

## A Novel Adaptive Controller for Robot Manipulators Based on Active Inference

Pezzato, Corrado; Ferrari, Riccardo M.G.; Hernández, Carlos

**DOI**

[10.1109/LRA.2020.2974451](https://doi.org/10.1109/LRA.2020.2974451)

**Publication date**

2020

**Document Version**

Final published version

**Published in**

IEEE Robotics and Automation Letters

**Citation (APA)**

Pezzato, C., Ferrari, R. M. G., & Hernández, C. (2020). A Novel Adaptive Controller for Robot Manipulators Based on Active Inference. *IEEE Robotics and Automation Letters*, 5(2), 2973-2980. <https://doi.org/10.1109/LRA.2020.2974451>

**Important note**

To cite this publication, please use the final published version (if applicable). Please check the document version above.

**Copyright**

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

**Takedown policy**

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

***Green Open Access added to TU Delft Institutional Repository***

***'You share, we take care!' – Taverne project***

**<https://www.openaccess.nl/en/you-share-we-take-care>**

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

# A Novel Adaptive Controller for Robot Manipulators Based on Active Inference

Corrado Pezzato , Riccardo Ferrari, and Carlos Hernández Corbato 

**Abstract**—More adaptive controllers for robot manipulators are needed, which can deal with large model uncertainties. This letter presents a novel active inference controller (AIC) as an adaptive control scheme for industrial robots. This scheme is easily scalable to high degrees-of-freedom, and it maintains high performance even in the presence of large unmodeled dynamics. The proposed method is based on active inference, a promising neuroscientific theory of the brain, which describes a biologically plausible algorithm for perception and action. In this work, we formulate active inference from a control perspective, deriving a model-free control law which is less sensitive to unmodeled dynamics. The performance and the adaptive properties of the algorithm are compared to a state-of-the-art model reference adaptive controller (MRAC) in an experimental setup with a real 7-DOF robot arm. The results showed that the AIC outperformed the MRAC in terms of adaptability, providing a more general control law. This confirmed the relevance of active inference for robot control.

**Index Terms**—Biologically-inspired robots, adaptive control of robotic systems, industrial robots, active inference, free-energy principle.

## I. INTRODUCTION

TRADITIONAL control approaches for industrial manipulators rely on an accurate model of the plant. However, there is an increasing demand in industry for robot controllers that are more flexible and adaptive to run-time variability. Often, robot manipulators are placed in dynamically changing surrounding, and they are subject to noisy sensory input and unexpected events. In these new applications, obtaining such a model is a major problem. For example, in pick and place tasks, the dynamics of the robot manipulators can change unpredictably while handling unknown objects. Recent research has focused on the use of machine learning methods to obtain accurate inverse dynamic models [1], [2]. In general, learning models using Neural Networks (NN) requires experts for defining the best topology for a particular problem [3]. Even though it is possible to exploit the physical knowledge of the system to simplify and improve the learning performance [4], the need of large amount

Manuscript received September 10, 2019; accepted January 19, 2020. Date of publication February 17, 2020; date of current version March 2, 2020. This letter was recommended for publication by Associate Editor Prof. S. Weiss and Editor Prof. P. Rocco upon evaluation of the reviewers' comments. This work was supported by Ahold Delhaize. (Corresponding author: Corrado Pezzato.)

Corrado Pezzato and Carlos Hernández Corbato are with the Cognitive Robotics Department, Delft University of Technology, 2628 CD Delft, The Netherlands (e-mail: c.pezzato@tudelft.nl; c.h.corbato@tudelft.nl).

Riccardo Ferrari is with the Department of Systems and Control, Delft University of Technology, 2628 CD Delft, The Netherlands (e-mail: riccardo.mg.ferrari@gmail.com).

This letter has supplementary downloadable material available at <http://ieeexplore.ieee.org>, provided by the authors.

Digital Object Identifier 10.1109/LRA.2020.2974451

of training data and several iterations for learning, still remains a problem and hard to generalise [5], [6]. Controllers that can dynamically adapt are required, but existing solutions in adaptive control either need an accurate model, or are difficult to tune and scale to higher number of DOFs. In this letter, we present a novel adaptive controller for robot manipulators, inspired by a recent theory of the brain, which does not require accurate plant dynamics, and that is less sensitive to large parameters variation.

The proposed control scheme is based on the general free-energy principle proposed by Karl Friston [7], and redefined in engineering terms [8], [9]. The main idea at the basis of Friston's neuroscientific theory, is that the brain's cognition and motor control functions could be described in terms of energy minimization. It is supposed [10] that humans have a set of sensory data and a specific internal model to characterize how the sensory data could have possibly been generated. Then, given this generative model, the causes of sensory data are inferred. Usually, the environment acts on humans to produce sensory impression, and humans can act on the environment to change it. In this view, the motor control of human body can be considered as the fulfillment of a prior expectation about proprioceptive sensations [11]. Although the general active inference framework is mathematically well defined, its application to robotics remains a challenge. Active inference has mainly been applied to neuronal simulations (for handwriting [7] for instance), supposing to know the true dynamical process. However, this is not the case in robotics. Even if the neuronal simulations are a strong proof of concept for the neuroscientific theory, in the present form their extension to realistic robotic scenarios [12], [13] does not provide advantages over other classical controllers. The main problems are the computational load and the definition of meaningful generative models. With our work we overcome these limitations, using active inference to derive a model-free control law. Instead of modeling the true unknown dynamical process, we define a reference model that active inference has to follow. The main contributions of this letter are twofold:

- Derivation of an online active inference control law for a generic  $n$ -DOF robot manipulator in joint space.
- Comparison of the adaptability of the AIC with a state-of-the-art model reference adaptive controller.

The contributions have been experimentally validated in a 7-DOF collaborative industrial manipulator.

## A. Related Work

At present, the use of active inference for robot control is still limited. In [12], the authors simulated a PR2 robot controlled in Cartesian space for a reaching task. The solution was offline, computationally expensive, open-loop, and it relied on an additional position controller. This makes the approach not

suitable for online tasks. A recent MSc thesis [13], based on [12], derived an offline closed-loop scheme of active inference. The feedforward torque commands for a simulated 7-DOF manipulator are computed offline, relying on additional controllers for feedback control. The scheme failed to control the robot in presence of gravity since the feedforward torques did not include the gravitational effect. Both [12] and [13] were based on the Statistical Parametric Mapping (SPM) by Friston. This toolbox is suitable for several offline applications, but it is too computationally heavy for online control. In [13], each iteration is reported to take about one second. Another recent work [14] formalised the use of the free-energy for static state estimation, using a real UR5 robot arm equipped with proprioceptive and visual sensors. Even though the results of the state estimation were promising, no control actions were included. The same authors presented in [15] the body estimation and control in the joint space of a simulated 2-DOF robot arm through active inference. This solution included state-of-the-art regressors to estimate online the generative models. However, during the simulations, the estimation of the acceleration was unreliable and substituted with the ground truth. Regardless of the fact that only forward dynamics models had to be learned, the authors pointed out how this approach is not simpler compared with classical inverse dynamics techniques. In a parallel, related work on active inference [16], the authors successfully controlled a real 3-DOF robot arm using velocity commands. In our approach we formulate an AIC for online closed loop control of industrial robots, using low-level torque commands. We also provide a comparison with a state-of-the-art adaptive controller, and insights for design and tuning. On the other hand, the adaptive control branch of control theory [17], offers solutions to deal with manipulators subject to parameters variation and abrupt changes in the dynamics. Within adaptive controllers, two main categories can be identified: the model reference adaptive systems, and the self-tuning regulators [18]. The first technique being studied for robot manipulators was the model reference adaptive control (MRAC) [19]. The idea behind this technique is to derive a control signal to be applied to the robot actuators which will force the system to behave as specified by a chosen reference model. Furthermore, the adaptation law is designed to guarantee stability using either Lyapunov theory or hyperstability theory [20]. The other most common approach for robot control is the self-tuning adaptive control [21], [22]. The main difference between this technique and the MRAC is that the self-tuning approach represents the robot as a linear discrete-time model and it estimates online the unknown parameters, substituting them in the control law. Adaptive control of robot manipulators is required in presence of uncertain dynamics and varying payloads, however, the complexity of the controller usually increases with increasing number of DOFs. Among all the possible adaptive controllers, in this letter we choose the MRAC with hyperstability theory [20] for comparison. This choice is motivated by the fact that this approach provides adaptability to abrupt changes in the robot dynamics, and it does not require the kinematic or dynamic description of the manipulator, similarly to the AIC.

