Lagrangian Inspired Polynomial Estimator for black-box learning and control of underactuated systems

Giulio Giacomuzzo GIACOMUZZO@DEI.UNIPD.IT
University of Padova, Padova, Italy

Riccardo Cescon RICCARDO.CESCON@EPFL.CH
École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland

Diego Romeres ROMERES@MERL.COM
Mitsubishi Electric Research Lab, Cambridge, MA, USA

Ruggero Carli CARLIRUG@DEI.UNIPD.IT
University of Padova, Padova, Italy

Alberto Dalla Libera DALLALIBER@DEI.UNIPD.IT
University of Padova, Padova, Italy

Abstract

The Lagrangian Inspired Polynomial (LIP) estimator Giacomuzzo et al. (2023) is a black-box estimator based on Gaussian Process Regression, recently presented for the inverse dynamics identification of Lagrangian systems. It relies on a novel multi-output kernel that embeds the structure of the Euler-Lagrange equation. In this work, we extend its analysis to the class of underactuated robots. First, we show that, despite being a black-box model, the LIP allows estimating kinetic and potential energies, as well as the inertial, Coriolis, and gravity components directly from the overall torque measures. Then we exploit these properties to derive a two-stage energy-based controller for the swing-up and stabilization of balancing robots. Experimental results on a simulated Pendubot confirm the feasibility of the proposed approach.

Keywords: Learning for Control, Inverse Dynamics Identification, Gaussian Process Regression

1. Introduction

In recent years, learning for control in complex robotics applications has gained increasing interest. Several machine learning methods have been presented, both for modeling and direct synthesis of the controller. In this context, under-actuated robots (UR) are an important class of systems. UR are mechanical systems characterized by fewer control inputs than degrees of freedom (DOF). Such systems are ubiquitous in robotics: examples are manipulators with passive joints, autonomous bicycles and motorcycles, bipedal robots, and most of aerospace and marine vehicles. To control such systems, for instance, several Reinforcement Learning algorithms have been explored. These algorithms aim at automatically learning a control law, see for example Lillicrap et al. (2015) for a model-free approach or Deisenroth and Rasmussen (2011); Amadio et al. (2022) for model-based solutions. Although effective, generally these methods ignore previous contributions on UR from control theory. As a result, they typically require a huge amount of interactions with the system and do not provide any performance guarantees.

A viable alternative consists of combining model learning with classic model-based control methods. Within this approach, a model of the system dynamics is derived from experimental data using machine learning techniques without the need for manually deriving accurate models of the

dynamics, which is the main drawback of classic model-based methods. Then, the learned model is used inside a model-based controller, exploiting results from control theory.

Typically, such control strategies rely on the inverse dynamics model, which relates torques to the robot trajectories. The problem of learning the inverse dynamics from data has been extensively explored in the literature Nguyen-Tuong and Peters (2010); Romeres et al. (2016); Camoriano et al. (2016); Dalla Libera and Carli (2020). An interesting class of solutions is represented by black-box methods, which propose to learn the inverse dynamics map by means of universal approximators, without any knowledge about the underlying system kinematics and dynamics Gijsberts and Metta (2011); Polydoros et al. (2015); Schreiter et al. (2015); Rueckert et al. (2017). However, learning inverse dynamics models for the control of UR is particularly challenging. Indeed, most of the strategies proposed for UR control require models that return the different inverse dynamics components and the system energy, see Spong and Praly (2005). Typically, black-box models do not provide such information. Underactuation further exacerbates these problems: torques of the under-actuated dimensions are constant signals equal to zero, leading to an ill-posed estimation problem.

Physics-Informed (PI) models are a class of black-box solutions that could mitigate the afore-mentioned problem Rezaei-Shoshtari et al. (2019); Lutter et al. (2019); Karniadakis et al. (2021). PI methods have been recently proposed to improve the data efficiency and generalization of black-box estimators. Instead of learning the inverse dynamics in a completely unstructured manner, PI methods embed knowledge from physics as priors in the model structure. As a result, learned models not only show better accuracy but also provide insights about relevant physical properties. For instance, in Lutter et al. (2019), the authors proposed Deep Lagrangian Networks (DeLaNs), a neural networks model with structure inspired by Euler-Lagrange (EL) equations. Similar ideas have been explored in Ayed et al. (2019), Greydanus et al. (2019), Cranmer et al. (2020).

In the context of inverse dynamics identification of fully-actuated manipulators, we proposed a PI solution based on Gaussian Process Regression (GPR), named LIP estimator Giacomuzzo et al. (2023). The LIP estimator derives a multioutput kernel of the torques by encoding EL equations. In this work, we extend the analysis of the LIP estimator to the class of UR and show its applicability in traditional model-based control strategies for UR.

Contribution: Our contribution is twofold. First, we review the recently proposed LIP estimator and extend its analysis in the context of UR, showing its ability to learn the inverse dynamics map, its different components, and the kinetic and potential energies. Second, we exploit the LIP to derive an energy-based controller for the stabilization of balancing UR. The controller is based on two modules: an energy-based controller for the swing-up, and an LQR which stabilizes the system once it has reached the unstable equilibrium. The energy-based controller performs a partial feedback linearization on the actuated system and a regulation of the energy to steer the non-actuated system to a trajectory passing through the unstable equilibrium. Once the system is sufficiently close to the target, the control is switched to the LQR. We show that the LIP model is suitable to implement this kind of controller since it returns the inertia matrix, the Coriolis and gravity torques, energy estimates, and the linearization of the system dynamics required by the LQR. This allows performing the swing-up and stabilization of the balancing robot in a completely black-box fashion. The method is validated on a simulated setup involving a Pendubot Spong and Block (1995). The LIP method is compared with DeLaNs and with the true analytical model, which has been included as a baseline. Results confirm the feasibility of the proposed approach.

The paper unfolds as follows. In Section 2, we review the modeling and control of UR and we define the inverse dynamics identification problem. Then, in Section 3, we present the LIP

estimator and show how to derive the balancing controller. Section 4 reports the results of the performed experiments, while Section 5 draws the conclusions.

