

Learning Control from Raw Position Measurements

Amadio, Fabio; Dalla Libera, Alberto; Nikovski, Daniel N.; Carli, Ruggero; Romeres, Diego

TR2023-060 June 01, 2023

Abstract

We propose a Model-Based Reinforcement Learning (MBRL) algorithm named VF-MC-PILCO, specifically designed for application to mechanical systems where velocities cannot be directly measured. This circumstance, if not adequately considered, can compromise the success of MBRL approaches. To cope with this problem, we define a velocity-free state formulation which consists of the collection of past positions and inputs. Then, VF-MC-PILCO uses Gaussian Process Regression to model the dynamics of the velocity-free state and optimizes the control policy through a particle-based policy gradient approach. We compare VF-MC-PILCO with our previous MBRL algorithm, MC-PILCO4PMS, which handles the lack of direct velocity measurements by modeling the presence of velocity estimators. Results on both simulated (cart-pole and UR5 robot) and real mechanical systems (Furuta pendulum and a ball-and-plate rig) show that the two algorithms achieve similar results. Conveniently, VF-MC-PILCO does not require the design and implementation of state estimators, which can be a challenging and time-consuming activity to be performed by an expert user.

American Control Conference (ACC) 2023

© 2023 MERL. This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

Learning Control from Raw Position Measurements

Fabio Amadio¹, Alberto Dalla Libera¹, Daniel Nikovski², Ruggero Carli¹, Diego Romeres²

Abstract—We propose a Model-Based Reinforcement Learning (MBRL) algorithm named VF-MC-PILCO, specifically designed for application to mechanical systems where velocities cannot be directly measured. This circumstance, if not adequately considered, can compromise the success of MBRL approaches. To cope with this problem, we define a velocity-free state formulation which consists of the collection of past positions and inputs. Then, VF-MC-PILCO uses Gaussian Process Regression to model the dynamics of the velocity-free state and optimizes the control policy through a particle-based policy gradient approach. We compare VF-MC-PILCO with our previous MBRL algorithm, MC-PILCO4PMS, which handles the lack of direct velocity measurements by modeling the presence of velocity estimators. Results on both simulated (cart-pole and UR5 robot) and real mechanical systems (Furuta pendulum and a ball-and-plate rig) show that the two algorithms achieve similar results. Conveniently, VF-MC-PILCO does not require the design and implementation of state estimators, which can be a challenging and time-consuming activity to be performed by an expert user.

I. INTRODUCTION

Model-Based Reinforcement Learning (MBRL) [1] proved to be a promising strategy to overcome the challenges of delivering Reinforcement Learning (RL) [2] solutions to real-world problems. In fact, standard Model-Free RL algorithms usually require a massive amount of interaction with the systems to solve the assigned task. This requirement might be unfeasible in many real-world applications, e.g. control of mechanical systems and robotics, due to the limited time available and the risk of damaging the devices involved in such a long training phase. On the other hand, MBRL uses the collected data to train a predictive model of the environment and updates the policy based on model simulations. With this strategy, we are able to extract more valuable information from the available data and increase data efficiency [3].

Nevertheless, the effectiveness of MBRL methods strongly depends on how accurately the trained model can simulate the behavior of the environment. For this reason, it is necessary to adopt stochastic models in order to capture uncertainty about predictions. Different classes of models have been employed, from Gaussian Processes (GPs) [4] in [5], [6], [7], [8], to ensembles of probabilistic deep neural networks in [9], [10].

The application of this kind of method to real-world environments is affected by another major problem: the full state of a real system is often only partially measurable.

For instance, when dealing with mechanical systems, joint positions can be measured by means of proper sensors, e.g. encoders, while velocities can only be estimated from the history of sampled positions. In our previous work [11], we proposed an MBRL algorithm, called MC-PILCO4PMS, specifically tailored to deal with Partially Measurable Systems and take correctly into account the presence of online and offline state observers. It proved able to robustly learn from scratch how to control mechanical systems, in both simulated and real environments even when the velocity is not directly measurable. However, the tuning of accurate filters and state estimators could be particularly challenging and time-consuming for systems affected by significant noise.

In this work, we present an alternative approach, called VF-MC-PILCO, that completely circumvents the necessity of estimating velocities, by working only with the history of measured positions and applied control actions. We adopted a Velocity-Free (VF) predictive model, similar to the one proposed in [12], [13], together with a control policy whose input depends only on positions. Compared to the works in [12], [13], which were focused on the modeling part, in this work we propose a complete VF solution to the RL problem. The proposed method is first tested in two simulated systems with an increasing number of DoF i.e., a cart-pole and a 6 DoF UR5 robot. Then, VF-MC-PILCO is tested in two real systems with an increasing level of difficulty for the velocity estimation i.e., a Furuta pendulum equipped with encoders and a ball-and-plate system equipped with an external camera to infer the ball positions. VF-MC-PILCO correctly solved all the tasks with performance similar to the one obtained by MC-PILCO4PMS. To solve such tasks, MC-PILCO4PMS must accurately reproduce the online filter employed inside the policy optimization algorithm and implement an offline filtering procedure for estimating the velocities needed for modeling. On the contrary, VF-MC-PILCO, by working directly with raw measurements, presents an alternative way that requires less effort and expertise from the user, obtaining similar performance despite the presence of significant noise. The comparisons are carried against MC-PILCO4PMS because this algorithm was shown to outperform other s.o.t.a. MBRL algorithms in [11].

