

Active inference body perception and action for humanoid robots

Guillermo Oliver¹, Pablo Lanillos^{1,2}, Gordon Cheng¹

Abstract—Providing artificial agents with the same computational models of biological systems is a way to understand how intelligent behaviours may emerge. We present an *active inference* body perception and action model working for the first time in a humanoid robot. The model relies on the free energy principle proposed for the brain, where both perception and action goal is to minimise the prediction error through gradient descent on the variational free energy bound. The body state (latent variable) is inferred by minimising the difference between the observed (visual and proprioceptive) sensor values and the predicted ones. Simultaneously, the action makes sensory data sampling to better correspond to the prediction made by the inner model. We formalised and implemented the algorithm on the iCub robot and tested in 2D and 3D visual spaces for online adaptation to visual changes, sensory noise and discrepancies between the model and the real robot. We also compared our approach with classical inverse kinematics in a reaching task, analysing the suitability of such a neuroscience-inspired approach for real-world interaction. The algorithm gave the robot adaptive body perception and upper body reaching with head object tracking (toddler-like), and was able to incorporate visual features online (in a closed-loop manner) without increasing the computational complexity. Moreover, our model predicted involuntary actions in the presence of sensorimotor conflicts showing the path for a potential proof of active inference in humans.

Index Terms—Active inference, free energy optimization, Bio-inspired Perception, Predictive coding, Humanoid robots, iCub.

I. INTRODUCTION

The medical doctor and physicist Hermann von Helmholtz described visual perception as an unconscious mechanism that infers the world [1]. In other words, the brain has generative models that reconstruct the world from partial information. Nowadays, there is a scientific mainstream that describes the inner workings of the brain as those of a Bayesian inference machine [2], [3]. This approach supports that we can adjust the cues (visual, proprioceptive, tactile, etc.) contribution to our interpretation in a Bayesian optimal way considering sensors and motor uncertainties [4]. This implies that the brain is able to encode uncertainty not only for perception but also for acting in the world.

Optimal feedback control was proposed for modelling motor coordination under uncertainty [5]. Alternatively, *active inference* [6], defended that both perception and action are two sides of the same process: the free energy principle.

¹Institute for Cognitive Systems, Technical University of Munich, Arcisstrasse 21, 80333 Munich, Germany.

²Donders Institute for Brain, Cognition and Behaviour, the Netherlands.

This work has been supported by SELFCEPTION project EU Horizon 2020 Programme, grant agreement n. 741941 and the EU's Erasmus+ Programme. Supplementary video: <https://youtu.be/rdbbmwo4TY4>, code: [tobereleased](#).

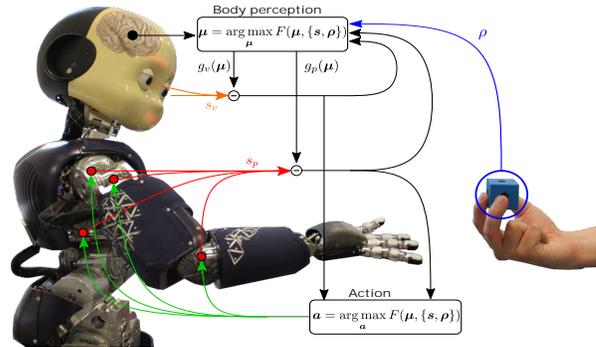


Fig. 1. Body perception and action via *active inference*. The robot infers its body (e.g., joint angles) by minimising the prediction error: discrepancy between the sensors (visual s_v and joint s_p) and their expected values ($g_v(\mu)$ and $g_p(\mu)$). In the presence of error it changes the perception of its body μ and generates an action a to reduce this discrepancy. Both are computed by optimising the free energy bound F . In a reaching task, the object is a causal variable ρ that acts as a perceptual attractor, producing an error in the desired sensory state and promoting an action towards the goal. The equilibrium point $\dot{\mu} = 0$ appears when the hand reaches the object. The head keeps the object in its visual field, improving the reaching performance.

This principle accounts for perception, action and learning through the minimization of *surprise*, i.e., the discrepancy between the current state and the predicted or desired one. According to this approach, free energy is a way of quantifying surprise and it can be optimised by changing the current beliefs (perception) or by acting on the environment (action) to adjust the difference between reality and prediction [7]. Active inference would allow living beings to infer and simultaneously adapt their body to uncertain environments [8], something just partially solved in robotics.

Hence, inspired by this principle, we present a robust and adaptive body perception and action mathematical model based on *active inference*, for a real humanoid robot (Fig. 1). Our approach describes artificial body perception and action, similar to biological systems, as a flexible and dynamic process that approximates the robot body latent state using the error between the expected and the observed sensory information. The proposed model enabled the robot to have adaptive body perception and to perform upper-body reaching behaviours even under high levels of sensor noise and discrepancies between the model and the real robot. We further compared our approach to standard inverse kinematics and analysed the relevance and the challenges of predictive coding interpretation solved through variational free energy methods for humanoid robotics. For reproducibility, the code, with the parameters used, is publicly available: [tobereleased](#).

A. Related work

Multisensory body perception has been widely studied in the literature, from the engineering point of view, and enables robots to estimate its body state by combining joint information with other sensors such as images and tactile cues. Bayesian estimation (e.g., particle filters) has been proved to achieve robust and accurate model-based robot arm tracking [9] even under occlusion [10], or for simultaneous joint offset parameters adaptation and hand pose estimation [11]. Furthermore, integrated visuomotor processes enabled humanoid robots to learn motor representations for robust reaching [12], [13], visuotactile motor representations for reaching and avoidance behaviours [14] or enabling visual robot body distinction for aiding the discovering of movable objects [15]. Free energy optimization for multisensory perception in a real robot was firstly presented in [16] working as an approximate Bayesian filter estimation. The robot was able to perceive its arm location fusing visual, proprioceptive and tactile information.

Active inference (under the free energy principle) managed to include the action within the multisensory perception process, explained as a classical spinal reflex arc pathway triggered by prediction errors. It has been mainly studied in theoretical or simulated conditions. Friston et al. [7] presented this neuroscience-inspired model as a generalization of the dynamic expectation-maximization algorithm. Recently, a few works analysed this approach for robotics in simulated conditions: one degree of freedom vehicle [17], a PR2 robot arm [18], and two degrees of freedom robot arm with generative function learning [19]. However, these studies did not validate the method on real robots bypassing the true difficulties, such as non-linear dynamics due to friction, backlash, unmodelled physical constraints, uncertainties in the dynamic model, sensor and model errors, etc.

In this work, we provide the active inference perception and action mathematical construct for the humanoid robot iCub upper-body working in different real setups. Thus, we show the plausibility of real implementations of the free energy principle in artificial agents, analysing its challenges and limitations. Furthermore, we introduce a novel way to compute the mapping from action to sensory consequences (partial derivatives of the sensory input with respect to the action) exploiting the properties of velocity control, and simplifying the calculus of the action within the free energy optimization framework.