2. Background and problem statement

In this section, we first provide the inverse dynamics model of UR under the rigid body assumption. Then we review the class of energy-based controllers we consider in this work.

2.1. Inverse dynamics

Consider a mechanical system with nDOF and let $q \in \mathbb{R}^n$ be the vector of generalized coordinates. We assume the system to have $m < n$ control inputs, each of which actuates a single DOF. We partition the vector $q \in \mathbb{R}^n$ as $q^T = [q_1^T, q_2^T]$, where $q_1 \in \mathbb{R}^m$ and $q_2 \in \mathbb{R}^{n-m}$ refer respectively to the actuated and the non-actuated DOFs. Under the rigid body assumption, the inverse dynamics of the system can be derived from the Euler-Lagrange equations (see Spong and Praly (2005)) as

$$\underbrace{\begin{bmatrix} M_{11}(q) & M_{12}(q) \\ M_{21}(q) & M_{22}(q) \end{bmatrix} \begin{bmatrix} \ddot{q}_1 \\ \ddot{q}_2 \end{bmatrix}}_{m(q,\ddot{q})} + \underbrace{\begin{bmatrix} c_1(q,\dot{q}) \\ c_2(q,\dot{q}) \end{bmatrix}}_{c(q,\dot{q})} + \underbrace{\begin{bmatrix} g_1(q) \\ g_2(q) \end{bmatrix}}_{g(q)} = \underbrace{\begin{bmatrix} \tau_1 \\ 0 \end{bmatrix}}_{\tau},$$

(1)

where

$$M(q) = \begin{bmatrix} M_{11}(q) & M_{12}(q) \\ M_{21}(q) & M_{22}(q) \end{bmatrix}$$

is the symmetric, positive definite inertia matrix, $m(q, \ddot{q})$ represents the inertial torque, $c(q, \dot{q})$ accounts for the Coriolis and centripetal torques, $g(q)$ represents the gravity contribution while $\tau_1 \in \mathbb{R}^m$ is the vector of generalized torques produced by the m actuators. For the sake of brevity, in the following, we will explicitly point out the dependency on $q, \dot{q}$ only when necessary.

The inverse dynamics identification problem consists of estimating the map in eq. (1) that relates $\tilde{x} = (q, \dot{q}, \ddot{q})$ and the torques τ from a set of noisy measures. Black-box solutions treat the inverse dynamics as an unknown function and, generally, rely on universal approximators to estimate the function from experimental data. Several solutions, including our approach, adopt GPR Williams and Rasmussen (2006), which is a framework for Bayesian inference widely used in machine learning and robotics.

Remark 1 *The standard approach when using GPR consists of considering the different torque components independently and solving n independent regression problems, one for each generalized coordinate. In UR, torques of the under-actuated dimensions are constant signals equal to zero, leading to an ill-posed estimation problem if the components are learned independently of each other, which prevents the possibility of deriving any inverse dynamics model useful for model-based control strategies.*

2.2. Model-based balancing control of underactuated systems

In what follows, we focus our attention on a particular class of robots described by eq. (1), known as balancing systems. Common examples of balancing systems are the Cartpole, the Furuta Pendulum, the Acrobot, and the Pendubot. For such systems, the typical control challenge requires swinging up and balancing the robot in the unstable equilibrium point, hereafter denoted by $x_\star = [q_\star^T, \dot{q}_\star^T]^T$ with

$\dot{\boldsymbol{q}}_\star = 0$. In the remainder of this section, we review one of the most common approaches presented in the literature, which combines an energy-based controller for the swing-up and an LQR for the balancing.

2.2.1. ENERGY-BASED SWING-UP CONTROLLER

We review the energy-based swing-up controller presented in Spong (1996); Kolesnichenko and Shiriaev (2002). The first step consists of a partial feedback linearization. From eq. (1) we isolate the dynamics of the actuated and non-actuated subsystems respectively as

$$M_{11}\ddot{\boldsymbol{q}}_1 + M_{12}\ddot{\boldsymbol{q}}_2 + c_1 + g_1 = \boldsymbol{\tau}_1, \tag{2}$$

$$M_{21}\ddot{\boldsymbol{q}}_1 + M_{22}\ddot{\boldsymbol{q}}_2 + c_2 + g_2 = 0. \tag{3}$$

We can solve eq. (3) for $\ddot{\boldsymbol{q}}_2$ as $\ddot{\boldsymbol{q}}_2 = -M_{22}^{-1}(M_{21}\ddot{\boldsymbol{q}}_1 + c_2 + g_2)$ since M_{22} is invertible (given that $M > 0$). Substituting the resulting expression into eq. (2) leads to

$$\bar{M}_{11}\ddot{\boldsymbol{q}}_1 + \bar{c}_1 + \bar{g}_1 = \boldsymbol{\tau}_1, \tag{4}$$

with $\bar{M}_{11} = M_{11} - M_{12}M_{22}^{-1}M_{21}$, $\bar{c}_1 = c_1 - M_{12}M_{22}^{-1}c_2$ and $\bar{g}_1 = g_1 - M_{12}M_{22}^{-1}g_2$.

A feedback linearizing controller for eq. (4) can be defined as

$$\boldsymbol{\tau}_1 = \bar{M}_{11}u + \bar{c}_1 + \bar{g}_1, \tag{5}$$

where u is a design parameter. Choosing $\boldsymbol{\tau}_1$ as in eq. (5) leads to the following closed-loop system

$$\ddot{\boldsymbol{q}}_1 = u \tag{6}$$

$$M_{22}\ddot{\boldsymbol{q}}_2 + c_2 + g_2 = -M_{21}u \tag{7}$$

We obtained a linear second-order dynamics for the actuated subsystem. Selecting u according to

$$u = -k_1\boldsymbol{q}_1 - k_2\dot{\boldsymbol{q}}_1 + \bar{u}, \quad k_1, k_2 > 0, \tag{8}$$

makes the linear subsystem in eq. (6) asymptotically stable for $\bar{u} = 0$. The remaining design problem is the choice of $\bar{u}$, which can be used to stabilize the non-actuated dynamics in eq. (7). Most of the existing literature proposes to design $\bar{u}$ based on energy concepts, penalizing the mismatch w.r.t. the system energy at the desired configuration. This leads to control laws of the form

$$\bar{u} = f_e(\mathcal{T}, \mathcal{V}), \tag{9}$$

where $\mathcal{T}(\boldsymbol{q}, \dot{\boldsymbol{q}}) = \frac{1}{2}\dot{\boldsymbol{q}}^T M(\boldsymbol{q})\dot{\boldsymbol{q}}$ is the kinetic energy, while $\mathcal{V}(\boldsymbol{q})$ denotes the potential energy. However, the choice of f_e depends on the system of interest. We will clarify our choice in the experimental section after introducing our setup. Examples of energy-compensating terms for common balancing robots can be found in Spong (1996), Åström and Furuta (2000), Xin et al. (2013).