## B. Paper Structure

The paper is organised as follows: In Sec. II we present the free-energy principle and active inference in control engineering terms. In Sec. III we derive a novel AIC for a 7-DOF

robot manipulator, and we explain the model assumptions and simplifications. In Sec. IV the MRAC is presented for comparison. In Sec. V we compare the adaptability of the AIC and MRAC in a simulated pick and place task, validating the results in the real setup. We also discuss the advantages of our AIC and the open questions. Finally, Sec. VI provides a summary and directions for future work.

## II. THE ACTIVE INFERENCE FRAMEWORK

In this section we report the free-energy principle and active inference from [8], [10], rewriting only the necessary concepts in control terms, to understand the derivation of our novel AIC in Sec. III.

### A. The Free-Energy Principle

The free-energy principle is formulated in terms of Bayesian inference [23]. In this view, body perception for state estimation is framed using Bayes rule:

$$p(\mathbf{x}|\mathbf{y}) = \frac{p(\mathbf{y}|\mathbf{x})p(\mathbf{x})}{p(\mathbf{y})} \quad (1)$$

where  $p(\mathbf{x}|\mathbf{y})$  is the probability of being in the  $n$ -dimensional state  $\mathbf{x}$  given the current  $m$ -dimensional sensory input  $\mathbf{y}$ . Instead of exactly inferring the posterior, which often involves intractable integrals, an auxiliary probability distribution  $r_d(\mathbf{x})$ , called recognition density, is introduced. By minimizing the Kullback-Leibler divergence ( $D_{KL}$ ) between the true posterior  $p(\mathbf{x}|\mathbf{y})$  and  $r_d(\mathbf{x})$ , the most probable state given a sensory input is inferred [8].  $D_{KL}$  is defined as:

$$D_{KL}(r_d(\mathbf{x})||p(\mathbf{x}|\mathbf{y})) = \int r_d(\mathbf{x}) \ln \frac{r_d(\mathbf{x})}{p(\mathbf{x}|\mathbf{y})} d\mathbf{x} = \mathcal{F} + \ln p(\mathbf{y}) \quad (2)$$

In the equation above, the scalar  $\mathcal{F}$  is the so called *free-energy*. By minimizing  $\mathcal{F}$ ,  $D_{KL}$  is also minimized and the recognition density approaches the true posterior. According to the Laplace approximation [24], the controller only parametrises the sufficient statistics (e.g. mean and variance) of the recognition density.  $r_d(\mathbf{x})$  is assumed Gaussian and sharply peaked at its mean value  $\boldsymbol{\mu}$ . This approximation allows to simplify the expression for  $\mathcal{F}$  which results:

$$\mathcal{F} \approx -\ln p(\boldsymbol{\mu}, \mathbf{y}) \quad (3)$$

The mean  $\boldsymbol{\mu}$  is the internal belief about the true states  $\mathbf{x}$ . Minimizing  $\mathcal{F}$ , the controller is continuously adapting the internal belief  $\boldsymbol{\mu}$  about the states  $\mathbf{x}$  based on the current sensory input  $\mathbf{y}$ .

### B. Free-Energy Equation

Equation (3) is still general and it has to be further specified to numerically evaluate  $\mathcal{F}$ . To do so, the joint probability  $p(\boldsymbol{\mu}, \mathbf{y})$  has to be defined. This is done by introducing two generative models, one to predict the sensory data  $\mathbf{y}$ , according to the current belief  $\boldsymbol{\mu}$ , and another to describe the dynamics of the evolution of the belief  $\boldsymbol{\mu}$ .

1) *Generative Model of the Sensory Data*: The sensory data is modeled using the following expression [8]:

$$\mathbf{y} = \mathbf{g}(\boldsymbol{\mu}) + \mathbf{z} \quad (4)$$

where  $\mathbf{g}(\boldsymbol{\mu})$  represents the non-linear mapping between sensory data and states of the environment, and  $\mathbf{z}$  is Gaussian noise

$z \sim (\mathbf{0}, \Sigma_y)$ . The covariance matrix  $\Sigma_y$  also represents the controller's confidence about each sensory input.

2) *Generative Model of the State Dynamics*: In presence of time varying states  $\mathbf{x}$ , the controller has to encode a dynamic generative model of the evolution  $\boldsymbol{\mu}'$  of the belief  $\boldsymbol{\mu}$ . This generative model is defined as [8]:

$$\frac{d\boldsymbol{\mu}}{dt} = \boldsymbol{\mu}' = \mathbf{f}(\boldsymbol{\mu}) + \mathbf{w} \quad (5)$$

where  $\mathbf{f}$  is a generative function dependant on the belief about the states  $\boldsymbol{\mu}$  and  $\mathbf{w}$  is Gaussian noise  $\mathbf{w} \sim (\mathbf{0}, \Sigma_\mu)$ .

3) *Generalised Motions*: To describe the dynamics of the states, or better the belief about these dynamics, we have to introduce the concept of generalised motions [25]. Generalised motions are used to represent the states of a dynamical system, using increasingly higher order derivatives of the states of the system itself. They apply to sensory inputs as well, meaning that the generalised motions of a position measurement, for example, correspond to its higher order temporal derivatives (velocity, acceleration, and so on). The use of generalised motions allows a more accurate description of the system's states. More precisely, the generalised motions  $\tilde{\boldsymbol{\mu}}$  of the belief under local linearity assumptions [24] are, up to the second order:

$$\begin{aligned} \boldsymbol{\mu}' &= \boldsymbol{\mu}^{(1)} = \mathbf{f}(\boldsymbol{\mu}) + \mathbf{w} \\ \boldsymbol{\mu}'' &= \boldsymbol{\mu}^{(2)} = \frac{\partial \mathbf{f}}{\partial \boldsymbol{\mu}} \boldsymbol{\mu}' + \mathbf{w}' \end{aligned} \quad (6)$$

In general, we indicate the generalised motions of the states up to order  $n_d$ .<sup>1</sup> as  $\tilde{\boldsymbol{\mu}} = [\boldsymbol{\mu}, \boldsymbol{\mu}', \boldsymbol{\mu}'', \boldsymbol{\mu}''', \dots, \boldsymbol{\mu}^{(n_d)}]$ .