B. Paper organization

First, in Sec. II we explain the general mathematical free energy optimization for perception and action. Afterwards, Sec. III describes the iCub physical model and in Sec. IV and V we detail the active inference computational model applied for object reaching and head object tracking. Finally, Sec. VI shows the results for several experiments and Sec. VII analyses the relevance of implementing active inference body perception in humanoids, for both human and robotics science.

II. FREE ENERGY OPTIMIZATION

A. Bayesian inference

According to the Bayesian inference model for the brain, the body configuration x is inferred using available sensory data s by applying Bayes' theorem: $p(x|s) = p(s|x)p(x)/p(s)$. Where this *posterior probability*, corresponding to the probability of body configuration x given the observed data s , is obtained as a consequence of three antecedents: (1) *likelihood*, $p(s|x)$, or compatibility of the observed data s with the body configuration x , (2) *prior probability*, $p(x)$, or current belief about the configuration before receiving the sensory data s , and (3) *marginal likelihood*, $p(s)$, a normalization term which corresponds to the marginalization of the likelihood of receiving sensory data s regardless of the configuration.

The goal is to find the value of x which maximises $p(x|s)$, because it is the most likely value for the real-world body configuration x taking into account the sensory data obtained s . This direct method becomes intractable for large state spaces [20]. In this case, marginalization over all the possible body and world states explodes computationally.

B. Free energy principle with the Laplace approximation

The *free energy principle* [6] provides a tractable solution to this obstacle, where, instead of calculating the marginal likelihood, the idea is to minimise the *Kullback-Leibler divergence* [21] between a reference distribution $q(x)$ and the real $p(x|s)$, in order for the first to be a good approximation of the second.

$$\begin{aligned} D_{KL}(q(x)||p(x|s)) &= \int q(x) \ln \frac{q(x)}{p(x|s)} dx \\ &= \int q(x) \ln \frac{q(x)}{p(s,x)} dx + \ln p(s) = F + \ln p(s) \geq 0 \end{aligned} \quad (1)$$

Minimising the first term, F , effectively minimises the difference between these two densities, with only the marginal likelihood remaining. Unlike the whole expression of the KL-divergence, the first term can be evaluated because it depends on the reference distribution and the knowledge about the environment we can assume the agent has, $p(s, x) = p(s|x)p(x)$. This term is defined as *variational free energy* [22].

According to the free energy optimization theory, there are two ways to minimise *surprise*, which accounts for the discrepancy between the current state and the predicted or desired one (prediction error): changing the belief (*perceptual inference*) or acting on the world (*active inference*). Perceptual inference and active inference optimise the value of variational free energy, while active inference also optimises the value of the marginal likelihood by acting on the environment and changing the sensory data s .

Under the Laplace approximation instead of computing the whole distribution we can track the mode [23] and simplify the calculus of F . Thus, the variational density $q(x)$ is assumed to have Gaussian form $\mathcal{N}(x|\mu, \Sigma)$. Defining the Laplace-encoded energy [22] as $L(s, x) = -\ln p(s, x)$, F can be approximated to:

$$F \approx L(s, \boldsymbol{\mu}) - \left[\frac{1}{2} \ln |\Sigma| + n \ln 2\pi \right] \quad (2)$$

where $\boldsymbol{\mu}$ conceptually describes the internal belief or state, s is the sensory input and n is the size of $\boldsymbol{\mu}$ vector.

C. Perceptual and active inference

Perceptual inference is the process of updating the inner model belief to best account for sensory data, minimising the prediction error. For instance for body perception, the agent must infer the most-likely or optimal value for the body configuration or state $\boldsymbol{\mu}$. This optimal value is the one that minimises free energy. To solve it, we can compute it by gradient descent over the free energy term. For static systems, this update is performed directly as [16]: $\dot{\boldsymbol{\mu}} = -\frac{\partial F}{\partial \boldsymbol{\mu}}$. In dynamic systems, time derivatives should be considered, the state variable is now a vector $\boldsymbol{\mu}$:

$$\dot{\boldsymbol{\mu}} = \mathbf{D}\boldsymbol{\mu} - \frac{\partial F}{\partial \boldsymbol{\mu}} \quad (3)$$

where the \mathbf{D} is the block-matrix derivative operator. When free energy is minimised, the value of its derivative is $\frac{\partial F}{\partial \boldsymbol{\mu}} = 0$, and the system is at equilibrium.

Active inference [7], is the extension of perceptual inference to the relationship between sensors and actions, taking into account that actions can change the world to make sensory data more accurate with predictions made by the inner model. The action plays a core role in the optimization and improves the approximation of the real distribution, therefore reducing the prediction error by minimising free energy. It also acts on the marginal likelihood by changing the real configuration which modifies the sensory data s to obtain new data that is more in concordance with the agent's belief. The optimal value is the one which minimises free energy, and again a gradient descent approach will be taken to update its value:

$$\dot{a} = -\frac{\partial F}{\partial a} \quad (4)$$

III. ROBOT PHYSICAL MODEL

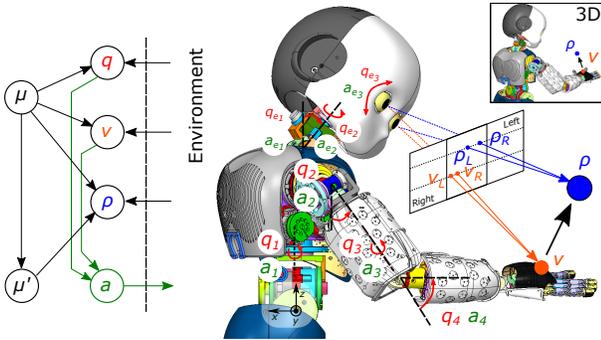


Fig. 2. Model description. (Left) Generative model of the robot. (Right) Notation and robot configuration. Shown variables: internal variables $\boldsymbol{\mu}$ and $\boldsymbol{\mu}'$, joint angle position \mathbf{q} , joint actions \mathbf{a} , and 3D location of the end-effector \mathbf{v} and causal variables $\boldsymbol{\rho}$ along with their projections in 2D visual space.

The robot iCub [24] (v1.4) is a humanoid powered by electric motors and driven by tendons. It is divided into several kinematic chains defined through homogeneous transformation matrices using Denavit-Hartenberg convention. We focused on two kinematic chains, those with the end-effector being the hands (without fingers), the head, the eyes and the torso.

As depicted in Fig. 2, each arm model is defined as a three degree of freedom system: *shoulder_roll*, *shoulder_yaw* and *r_elbow*. To provide improved reaching capabilities, an additional degree of freedom is considered for the torso: *torso_yaw*. Left and right eye cameras observe the end-effector position and the world around it, providing stereo vision. The joints considered for the head are: *neck_pitch* and *neck_yaw*. Moreover, vertical synchronised motion is performed in both eyes using: *eyes_tilt*.