Remark 2 *As highlighted in Spong (1996), the controller presented in this section is not stabilizing the system to a fixed point but only to a manifold. For this reason, in applications such as the swing-up of balancing robots, the control must switch to another controller achieving local asymptotic stability to the equilibrium.*

2.2.2. LINEAR QUADRATIC REGULATOR

To stabilize the system at the equilibrium $\boldsymbol{x}_\star$ we resort to a Linear Quadratic Regulator (LQR). First, we provide a state space description for the system with dynamics expressed in eq. (1). Let the system state $\boldsymbol{x}$ be such that $\boldsymbol{x}^T = [\boldsymbol{q}^T, \dot{\boldsymbol{q}}^T]$. The state evolution can be derived from eq. (1) as

$$\dot{\boldsymbol{x}} = \begin{bmatrix} \dot{\boldsymbol{q}} \\ -M^{-1}(\boldsymbol{q})[\boldsymbol{c}(\boldsymbol{q},\dot{\boldsymbol{q}}) + \boldsymbol{g}(\boldsymbol{q})] \end{bmatrix} + \begin{bmatrix} 0_n \\ M^{-1}(\boldsymbol{q}) \end{bmatrix} \boldsymbol{\tau} = f(\boldsymbol{x}) + g(\boldsymbol{x}, \boldsymbol{\tau}) \tag{10}$$

Now we linearize the non-linear system in eq. (10) around $\boldsymbol{x}_\star^T$. Moreover, let $\boldsymbol{\tau}_\star$ be the reference input at the equilibrium. Applying a first-order Taylor expansion, the system dynamics around $(\boldsymbol{x}_\star, \boldsymbol{\tau}_\star)$ can be approximated as

$$\dot{\boldsymbol{x}} \approx A(\boldsymbol{x} - \boldsymbol{x}_\star) + B(\boldsymbol{\tau} - \boldsymbol{\tau}_\star). \tag{11}$$

Recalling that at the equilibrium $\dot{\boldsymbol{q}}_\star = 0$, matrices A and B are

$$A = \frac{\partial f}{\partial \boldsymbol{x}}^T \bigg|_{\boldsymbol{x}_\star, \boldsymbol{\tau}_\star} = \begin{bmatrix} 0 & I \\ -M^{-1}(\boldsymbol{q})\frac{\partial g(\boldsymbol{q})}{\partial \boldsymbol{q}} & M^{-1}(\boldsymbol{q})C(\boldsymbol{q},\dot{\boldsymbol{q}}) \end{bmatrix}\bigg|_{\boldsymbol{x}_\star, \boldsymbol{\tau}_\star} = \begin{bmatrix} 0 & I \\ -M^{-1}(\boldsymbol{q}_\star)\frac{\partial g(\boldsymbol{q}_\star)}{\partial \boldsymbol{q}} & 0 \end{bmatrix} \tag{12}$$

and

$$B = \frac{\partial g}{\partial \boldsymbol{\tau}}^T \bigg|_{\boldsymbol{x}_\star, \boldsymbol{\tau}_\star} = \begin{bmatrix} 0 \\ M^{-1}(\boldsymbol{q}_\star) \end{bmatrix}, \tag{13}$$

where $C(\boldsymbol{q}, \dot{\boldsymbol{q}}) \in \mathbb{R}^{n \times n}$ is the skew-symmetric Coriolis matrix, such that $\boldsymbol{c}(\boldsymbol{q}, \dot{\boldsymbol{q}}) = C(\boldsymbol{q}, \dot{\boldsymbol{q}})\dot{\boldsymbol{q}}$.

Then, we introduce the infinite horizon control problem on the linearized system in (11), namely,

$$\operatorname*{argmin}_{\boldsymbol{\tau}} \int_0^\infty (\boldsymbol{x} - \boldsymbol{x}_\star)^T Q(\boldsymbol{x} - \boldsymbol{x}_\star) + (\boldsymbol{\tau} - \boldsymbol{\tau}_\star)^T R(\boldsymbol{\tau} - \boldsymbol{\tau}_\star)dt,$$

which leads to the control input $\boldsymbol{\tau} = -K\boldsymbol{x}$, with K being the optimal control gain, which can be computed by solving the *continuous time algebraic Riccati equation*.

3. Lagrangian Inspired Polynomial Estimator for modeling and control of under-actuated systems

In this section, we briefly review the proposed inverse dynamics LIP estimator for the control of UR. Then, we detail the application of the LIP model to the UR control described in Section 2.

3.1. LIP estimator

The LIP estimator is based on GPR Williams and Rasmussen (2006), which is a framework for Bayesian inference widely used in machine learning and robotics applications. Generally, GPR solutions for inverse dynamics identification model each torque component $\tau^i(\tilde{\boldsymbol{x}})$ as a Gaussian Process (GP) by assuming $\tau^i(\tilde{\boldsymbol{x}})$s are independent given $\tilde{\boldsymbol{x}}$, and then apply standard GPR inference. As discussed in the previous section, this approach is not effective for UR. The LIP estimator follows an alternative strategy and defines the kinetic and potential energies as two independent GPs. Then, it derives a multi-output kernel of the torques by exploiting EL equations. In this way, the inverse dynamics problem is well defined also in the UR setup.