Similarly, the generalised motions of the sensory input are:

$$\begin{aligned} \mathbf{y} &= \mathbf{y}^{(0)} = \mathbf{g}(\boldsymbol{\mu}) + \mathbf{z} \\ \mathbf{y}' &= \mathbf{y}^{(1)} = \frac{\partial \mathbf{g}}{\partial \boldsymbol{\mu}} \boldsymbol{\mu}' + \mathbf{z}' \end{aligned} \quad (7)$$

We indicate the generalised motions of the sensory input up to order  $n_d$  as  $\tilde{\mathbf{y}} = [\mathbf{y}, \mathbf{y}', \mathbf{y}'', \mathbf{y}''', \dots, \mathbf{y}^{(n_d)}]$ .

4) *General Free-Energy Expression*: With the extra theoretical knowledge about the generalised motions, we can define an expression for the free-energy for a multivariate case in a dynamically changing environment:

$$\mathcal{F} = -\ln p(\tilde{\boldsymbol{\mu}}, \tilde{\mathbf{y}}) \quad (8)$$

The joint probability  $p(\tilde{\boldsymbol{\mu}}, \tilde{\mathbf{y}})$  has to be specified. According to [8] and to the definitions previously given, the noise at each dynamical order is considered uncorrelated. Then, according to the generalised sensory input, the sensory data at a particular order relates only with the states at the same dynamical order. Similarly, for the state dynamics, the state at a certain dynamical order are related only with those which are one order below. Then, using the chain rule, it results:

$$p(\tilde{\boldsymbol{\mu}}, \tilde{\mathbf{y}}) = \prod_{i=0}^{n_d-1} p(\mathbf{y}^{(i)} | \boldsymbol{\mu}^{(i)}) p(\boldsymbol{\mu}^{(i+1)} | \boldsymbol{\mu}^{(i)}) \quad (9)$$

<sup>1</sup>Generalised motions can extend up to infinite order but the noise at high orders is predominant, thus we can limit the chosen order to  $n_d$  [26].

Using the Laplace assumption, and thus considering Gaussian distributed probability densities, we can write:

$$\begin{aligned} p(\boldsymbol{\mu}^{(i+1)} | \boldsymbol{\mu}^{(i)}) &= \frac{1}{|\Sigma_{\boldsymbol{\mu}^{(i)}}| \sqrt{2\pi}} \exp \left\{ -\frac{1}{2} \boldsymbol{\varepsilon}_\mu^{(i)\top} \Sigma_{\boldsymbol{\mu}^{(i)}}^{-1} \boldsymbol{\varepsilon}_\mu^{(i)} \right\} \\ p(\mathbf{y}^{(i)} | \boldsymbol{\mu}^{(i)}) &= \frac{1}{|\Sigma_{\mathbf{y}^{(i)}}| \sqrt{2\pi}} \exp \left\{ -\frac{1}{2} \boldsymbol{\varepsilon}_y^{(i)\top} \Sigma_{\mathbf{y}^{(i)}}^{-1} \boldsymbol{\varepsilon}_y^{(i)} \right\} \end{aligned} \quad (10)$$

where  $\boldsymbol{\varepsilon}_y^{(i)} = (\mathbf{y}^{(i)} - \mathbf{g}^{(i)}(\boldsymbol{\mu}))$  and  $\boldsymbol{\varepsilon}_\mu^{(i)} = (\boldsymbol{\mu}^{(i+1)} - \mathbf{f}^{(i)}(\boldsymbol{\mu}))$  are respectively the sensory and state model prediction errors. Furthermore it holds:

$$\mathbf{g}^{(i)} = \frac{\partial \mathbf{g}}{\partial \boldsymbol{\mu}} \boldsymbol{\mu}^{(i)}, \quad \mathbf{f}^{(i)} = \frac{\partial \mathbf{f}}{\partial \boldsymbol{\mu}} \boldsymbol{\mu}^{(i)}, \quad \mathbf{g}^{(0)} = \mathbf{g}, \quad \mathbf{f}^{(0)} = \mathbf{f} \quad (11)$$

Substituting (9) in (8) leads to:

$$\mathcal{F} = - \sum_{i=0}^{n_d-1} \left[ \ln p(\mathbf{y}^{(i)} | \boldsymbol{\mu}^{(i)}) + \ln p(\boldsymbol{\mu}^{(i+1)} | \boldsymbol{\mu}^{(i)}) \right] \quad (12)$$

Finally, according to (10),  $\mathcal{F}$  can be expressed up to a constant as a weighted sum of squared prediction errors:

$$\mathcal{F} = \frac{1}{2} \sum_{i=0}^{n_d-1} \left[ \boldsymbol{\varepsilon}_y^{(i)\top} \Sigma_{\mathbf{y}^{(i)}}^{-1} \boldsymbol{\varepsilon}_y^{(i)} + \boldsymbol{\varepsilon}_\mu^{(i)\top} \Sigma_{\boldsymbol{\mu}^{(i)}}^{-1} \boldsymbol{\varepsilon}_\mu^{(i)} \right] + K \quad (13)$$

where  $n_d$  is the number of generalised motions chosen and  $K$  is a constant term resulting from the substitution. The minimisation of this expression can be done by refining the internal belief, thus performing state estimation, but also computing the control actions to fulfill the prior expectations and achieve a desired motion. The constant term  $K$  is neglected in the sequel since it plays no role into the minimisation problem. The next two subsections describe the approach proposed by Friston [10], [27] to minimise  $\mathcal{F}$ , using gradient descent.

### C. Belief Update for State Estimation

The belief update law for state estimation is determined from the gradient of the free-energy, with respect to each generalised motion [8], [25]:

$$\dot{\tilde{\boldsymbol{\mu}}} = \frac{d}{dt} \tilde{\boldsymbol{\mu}} - \kappa_\mu \frac{\partial \mathcal{F}}{\partial \tilde{\boldsymbol{\mu}}} \quad (14)$$

The learning rate  $\kappa_\mu$ , can be seen from a control perspective as a tuning parameter for the state update.

### D. Control Actions

In the free-energy principle the control actions play a fundamental role in the minimisation process. In fact, the control input  $\mathbf{u}$  allows to steer the system to a desired state while minimising the prediction errors in  $\mathcal{F}$ . This is done using gradient descent. Since the free-energy is not a function of the control actions directly, but the actions  $\mathbf{u}$  can influence  $\mathcal{F}$  by modifying the sensory input, we can write [8]:

$$\frac{\partial \mathcal{F}(\tilde{\boldsymbol{\mu}}, \tilde{\mathbf{y}}(\mathbf{u}))}{\partial \mathbf{u}} = \frac{\partial \tilde{\mathbf{y}}(\mathbf{u})}{\partial \mathbf{u}} \frac{\partial \mathcal{F}(\tilde{\boldsymbol{\mu}}, \mathbf{y}(\mathbf{u}))}{\partial \tilde{\mathbf{y}}(\mathbf{u})} \quad (15)$$

Dropping the dependencies for a more compact notation, the dynamics of the control actions can be written as:

$$\dot{\mathbf{u}} = -\kappa_a \frac{\partial \tilde{\mathbf{y}}}{\partial \mathbf{u}} \frac{\partial \mathcal{F}}{\partial \tilde{\mathbf{y}}} \quad (16)$$

where  $\kappa_a$  is the tuning parameter to be chosen.