IV. ACTIVE INFERENCE COMPUTATIONAL MODEL FOR ICUB REACHING TASK

A. Problem formulation

The body configuration, or internal variables, is defined as the joint angles $\mathbf{g}_p(\boldsymbol{\mu}) = \boldsymbol{\mu}$. The estimated states $\boldsymbol{\mu} \in \mathbb{R}^4$ are the belief the agent has about the joint angle position and the action $\mathbf{a} \in \mathbb{R}^4$ is the angular velocity of those same joints. Due to the fact that we use a velocity control for the joints, first-order dynamics must also be considered $\boldsymbol{\mu}' \in \mathbb{R}^4$.

Sensory data will be obtained through several input sensors. Using binocular disparity between the images obtained from both eyes, stereo vision reconstruction is performed to obtain the 3D position of the end-effector, $\mathbf{s}_v \in \mathbb{R}^3$. All joint motors have encoders which provide joint angle position, $\mathbf{s}_p \in \mathbb{R}^4$.

$$\mathbf{s}_p = \begin{bmatrix} q_1 \\ q_2 \\ q_3 \\ q_4 \end{bmatrix} \quad \mathbf{s}_v = \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix} \quad (5)$$

The likelihood $p(s|\boldsymbol{\mu})$ is made up of proprioception functions in terms of the current body configuration, while the prior $p(\boldsymbol{\mu})$ will take into account the dynamic model of the agent that describes how this internal state changes with time. Adapting the Laplace-encoded energy for our robot model and assuming independence between different modalities sensory readings:

$$\ln p(s, \boldsymbol{\rho}, \boldsymbol{\mu}) = \ln p(\mathbf{s}_p|\boldsymbol{\mu})p(\mathbf{s}_v|\boldsymbol{\mu})p(\boldsymbol{\mu}'|\boldsymbol{\mu}, \boldsymbol{\rho}) \quad (6)$$

B. Free energy optimization

In order to define the conditional densities for each of the terms, we should define the expressions for the sensory data. Joint angle position, \mathbf{s}_p , is obtained directly from the joint angle sensors. Let us assume that the input is noisy and follows a normal distribution with mean at the internal value $\boldsymbol{\mu}$ and variance Σ_{s_p} . End-effector 2D or 3D position, \mathbf{s}_v , is defined by a non-linear function dependent on the body configuration and obtained using the forward model of the arm. Again, the input is noisy and follows a normal distribution with mean at the value of this function $\mathbf{g}_v(\boldsymbol{\mu}) \in \mathbb{R}^{n_v}$ and variance Σ_{s_v} , where n_v is the visual dimensions. The dynamic model for n_p latent

variables (joints angles) is determined by a function which depends on both the current state $\boldsymbol{\mu}$ and the causal variables $\boldsymbol{\rho}$ (e.g. 3D position of the object to be reached), with a noisy input following a normal distribution with mean at the value of this function $\boldsymbol{f}(\boldsymbol{\mu}, \boldsymbol{\rho}) \in \mathbb{R}^{n_p}$ and variance Σ_{s_μ} . Therefore, likelihood functions in Eq. (6) are defined as:

$$p(\boldsymbol{s}_p|\boldsymbol{\mu}) = \prod_{i=1}^{n_p} \mathcal{N}(q_i|\mu_i, \Sigma_p) \quad p(\boldsymbol{s}_v|\boldsymbol{\mu}) = \prod_{i=1}^{n_v} \mathcal{N}(v_i|g_{v_i}(\boldsymbol{\mu}), \Sigma_v)$$

$$p(\boldsymbol{\mu}'|\boldsymbol{\mu}, \boldsymbol{\rho}) = \prod_{i=1}^{n_p} \mathcal{N}(\mu'_i|f_i(\boldsymbol{\mu}, \boldsymbol{\rho}), \Sigma_{s_\mu}) \quad (7)$$

The optimization of the Laplace-encoded energy is performed using gradient descent (Eq. (3) and (4)). The dependency of F with respect to the vector of internal variables $\boldsymbol{\mu}$ can be calculated using the chain rule on the functions that depend on those internal variables. The dependency of F with respect to the vector of actions \boldsymbol{a} is calculated considering that the only magnitudes directly affected by action are the values obtained from the sensors: $\frac{\partial F}{\partial \boldsymbol{a}} = \frac{\partial s}{\partial \boldsymbol{a}} \frac{\partial F}{\partial s}$ - see Sec. IV-D.

$$-\frac{\partial F}{\partial \boldsymbol{\mu}} = \frac{1}{\Sigma_{s_p}} (\boldsymbol{s}_p - \boldsymbol{\mu}) + \frac{1}{\Sigma_{s_v}} \frac{\partial g_v(\boldsymbol{\mu})^T}{\partial \boldsymbol{\mu}} (\boldsymbol{s}_v - \boldsymbol{g}(\boldsymbol{\mu}))$$

$$+ \frac{1}{\Sigma_{s_\mu}} \frac{\partial \boldsymbol{f}(\boldsymbol{\mu}, \boldsymbol{\rho})^T}{\partial \boldsymbol{\mu}} (\boldsymbol{\mu}' - \boldsymbol{f}(\boldsymbol{\mu}, \boldsymbol{\rho})) \quad (8)$$

$$-\frac{\partial F}{\partial \boldsymbol{a}} = - \left(\frac{1}{\Sigma_{s_p}} \frac{\partial \boldsymbol{s}_p^T}{\partial \boldsymbol{a}} (\boldsymbol{s}_p - \boldsymbol{\mu}) + \frac{1}{\Sigma_{s_v}} \frac{\partial \boldsymbol{s}_v^T}{\partial \boldsymbol{a}} (\boldsymbol{s}_v - \boldsymbol{g}_v(\boldsymbol{\mu})) \right) \quad (9)$$

The agent also infers the first-order dynamics of the body $\boldsymbol{\mu}'$ and updates their values using a gradient descent formulation. The dependency of F with respect to $\boldsymbol{\mu}'$ is limited to the influence of the dynamic model.

$$-\frac{\partial F}{\partial \boldsymbol{\mu}'} = -\frac{1}{\Sigma_{s_\mu}} (\boldsymbol{\mu}' - \boldsymbol{f}(\boldsymbol{\mu}, \boldsymbol{\rho})) = \frac{1}{\Sigma_{s_\mu}} (\boldsymbol{f}(\boldsymbol{\mu}, \boldsymbol{\rho}) - \boldsymbol{\mu}') \quad (10)$$

Finally, the differential equations for $\boldsymbol{\mu}$, $\boldsymbol{\mu}'$ and \boldsymbol{a} are:

$$\dot{\boldsymbol{\mu}} = \boldsymbol{\mu}' - \frac{\partial F}{\partial \boldsymbol{\mu}} \quad \dot{\boldsymbol{\mu}'} = -\frac{\partial F}{\partial \boldsymbol{\mu}'} \quad \dot{\boldsymbol{a}} = -\frac{\partial F}{\partial \boldsymbol{a}} \quad (11)$$

A first-order Euler integration method is applied to update the values of $\boldsymbol{\mu}$, $\boldsymbol{\mu}'$ and \boldsymbol{a} in each iteration. E.g., $\mu_{i+1} = \mu_i + \dot{\mu}_i \Delta t$, where $\Delta t = T$ is the period of execution of the updating cycle.