The remainder of this paper is structured as follows. In Sec. II, we formulate the problem we aim to solve, as well as describe the use of GPs for modeling. In Sec. III, we detail the proposed algorithm, VF-MC-PILCO. Section IV illustrates the validation conducted on the simulated cart-pole benchmark. Section V reports the experiment performed on a simulated UR5 robot to test VF-MC-PILCO capacity of handling systems up to 6 DoF. Section VI shows the results

¹ Fabio Amadio, Alberto Dalla Libera, and Ruggero Carli are with the Department of Information Engineering, University of Padova, Via Gradenigo 6/B, 35131 Padova, Italy [fabio.amadio@phd.unipd.it, dallalibera@dei.unipd.it, carlirug@dei.unipd.it].

³ Diego Romeres and Daniel Nikovski are with Mitsubishi Electric Research Laboratories (MERL), Cambridge, MA 02139 [romeres@merl.com, nikovski@merl.com].

of the experiments with the two real mechanical systems. Finally, we draw conclusions in Section VII.

II. BACKGROUND

In this section, we first introduce the problem of MBRL on real mechanical systems. Then, we briefly discuss how Gaussian Process Regression (GPR) is usually used for modeling purposes.

A. Problem Formulation

Consider a mechanical system with d_q degrees of freedom, and denote by \mathbf{x}_t its state at time t . Typically, \mathbf{x}_t is defined as $\mathbf{x}_t = [\mathbf{q}_t^T, \dot{\mathbf{q}}_t^T]^T$, where $\mathbf{q}_t \in \mathbb{R}^{d_q}$ and $\dot{\mathbf{q}}_t \in \mathbb{R}^{d_q}$ are, respectively, the vector of joint positions and velocities. Assume that only \mathbf{q}_t can be directly measured, whereas $\dot{\mathbf{q}}_t$ is not directly measurable, but instead must be estimated. We argue that this is a common scenario as mechanical systems are often equipped with position sensors such as encoders, but lack velocity sensors. More accurate velocity estimates can be obtained from the history of position measurements, exploiting both past as well as future samples. These kinds of estimation techniques are intrinsically acausal, hence they can be performed offline and used only for modeling. Thus, controllers must rely on causal online estimates, which are usually less accurate and affected by delays.

The discrete-time dynamics of the system is given by $\mathbf{x}_{t+1} = f(\mathbf{x}_t, \mathbf{u}_t) + \mathbf{w}_t$, where $f(\cdot)$ is an unknown transition function, $\mathbf{u}_t \in \mathbb{R}^{d_u}$ represents the control action, and $\mathbf{w}_t \sim \mathcal{N}(0, \Sigma_w)$ models uncertainty. RL algorithms aim to learn how to accomplish a given task based on interaction data. The task is encoded by a cost function $c(\mathbf{x}_t)$, defined to characterize the immediate penalty associated with being in state \mathbf{x}_t . Control actions are chosen from a policy $\mathbf{u} = \pi_\theta(\mathbf{x})$, parameterized by θ . Then, the objective is to find the policy that minimizes the expected cumulative cost over T time steps, starting from the initial state distribution $p(\mathbf{x}_0)$, i.e.,

$$J(\theta) = \sum_{t=0}^T \mathbb{E}_{\mathbf{x}_t} [c(\mathbf{x}_t)]. \quad (1)$$

An MBRL algorithm consists, generally, of the succession of several trials, i.e. attempts to solve the desired task, and each trial is structured in three main phases:

- *Model Learning*: data collected during previous interactions are used to learn a model of the system dynamics. At the beginning, first data are collected by applying an exploratory policy, e.g. random control actions;
- *Policy Update*: the control policy is optimized in order to minimize an estimate of the cost $J(\theta)$ obtained by exploiting the trained model;
- *Policy Execution*: the updated control policy is applied to the system and interaction data are stored.

In order to comply with the common conditions in real mechanical systems described above, we propose an MBRL algorithm to control mechanical systems without assuming to have neither measurements nor estimations of velocities.