### III. ROBOT ARM CONTROL WITH ACTIVE INFERENCE

In this section we derive the first model-free, computationally lightweight, online torque controller for joint space control using active inference. The established theory of Sec. II is adapted to define a novel control scheme for a generic  $n$ -DOF manipulator. The challenging problem of finding suitable generative models  $\mathbf{f}(\cdot)$  and  $\mathbf{g}(\cdot)$ , and the relation  $\partial \tilde{\mathbf{y}} / \partial \mathbf{u}$  in such a complex scenario is solved.

*Assumption 1:* The robot manipulator is equipped with position and velocity sensors, which respectively provide the two variables  $\mathbf{y}_q, \mathbf{y}_{\dot{q}} \in \mathbb{R}^n$ .

*Assumption 2:* Since only the position and velocity measurements are available, we will consider the generalised motions up to order two, so  $n_d = 2$ .

*Assumption 3:* The Gaussian noise affecting the different sensory channels is supposed uncorrelated [8], [24]. The covariance matrices for sensory input and state belief are:

$$\Sigma_{\mathbf{y}(0)} = \sigma_q I_n, \quad \Sigma_{\mathbf{y}(1)} = \sigma_{\dot{q}} I_n, \quad (17)$$

$$\Sigma_{\boldsymbol{\mu}(0)} = \sigma_{\boldsymbol{\mu}} I_n, \quad \Sigma_{\boldsymbol{\mu}(1)} = \sigma_{\boldsymbol{\mu}'} I_n \quad (18)$$

where we supposed that the controller associates four different variances to describe its confidence about sensory input and internal belief.

*Assumption 4:* The states of the environment  $\mathbf{x}$  are set as the joint positions of the robot arm. Doing so, we can control the robot arm in joint space through free-energy minimization, and simplify the equations for states update and control actions.

#### A. Generative Models and $\mathcal{F}$ for a Robot Manipulator

In order to numerically evaluate the free-energy as in (13), the two functions  $\mathbf{g}(\boldsymbol{\mu})$  and  $\mathbf{f}(\boldsymbol{\mu})$  have to be chosen.

1) *Generative Model of the Sensory Data:*  $\mathbf{g}(\boldsymbol{\mu})$  indicates the relation between the sensed values and the states. Since we chose the states to be the joint positions and the sensory data provides directly the noisy values  $\mathbf{y}_q$  and  $\mathbf{y}_{\dot{q}}$ , it holds:

$$\mathbf{g}_q(\boldsymbol{\mu}) = \boldsymbol{\mu}, \quad \partial \mathbf{g}_q / \partial \boldsymbol{\mu} = \mathbf{I} \quad (19)$$

2) *Dynamic Generative Model of the World:* Instead of modelling the true dynamics of the manipulator, we propose to define a reference model to specify the desired behaviour of the robot [8]. In particular, the world dynamics are chosen such that the robot is steered to a desired position  $\boldsymbol{\mu}_d$ . In other words, the controller believes that the states will evolve in such a way that they will reach the goal  $\boldsymbol{\mu}_d$  with the dynamics of a first order system with unitary time constant:

$$\mathbf{f}(\boldsymbol{\mu}) = \boldsymbol{\mu}_d - \boldsymbol{\mu} \quad (20)$$

The value  $\boldsymbol{\mu}_d$  is a constant  $\in \mathbb{R}^n$  corresponding to the desired set-point for the joints of the manipulator. Substituting (19) and (20) in (7) and (6), it results:

$$\begin{cases} \boldsymbol{\mu}' &= \boldsymbol{\mu}_d - \boldsymbol{\mu} + \mathbf{w} \\ \boldsymbol{\mu}'' &= -\boldsymbol{\mu}' + \mathbf{w}' \end{cases} \begin{cases} \mathbf{y}_q &= \boldsymbol{\mu} + \mathbf{z} \\ \mathbf{y}_{\dot{q}} &= \boldsymbol{\mu}' + \mathbf{z}' \end{cases} \quad (21)$$

According to (21) and (13), the free-energy expression for a generic robot manipulator under the assumptions given is:

$$\begin{aligned} \mathcal{F} &= \frac{1}{2} (\mathbf{y}_q - \boldsymbol{\mu})^\top \Sigma_{\mathbf{y}(0)}^{-1} (\mathbf{y}_q - \boldsymbol{\mu}) \\ &+ \frac{1}{2} (\mathbf{y}_{\dot{q}} - \boldsymbol{\mu}')^\top \Sigma_{\mathbf{y}(1)}^{-1} (\mathbf{y}_{\dot{q}} - \boldsymbol{\mu}') \\ &+ \frac{1}{2} (\boldsymbol{\mu}' + \boldsymbol{\mu} - \boldsymbol{\mu}_d)^\top \Sigma_{\boldsymbol{\mu}(0)}^{-1} (\boldsymbol{\mu}' + \boldsymbol{\mu} - \boldsymbol{\mu}_d) \\ &+ \frac{1}{2} (\boldsymbol{\mu}'' + \boldsymbol{\mu}')^\top \Sigma_{\boldsymbol{\mu}(1)}^{-1} (\boldsymbol{\mu}'' + \boldsymbol{\mu}') \end{aligned} \quad (22)$$

#### B. Belief Update and State Estimation for a Manipulator

According to the free-energy principle, the states of the robot manipulator can be estimated using a gradient descent scheme. Applying (14), having defined  $\mathcal{F}$  as in (22), leads to the following state update law:

$$\begin{aligned} \dot{\boldsymbol{\mu}} &= \boldsymbol{\mu}' + \kappa_{\boldsymbol{\mu}} \Sigma_{\mathbf{y}(0)}^{-1} (\mathbf{y}_q - \boldsymbol{\mu}) - \kappa_{\boldsymbol{\mu}} \Sigma_{\boldsymbol{\mu}(0)}^{-1} (\boldsymbol{\mu}' + \boldsymbol{\mu} - \boldsymbol{\mu}_d) \\ \dot{\boldsymbol{\mu}'} &= \boldsymbol{\mu}'' + \kappa_{\boldsymbol{\mu}} \Sigma_{\mathbf{y}(1)}^{-1} (\mathbf{y}_{\dot{q}} - \boldsymbol{\mu}') - \kappa_{\boldsymbol{\mu}} \Sigma_{\boldsymbol{\mu}(0)}^{-1} (\boldsymbol{\mu}' + \boldsymbol{\mu} - \boldsymbol{\mu}_d) \\ &\quad - \kappa_{\boldsymbol{\mu}} \Sigma_{\boldsymbol{\mu}(1)}^{-1} (\boldsymbol{\mu}'' + \boldsymbol{\mu}') \\ \dot{\boldsymbol{\mu}''} &= -\kappa_{\boldsymbol{\mu}} \Sigma_{\boldsymbol{\mu}(1)}^{-1} (\boldsymbol{\mu}'' + \boldsymbol{\mu}') \end{aligned} \quad (23)$$

Note that  $\kappa_{\boldsymbol{\mu}}$  is the tuning parameter for state estimation.