C. Perceptual attractor dynamics

The reaching goal is defined in the dynamics of the model by introducing a perceptual attractor $\boldsymbol{\rho}$ in the 2D or 3D space:

$$\boldsymbol{A}(\boldsymbol{\mu}, \boldsymbol{\rho}) = \rho_4 \left(\begin{bmatrix} \rho_1 \\ \rho_2 \\ \rho_3 \end{bmatrix} - \begin{bmatrix} g_{v_1}(\boldsymbol{\mu}) \\ g_{v_2}(\boldsymbol{\mu}) \\ g_{v_3}(\boldsymbol{\mu}) \end{bmatrix} \right) \quad (12)$$

where $\rho_{1:3}$ defines the object location in the visual space and ρ_4 describes the gain of the attractor.

Internal variable dynamics are then defined in terms of the attractor: $\boldsymbol{f}(\boldsymbol{\mu}, \boldsymbol{\rho}) = \boldsymbol{T}(\boldsymbol{\mu})\boldsymbol{A}(\boldsymbol{\mu}, \boldsymbol{\rho})$. $\boldsymbol{T}(\boldsymbol{\mu})$ is the function that

transforms the attractor vector from target space (3D) to joint space. The system is velocity controlled, therefore target space is a linear velocity vector and joint space is angular velocity.

When the Jacobian matrix is rectangular (e.g., 3D space and 4 DOF arm joint space), we used the generalised inverse (Moore-Penrose pseudoinverse) as the mapping function: $\boldsymbol{T}(\boldsymbol{\mu}) = \boldsymbol{J}^+(\boldsymbol{\mu})$. This matrix is calculated using the singular-value decomposition (SVD), where $\boldsymbol{J}^+ = \boldsymbol{V}\boldsymbol{\Sigma}^+\boldsymbol{U}^T$.

D. Active inference

Action is set to be an angular velocity magnitude, which corresponds with angular joint velocity in the latent space. In order to compute Eq. (9) we need the mapping between the sensor values and the velocities of the joints, expressed in the partial derivatives $\frac{\partial \boldsymbol{s}_p}{\partial \boldsymbol{a}}$ and $\frac{\partial \boldsymbol{s}_v}{\partial \boldsymbol{a}}$.

We assume that the control action, a , is updated for every cycle, and therefore for each interval of time between cycles it has the same value. For each period (cycle time between updates), the equation of uniform circular motion is satisfied for each joint position. If this equation is discretised, for each sampling moment, which are T seconds apart, the joint value will be updated in the following way: $q_{i+1} = q_i + a_i T$. The dependency of the joint angle position with respect to the control action is therefore defined.

The partial derivatives of joint position \boldsymbol{s}_p with respect to action, considering there is no cross-influence or coupling between joint velocities and that q_i and its expected μ_i should converge at equilibrium, are given by the following expression:

$$\frac{\partial q_i}{\partial a_j} = \frac{\partial \mu_i}{\partial a_j} = \begin{cases} T & i = j, \\ 0 & \text{otherwise.} \end{cases} \quad (13)$$

If the dependency of joint position with respect to action is known, we can use the chain rule to calculate the dependency for the visual sensor \boldsymbol{s}_v . Considering that the values of $g_{v_i}(\boldsymbol{\mu})$ should also converge to v_i at equilibrium, the partial derivatives are given by the following expression:

$$\frac{\partial v_i}{\partial a_j} = \frac{\partial v_i}{\partial q_j} \frac{\partial q_j}{\partial a_j} = \frac{\partial g_{v_i}(\boldsymbol{\mu})}{\partial \mu_j} \frac{\partial \mu_j}{\partial a_j} \quad (14)$$

V. ACTIVE INFERENCE COMPUTATIONAL MODEL FOR ICUB HEAD OBJECT TRACKING TASK

We extend the arm reaching model for the head to obtain an object tracking motion behaviour. The goal of this task is to maintain the object in the centre of the visual field, thus increasing its reaching working range capabilities.

Sensory data and proprioception for the head is defined by internal variables beliefs $\boldsymbol{\mu}_e \in \mathbb{R}^3$, actions $\boldsymbol{a}_e \in \mathbb{R}^3$, and first-order dynamics $\boldsymbol{\mu}'_e \in \mathbb{R}^3$ and because the end-effector are the eyes, there is only joint angle position, $\boldsymbol{s}_e \in \mathbb{R}^3$.

To obtain the desired motion, an attractor is defined towards the centre of the left eye image (c_x, c_y). Attractor position (ρ_5, ρ_6) is read from the projected 2D visual input of the left eye and dynamically updates with head motion.

$$\boldsymbol{A}_e(\boldsymbol{\mu}_e, \boldsymbol{\rho}) = \rho_4 \left(\begin{bmatrix} c_x \\ c_y \end{bmatrix} - \begin{bmatrix} \rho_5 \\ \rho_6 \end{bmatrix} \right) \quad (15)$$