Our framework models $\mathcal{T}$ and $\mathcal{V}$ as independent GPs, namely $\mathcal{T} \sim \mathcal{GP}\left(0, k^{\mathcal{T}}(\cdot, \cdot)\right)$ and $\mathcal{V} \sim \mathcal{GP}\left(0, k^{\mathcal{V}}(\cdot, \cdot)\right)$, where $k^{\mathcal{T}}$ and $k^{\mathcal{V}}$ are the kernels that defines the covariance of $\mathcal{T}$ and $\mathcal{V}$. For instance, let $\tilde{x}$ and $\tilde{x}'$ be two input locations, then the covariance between the values of $\mathcal{T}$ at $\tilde{x}$ and $\tilde{x}'$ is $E\left[\mathcal{T}(\tilde{x})\mathcal{T}(\tilde{x}')\right] = k^{\mathcal{T}}(\tilde{x}, \tilde{x}')$. For convenience, we introduce also $K_{XX'}^{\mathcal{T}}$, that is the matrix that collects $k^{\mathcal{T}}$ evaluated at $X = \{\tilde{x}_1, \ldots, \tilde{x}_N\}$, $X' = \{\tilde{x}'_1, \ldots, \tilde{x}'_M\}$:

$$K_{XX'}^{\mathcal{T}} = \begin{bmatrix} k^{\mathcal{T}}(\tilde{x}_1, \tilde{x}'_1) & \ldots & k^{\mathcal{T}}(\tilde{x}_1, \tilde{x}'_M) \\ \vdots & \vdots & \vdots \\ k^{\mathcal{T}}(\tilde{x}_N, \tilde{x}'_1) & \ldots & k^{\mathcal{T}}(\tilde{x}_N, \tilde{x}'_M) \end{bmatrix}.$$

Similar definitions hold for $k^{\mathcal{V}}$ and all the other kernels we will introduce. The LIP estimator defines $k^{\mathcal{T}}$ and $k^{\mathcal{V}}$ relying on a polynomial formulation that extends Dalla Libera and Carli (2020). Due to space constraint, we refer the interested reader to Giacomuzzo et al. (2023) for further details.

Notice that, (i) once we define $\mathcal{T}$ and $\mathcal{V}$ as zero-mean GPs, for the properties of GPs also the Lagrangian function $\mathcal{L} = \mathcal{T} - \mathcal{V}$ is a zero-mean GP with kernel $k^{\mathcal{L}}(\cdot, \cdot) = k^{\mathcal{T}}(\cdot, \cdot) + k^{\mathcal{V}}(\cdot, \cdot)$, namely $\mathcal{L} \sim \mathcal{GP}\left(0, k^{\mathcal{L}}(\cdot, \cdot)\right)$. Furthermore, (ii) under rigid body assumptions, each $\tau^i(\tilde{x})$ is described by a linear differential equation of $\mathcal{L}$, namely, $\tau^i = \frac{d}{dt}\left(\frac{\partial \mathcal{L}}{\partial \dot{q}^i}\right) - \frac{\partial \mathcal{L}}{\partial q^i}$. Expanding explicit derivations w.r.t. time, we obtain

$$\tau^i = \sum_{j=1}^{n} \left(\frac{\partial^2 \mathcal{L}}{\partial \dot{q}^i \partial \dot{q}^j} \ddot{q}^j + \frac{\partial^2 \mathcal{L}}{\partial \dot{q}^i \partial q^j} \dot{q}^j \right) - \frac{\partial \mathcal{L}}{\partial q^i} =: \mathcal{G}_i \mathcal{L},$$

where $\mathcal{G}_i$ is the linear operator that maps $\mathcal{L}$ in the linear differential equation of τ^i. Finally, (iii) GPs are closed w.r.t. linear operators Särkkä (2011), namely, if f is a zero-mean GP with kernel $k^f(\cdot, \cdot)$ and $g(\tilde{x}) = \mathcal{G}f(\tilde{x})$, then also $g(\tilde{x})$ is a zero-mean GP with kernel $k^g(\tilde{x}, \tilde{x}') = \mathcal{G}\mathcal{G}'k^f(\tilde{x}, \tilde{x}')$. The last expression means that $k^g(\tilde{x}, \tilde{x}')$ is obtained by applying two times the operator $\mathcal{G}$ to $k^f(\tilde{x}, \tilde{x}')$, first w.r.t. the input $\tilde{x}$ then w.r.t. $\tilde{x}'$.

Based on (i), (ii), and (iii), we conclude that torques are a zero-mean GP, with covariance defined by a multi-output kernel $k^\tau(\tilde{x}, \tilde{x}') \in \mathbb{R}^{n \times n}$ that encodes the EL equations, and is expressed as

$$k^\tau(\tilde{x}, \tilde{x}') = \begin{bmatrix} \mathcal{G}_1 \mathcal{G}'_1 k^{\mathcal{L}}(\tilde{x}, \tilde{x}') & \ldots & \mathcal{G}_1 \mathcal{G}'_n k^{\mathcal{L}}(\tilde{x}, \tilde{x}') \\ \vdots & \ddots & \vdots \\ \mathcal{G}_n \mathcal{G}'_1 k^{\mathcal{L}}(\tilde{x}, \tilde{x}') & \ldots & \mathcal{G}_n \mathcal{G}'_n k^{\mathcal{L}}(\tilde{x}, \tilde{x}') \end{bmatrix}. \tag{14}$$

To derive (14) we applied the multi-output version of property (iii). For a more detailed derivation, we refer the interested reader to Giacomuzzo et al. (2023).