#### C. Control Actions for a Robot Manipulator

The final step in order to be able to steer the joints of a robot manipulator to a desired value  $\boldsymbol{\mu}_d$ , is the definition of the control actions.

1) *General Considerations:* The general actions update is expressed by (16). The partial derivatives of (22) with respect to the generalised sensory input are given by:

$$\frac{\partial \mathcal{F}}{\partial \mathbf{y}_q} = \Sigma_{\mathbf{y}(0)}^{-1} (\mathbf{y}_q - \boldsymbol{\mu}), \quad \frac{\partial \mathcal{F}}{\partial \mathbf{y}_{\dot{q}}} = \Sigma_{\mathbf{y}(1)}^{-1} (\mathbf{y}_{\dot{q}} - \boldsymbol{\mu}') \quad (24)$$

Having said that, the actions update is expressed as:

$$\dot{\mathbf{u}} = -\kappa_a \left[ \frac{\partial \mathbf{y}_q}{\partial \mathbf{u}} \Sigma_{\mathbf{y}(0)}^{-1} (\mathbf{y}_q - \boldsymbol{\mu}) + \frac{\partial \mathbf{y}_{\dot{q}}}{\partial \mathbf{u}} \Sigma_{\mathbf{y}(1)}^{-1} (\mathbf{y}_{\dot{q}} - \boldsymbol{\mu}') \right] \quad (25)$$

Active inference requires then to define the change in the sensory input with respect to the control actions, namely  $\partial \mathbf{y}_q / \partial \mathbf{u}$  and  $\partial \mathbf{y}_{\dot{q}} / \partial \mathbf{u}$ . This is usually a hard forward dynamic problem, which constituted a major complication in past control strategies. One approach to compute these relations is through online learning using high-dimensional space regressors. However, this increases the complexity of the overall scheme and can produce unreliable results, as shown by the authors in [15]. In this letter we propose to approximate the partial derivatives relying on the high adaptability of the active inference controller against unmodeled dynamics, as suggested in the conclusive remarks in [15].

2) *Approximation of the True Relation Between  $\mathbf{u}$  and  $\tilde{\mathbf{y}}$ :* Let us first analyse the structure of the partial derivative matrices in (25). The control action is a vector of  $n$  torques applied to the  $n$  joints of the robot manipulator. Each torque has a direct effect only on the corresponding joint to which it is applied. This allows us to conclude that  $\partial \mathbf{y}_q / \partial \mathbf{u}$  and  $\partial \mathbf{y}_{\dot{q}} / \partial \mathbf{u}$  are diagonal matrices.

**Algorithm 1: AIC for Robot Control.***Initialization*

$Par \leftarrow \sigma_q, \sigma_{\dot{q}}, \sigma_\mu, \sigma_{\mu'}, \kappa_\mu, \kappa_a$      $\triangleright$  Set AIC parameters  
 $\boldsymbol{\mu} = \mathbf{y}_q \in \mathbb{R}^n$      $\triangleright$  Initialise belief  
 $\boldsymbol{\mu}' = \mathbf{y}_{\dot{q}} \in \mathbb{R}^n$   
 $\boldsymbol{\mu}'' = \mathbf{0} \in \mathbb{R}^n$   
 $\mathbf{u} = \mathbf{0} \in \mathbb{R}^n$      $\triangleright$  Initialise torque commands  
 $\boldsymbol{\mu}_d \in \mathbb{R}^n$      $\triangleright$  Set prior, desired goal

*Control Loop*

$\mathbf{y}_q, \mathbf{y}_{\dot{q}}$      $\triangleright$  At high frequency  
 $\triangleright$  Retrieve sensory input  
 $\dot{\tilde{\boldsymbol{\mu}}} = \frac{d}{dt}\tilde{\boldsymbol{\mu}} - \kappa_\mu \frac{\partial \mathcal{F}}{\partial \tilde{\boldsymbol{\mu}}}$      $\triangleright$  Belief dynamics (14)  
 $\tilde{\boldsymbol{\mu}} = \tilde{\boldsymbol{\mu}} + \Delta_t \dot{\tilde{\boldsymbol{\mu}}}$      $\triangleright$  Belief update, integration  
 $\dot{\mathbf{u}} = -\kappa_a \frac{\partial \mathcal{F}}{\partial \mathbf{u}} \frac{\partial \mathcal{F}}{\partial \tilde{\boldsymbol{\mu}}}$      $\triangleright$  Action dynamics (16)  
 $\mathbf{u} = \mathbf{u} + \Delta_t \dot{\mathbf{u}}$      $\triangleright$  Action update, integration  
**return**  $\mathbf{u}$      $\triangleright$  Commanded torque

Furthermore, considering the second Newton's law, the total torque applied to a rotational joint equals the moment of inertia times the angular acceleration. The diagonal terms of the partial derivatives matrices are then time varying positive values which depend on the current robot configuration. In other words, this means that a positive torque applied to a joint will always result in a positive contribution for both position and velocity of that specific joint. In this control scheme we propose to approximate the true positive time-varying relation with a positive constant, making use of the learning rate  $\kappa_a$  as tuning parameter to achieve a sufficiently fast actions update. The control update law is finally given by:

$$\dot{\mathbf{u}} = -\kappa_a \left[ C_q \Sigma_{y(0)}^{-1} (\mathbf{y}_q - \boldsymbol{\mu}) + C_{\dot{q}} \Sigma_{y(1)}^{-1} (\mathbf{y}_{\dot{q}} - \boldsymbol{\mu}') \right] \quad (26)$$

$$\frac{\partial \mathbf{y}_q}{\partial \mathbf{u}} \approx C_q, \quad \frac{\partial \mathbf{y}_{\dot{q}}}{\partial \mathbf{u}} \approx C_{\dot{q}} \quad (27)$$

The positive definite diagonal constant matrices  $C_q$ ,  $C_{\dot{q}}$  are then set to the identity, meaning that we only encode the sign of the relation between  $\mathbf{u}$  and the change in  $\tilde{\mathbf{y}}$ .

3) *Tuning Parameters AIC*: The tuning parameters for the active inference controller are:

- $\sigma_q, \sigma_{\dot{q}}, \sigma_\mu, \sigma_{\mu'}$ : the standard deviations representing the confidence of the controller regarding its sensory input and internal belief about the states;
- $\kappa_\mu, \kappa_a$ : the learning rates for state update and control actions respectively.

Algorithm 1 reports the pseudo-code of our AIC. For state and actions update, first-order Euler integration is used.