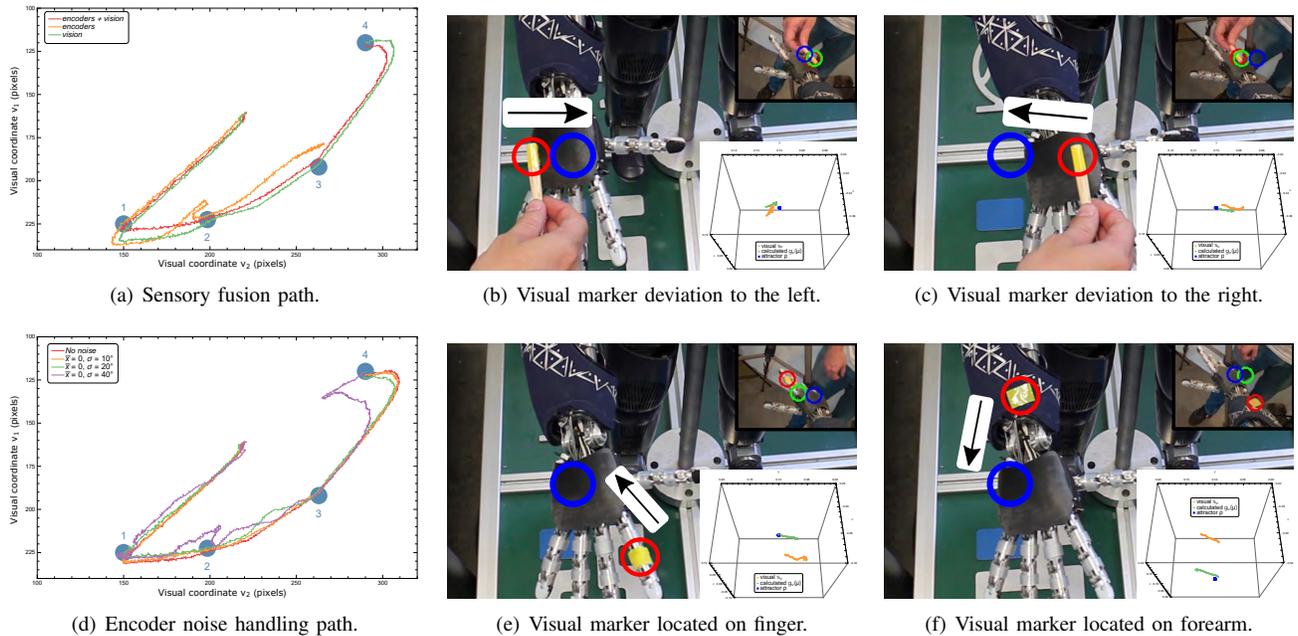


Fig. 3. (Left) End-effector averaged paths in the right arm object reaching for (a) sensory fusion and (d) under noisy encoder sensors. Attractor position is represented by a blue disk. (Right) Results of adaptive behaviour using sensory fusion. If the visual marker position is modified, the algorithm produces a reaching action that tries to approach this new position to the attractor location. On the top right of each image the left eye camera is shown. Attractor position is shown in blue, visual perception in red and calculated position in green. The resulting direction of motion obtained from the algorithm is shown as an arrow.

Internal variable dynamics are then defined in terms of the attractor as: $f_e(\mu_e, \rho) = T_e(\mu_e)A_e(\mu_e, \rho)$, $f_e(\mu_e, \rho) \in \mathbb{R}^3$. With two pixel coordinates and three degrees of freedom, the pseudoinverse of the visual Jacobian matrix is used as the mapping matrix in the visual space: $T_e(\mu_e) = J_v^+(\mu_e)$.

VI. RESULTS

We evaluated the proposed approach for humanoid upper-body reaching in 2D and 3D visual space scenarios. The proposed algorithm, in order to reduce the discrepancy between the current sensory input and the desired state, produces a reaching behaviour with the arm and torso of the robot while the head tracks the object to maintain it in its visual field.

Several experiments were performed: (1) *proof of concept*, right arm reaching towards a series of 2D locations in the visual plane with different levels of noise and sensor sources; (2) *adaptation*, the robot adapts its reaching behaviour during changes on the visual feature location; (3) *comparison*, motion from the active inference algorithm is compared to inverse kinematics; and (4) *dynamics evaluation*, body perception and action variables are analysed during an arm reaching with active head towards a moving object. Video footage of the experiments and some extras can be found in the supplementary video: <https://youtu.be/rdbbmwo4TY4>.

The iCub robot is placed in a controlled environment and an easily recognizable object is used as a perceptual attractor to produce the movement of the robot. In 3D space, the values of the causal variables ρ_1, ρ_2, ρ_3 are the coordinates of the 3D position of the object obtained using stereo vision, while in 2D space, the values used for the attractor in (12) are the object pixel location at the left eye. The value of ρ_4 is a weighting

factor to adjust the power of the attractor. The right arm end-effector has a visual marker and its 3D position is obtained as the values of v_1, v_2 and v_3 , along with its projections in camera space. The relevant parameters of the algorithm are: encoder sensor variance Σ_{s_p} , visual perception variance Σ_{s_v} , attractor dynamics variance Σ_{s_μ} and action gains k_a . These parameters were tuned empirically with their physical meaning in mind and remained constant during the experiments unless stated otherwise. Their values can be found in the code repository.

A. Sensory fusion and robustness under noisy sensors

The first experiment was designed to test sensory fusion (Fig. 3(a)) and the impact of noisy sensors on the algorithm (Fig. 3(d)). The robot worked in 2D visual space using the left-eye camera (monocular vision) and the torso motion is locked. The robot had to reach 10 times four different static locations in the visual field with the right arm. A location is considered to be reached when the visual position is inside a disk with a radius of five pixels centred at the location position.

Figure 3(a) shows the averaged trajectory of the visual marker with different sensor modalities (considering only intrinsic noise from the sensors and model errors): visual, joint angle encoders, and both together. Even though the model has been verified and the camera calibrated, there was a difference between the forward model and the real robot, due to possible deviations in parameters and motor backlash, which implies that the robot had to use visual information to correct its real position. Employing joint angle encoders and vision provided the best behaviour for the reaching of the fixed locations in the visual field, achieving all positions in the shortest fashion. Visual perception reached all the positions but it did not follow