Once k^τ is defined, we can compute torque estimates following standard GPR. Let X be a set of N training input locations, and $y = \left[y_1^T, \ldots, y_N^T\right]^T$ the respective torque measurements, with $y_i \in \mathbb{R}^n$ equal to the torque measures at input $\tilde{x}_i$. The LIP torque estimate in a general input location $\tilde{x}$ is

$$\hat{\tau}(\tilde{x}) = K_{\tilde{x}X}^\tau \left(K_{XX}^\tau + \Sigma_e\right)^{-1} y, \tag{15}$$

where Σ_e is a regularization parameter that accounts for the additive Gaussian noise modeled by GPR. For each dimension, we assumed independent and identically distributed noise, thus obtaining a block diagonal matrix with equal diagonal blocks, namely $\Sigma_e = \text{diag}(\Sigma_{e_1}, \ldots, \Sigma_{e_N})$, with $\Sigma_{e_1} = \Sigma_{e_2} = \ldots = \Sigma_{e_N} = \text{diag}(\sigma_{e_1}^2, \ldots, \sigma_{e_n}^2)$, where $\sigma_{e_i}^2$ is the variance of the τ^i measures.

3.2. LIP for control

In this section, we describe how to estimate the kinetic and potential energies, the inertial, Coriolis, and gravity vector as well as $\frac{\partial g}{\partial q}$ required by the control laws presented in Section 2.

3.2.1. KINETIC AND POTENTIAL ENERGIES

The kinetic energy $\mathcal{T}$ and potential energy $\mathcal{V}$ are required to implement the energy based control law described by eq. (9). The LIP model provides a principled way to estimate them from the torque measurements y. Indeed, within the LIP framework, $\mathcal{T}$, $\mathcal{V}$ and τ are jointly Gaussian distributed, since the prior of τ is derived by applying the linear operator $\mathcal{G}$ to the kinetic and potential GPs $\mathcal{T}$ and $\mathcal{V}$. The covariances between $\mathcal{T}$ and τ and between $\mathcal{V}$ and τ at general input locations $\tilde{x}$ and $\tilde{x}'$ are

$$\mathrm{Cov}[\mathcal{T}(\tilde{x}), \tau(\tilde{x}')] = \mathrm{Cov}[\mathcal{T}(\tilde{x}), \mathcal{G}\mathcal{L}(\tilde{x}')] =: k^{\mathcal{T}\tau}(\tilde{x}, \tilde{x}'), \tag{16a}$$

$$\mathrm{Cov}[\mathcal{V}(\tilde{x}), \tau(\tilde{x}')] = \mathrm{Cov}[\mathcal{V}(\tilde{x}), \mathcal{G}\mathcal{L}(\tilde{x}')] =: k^{\mathcal{V}\tau}(\tilde{x}, \tilde{x}'). \tag{16b}$$

Recalling that we model $\mathcal{T}$ and $\mathcal{V}$ as independent GPs, and in view of the properties of GPs under linear operators, we obtain

$$k^{\mathcal{T}\tau}(\tilde{x}, \tilde{x}') = \mathrm{Cov}[\mathcal{T}(\tilde{x}), \mathcal{G}\mathcal{T}(\tilde{x}')] = \left[\mathcal{G}'k^{\mathcal{T}}(\tilde{x}, \tilde{x}')\right]^T = [\mathcal{G}'_1 k^{\mathcal{T}}(\tilde{x}, \tilde{x}'), \dots, \mathcal{G}'_n k^{\mathcal{T}}(\tilde{x}, \tilde{x}')] \quad \text{(17a)}$$

$$k^{\mathcal{V}\tau}(\tilde{x}, \tilde{x}') = \mathrm{Cov}[\mathcal{V}(\tilde{x}), \mathcal{G}\mathcal{V}(\tilde{x}')] = \left[\mathcal{G}'k^{\mathcal{V}}(\tilde{x}, \tilde{x}')\right]^T = [\mathcal{G}'_1 k^{\mathcal{V}}(\tilde{x}, \tilde{x}'), \dots, \mathcal{G}'_n k^{\mathcal{V}}(\tilde{x}, \tilde{x}')] \quad \text{(17b)}$$

From the posterior distributions of $\mathcal{T}$ and $\mathcal{V}$ given y we can obtain an estimate of the energies at arbitrary input location $\tilde{x}$ as

$$\hat{\mathcal{T}}(\tilde{x}) = K_{\tilde{x}X}^{\mathcal{T}\tau}(K_{XX}^{\tau} + \Sigma_e)^{-1}y, \quad \hat{\mathcal{V}}(\tilde{x}) = K_{\tilde{x}X}^{\mathcal{V}\tau}(K_{XX}^{\tau} + \Sigma_e)^{-1}y, \tag{18}$$

where the covariance matrices $K_{\tilde{x}X}^{\mathcal{T}\tau} \in \mathbb{R}^{1 \times nN}$ and $K_{\tilde{x}X}^{\mathcal{V}\tau} \in \mathbb{R}^{1 \times nN}$ are obtained as $K_{\tilde{x}X}^{\mathcal{T}\tau} = [k^{\mathcal{T}\tau}(\tilde{x}, x_1), \dots, k^{\mathcal{T}\tau}(\tilde{x}, \tilde{x}_N)]$ and $K_{\tilde{x}X}^{\mathcal{V}\tau} = [k^{\mathcal{V}\tau}(\tilde{x}, \tilde{x}_1), \dots, k^{\mathcal{V}\tau}(\tilde{x}, \tilde{x}_N)]$.

3.2.2. TORQUE COMPONENTS

The energy based control law in eq. (5) requires to estimate m, c and g, while the LQR in Sec. 2.2.2 requires the inverse of the inertia matrix M as well as the term $\frac{\partial g}{\partial q}$. These quantities are derived similarly to how we computed the energies in the previous section.