## IV. MODEL REFERENCE ADAPTIVE CONTROLLER

The controller chosen for comparison is an MRAC. This adaptive controller allows to obtain decoupled joint dynamics, forcing every single joint  $i = 1, \dots, n$  to respond as a second order linear system with transfer function:

$$G_i(s) = \frac{\omega_i^2}{s^2 + 2\zeta\omega_i s + \omega_i^2} q_{ri}(s) \quad (28)$$

The control architecture is taken from [20], where the control is specified in terms of feedforward and feedback adaptive gain matrices. These time-varying gain matrices are adjusted by means of adaptation laws to guarantee closed loop stability in




Fig. 1. Simulated and real robot for pick and place cycle.

case of large parameters perturbations. Supposing zero initial conditions for the gains, and neglecting the derivative terms as described in [20], it holds:

$$K_0(t) = E_{01} \bar{\mathbf{q}}_e(t) \mathbf{q}(t)^\top + E_{02} \int_0^T \bar{\mathbf{q}}_e(\tau) \mathbf{q}(\tau) d\tau \quad (29)$$

$$K_1(t) = E_{11} \bar{\mathbf{q}}_e(t) \dot{\mathbf{q}}(t)^\top + E_{12} \int_0^T \bar{\mathbf{q}}_e(\tau) \dot{\mathbf{q}}(\tau) d\tau \quad (30)$$

$$Q_0(t) = F_{01} \bar{\mathbf{q}}_e(t) \mathbf{q}_r(t)^\top + F_{02} \int_0^T \bar{\mathbf{q}}_e(\tau) \mathbf{q}_r(\tau) d\tau \quad (31)$$

$$Q_1(t) = F_{11} \bar{\mathbf{q}}_e(t) \dot{\mathbf{q}}_r(t)^\top + F_{12} \int_0^T \bar{\mathbf{q}}_e(\tau) \dot{\mathbf{q}}_r(\tau) d\tau \quad (32)$$

$$f(t) = \alpha_1 \bar{\mathbf{q}}_e(t) + \alpha_2 \int_0^T \bar{\mathbf{q}}_e(\tau) d\tau \quad (33)$$

The variables  $\mathbf{q}_r$  and  $\dot{\mathbf{q}}_r$  are the desired references to track. The diagonal matrices  $E_{jk}$  and  $F_{jk} \in \mathbb{R}^{n \times n}$ , and the vector  $\alpha_k \in \mathbb{R}^n$  with  $j = \{0, 1\}$  and  $k = \{1, 2\}$ , are the tuning parameters for the proportional-integral adaptation law. The term  $\bar{\mathbf{q}}_e$  is called modified joint angle error vector [20]:

$$\bar{\mathbf{q}}_e = P_2 [\mathbf{q}_r(t) - \mathbf{q}(t)] + P_3 [\dot{\mathbf{q}}_r(t) - \dot{\mathbf{q}}(t)] \quad (34)$$

with  $P_2$  and  $P_3$  diagonal weighting matrices. The MRAC, similarly to the AIC, does not need the dynamic description of the robot manipulator, and it is scalable to high DOF. However, the number of the tuning parameters increases with the degrees of freedom, unlike for the AIC.

## V. EXPERIMENTAL EVALUATION

This section presents the performance comparison between AIC and MRAC. To analyse the adaptability of the algorithms against unmodeled dynamics, the controllers are tuned in simulation using an approximated model of the robot, and then transferred to the real system. The tests performed are based on a pick and place cycle using the Franka Emika Panda 7-DOF robot manipulator, as in Fig. 1, with different payloads.

## A. Remarks About the Tuning Procedure for the Controllers

Before presenting the simulations and experimental results, we provide some observations regarding the number of parameters and the different tuning procedures for the AIC and MRAC.

1) *Number of Tuning Parameters*: The number of tuning parameters for the MRAC equals the number of DOFs times the number of weighting terms. According to Sec. IV, this results in  $17 \times n$  parameters to be tuned. Regarding the AIC, instead, the number of tuning parameters is independent from the DOFs

and it equals 6, following the formulation presented in Sec. III. The lower number of parameters resulted in an overall easier tuning procedure for the active inference controller. As a final remark, to modify the behaviour of the step response for the AIC, such as rise time and settling time, one should change the internal reference model  $f(\mu)$  instead of fine tuning the controller's parameters.

2) *AIC Tuning Procedure*: To obtain a satisfactory response for the AIC, we performed the following steps: 1) We set the controller confidence about sensory input and internal belief to one; 2) We disabled the control actions and incremented the learning rate  $\kappa_\mu$  until the state estimation in a static situation was fast enough; 3) We included the control actions and increased the learning rate  $\kappa_a$  until the robot was steered to the desired position, showing significant oscillations; 4) We dampened the oscillatory behaviour decreasing the sensory confidence about the most noisy sensors and the internal belief about velocities.

### B. Simulations With Approximated Model

The performance of AIC and MRAC in simulation are now presented. The task is a pick and place cycle where the desired joint values are chosen such that the arm simulates the pick and place of an object from one bin to the other, positioning the end-effector in A, B or C, see Fig 1. This is achieved giving every 6 [s] a set-point in joint space following the sequence:  $q_A$ ,  $q_B$ ,  $q_C$ ,  $q_B$ ,  $q_A$ , where:

- $q_A = [1, 0.5, 0, -2, 0, 2.5, 0]$  [rad]
- $q_B = [0, 0.2, 0, -1, 0, 1.2, 0]$  [rad]
- $q_C = [-1, 0.5, 0, -1.2, 0, 1.6, 0]$  [rad]

The controllers have been tuned using a considerably inaccurate model of the robot arm on purpose. The links have been approximated as cuboids, and 20% random uncertainty in each link's mass has been assumed. This will allow to evaluate later on the adaptability performance while applying the controllers to the real manipulator. The joint values and control actions using AIC and MRAC, are depicted in Fig. 2. Note that, for the MRAC, saturation of the control input at  $\pm 85$  Nm is reached for some of the joints, after providing the new goal position.

### C. Experiments on the Real Setup

The same controllers tuned in simulation using the approximated model of the 7-DOF robot arm are now applied to control the real manipulator. Two tests are performed: first, the pick and place cycle of the previous section is repeated in the real robot, without re-tuning the controllers. Second, the AIC and MRAC are re-tuned in the real robot and used to pick and place different objects. The real setup is controlled using a standard laptop running Ubuntu 16.04 with RT kernel, 8-cores Intel i7-4710MQ 2.50 GHz.

1) *Pick and Place Cycle on the Real Robot*: We applied the MRAC and AIC from simulation to the real 7-DOF Franka Emika Panda. It is important to notice that, besides having different physical parameters, the real setup is already gravity compensated. The AIC and MRAC are simply applied on top of this intrinsic controller. This is already a considerable change in the system's dynamics, but to further increase the level of uncertainties, an end-effector is attached to the robot. From a modeling point of view, the system used for tuning the controllers in simulation is completely different from the real one. Usually, a controller tuned in simulation will not directly work on a real setup, especially if the initial model was not

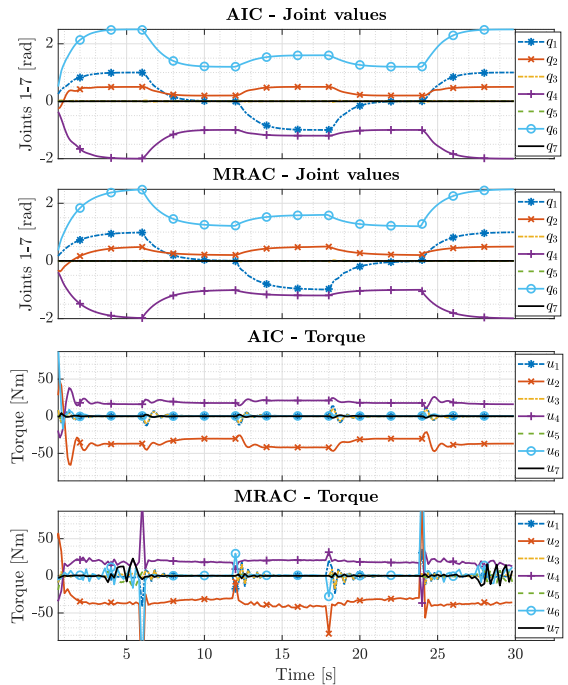


Fig. 2. Response and control actions for the 7-DOF robot arm controlled through AIC and MRAC with approximated dynamics.