B. GPR for Model Learning

Given a data set of state-action pairs measured during the interactions with the system, it is possible to use GPR to train a probabilistic model that approximates the unknown transition function $f(\cdot)$. A common strategy in the literature [14], [15] is to model the evolution of each state dimension with a distinct zero mean GP. Let us indicate with $x^{(i)}$ the i -th component of the state, for $i \in \{1, \dots, d_x\}$ (d_x is the dimension of the state vector), and define $\tilde{\mathbf{x}}_t = [\mathbf{x}_t^T, \mathbf{u}_t^T]^T$. The i -th GP takes $\tilde{\mathbf{x}}_t$ as input, and predicts $x_{t+1}^{(i)} - x_t^{(i)}$. The GPs are completely characterized by their kernel functions that represent our belief on the a priori covariance. A common choice is the Squared Exponential (SE) kernel,

$$k_{SE}(\tilde{\mathbf{x}}_{t_j}, \tilde{\mathbf{x}}_{t_k}) := \lambda^2 e^{-\|\tilde{\mathbf{x}}_{t_j} - \tilde{\mathbf{x}}_{t_k}\|_{\Lambda}^2}. \quad (2)$$

Given a data set of state-action pairs \mathcal{D} , the GPs provide a closed form expression of $p(\mathbf{x}_{t+1} | \tilde{\mathbf{x}}_t, \mathcal{D})$, the posterior distribution of the estimated state at time $t+1$. For further details about GPR and its application to dynamical system modeling, the readers can refer to [4].

III. VELOCITY-FREE MC-PILCO

Here we present the algorithm VF-MC-PILCO (*Velocity Free Monte Carlo Probabilistic Inference for Learning Control*), whose objective is to solve the problem defined in Sec. II-A without the need of performing any kind of velocity estimation. In fact, tuning effective estimators may be a tedious and complex operation, especially in the presence of high measurement noise. This might significantly compromise the MBRL algorithm, if not duly considered. VF-MC-PILCO circumvents these issues by adopting a VF formulation. Inspired by [12], [13], we consider a VF model of the system dynamics, given in the following general form

$$\mathbf{q}_{t+1} = f_{\text{dt}}(\mathbf{q}_t, \mathbf{q}_{t-1}, \dots, \mathbf{q}_{t-m_q}, \mathbf{u}_t, \dots, \mathbf{u}_{t-m_u}). \quad (3)$$

The joint positions at the next time step are predicted based on the history of the past positions, from t up to $t-m_q$, and the history of applied control actions, from t up to $t-m_u$. Let m_q and m_u be called, respectively, the position memory and the control memory of the VF model. In this new VF framework, it is convenient to redefine the state of the system as $\mathbf{x}_t = [\mathbf{q}_t^T, \dots, \mathbf{q}_{t-m_q}^T, \mathbf{u}_{t-1}^T, \dots, \mathbf{u}_{t-m_u}^T]^T$.

In the following, we present the model learning and policy update phases of the VF-MC-PILCO algorithm, detailing how they have been adapted to the new VF formulation.

A. VF Model Learning

We employ the GPR framework of Sec. II-B, but instead of considering a full state representation with velocities, we train a VF GP model of form (3). Let us denote with $q_t^{(i)}$ the position of the i -th joint at time t , and define $\Delta_{q_t}^{(i)} = q_{t+1}^{(i)} - q_t^{(i)}$, for $i \in \{1, \dots, d_q\}$. The evolution of $\Delta_{q_t}^{(i)}$ for all i is modeled using a distinct GP, whose input depends upon $[\mathbf{q}_t^T, \dots, \mathbf{q}_{t-m_q}^T, \mathbf{u}_t^T, \mathbf{u}_{t-1}^T, \dots, \mathbf{u}_{t-m_u}^T]^T$. Trivially, the transition functions of $\mathbf{q}_{t-1}, \dots, \mathbf{q}_{t-m_q}, \mathbf{u}_{t-1}, \dots, \mathbf{u}_{t-m_u}$ are deterministic and known.

Experimentally, we found it beneficial in terms of data efficiency (details in Sec. IV-C) to rearrange GP input as

$$\tilde{\mathbf{x}}_t = \left[\mathbf{q}_t^T, \Delta_{\mathbf{q}_{t-1}}^T, \dots, \Delta_{\mathbf{q}_{t-m_q}}^T, \mathbf{u}_t^T, \dots, \mathbf{u}_{t-m_u}^T \right]^T, \quad (4)$$

where $\Delta_{\mathbf{q}_{t-i}} = (\mathbf{q}_{t-i+1} - \mathbf{q}_{t-i})$, for $i = 1, \dots, m_q$, following the same notation used before when defining the GP targets.

In this way, we are providing the model with additional information about the rates of change observed inside the considered past position memory interval m_q . Depending on the considered application, it may be convenient to further modify the GP input vector w.r.t. (4), in order to exploit particular characteristics of the considered quantities. For instance, we applied a *sin-cos* expansion to angular quantities during some of the experiments presented in the next sections.