First, the inertia matrix is estimated component wise. The element in position ij of M is

$$M^{ij}(q) = \frac{\partial^2 \mathcal{T}(q, \dot{q})}{\partial \dot{q}^i \partial \dot{q}^j} =: \mathcal{G}_{M^{ij}}\mathcal{T}(\tilde{x}),$$

where we introduced the linear operator $\mathcal{G}_{M^{ij}}$. The covariance between $M_{i,j}$ and τ at general input locations $\tilde{x}$ and $\tilde{x}'$ can be computed as

$$\mathrm{Cov}\left[M^{ij}(x), \tau(\tilde{x}')\right] = \mathrm{Cov}\left[\mathcal{G}_{M^{ij}}\mathcal{T}(\tilde{x}), \mathcal{G}\mathcal{T}(\tilde{x}')\right] = \mathcal{G}_{M^{ij}}\mathcal{G}k^{\mathcal{T}}(\tilde{x}, \tilde{x}') =: k^{M^{ij}\tau}(\tilde{x}, \tilde{x}').$$

Accordingly, we can estimate M^{ij} at any input location $\tilde{x}$ as $\hat{M}^{ij}(\tilde{x}) = K_{\tilde{x}X}^{M^{ij}\tau}(K_{XX}^{\tau} + \Sigma_e)^{-1}y$. with $K_{\tilde{x}X}^{M^{ij}\tau} = [k^{M^{ij}\tau}(\tilde{x}, \tilde{x}_1), \dots, k^{M^{ij}\tau}(\tilde{x}, \tilde{x}_N)]$. Then, from the estimate of the inertia matrix we can derive an estimate of the inertial torque component as $\hat{m}(\tilde{x}) = \hat{M}(\tilde{x})\ddot{q}$.

Next, we can estimate the gravity contribution $\boldsymbol{g}$. Recall that the i-th component of the vector $\boldsymbol{g}$ is defined as $g^i(\boldsymbol{q}) = \frac{\partial \mathcal{V}(\boldsymbol{q})}{\partial q^i}$. The covariance between g^i and $\boldsymbol{\tau}$ at general input locations $\tilde{\boldsymbol{x}}$ and $\tilde{\boldsymbol{x}}'$ is

$$\text{Cov}\left[g^i(\boldsymbol{x}), \boldsymbol{\tau}(\tilde{\boldsymbol{x}}')\right] = \text{Cov}\left[\frac{\partial \mathcal{V}(\tilde{\boldsymbol{x}})}{\partial q^i}, \mathcal{G}\mathcal{V}(\tilde{\boldsymbol{x}}')\right] = \frac{\partial}{\partial q^i}\mathcal{G}'k^{\mathcal{V}}(\tilde{\boldsymbol{x}}, \tilde{\boldsymbol{x}}') =: k^{g^i\tau}(\tilde{\boldsymbol{x}}, \tilde{\boldsymbol{x}}').$$

Accordingly g^i can be estimated at any input location $\tilde{\boldsymbol{x}}$ as $\hat{g}(\tilde{\boldsymbol{x}}) = K_{\tilde{\boldsymbol{x}}X}^{g^i\tau}(K_{XX}^\tau + \Sigma_e)^{-1}\boldsymbol{y}$, with $K_{\tilde{\boldsymbol{x}}X}^{g^i\tau} = [k^{g^i\tau}(\tilde{\boldsymbol{x}}, \tilde{\boldsymbol{x}}_1), \dots, k^{g^i\tau}(\tilde{\boldsymbol{x}}, \tilde{\boldsymbol{x}}_N)]$.

Given the estimates of $\boldsymbol{m}$ and $\boldsymbol{g}$, it is possible to obtain an estimate of the Coriolis and centripetal contribution $\boldsymbol{c}$ as $\hat{\boldsymbol{c}}(\tilde{\boldsymbol{x}}) = \hat{\boldsymbol{\tau}}(\tilde{\boldsymbol{x}}) - \hat{\boldsymbol{m}}(\tilde{\boldsymbol{x}}) - \hat{\boldsymbol{g}}(\tilde{\boldsymbol{x}})$.

Finally, the matrix $\frac{\partial \boldsymbol{g}}{\partial \boldsymbol{q}} \in \mathbb{R}^{n \times n}$ required in the computation of matrix A in eq. (12), is estimated following the same procedure adopted for the inertia matrix. Its element in position ij is given by

$$\left[\frac{\partial \boldsymbol{g}(\boldsymbol{q})}{\partial \boldsymbol{q}}\right]_{ij} = \frac{\partial^2 \mathcal{V}(\boldsymbol{q})}{\partial q^i \partial q^j} =: \mathcal{G}_{G^{ij}}\mathcal{V}(\boldsymbol{q}) =: G^{ij}(\tilde{\boldsymbol{x}}).$$

Then, the covariance between $G^{ij}(\tilde{\boldsymbol{x}})$ and the torque vector $\boldsymbol{\tau}$ at general input locations $\tilde{\boldsymbol{x}}$ and $\tilde{\boldsymbol{x}}'$ is

$$\text{Cov}\left[G^{ij}(\tilde{\boldsymbol{x}}), \boldsymbol{\tau}(\tilde{\boldsymbol{x}}')\right] = \text{Cov}\left[\mathcal{G}_{G^{ij}}\mathcal{V}(\boldsymbol{q}), \mathcal{G}\mathcal{V}(\tilde{\boldsymbol{x}}')\right] = \mathcal{G}_{G^{ij}}\mathcal{G}'k^{\mathcal{V}}(\tilde{\boldsymbol{x}}, \tilde{\boldsymbol{x}}') =: k^{G^{ij}\tau}(\tilde{\boldsymbol{x}}, \tilde{\boldsymbol{x}}').$$

Finally, the estimate at a general input location $\tilde{\boldsymbol{x}}$ is computed as $\hat{G}^{ij}(\tilde{\boldsymbol{x}}) = K_{\tilde{\boldsymbol{x}}X}^{G^{ij}\tau}(K_{XX}^\tau + \Sigma_e)^{-1}\boldsymbol{y}$, with $K_{\tilde{\boldsymbol{x}}X}^{G^{ij}\tau} = [k^{G^{ij}\tau}(\tilde{\boldsymbol{x}}, \tilde{\boldsymbol{x}}_1), \dots, k^{G^{ij}\tau}(\tilde{\boldsymbol{x}}, \tilde{\boldsymbol{x}}_N)]$.

4. Experimental results

We chose the Pendubot as a benchmark for testing the LIP estimator performance on the estimation and control of a UR. The Pendubot is a double-inverted pendulum composed of two links with only the first one actuated. In this example the links lengths are $l_1 = l_2 = 20\,cm$ and masses are $m_1 = m_2 = 0.3\,kg$ respectively. The task consists of swinging up and controlling the system in the unstable equilibrium point. To carry out the task, we first derive an inverse dynamics model using the LIP estimator, then we implement the controller presented in Section 2 with the learned model. We implemented all the simulations and controllers in Python and with the auxiliary library PyTorch, Paszke et al. (2019).

