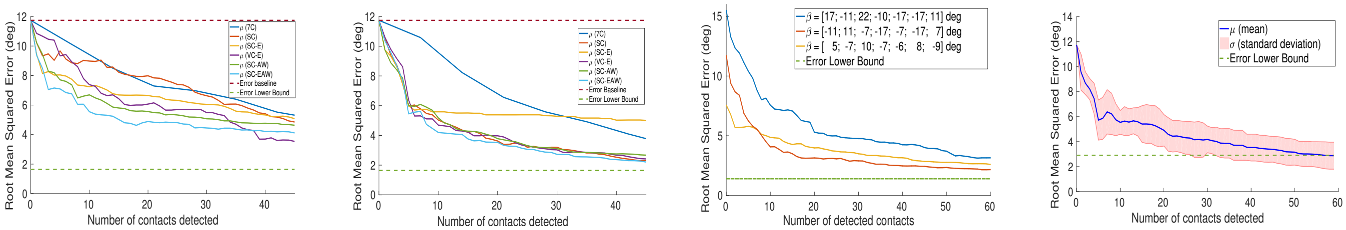

(a) Comparison between data incorporation methods for contacts one surface, using EKF.

(b) Comparison between all data incorporation methods using contacts over three surfaces.

(c) 3 different offsets sets estimation using the SC-EAW method and contacts on 3 surfaces.

(d) Slow changing $\boldsymbol{\beta}$ estimation using EKF (**SC-EAW**) and NLS settings.

Fig. 2. $\boldsymbol{\beta}$ estimation performance of all data incorporation methods with contacts on either one or three surfaces. The error lower bound is the estimation performed by the Non-Linear Least squares (NLS) algorithm. This method uses the information of 45 contact events, *i.e.*, a batch estimation method.

= 2.80 deg (Fig. 2(d)), corresponding to an improvement of 76% relative to the initial true offsets values. This experiment shows that the system is capable of coping with small changes in $\boldsymbol{\beta}$ throughout time, achieving a low estimation error after 60 contacts, keeping with a low standard deviation ($\sigma[°]$ = 0.73 deg), and even outperforming a batch estimation baseline.

VI. CONCLUSIONS AND FUTURE WORK

We devised a novel approach for online body schema adaptation based on proprioceptive and contact sensing, implemented on the iCub humanoid robot. The robot arm offsets (*i.e.*, model inaccuracies) are estimated with an EKF fed with contact constraints obtained during the execution of touch movements. Overall, our experiments prove our strategy successful in correcting model inaccuracies in real-time. Moreover, results show that touching surfaces with three different orientations is more effective than touching only one surface: model inaccuracies can be reduced up to 80% by performing contacts on 3 different surfaces, and up to 63% for contacts on a single surface.

Clearly, these simulations must be taken as a preliminary, though promising, result. To realize the proposed strategy in the real-world, the non-trivial problem of estimating the pose of the touched surfaces from on-board vision sensing should be solved. However, it could also be assumed that the pose of a few surfaces within a certain environment is known by the robot in advance: imagine for example a service robot in an indoor space, that could touch the ground and the walls. Then, it would be interesting to relate this strategy to other ones, *e.g.*, based on vision or self-touch, both in robots and in humans: could they be effectively combined? Do they appear in the developing child, maybe at different stages? Indeed, toddlers do interact with a number of familiar surfaces during development (*e.g.*, the walls of the crib, the floor at home, the eating table of the booster seat), and in some cases these surfaces have an almost constant orientation with respect to the body: do they provide stable (and important) references for the development and continuous self-calibration of the body schema?

REFERENCES

[1] G. Berlucchi and S. Aglioti, "The body in the brain: neural bases of corporeal awareness," *TINS*, no. 20(12), pp. 560–564, 1997.

[2] P. J. Marshall and A. N. Meltzoff, "Body maps in the infant brain," *Trends Cogn. Sci.*, no. 19(9), pp. 499–505, 2015.