B. VF Particle-Based Policy Gradient

The GP-based VF predictive model of Sec. III-A is now employed to optimize the policy parameters θ following a particle-based policy gradient strategy. VF-MC-PILCO computes $\hat{J}(\theta)$, an approximation of $J(\theta)$ in (1) exploiting the posterior distribution $p(\mathbf{x}_{t+1}|\tilde{\mathbf{x}}_t, \mathcal{D})$ defined by the GPs. Finally, it updates the θ with a gradient-based procedure.

The computation of $\hat{J}(\theta)$ entails the simulation of the effects of the policy π_θ on M independent state particles by cascading the one-step-ahead stochastic predictions. In particular, let $\mathbf{q}_t^{(m)}$, for $m = 1, \dots, M$, represent the position of the M state particles simulated by the VF GP model. Starting positions are sampled from a given distribution $p(\mathbf{q}_0)$. We assume that the system is not moving at $t = 0$, i.e., $\mathbf{q}_0^{(m)} = \mathbf{q}_{-1}^{(m)} = \dots = \mathbf{q}_{-m_q}^{(m)}$. At each time step t , in order to simulate the presence of measurement error, we corrupt the particle positions $\mathbf{q}_t^{(m)}$ with a fictitious noise $e_t^{(m)}$, e.g. a zero mean Gaussian i.i.d. noise, obtaining a set of simulated measurements $\bar{\mathbf{q}}_t^{(m)} = \mathbf{q}_t^{(m)} + e_t^{(m)}$. Then, for each particle, the policy π_θ selects the next control actions $\mathbf{u}_t^{(m)}$ according to the history of the simulated measurements, $\tilde{\mathbf{x}}_t^{(m)} = [\bar{\mathbf{q}}_t^{(m)T}, \dots, \bar{\mathbf{q}}_{t-m_q}^{(m)T}]^T$. Finally, the M positions at the next time step, $t + 1$, are simulated by forward sampling from the distributions derived by the VF GP model $p(\mathbf{x}_{t+1}^{(m)}|\tilde{\mathbf{x}}_t^{(m)}, \mathcal{D})$ (for $m = 1 \dots M$) with $\tilde{\mathbf{x}}_t^{(m)}$ defined for each particle m as in (4). This procedure is iterated for T time steps, obtaining M different particle trajectories $\{\{\mathbf{x}_t^{(m)}\}_{m=1}^M\}_{t=0}^T$, that simulate the results of the policy. The particle generation procedure is depicted in the block scheme of Fig. 1. The sample mean of the costs incurred by the different particles provides an estimate of the expected cumulative cost, namely

$$\hat{J}(\theta) = \sum_{t=0}^T \left(\frac{1}{M} \sum_{m=1}^M c(\mathbf{x}_t^{(m)}) \right). \quad (5)$$

The computational graph resulting from (5) allows us to compute $\nabla_\theta \hat{J}(\theta)$, i.e., the gradient of $\hat{J}(\theta)$ w.r.t. θ , through backpropagation, exploiting the reparametrization trick [16], [17] to propagate the gradient through the stochastic operations. Finally, a stochastic gradient descent algorithm, e.g. Adam [18], can exploit the estimated gradient to update θ .

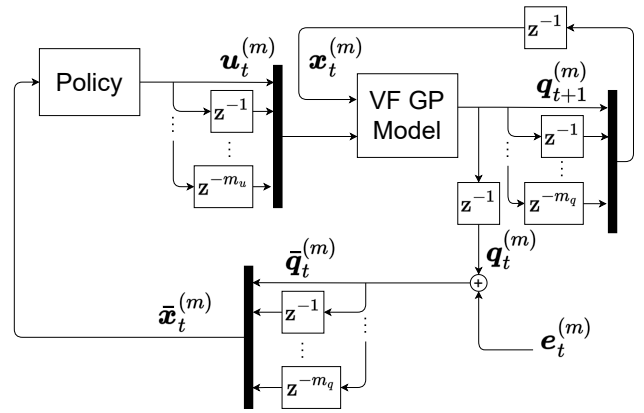


Fig. 1: VF-MC-PILCO particles generation block schemes.

C. Policy structure

We considered an RBF network policy with outputs limited by a hyperbolic tangent function, properly scaled. We call this function *squashed-RBF-network*, and it is expressed as

$$\pi_\theta(\mathbf{x}^*) = u_{max} \tanh \left(\frac{1}{u_{max}} \sum_{i=1}^{n_b} w_i e^{-\|\mathbf{a}_i - \mathbf{x}^*\|_{\Sigma_\pi}^2} \right), \quad (6)$$

The input vector of the policy is defined as

$$\mathbf{x}_t^* = \left[\mathbf{q}_t^T, \Delta_{\mathbf{q}_{t-1}}^T, \dots, \Delta_{\mathbf{q}_{t-m_q}}^T \right]^T, \quad (7)$$

where we are providing the policy with the same consecutive differences of position measures used for GP input in (4).