4.1. Estimation evaluation

First, we test the ability to estimate the torque $\boldsymbol{\tau}$ and the components $\mathbf{m}(\boldsymbol{q}, \ddot{\boldsymbol{q}}), \mathbf{c}(\boldsymbol{q}, \dot{\boldsymbol{q}}), \mathbf{g}(\boldsymbol{q})$ and the potential and kinetic energies on the simulated Pendubot. The trajectories used for training and test were collected simulating eq. (1). The input torques are different realizations of filtered white Gaussian noise with variance 0.5 and cutoff frequency of the filter 1 Hz. Test trajectories are significantly different from the training trajectories, because of high accelerations that are not present in the training set. The hyperparameters of the LIP estimator were optimized by marginal likelihood maximization. As a baseline, the LIP model is compared with the DeLaN model and we considered the same structure as in Lutter et al. (2019) and used the Adam optimizer. The accuracy when predicting torques and energies is evaluated by means of Mean Squared Error (MSE) in Fig. 1. We also show the estimated trajectories on the test dataset in Fig. 2. The proposed LIP

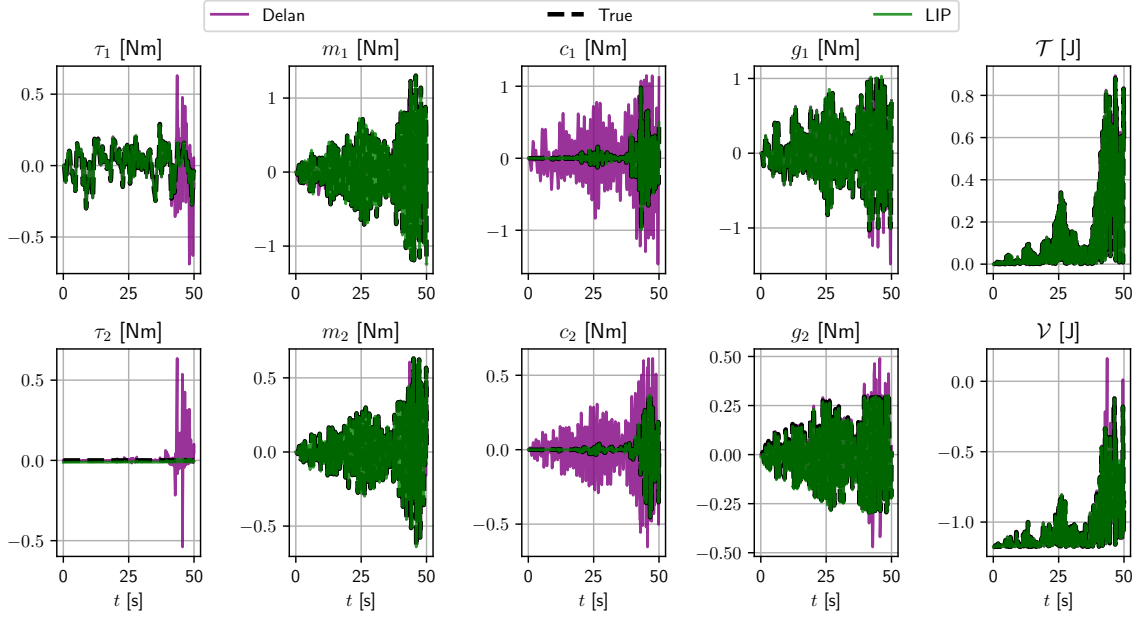Figure 1: MSE of torque components and energies of the LIP (green) and DeLaN (red).

Figure 2: Torques and energies and estimates. The dashed blue line represents the ground truth, LIP and DeLaN are in green and red, respectively.

model outperforms the NN-based torque estimator and can also well estimate all the components of the torque and the energies. Instead, even though the DeLaN produces reasonable estimates of torque and energies completely misses the decomposition of inertia and Coriolis components. The ability of the LIP model to well capture all the different components is fundamental for model-based controllers like the energy-based one presented in Section 2.

4.2. Control performance evaluation

The goal of this section is to show that using the black-box LIP model we can perform a complete swing up and balancing of the Pendubot without requiring the knowledge of the model. We implemented the general control scheme described in eq. (5) and eq. (8), with $k_1 := k_p/k_d$ and $k_2 := k_v/k_d$. Regarding the choice of $\bar{u}$, we implemented the same strategy as in Xin et al. (2013), namely

$$\bar{u} = \frac{k_E(E - E^\star)\Xi}{k_d(k_d M_{22} + k_E(E - E^\star)|M|)},$$

9

Figure 3: Time evolution of the generalized coordinates q_1 and q_2. In dashed black is the real controller, in green the LIP-based one, and in orange the nominal one.

with $\Xi = |M|(k_p q_1 + k_v \dot{q}_1) + k_d(M_{12}(c_2 + g_2) - M_{22}(c_1 + g_1))$. The values chosen for the controller's gains are $k_E = 0.8$, $k_p = 0.31$, $k_d = 0.029$ and $k_v = 0.064$. The energy-based controller is switched to the LQR when the Pendubot reaches a neighborhood of the vertical equilibrium, described by the condition $|q_1| + |q_2| + 0.1|\dot{q}_1| + 0.1|\dot{q}_2| < \epsilon$, with $\epsilon = 0.6$.

The feedback matrix K of the LQR controller is computed using the *lqr* routine of the *Python Control Systems Toolbox* with $Q = I$ and $R = 100$, obtaining $K = [8.2\ 8.2\ 1.8\ 1.2]$. The response of the LIP-based controller is compared against the response given by the true model and a nominal model. In the latter, we considered uncertain knowledge of the system's parameters, assuming that $m_1 = 0.315\,kg$, $m_2 = 0.285\,kg$, $l_1 = 20.4\,cm$, and $l_1 = 19.2\,cm$. The DeLaN estimator has not been considered for this task, since it does not provide the possibility to implement the LQR controller. Fig. 3 depicts the time evolution of the joints' variables. We can notice that the LIP based controller is capable of swinging up and balancing the system and that the trajectory produced is almost identical to the one we obtain using the true model. With the nominal model we are still able to achieve the same goal, however the performances of this controller are worse than the data-driven ones. The swing-up time needed to reach a neighborhood of the vertical position and switch to the LQR is higher. Also, the stabilizing controller takes more time, with a larger overshoot.