accurate. This was indeed the case for the MRAC which, when transferred to the real robot, could not control the setup leading to an immediate safety stop. Nonetheless, this was not the case for our novel AIC: its strong capabilities to cope with unmodeled dynamics allowed to transfer the controller from the simulation to the real setup without re-tuning. For clarity, we only report the response of the AIC during the initial part of the pick and place cycle ( $q_B \rightarrow q_A \rightarrow q_B$ ) in Fig. 3. Joint 7 is not reported to limit redundant information, since no motion was required. As can be seen the AIC can successfully control the manipulator, however, the effect of the large uncertainties introduced for the tuning, resulted in some initial jittering, especially in joint 6.<sup>2</sup> In other words, the AIC tuned in simulation resulted too aggressive for the real robot. This is because in simulation the AIC had to compensate also for gravity, thus a faster torque update was required. The learning rate  $\kappa_a$  is the same for every joint but the jittering effect is mostly visible in joint 6. This is because in the last part of the kinematic chain, the resulting inertia acting on a joint is lower, and so it is its reluctance to changes in velocities. To completely remove the jittering, one can simply reduce the learning rate  $\kappa_a$  to lower the torque update rate. The AIC and MRAC have been tested against large external disturbances such as a human pushing the robot during motion. AIC resulted more compliant than MRAC, showing at the same time a faster and less oscillatory disturbance rejection.

2) *Pick and Place With Different Payloads*: In order to use the MRAC on the real robot, a severe re-tuning of 63 parameters had to be performed, to stabilise the response due to the large unmodeled dynamics. For the AIC,  $\kappa_a$  has been reduced to eliminate the jittering, as well as  $\sigma_q$ ,  $\sigma_{\dot{q}}$  to give more importance to the measurements and further reduce oscillations. The two controllers are used to perform a pick and place of an almost empty water bottle ( $\approx 0.1$  [kg]) and a full water bottle

<sup>2</sup>[https://youtu.be/Vsb0MzOp\\_TY](https://youtu.be/Vsb0MzOp_TY)






Fig. 3. AIC on real setup without re-tuning from simulation. Focus on initial part of the pick and place cycle ( $q_B \rightarrow q_A \rightarrow q_B$ ) to highlight jittering.

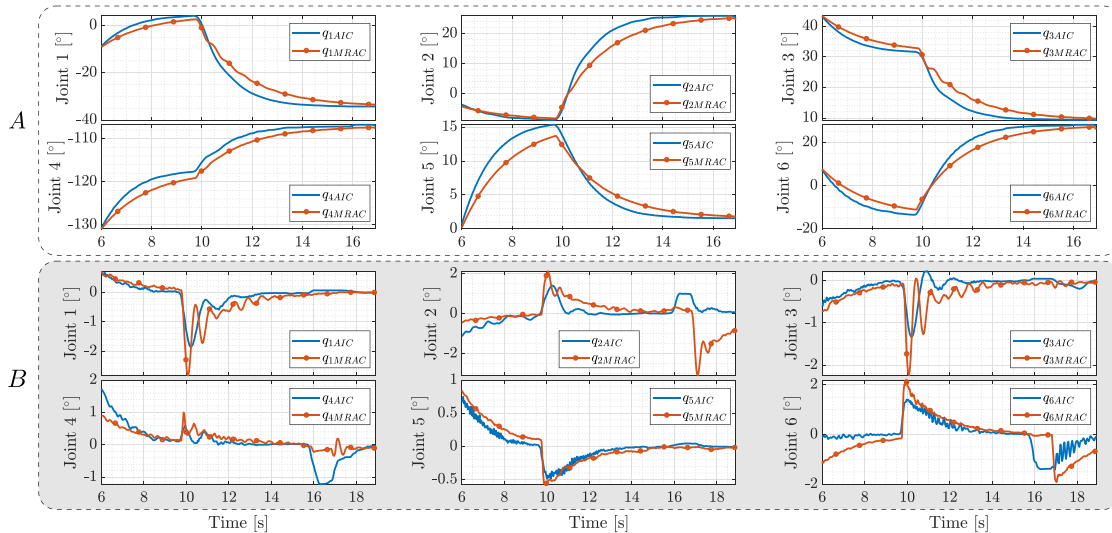


Fig. 4. (A) Lift and place of empty bottle. (B) Difference of trajectories between empty and filled bottle during lift, place, and release.

( $\approx 0.7$  [kg]), as in Fig. 1. In Fig. 4 A we show the responses of AIC and MRAC in case of the almost empty bottle, during *lifting* and *placing*. As can be seen, the AIC presents a faster convergence to the set-point, as well as smoother trajectories with less oscillations. To achieve a satisfactory response, we had to increase the stiffness of the MRAC, while the AIC could be kept compliant. Furthermore, in Fig. 4 B we show the difference of the trajectories in joint space between the case with empty and full bottle, considering *lifting*, *placing* and *releasing*. Both controllers adapt to the heavier payload, making the trajectory converge to the one with lightweight bottle. AIC behaves similarly to the MRAC, yet it presents considerably less oscillations which reflected in smoother placing of the heavy object. The bigger error appearing at around 16 [s] is due to the releasing of the heavy object. The effect is more visible in the AIC since the robot is more compliant. In a sense, the AIC behaves similarly to a human arm, when an unexpected weight is dropped. This is an additional evidence of the bio-inspired character of the controller. The AIC can also be tuned to be stiffer if this effect is not desired.

#### D. Discussion and Implementation Notes

Our novel AIC showed high adaptability, allowing to transfer from simulation to real robot without re-tuning. Furthermore, the AIC showed superior performance with respect to the MRAC in pick and place scenarios. The AIC is compliant while allowing to compensate for large perturbations. However, even though there is a strong evidence of stability and robustness of the AIC for a complex non-linear system, finding a formal stability proof

is still an open question. Similarly to a linear case, one should determine a set of learning rates which guarantees convergence. Intuitively, active inference is a gradient descent on a quadratic and convex function thus, for some set of learning rates, the algorithm should converge to the global minima. A possible approach to a formal proof is to use Lyapunov theory as for the back-propagation algorithm in neural networks. Active inference is, in a sense, back-propagating the sensitivity of the control input with respect to the free-energy, to minimise  $\mathcal{F}$ . Properly addressing this proof mathematically would require a deep analysis which is out of the scope of the current paper. Another remark relates to the computational load of AIC. According to Algorithm 1, our novel AIC has a computational complexity of  $\mathcal{O}(n)$  where  $n$  is the number of DOFs. Given the structure of the generative models and covariance matrices chosen, the AIC reduces to 16 sums of vectors and 15 scalar-vector multiplications with  $n$ -dimensional vectors. On the other hand, the complexity of the MRAC is  $\mathcal{O}(n^3)$ . Another optimised computed torque algorithm such as LGP [28], which relies on learning dynamical models, has a cost of  $\mathcal{O}(N^2)$  for online learning, where  $N$  is the number of data points (i.e.  $N \approx 300$ ). Finally, the Franka Emika Panda requires the control signals to be ready within 300 [ $\mu$ s] to guarantee a functioning frequency of 1 [kHz]: Our AIC can perform at such a high loop rate without any package loss; is straightforward to implement; and extremely simple to tune. The source code for simulations<sup>3</sup> and experiments<sup>4</sup> is freely available on GitHub.