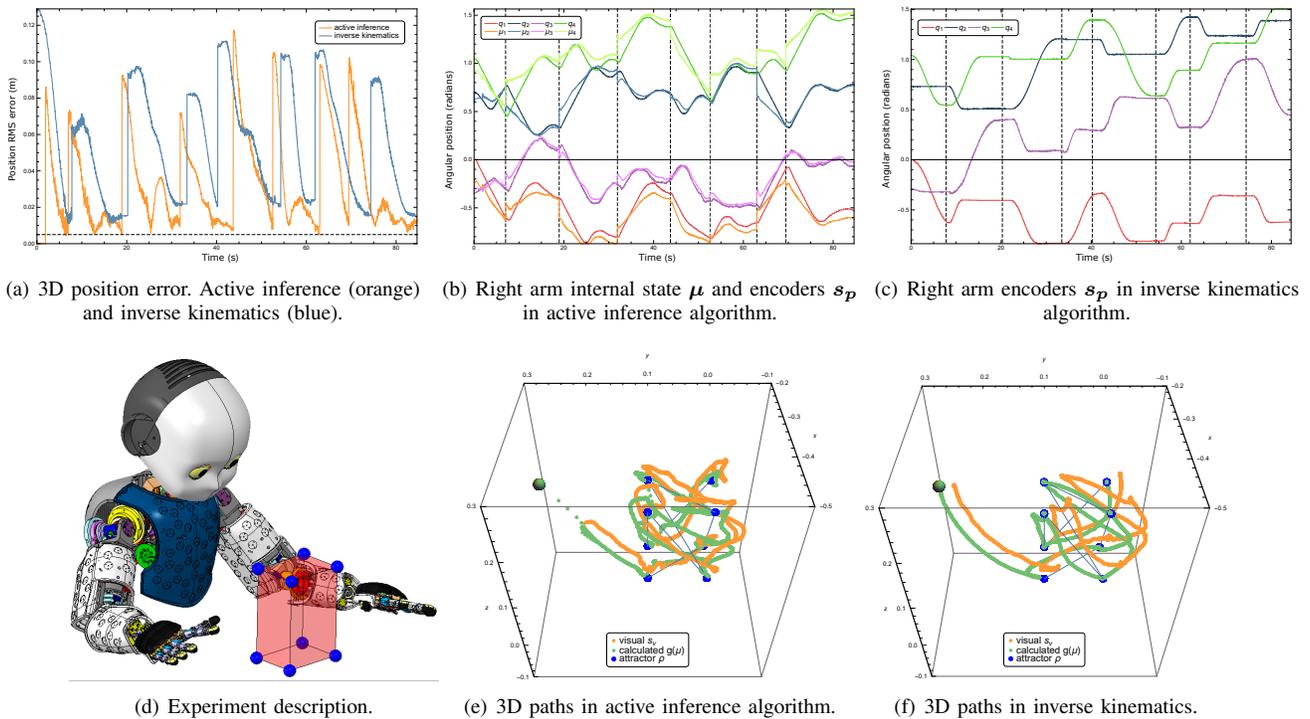


Fig. 4. Comparison between active inference and inverse kinematics. (Top left) RMS 3D position error between the real visual position and the attractor position. (Top centre and right) Internal state and encoder values of the right arm for each of the algorithms. (Bottom left) Experiment description showing the initial arm position and the rectangular prism. (Bottom centre and right) 3D paths taken for each of the algorithms.

the optimum path while using only the encoder values failed to reach all locations.

Secondly, in order to test the robustness against sensory noise, we injected noise in the robot motors encoders s_p . Four conditions were tested: Gaussian noise with a zero mean and with standard deviation of 0° (control test), 10° , 20° and 40° . Again, averaged path is shown (Fig. 3(d)). Trajectories with no noise and $\sigma = 10^\circ$ achieved very similar results, with the first one achieving the objectives slightly faster. Extreme case with $\sigma = 40^\circ$ caused oscillations and erroneous approximation trajectories that produced significant delays in the reaching of the target locations. Intuitively, with low sensor variance $\Sigma_{s_p} < 0.5$ the system was not able to perform the task with high sensor noise.

B. Online adaptive behaviour

We tested the online perception and action loop when there was a change in the visual feature location. The visual feature was in the 3D space, and all degrees of freedom and binocular vision were considered. Figure 3(b,c,e,f) shows the behaviour of the arm trying to correct the forward model prediction with the altered location of the visual marker. We set the perceptual attractor position to a constant fixed point in 3D space. The end-effector visual marker was removed and replaced with a similar marker located on a tool that was operated by a human. When the marker was moved sideways (Fig. 3(b) and 3(c)), the algorithm generated an action that moved the end-effector towards the opposite direction to correct the deviation. The internal belief was also updated to find an equilibrium point.

Similar behaviour was obtained when the marker was located on one of the fingers (Fig. 3(e)) and on the forearm (Fig. 3(f)). In this case, the visual discrepancy with the forward model is much more noticeable and the algorithm is not able of fully reaching the visual marker position to the attractor, but the same opposite action behaviour as in the previous case is obtained.

This experiment shows that in the case of visual sensorimotor conflict a perception body alteration and action to reduce the uncertainty is generated. In essence, this is a simplified body illusion that only uses one visual feature. Thus, this is extremely relevant as the model of the robot is predicting drifts on the hand localization [25] but also actions.

C. Comparison of active inference and inverse kinematics

We compared our algorithm with the inverse kinematics approach in a reaching task using the right arm, the torso and the head. This comparison was performed by defining a set of 3D waypoints matching the vertices of a rectangular prism that the robot must reach, as described in Fig. 4(d).

The inverse kinematics module from iCub¹ was used sequentially to solve each of the eight vertices reaching problems. The robot joints were position-controlled to the desired state and joint speed was adjusted in order for the motion to have the same time duration as the active inference test. Head angles were positioned in a way such that the perceptual

¹iKin uses the nonlinear optimization package *Ipopt* to obtain the joint angle position to arrive at a certain 3D point in space taking a certain angle configuration as the starting point.

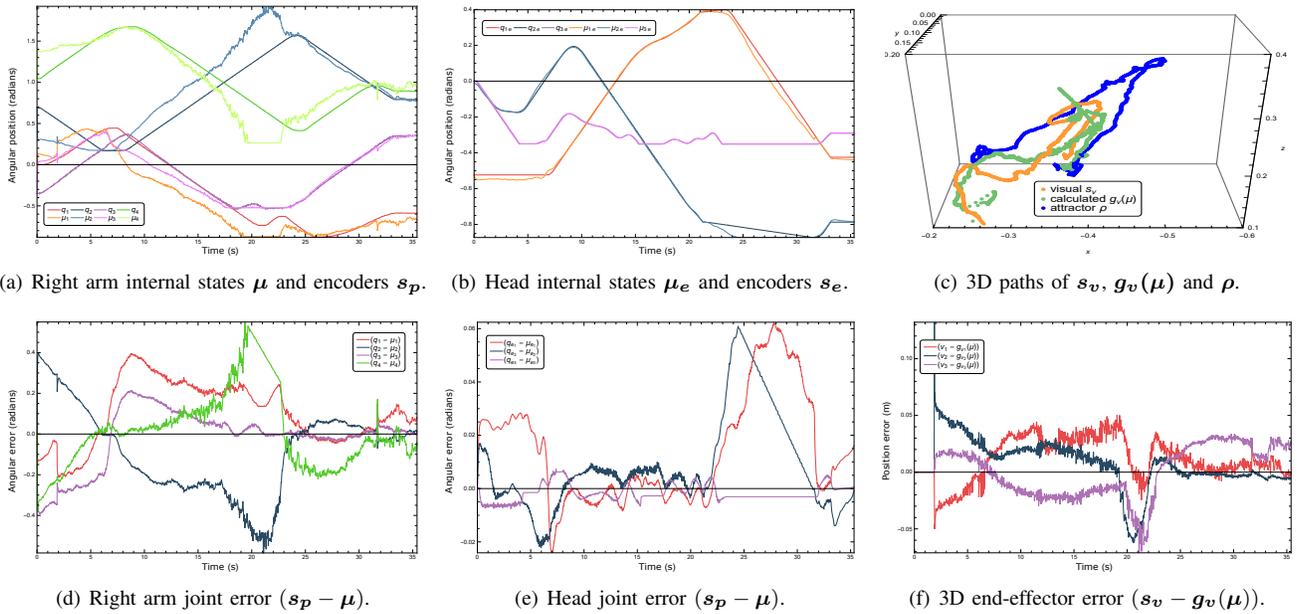


Fig. 5. System dynamics for reaching a moving object in 3D. Internal states are driven towards the perceptual attractor position. The difference between the value of internal states and encoders (5(d)) and between 3D calculated position and visual location of end-effector (5(f)) drive the actions.

attractor was always inside the visual field. Our approach included in addition to the arm and the torso, also active head object tracking. In this case, the head started looking front and the arm was outside of the visual field.