5. Conclusions and future works

In this paper, we extended our GP-based PI inverse dynamics estimator, named Lagrangian Inspired Polynomial estimator, to modeling and control of UR. Differently from standard GP-based black-box models, the LIP estimator, besides deriving an inverse dynamics model in a black-box fashion, can also estimate the different dynamics components and potential and kinetic energies. Thanks to these properties, the resulting model can rely on model-based controllers from control theory to solve control tasks for UR. The proposed strategy has been compared in simulation with state-of-art solutions, confirming the effectiveness of the proposed strategy.

Acknowledgements

References

Fabio Amadio, Alberto Dalla Libera, Riccardo Antonello, Daniel Nikovski, Ruggero Carli, and Diego Romeres. Model-based policy search using monte carlo gradient estimation with real systems application. *IEEE Transactions on Robotics*, 38(6):3879–3898, 2022.

Ibrahim Ayed, Emmanuel de Bézenac, Arthur Pajot, Julien Brajard, and Patrick Gallinari. Learning dynamical systems from partial observations. *arXiv preprint arXiv:1902.11136*, 2019.

R. Camoriano, S. Traversaro, L. Rosasco, G. Metta, and F. Nori. Incremental semiparametric inverse dynamics learning. In *ICRA*, pages 544–550, 2016.

Miles Cranmer, Sam Greydanus, Stephan Hoyer, Peter Battaglia, David Spergel, and Shirley Ho. Lagrangian neural networks. *arXiv preprint arXiv:2003.04630*, 2020.

A. Dalla Libera and R. Carli. A data-efficient geometrically inspired polynomial kernel for robot inverse dynamic. *IEEE Robotics and Automation Letters*, 5(1):24–31, 2020. doi: 10.1109/LRA.2019.2945240.

Marc Peter Deisenroth and Carl Edward Rasmussen. Pilco: A model-based and data-efficient approach to policy search. In *ICML*, 2011.

Giulio Giacomuzzo, Alberto Dalla Libera, Diego Romeres, and Ruggero Carli. A black-box physics-informed estimator based on gaussian process regression for robot inverse dynamics identification. *arXiv preprint arXiv:2310.06585*, 2023.

Arjan Gijsberts and Giorgio Metta. Incremental learning of robot dynamics using random features. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 951–956, 2011.

Samuel Greydanus, Misko Dzamba, and Jason Yosinski. Hamiltonian neural networks. *Advances in neural information processing systems*, 32, 2019.

George Em Karniadakis, Ioannis G Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3(6):422–440, 2021.

Olga Kolesnichenko and Anton S Shiriaev. Partial stabilization of underactuated euler–lagrange systems via a class of feedback transformations. *Systems & Control Letters*, 45(2):121–132, 2002.

Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.

Michael Lutter, Christian Ritter, and Jan Peters. Deep lagrangian networks: Using physics as model prior for deep learning. *arXiv preprint arXiv:1907.04490*, 2019.

Duy Nguyen-Tuong and Jan Peters. Using model knowledge for learning inverse dynamics. In *2010 IEEE International Conference on Robotics and Automation*, pages 2677–2682, 2010. doi: 10.1109/ROBOT.2010.5509858.

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, 2019.

Athanasios S. Polydoros, Lazaros Nalpantidis, and Volker Krüger. Real-time deep learning of robotic manipulator inverse dynamics. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3442–3448, 2015. doi: 10.1109/IROS.2015.7353857.

S. Rezaei-Shoshtari, D. Meger, and I. Sharf. Cascaded gaussian processes for data-efficient robot dynamics learning. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6871–6877, 2019. doi: 10.1109/IROS40897.2019.8968107.

Diego Romeres, Mattia Zorzi, Raffaello Camoriano, and Alessandro Chiuso. Online semi-parametric learning for inverse dynamics modeling. In *2016 IEEE 55th Conference on Decision and Control (CDC)*, pages 2945–2950. IEEE, 2016.

Elmar Rueckert, Moritz Nakatenus, Samuele Tosatto, and Jan Peters. Learning inverse dynamics models in o(n) time with lstm networks. In *2017 IEEE-RAS 17th International Conference on Humanoid Robotics (Humanoids)*, pages 811–816, 2017. doi: 10.1109/HUMANOIDS.2017. 8246965.

Simo Särkkä. Linear operators and stochastic partial differential equations in gaussian process regression. In *International Conference on Artificial Neural Networks*, pages 151–158. Springer, 2011.

J. Schreiter, P. Englert, D. Nguyen-Tuong, and M. Toussaint. Sparse gaussian process regression for compliant, real-time robot control. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2586–2591, 2015. doi: 10.1109/ICRA.2015.7139547.

Mark W Spong. Energy based control of a class of underactuated mechanical systems. *IFAC Proceedings Volumes*, 29(1):2828–2832, 1996.

Mark W Spong and Daniel J Block. The pendubot: A mechatronic system for control research and education. In *Proceedings of 1995 34th IEEE conference on decision and control*, volume 1, pages 555–556. IEEE, 1995.

Mark W Spong and Laurent Praly. Control of underactuated mechanical systems using switching and saturation. In *Control using logic-based switching*, pages 162–172. Springer, 2005.

Christopher KI Williams and Carl Edward Rasmussen. *Gaussian processes for machine learning*, volume 2. MIT press Cambridge, MA, 2006.

Xin Xin, Seiji Tanaka, Jinhua She, and Taiga Yamasaki. New analytical results of energy-based swing-up control for the pendubot. *International Journal of Non-Linear Mechanics*, 52:110–118, 2013.

KJ Åström and K Furuta. Swinging up a pendulum by energy control. *Automatica*, 36(2):287–295, 2000.