<sup>3</sup>[https://github.com/cpezzato/panda\\_simulation](https://github.com/cpezzato/panda_simulation)

<sup>4</sup>[https://github.com/cpezzato/active\\_inference](https://github.com/cpezzato/active_inference)

## VI. CONCLUSION

In this letter we derived the first active inference torque controller for online joint space control of robot manipulators. Our approach makes use of the alleged adaptability of active inference, to introduce simplifications for the generative models, obtaining a model-free scheme which is less sensitive to unmodeled dynamics, is easily scalable to high DOF and is computationally inexpensive. With the proposed controller structure we overcame the complexity barrier of previous approaches, making possible control loops at high frequency with active inference. Simulations and experiments in a real setup with a 7-DOF robot arm showed that our AIC is suitable for tasks in which the dynamic model of the plant is unknown or subject to large changes. The performance of our novel AIC has been compared with that of a state-of-the-art MRAC, in different pick and place scenarios. The AIC shows better adaptability properties, allowing to transfer from simulation to real setup without re-tuning. In addition, the AIC resulted easier to tune and implement. With this work we confirmed the value of active inference to develop more adaptive control of robot manipulators. This is only the first step in this direction, future work should proof the closed-loop stability of active inference, define generative models to account for dynamic requirements and motion constraints, and be extended to other control modalities, such as control in Cartesian space or impedance control.

## ACKNOWLEDGMENT

The authors would like to thank Prof. Dr. Martijn Wisse for the helpful discussions, together with the whole group working on active inference at the Cognitive Robotics department.

## REFERENCES

- [1] S. Vijayakumar and S. Shaal, "Locally weighted projection regression: Incremental real time learning in high dimensional space," in *Proc. Int. Conf. Mach. Learn.*, 2000, pp. 1079–1086.
- [2] D. Nguyen-Tuong, J. Peters, and M. Seeger, "Local gaussian process regression for real time online model learning," in *Proc Neural Inf. Process. Syst.*, 2008, pp. 1193–1200.
- [3] M. Matteucci, "Elearn: Evolutionary learning of rich neural topologies," vol. 103, Tech. Rep. CMU-CALD-02, 2006. [Online]. Available: <https://apps.dtic.mil/dtic/tr/fulltext/u2/a456062.pdf>
- [4] F. Ledezma and S. Haddadin, "First-order-principles-based constructive network topologies: An application to robot inverse dynamics," in *Proc. IEEE-RAS 17th Int. Conf. Humanoid Robot. (Humanoids)*, 2017, pp. 438–445.
- [5] D. Keppler, F. Peters, N. Ratliff, and S. Shaal, "A new data source for inverse dynamics learning," in *Proc IEEE/RJS Conf. Intell. Robots Syst.*, 2017, pp. 4723–4730.
- [6] L. Jamone, B. Damas, and J. Santos-Victor, "Incremental learning of context-dependent dynamic internal models for robot control," in *Proc. IEEE Int. Symp. Intell. Control*, 2014, pp. 1336–1341.
- [7] K. J. Friston, J. Mattout, and J. Kilner, "Action understanding and active inference," *Biol. Cybern.*, vol. 104, no. 1–2, pp. 137–160, 2011.
- [8] C. Buckley, C. Kim, S. McGregor, and A. Seth, "The free energy principle for action and perception: A mathematical review," *J. Math. Psychol.*, vol. 81, pp. 55–79, 2017.
- [9] R. Bogacz, "A tutorial on the free-energy framework for modelling perception and learning," *J. Math. Psychol.* vol. 76, pp. 198–211, 2017.
- [10] K. J. Friston, "The free-energy principle: a unified brain theory?," *Nature Rev. Neuroscience*, vol. 11, no. 2, pp. 27–138, 2010.
- [11] K. J. Friston, J. Daunizeau, and S. Kiebel, "Action and behavior: a free-energy formulation," *Biol. Cybern.*, vol. 102, no. 3, pp. 227–260, 2010.
- [12] L. Pio-Lopez, A. Nizard, K. Friston, and G. Pezzulo, "Active inference and robot control: a case study," *J. Royal Soc. Interface*, vol. 13, pp. 1–12, 2016. [Online]. Available: <http://doi.org/10.1098/rsif.2016.0616>
- [13] A. C. Mercadé, "Robot manipulator control under the active inference framework," (*Unpublished MSc thesis*), TU Delft, pp. 1–72, 2018.
- [14] P. Lanillos and G. Cheng, "Adaptive robot body learning and estimation through predictive coding," in *Proc. Int. Conf. Intell. Robots Syst.*, 2018, pp. 4083–4090.
- [15] P. Lanillos and G. Cheng, "Active inference with function learning for robot body perception," in *Proc. Int. Workshop Continual Unsupervised Sensorimotor Learn.*, ICDL-Epirob, 2018, pp. 1–5.
- [16] G. Oliver, P. Lanillos, and G. Cheng, "Active inference body perception and action for humanoid robots," 2020, *arXiv:1906.03022v3*.
- [17] K. Astrom, "Theory and applications of adaptive control - a survey," *Automatica*, vol. 19, no. 5, pp. 471–486, 1983.
- [18] T. Hsia, "Adaptive control of robot manipulators - a review," in *Proc IEEE Int. Conf. Robot. Autom.*, 1986, pp. 183–189.
- [19] D. D. Zhang and B. Wei, "A review on model reference adaptive control of robotic manipulators," *Annu. Rev. Control*, vol. 43, pp. 188–198, 2017.
- [20] M. Tarokh, "Hyperstability approach to the synthesis of adaptive controllers for robot manipulators," in *Proc IEEE Int. Conf. Robot. Autom.*, 1991, pp. 2154–2159.
- [21] R. Walters and M. Bayoumi, "Application of a self-tuning pole-placement regulator to an industrial manipulator," in *Proc 21st IEEE Conf. Decis. Control*, 1991, pp. 323–329.
- [22] A. Koivo and T. Guo, "Adaptive linear controller for robotic manipulators," *IEEE Trans. Autom. Control*, vol. AC-28, no. 2, pp. 162–171, Feb. 1983.
- [23] D. Lindley, "Bayesian statistics, a review," *SIAM*, vol. 2, pp. 1–83, 1972.
- [24] 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.
- [25] K. Friston, K. Stephan, B. Li, and J. Daunizeau, "Generalised filtering," *Math. Problems Eng.*, vol. 2010, pp. 1–34, 2010.
- [26] K. Friston, "Hierarchical models in the brain," *PLoS Comput. Biol.*, vol. 4, no. 11, 2008, Art. no. e1000211.
- [27] K. Friston, J. Daunizeau, and S. Kiebel, "Reinforcement learning or active inference?," *PLoS one*, vol. 4, no. 7, 2009, Art. no. e6421.
- [28] D. Nguyen-Tuong and J. Peters, "Learning robot dynamics for computed torque control using local gaussian processes regression," in *Proc. Symp. Learn. Adaptive Behav. Robot. Syst.*, 2008, pp. 59–64.