[3] R. Joseph, "Fetal brain behavior and cognitive development," *Developmental Review*, vol. 20, pp. 81–98, 2000.

[4] C. von Hofsten, "An action perspective on motor development," *Trends Cogn. Sci.*, vol. 8, pp. 266–272, 2004.

[5] M. Hoffmann, H. Marques, A. Arieta, H. Sumioka, M. Lungarella, and R. Pfeifer, "Body schema in robotics: a review," *IEEE TAMD*, vol. 2, no. 4, pp. 304–324, 2010.

[6] P. Lanillos, E. Dean-Leon, and G. Cheng, "Enactive self: a study of engineering perspectives to obtain the sensorimotor self through enaction," in *IEEE ICDL-Epirob*, 2017.

[7] S. Ulbrich, V. de Angulo, T. Asfour, and C. Torras, "Rapid learning of humanoid body schemas with kinematic bezier maps," in *IEEE-RAS Intl. Conf. on Humanoid Robots*, 2009, pp. 431–438.

[8] L. Jamone, L. Natale, G. Metta, F. Nori, and G. Sandini, "Autonomous online learning of reaching behavior in a humanoid robot," *Intl. J. of Humanoid Robotics*, no. 9(3), pp. 1 250 017(1–26), 2012.

[9] L. Jamone, M. Brandao, L. Natale, K. Hashimoto, G. Sandini, and A. Takanishi, "Autonomous online generation of a motor representation of the workspace for intelligent whole-body reaching," *Rob. Auton. Syst.*, no. 64(4), pp. 556–567, 2014.

[10] L. Jamone, B. Damas, N. Endo, V. Santos, and A. Takanishi, "Incremental development of multiple tool models for robotic reaching through autonomous exploration," *Paladyn*, no. 3, pp. 113–127, 2013.

[11] L. Jamone, L. Natale, K. Hashimoto, G. Sandini, and A. Takanishi, "Learning task space control through goal directed exploration," in *2011 IEEE ROBIO*, pp. 702–708.

[12] M. Rolf, "Goal babbling with unknown ranges: A direction-sampling approach," in *2013 IEEE ICDL*, pp. 1–7.

[13] P. Vicente, L. Jamone, and A. Bernardino, "Online body schema adaptation based on internal mental simulation and multisensory feedback," *Front. robot. AI*, vol. 3, p. 7, 2016.

[14] ——, "Robotic hand pose estimation based on stereo vision and gpu-enabled internal graphical simulation," *J. Intell. Robotic Syst.*, vol. 83, no. 3-4, pp. 339–358, 2016.

[15] A. Roncone, M. Hoffmann, U. Pattacini, and G. Metta, "Automatic kinematic chain calibration using artificial skin: self-touch in the icub humanoid robot," in *2014 IEEE ICRA*, pp. 2305–2312.

[16] M. Klingensmith, C. Koval, S. Srinivasa, S. Pollard, and M. Kaess, "The manifold particle filter for state estimation on high-dimensional implicit manifolds," *2017 IEEE ICRA*, pp. 4673–4680.

[17] S. Jurado, R. Salinas, J. Cuevas, and J. Jiménez, "Automatic generation and detection of highly reliable fiducial markers under occlusion," *Pattern Recognition*, vol. 47, no. 6, pp. 2280–2292, 2014.

[18] J. Maye, H. Sommer, G. Agamennoni, R. Siegwart, and P. Furgale, "Online self-calibration for robotic systems," *Int. J. Robotics Res.*, vol. 35, no. 4, pp. 357–380, 2016.

[19] B. Stenlund and F. Gustafsson, "Avoiding windup in recursive parameter estimation," *Preprints of reglermöte 2002*, pp. 148–153, 2002.

[20] V. Tikhanoff, A. Cangelosi, P. Fitzpatrick, G. Metta, L. Natale, and F. Nori, "An open-source simulator for cognitive robotics research: the prototype of the icub humanoid robot simulator," in *Proceedings of the 8th workshop on PerMIS*. ACM, 2008, pp. 57–61.