The policy parameters are $\theta = \{\mathbf{w}, A, \Sigma_\pi\}$, where $\mathbf{w} = [w_1 \dots w_{n_b}]$ and $A = \{\mathbf{a}_1 \dots \mathbf{a}_{n_b}\}$ are, respectively, the weights and the centers of the n_b basis functions, while Σ_π is a diagonal matrix that determines their shapes. The maximum control u_{max} is constant and depends on the application. It is worth mentioning that VF-MC-PILCO is not restricted to this particular choice of policy function.

IV. SIMULATED EXPERIMENT: CART-POLE SWING-UP

As a preliminary validation, we tested VF-MC-PILCO on a simulated cart-pole swing-up task to analyze its performance under different setups. We compare the proposed approach with the s.o.t.a. MBRL algorithm specifically designed to deal with partial state measurability of real mechanical systems, MC-PILCO4PMS [11]. MC-PILCO4PMS follows a particle-based policy gradient framework similar to the one depicted in Sec. III-B, but, differently from the proposed VF-MC-PILCO, it works with velocity estimates by simulating not only the evolution of the system state but also the evolution of the estimated state, which entails modeling the measurement system and the implemented online filters. Notice that the implementation of MC-PILCO4PMS could be in some cases complex or time-consuming, due to its requirement to reproduce the online filtering procedure inside the policy optimization phase, and the need to adopt a different offline filter for model learning. This is the limitation that the proposed method aims to solve. Both algorithms have been implemented in Python¹, exploiting the PyTorch library [19].

¹<https://www.merl.com/research/license/MC-PILCO>

Now, let us briefly describe the characteristics of the simulated scenario. Let p_t and α_t be, respectively, the position of the cart and the angle of the pole at time step t , hence $\mathbf{q}_t = [p_t, \alpha_t]^T$. The target configurations corresponding to the pendulum swing-up are given by $p^{des} = 0$ [m] and $|\alpha^{des}| = \pi$ [rad]. The cart-pole starts from $\theta_0 = 0$ [rad] and $p_0 = 0$ [m]. The control action is the force applied to the cart, and the system is controlled at 30 [Hz]. We considered a Gaussian measurement noise with standard deviation of 10^{-3} [m] for positions and $2 \cdot 10^{-3}$ [rad] for angles.

The GPs of the VF model are equipped with the SE kernel described in (2). The policy adopted is a *squashed-RBF-network* policy with $n_b = 200$ basis functions and $u_{max} = 10$ [N]. The number of particles is set to $M = 400$ during policy optimization. In order to avoid singularities due to the angles, we replaced, in both the model inputs $\hat{\mathbf{x}}_t$ defined in (4) and policy input \mathbf{x}^* defined in (7), occurrences of α_t with $\sin(\alpha_t)$ and $\cos(\alpha_t)$. Exploration data were collected by random actions, obtained by filtering Gaussian white noise with cut-off frequency 1.5 [Hz]. The cost function is

$$c(\mathbf{x}_t) = 1 - \exp\left(-\left(\frac{|\alpha_t| - \pi}{l_\alpha}\right)^2 - \left(\frac{p_t}{l_p}\right)^2\right), \quad (8)$$

where $l_\alpha = 3$ and $l_p = 1$ define the shape of $c(\cdot)$. The absolute value on α_t is needed to allow different swing-up solutions to both the equivalent target pole angles π and $-\pi$.

The objective is to analyze different VF-MC-PILCO configurations and compare their performance with the results obtained by MC-PILCO4PMS, as a benchmark. We analyzed the results obtained in 50 distinct experiments, consisting of 5 trials of length 3 seconds, varying the random seed each time. In this way, it is possible to evaluate the robustness of the algorithm to different exploration trajectories and policy initialization, as well as to different noise realizations. In particular, we investigate the effects that different position and control memories, m_q and m_u , have on modeling and policy learning. We studied four different VF-MC-PILCO configurations, choosing the value of m_q between 1 and 2, and m_u between 0 and 1. In the following, we will refer to these different alternatives with the symbol $VF_{m_q}^{m_u}$.

A. Modeling results

We compared the accuracy of the different VF GP models by looking at the absolute values of the prediction errors observed on the data registered at the last trial in all 50 experiments. Models were trained using all the data collected up to that trial. Fig. 2 reports the results by means of box plots, showing median values, confidence intervals, and outliers. The results show that the presence of input history, \mathbf{u}_{t-1} , as part of the GP inputs is beneficial, and within this choice, see models with $m_u = 1$, the best results are obtained by $VF_{m_q=2}^{m_u=1}$. On the other hand, the greater errors and the significant number of outliers obtained by models with $m_u = 0$ seem to indicate that these kinds of setups are not fully capable of fitting the registered position changes. Also, it appears that using a longer position memory leads to an improvement of prediction accuracy only when the control memory is $m_u = 1$. We can




Fig. 2: Absolute p and α prediction errors obtained by different VF GP models at trial 5 in the simulated cart-pole experiments.

conclude that it seems beneficial to provide VF GP models with information about past control actions ($m_u = 1$) for fitting the system dynamics without relying on velocity.