Each location was considered to be reached once the calculated position of the visual marker was inside a sphere of $r = 0.5 \text{ cm}$ centred at the attractor location. Setting the 3D position reconstructed from the visual input as the ground truth for the location of the end-effector, the performance was assessed using the root mean square (RMS) error between this location and the target position of each vertex. Figure 4(a) shows the RMS error between the visual perception position s_v and the attractor position ρ . The dashed line marks an error of 0.5 cm . Because the active inference algorithm implements online sensory fusion (forward model and stereo vision) the performance of the reaching task was improved over the inverse kinematics approach. The minimum error obtained for the inverse kinematics algorithm was 1.1 cm , while the minimum error for the active inference algorithm was 0.42 cm . The mean error for the reaching of all 3D points in the inverse kinematics was $\bar{e}_{IK} = 1.69 \text{ cm}$, while the active inference algorithm obtained $\bar{e}_{AI} = 0.67 \text{ cm}$.

Figure 4(b) shows the internal state belief μ and the encoder values q of the active inference algorithm, while Fig. 4(c) shows the encoder values q for the inverse kinematics algorithm. The attractor position change is represented using a dashed line. In the active inference algorithm, each time the attractor position changes the believed state μ gets pushed towards the new position and there is a relatively big discrepancy between μ and q that gets reduced as the goal is reached.

Figures 4(e) and 4(f) show the paths taken by the active inference and the inverse kinematics algorithm. In both algorithms, the robot reached all targets, but there was a clear

difference in the discrepancy between the calculated and the visual perception position. Inverse kinematics paths show the error between the model and the physical robot (distance between the orange and the green line of Fig. 4(f)). Our approach reduced this distance when approaching the target location in a closed-loop manner (Fig. 4(e)), while inferring its real body configuration and using head object tracking.

D. System dynamics when reaching a moving object

We evaluated the algorithm in a reaching task of a moving object manually operated by a human. The experiment started with the object near the robot and it was progressively moved away and in an ascendant vertical direction, to finally be handed to the robot in an intermediate and lower location.

This experiment was performed both in 2D and 3D visual space. In 2D space, the setup used monocular vision from the left eye and the torso was blocked. Our model was tested for both arms and using active head object tracking. In 3D space, we enabled the torso motion but restricted the model for the right arm and the head. Figure 6 shows the vision of the robot and the free energy values for the head and the arm during the reaching experiment in 2D space. F is minimised modifying the arm perception and generating actions towards the object.

The resulting variable dynamics of the 3D space experiment are shown in Fig. 5. The initial position of the right arm lied outside of the visual plane (Fig. 5(f) showing zero error). When there was no visual input, the free energy optimization algorithm only relied on joint measurements and the forward model to produce the reaching motion until the right hand appears in the visual plane. Figures 5(a) and 5(b) show both the encoders measurements q and the estimated joint angle μ of the arm and head. Figure 5(c) shows how calculated $g_v(\mu)$ and real s_v 3D positions of the right arm end-effector

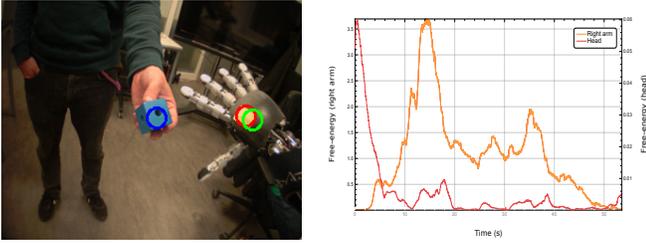


Fig. 6. Reaching a moving object in 2D. (Left) Robot vision: visual feature (red), predicted end-effector location (green), object (blue). (Right) Free energy optimization for the arm (orange) and the head (red).

follow the perceptual attractor ρ . Although the cameras were previously calibrated, noise was present in the stereo reconstruction of visual perception. This noise was more relevant in the depth estimation (first term in Fig. 5(f)). Stop action for this experiment was produced by the sense of touch. Contact in the hand pressure sensors triggers the grasping motion.

VII. ACTIVE INFERENCE IN REAL ARTIFICIAL AGENTS

Being able to develop robots with biologically plausible body perception and action models allow us to validate current brain-inspired computational models [26], [27]. However, there is a big gap between the theoretical mathematical construct and physical reality. This work reduces this gap and shows the plausibility of implementing the free energy principle [6] on artificial agents for real-world applications, joining together robotics and computational neuroscience. The reproduced robot behaviour, during the system dynamics study (Sec. VI-D), resembled a one-year-old infant. Just by prediction error minimization, in this case, through surprise minimization, the robot was able to perform human-like movements in a reaching and object tracking task. The model and the experiments can be reproduced and replicated by downloading the open-source code [tobereleased](#).

A. Relevant model predictions for human body perception

Interestingly, our model predicts the appearance of involuntary actions in a sensorimotor conflict situation, such as the rubber-hand illusion. In the presence of sensory errors (arm location expectation mismatch), we observed movements towards the visual feature of the arm. In previous work, we showed that the proprioceptive drift (mislocalization of the hand) using a similar free energy optimization algorithm could be modelled [25]. By adding the action term of active inference, forces towards minimising the prediction error should appear, i.e., in the same direction of the drift. This behaviour could be in charge of the online adaptation of movements in the presence of disturbances or not trained conditions, bypassing the voluntary movement pathway.

B. Technical challenges for a full-fledged body model

Closed-loop adaptive perception and action is an important advantage. The model is assumed to be an approximation of the reality and the minimization of the variational free

energy tackles both sensory noise and model bias. In fact, the equilibrium point where the algorithm converges reduces the differences between the model and the real observations. However, it is prone to local minima. As there are no task hierarchies, dual arm reaching produced redundant motions in both arms. A hierarchical version of the proposed approach, as proposed in [28], could circumvent this limitation by generating attractors for specialised behaviours.

The engineering analysis needed on iCub specific parameters and actuation was the biggest limitation for generalization. Sensor variances, action gains and velocity limits were experimentally tuned. In theory, variances and gains could be also optimised within the free energy framework [7]. Besides, we introduced how time discretization in a velocity control scheme can be used to simplify the action computation. In force controlled non-linear systems the mapping between the sensory consequence and the force applied must be learnt [19] or explicitly computed through the forward dynamics [18], considerably increasing the complexity of the optimization framework.

The proposed algorithm is scalable to integrate features from different modalities (i.e., adding a new term to the summation) and handles sensory fusion in both perception and action. Furthermore, the needed generative functions can be easily computed from the forward kinematics or be even replaced by learned ones, as shown in [16]. However, more complex sensory input features should be incorporated to be able to compare with other bio-inspired approaches such as deep reinforcement learning.