B. Policy learning results

In this section, we evaluate the performance of the control policies learned by the different VF-MC-PILCO setups and by MC-PILCO4PMS. Notice that MC-PILCO4PMS achieved results comparable to or better than other state-of-the-art GP-based MBRL algorithms, see [11]. The cumulative costs and success rates obtained at each trial in the 50 experiments are reported in Fig. 3. In the two plots, the cumulative cost is reported in terms of median values and confidence intervals defined by the 5-th and 95-th percentiles. As one would expect, the worse modeling results of $VF_{m_q=1}^{m_u=0}$ and $VF_{m_q=2}^{m_u=0}$ lead to an unsatisfactory policy learning. These VF-MC-PILCO setups manage to complete a successful swing-up only in, approximately, half of the cases. On the other hand, when using $m_u = 1$, VF-MC-PILCO is able to robustly find an optimal solution for the task by trial 5. In particular, the performance of $VF_{m_q=2}^{m_u=1}$ are almost equivalent to the results of MC-PILCO4PMS. This result confirms the effectiveness of the proposed method: with less information, as we are not manually tuning any velocity estimator, VF-MC-PILCO achieves state-of-the-art performance. For the user, this corresponds to less effort and a more general method without compromising significantly the performance.

C. Analysis of input vector structure

Before concluding this section, we would like to analyze the reasons behind the decision to use (4) and (7) as GP and policy input vectors, respectively. In this respect, we compared the results obtained by VF-MC-PILCO with position memory $m_q = 2$ and control memory $m_u = 1$ using two different structures for the input vectors. The first employs directly the history of positions and actions up to time step t as GP input, e.g., $\hat{\mathbf{x}}_t = [\mathbf{q}_t^T, \dots, \mathbf{q}_{t-m_q}^T, \mathbf{u}_t^T, \dots, \mathbf{u}_{t-m_u}^T]^T$, and the history of positions as policy input, e.g., $\mathbf{x}_t^* = [\mathbf{q}_t^T, \dots, \mathbf{q}_{t-m_q}^T]^T$. The second version is the one employed




Fig. 3: Cumulative costs registered during simulated cart-pole experiments from the four considered VF-MC-PILCO setups and the MC-PILCO4PMS benchmark (indicated by the shorthand PMS). The observed success rates are presented in the table below.

	Trial 1	Trial 2	Trial 3	Trial 4	Trial 5
$VF_{m_q=1}^{m_u=0}$	0%	14%	34%	46%	56%
$VF_{m_q=2}^{m_u=0}$	0%	10%	18%	28%	52%
$VF_{m_q=1}^{m_u=1}$	0%	8%	52%	86%	96%
$VF_{m_q=2}^{m_u=1}$	0%	20%	73%	93%	100%
PMS	0%	14%	82%	98%	96%

previously with GP input and policy input defined as in (4) and (7), respectively. To distinct the two implementations, we labeled the first as $VF_{m_q=2}^{m_u=1}$ naive, and the second $VF_{m_q=2}^{m_u=1}$ with position differences. We analyzed the results obtained by these two setups in 50 distinct experiments, consisting of 7 trials of length 3 seconds, varying the random seed each time. The obtained cumulative costs are reported in Fig. 4 in terms of median values and 5-95 percentile ranges.

It is clear that providing information about the rate of change of position measures, by using (4) and (7) as input vectors, greatly improves the data efficiency of VF-MC-PILCO algorithm. In fact, the $VF_{m_q=2}^{m_u=1}$ naive implementation (that uses directly the history of position and controls) shows a much slower convergence, reaching a 79% success rate only at trial 7. On the other hand, $VF_{m_q=2}^{m_u=1}$ with position differences is able to always find a solution by trial 5.

This result underlines the importance of the information carried out by the differences between consecutive measured positions. Without this information, the model needs more data to correctly capture the dynamics of the system relying only on positions. Through input vectors (4) and (7), we are able to provide the model with knowledge about a sort of velocity, without requiring any kind of filtering procedure.

V. SIMULATED EXPERIMENT: UR5 ROBOT CONTROL

The objective of this experiment is to test VF-MC-PILCO in a more complex system with a higher DoF. We used VF-MC-PILCO to learn a joint-space torque controller for a UR5, a robotic manipulator with 6 DoF, simulated in




Fig. 4: Cumulative costs registered by VF-MC-PILCO with different GP and policy input structures that include past positions either as consecutive differences ($\Delta_{q_{t-1}}, \dots, \Delta_{q_{t-m_q}}$) or directly ($q_{t-1}, \dots, q_{t-m_q}$).

MuJoCo [20], assuming to measure only joint angles and not velocities. Measurements are perturbed by the presence of white Gaussian noise with a standard deviation of 10^{-3} . Let us denote with $q_t \in \mathbb{R}^6$ the joint angles and with $u_t \in \mathbb{R}^6$ the applied torques. Our objective is to learn a VF control policy able to follow a desired trajectory $\{q_t^r\}_{t=1}^T$. Let $e_t = q_t^r - q_t$ denote the position error at time t . VF-MC-PILCO memories were set to $m_q = 2$ and $m_u = 1$, hence the VF state of the system at time step t is defined as $x_t = [q_t^T, q_{t-1}^T, q_{t-2}^T, u_{t-1}^T]^T$. The GP input vector was defined applying a *sin-cos* expansion of angular quantities as $\tilde{x}_t = [\sin(q_t)^T, \cos(q_t)^T, \Delta_{q_{t-1}}^T, \Delta_{q_{t-2}}^T, u_t^T, u_{t-1}^T]^T$.

The policy adopted was a multi-output *squashed-RBF-network* with $n_b = 400$ basis functions and $u_{max} = 1$ [N·m] for all the joints. $M = 200$ particles were used during optimization. The policy takes in input the vector $x_t^* = [\sin(q_t)^T, \cos(q_t)^T, \Delta_{q_{t-1}}^T, \Delta_{q_{t-2}}^T, e_t^T]^T$. Fig. 5 represents the overall control scheme.




Fig. 5: VF-MC-PILCO control scheme for the simulated UR5.

In this experiment, we considered a control horizon of 4 seconds with a sampling time of 0.02 seconds. The reference trajectory has been calculated to make the end-effector draw a circle in the X-Y operational space. The initial exploration used to initialize the VF GP model is provided by a poorly-tuned PD controller (for which we estimated velocity by backward differentiation). We used $M = 200$ of particles for gradient estimation and considered the following cost,

$$c(x_t) = 1 - \exp(-\|q_t^r - q_t\|^2).$$




Fig. 6: Average end-effector position tracking errors obtained during the exploratory phases and by the control policy learned at trial 3.




Fig. 7: Example of explorative and final end-effector trajectories.

The experiment was repeated 10 different times, varying the random seed and the initial exploration trajectories, obtained each time by using random PD gains, uniformly sampled from $K_P \sim \mathcal{U}(0.5, 2)$ and $K_D \sim \mathcal{U}(0.01, 0.2)$. VF-MC-PILCO managed to learn an effective control policy by the third trial in all the repetitions, with average positioning errors not superior to 2 [mm]. The average end-effector tracking errors obtained are reported in Fig. 6, where results are given by means of box plots. Fig. 7 shows an example of exploratory and final trajectories, taken from one of the conducted tests.

VI. EXPERIMENTS ON REAL MECHANICAL SYSTEMS

In this section, we report the results obtained by VF-MC-PILCO when applied to real systems. In particular, we experimented on two benchmark systems: a Furuta pendulum, and a ball-and-plate (Figure 8)². The objective is to compare the performance obtained by VF-MC-PILCO in these two setups with the results of MC-PILCO4PMS reported in [11].

²A video of the experiments on real mechanical systems is available at the following link <https://youtu.be/Hx3Y11B-6Tc>.



Fig. 8: (Left) Furuta pendulum. (Right) Ball-and-plate system.

A. Furuta pendulum

The Furuta pendulum [21] is a popular nonlinear control benchmark composed of two revolute joints and three links (see Fig. 8, left). It is an under-actuated system as only the horizontal joint is actuated by a DC servomotor. The two angles are measured by optical encoders with 4096 [ppr]. The control action is the motor voltage, and its maximum allowed value is 10 [V]. Let the pose at time step t be $\mathbf{q}_t = [\alpha_t^h, \alpha_t^v]^T$, where α_t^h is the angle of the horizontal joint and α_t^v the angle of the vertical joint attached to the pendulum. The objective is to learn how to swing-up the pendulum and stabilize it in the upward equilibrium ($\alpha_t^v = \pm\pi$ [rad]) with $\alpha_t^h = 0$ [rad], starting from $\mathbf{q}_0 = [0, 0]^T$. The trial length is 3 [s] and the system is controlled at 30 [Hz]. The cost is defined as

$$c(\mathbf{x}_t) = 1 - \exp\left(-\left(\frac{\alpha_t^h}{2}\right)^2 - \left(\frac{|\alpha_t^v| - \pi}{2}\right)^2\right) + c_b(\alpha_t^h), \quad (9)$$

with

$$c_b(\alpha_t^h) = \frac{1}{1 + \exp\left(-10\left(-\frac{3}{4}\pi - \alpha_t^h\right)\right)} + \frac{1}{1 + \exp\left(-10\left(\alpha_t^h - \frac{3}{4}\pi\right)\right)}.$$

The first part of the function in (9) aims at driving the two angles towards the target, while $c_b(\alpha_t^h)$ penalizes solutions that push the horizontal joint beyond a certain safety threshold.