VIII. CONCLUSIONS

This work presented a neuroscience-inspired closed-loop algorithm for body perception and action working in a humanoid robot. Our *active inference* model is the first one validated on a real robot for upper-body reaching and active head object tracking. Robots perceiving and interacting under the active inference approach is an important step for evaluating the plausibility of the model in biological systems as we systematically observe the behaviours and connect them to the mechanisms behind, and even make predictions. As a proof of concept, we showed its robustness to sensory noise and discrepancies between the robot model and the real body, its adaptive characteristic to visual online changes and its capacity to embed sensory fusion. Body perception was approached as a hidden Markov model and the unobserved body configuration/state was continuously approximated with visual and proprioceptive inputs by minimising the free energy bound. The action was modelled within perception as a reflex to reduce the prediction error. We tested the algorithm on the iCub robot in 2D and 3D using binocular vision. Results showed that our approach was capable of combining different independent sources of sensory information without an increase of computational complexity, displaying online adaptation capabilities and performing more accurate reaching behaviours than inverse kinematics

REFERENCES

- [1] H. v. Helmholtz, *Handbuch der physiologischen Optik*. L. Voss, 1867.

- [2] D. C. Knill and A. Pouget, "The bayesian brain: the role of uncertainty in neural coding and computation," *TRENDS in Neurosciences*, vol. 27, no. 12, pp. 712–719, 2004.
- [3] K. Friston, "A theory of cortical responses," *Philos Trans R Soc Lond B: Biological Sciences*, vol. 360, no. 1456, pp. 815–836, 2005.
- [4] T. Parr and K. J. Friston, "Uncertainty, epistemics and active inference," *Journal of The Royal Society Interface*, vol. 14, no. 136, p. 20170376, 2017.
- [5] E. Todorov and M. I. Jordan, "Optimal feedback control as a theory of motor coordination," *Nature neuroscience*, vol. 5, no. 11, p. 1226, 2002.
- [6] K. J. Friston, "The free-energy principle: a unified brain theory?" *Nature Reviews. Neuroscience*, vol. 11, pp. 127–138, 02 2010.
- [7] K. J. Friston, J. Daunizeau, J. Kilner, and S. J. Kiebel, "Action and behavior: a free-energy formulation," *Biological cybernetics*, vol. 102, no. 3, pp. 227–260, 2010.
- [8] M. Kirchhoff, T. Parr, E. Palacios, K. Friston, and J. Kiverstein, "The markov blankets of life: autonomy, active inference and the free energy principle," *Journal of The royal society interface*, vol. 15, no. 138, p. 20170792, 2018.
- [9] C. Fantacci, U. Pattacini, V. Tikhonoff, and L. Natale, "Visual end-effector tracking using a 3d model-aided particle filter for humanoid robot platforms," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 1411–1418.
- [10] C. Garcia Cifuentes, J. Issac, M. Wthrich, S. Schaal, and J. Bohg, "Probabilistic articulated real-time tracking for robot manipulation," *IEEE Robotics and Automation Letters*, vol. 2, 2017.
- [11] P. Vicente, L. Jamone, and A. Bernardino, "Online body schema adaptation based on internal mental simulation and multisensory feedback," *Frontiers in Robotics and AI*, vol. 3, p. 7, 2016.
- [12] C. Gaskett and G. Cheng, "Online learning of a motor map for humanoid robot reaching," in *2nd Int. Conf. on computational inte., robotics and autonomous systems, Singapore*, 2003.
- [13] 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," *Robotics and Autonomous Systems*, vol. 62, no. 4, pp. 556–567, 2014.
- [14] A. Roncone, M. Hoffmann, U. Pattacini, L. Fadiga, and G. Metta, "Peripersonal space and margin of safety around the body: learning visuo-tactile associations in a humanoid robot with artificial skin," *PLoS one*, vol. 11, no. 10, p. e0163713, 2016.
- [15] P. Lanillos, E. Dean-Leon, and G. Cheng, "Yielding self-perception in robots through sensorimotor contingencies," *IEEE Trans. on Cognitive and Developmental Systems*, no. 99, pp. 1–1, 2016.
- [16] P. Lanillos and G. Cheng, "Adaptive robot body learning and estimation through predictive coding," *Intelligent Robots and Systems (IROS), 2018 IEEE/RSJ Int. Conf. on*, 2018.
- [17] M. Baltieri and C. L. Buckley, "An active inference implementation of phototaxis," *2018 Conference on Artificial Life*, no. 29, pp. 36–43, 2017.
- [18] L. Pio-Lopez, A. Nizard, K. Friston, and G. Pezzulo, "Active inference and robot control: a case study," *J R Soc Interface*, vol. 13, 2016.
- [19] P. Lanillos and G. Cheng, "Active inference with function learning for robot body perception," *International Workshop on Continual Unsupervised Sensorimotor Learning, ICDL-Epirob*, 2018.
- [20] M. J. Beal and Z. Ghahramani, "The variational bayesian em algorithm for incomplete data," *Statistics*, 07 2002.
- [21] S. Kullback and R. A. Leibler, "On information and sufficiency," *The Annals of Mathematical Statistics*, vol. 22, no. 1, pp. 79–86, 03 1951.
- [22] C. L. Buckley, C. S. Kim, S. McGregor, and A. K. Seth, "The free energy principle for action and perception: A mathematical review," *Journal of Mathematical Psychology*, 2017.
- [23] K. Friston, J. Mattout, N. Trujillo-Barreto, J. Ashburner, and W. Penny, "Variational free energy and the laplace approximation," *Neuroimage*, vol. 34, no. 1, pp. 220–234, 2007.
- [24] G. Metta, G. Sandini, D. Vernon, L. Natale, and F. Nori, "The icub humanoid robot: An open platform for research in embodied cognition," *Performance Metrics for Intelligent Systems Workshop*, 01 2008.
- [25] N.-A. Hinz, P. Lanillos, H. Mueller, and G. Cheng, "Drifting perceptual patterns suggest prediction errors fusion rather than hypothesis selection: replicating the rubber-hand illusion on a robot," *arXiv preprint arXiv:1806.06809*, 2018.
- [26] P. Lanillos, E. Dean-Leon, and G. Cheng, "Enactive self: a study of engineering perspectives to obtain the sensorimotor self through enaction," in *Developmental Learning and Epigenetic Robotics, Joint IEEE Int. Conf. on*, 2017.
- [27] M. Hoffmann and R. Pfeifer, "Robots as powerful allies for the study of embodied cognition from the bottom up," in *The Oxford Handbook of 4E Cognition*, 2018.
- [28] J. Tani, "Learning to generate articulated behavior through the bottom-up and the top-down interaction processes," *Neural Networks*, vol. 16, no. 1, pp. 11–23, 2003.