In this scenario, we used position memory $m_q = 2$ and control memory $m_u = 1$. We equipped the VF GP model with an SE kernel and adopted a *squashed-RBF-network* with $n_b = 200$ basis functions as control policy. $M = 400$ particles were simulated during policy optimization. We replaced, in both GP inputs $\tilde{\mathbf{x}}_t$ and policy input \mathbf{x}^* , occurrences of α_t^h and α_t^v with their *sin-cos* expansion, as previously done in the simulated cart-pole case. The exploration trajectory has been obtained using as input a sum of ten cosine waves of random frequencies and the same amplitudes. The presence of quantization errors was simulated during particles generation by corrupting predicted angles with a uniform fictitious measurement noise $\mathcal{U}\left(\frac{-\pi}{4096}, \frac{\pi}{4096}\right)$ [rad].




Fig. 9: Real swing-up trajectory (bullets) and particles prediction (shaded lines) obtained by VF-MC-PILCO at trial 6 of the Furuta pendulum experiment.

VF-MC-PILCO learned how to swing-up the Furuta pendulum at trial 6, i.e. after 18 seconds of experience. That is the same result obtained by MC-PILCO4PMS when using the SE kernel. Hence, the VF approach showed no particular differences in terms of data efficiency when compared with an approach that makes use of velocity estimates. In Fig. 9, we report the successful swing-up performed by VF-MC-PILCO at trial 6, together with the particles predicted by the VF GP model, simulating the effects of the same control policy. Notice how the particles’ trajectories resemble almost perfectly the real behaviour of the two angles.

B. Ball-and-plate

The ball-and-plate system is composed of a square plate that can be tilted in two orthogonal directions, and a ball that is free to roll over it (see Fig. 8, right). A camera is placed on top of the system to track the ball and measure its position on the plate, with a precision of one millimeter. Let, at time t , (b_t^x, b_t^y) be the position of the center of the ball, while $\alpha_t^{(1)}$ and $\alpha_t^{(2)}$ are the angles of the two motors tilting the plate. Thus, $\mathbf{q}_t = [b_t^x, b_t^y, \alpha_t^{(1)}, \alpha_t^{(2)}]^T$. The drivers of the motors allow only position control and do not provide feedback about the motors’ angles. To keep track of them, we defined the control actions as the difference between two consecutive reference values sent to the motor, and we limited the maximum input to a sufficiently small value, i.e. 4 [deg], such that the motor controllers are able to reach the target within the sampling time. Then, as a first approximation, the reference angles, and the actual motor angles coincide, and we have $u_t^{(1)} = \alpha_{t+1}^{(1)} - \alpha_t^{(1)}$ and $u_t^{(2)} = \alpha_{t+1}^{(2)} - \alpha_t^{(2)}$. The objective of the experiment is to learn how to control the motor angles in order to stabilize the ball around the center of the plate. The trial length is 3 seconds, with a sampling frequency of 30 [Hz]. The cost function encoding the task is

$$c(\mathbf{x}_t) = 1 - \exp(-g(\mathbf{x}_t)), \quad \text{with}$$

$$g(\mathbf{x}_t) = \left(\frac{b_t^x}{0.15}\right)^2 + \left(\frac{b_t^y}{0.15}\right)^2 + \left(\alpha_t^{(1)}\right)^2 + \left(\alpha_t^{(2)}\right)^2.$$




Fig. 10: Example of GP targets in the ball-and-plate experiment, i.e. measured ball position changes in X and Y directions.

With regards to the VF model setup, we considered position memory $m_q = 2$ and control memory $m_u = 1$, and we replaced in both GP inputs $\tilde{\mathbf{x}}_t$ and policy input \mathbf{x}^* , the occurrences of $\alpha_t^{(1)}$ and $\alpha_t^{(2)}$ with their *sin-cos* expansion. Analogously to the previous MC-PILCO4PMS experiment, the kernel function of the VF GP model is given by the sum of a SE kernel that takes as input the whole GP input vector, and of a linear kernel that takes as input only the *sin-cos* expansion of angular quantities. The control policy is a *squashed-RBF-network* with $n_b = 400$ basis functions. Policy optimization involves the use of $M = 400$ particles.

The initial exploration is implemented in two trials, in which the control signals are two distinct noisy triangular waves. Mostly during exploration and initial trials, the ball might touch the borders of the plate. In those cases, we kept data up to the collision instant and discarded it thereafter. The presence of quantization errors was simulated during particles generation by corrupting predicted angles with a uniform fictitious measurement noise $\mathcal{U}(-\frac{0.001}{2}, \frac{0.001}{2})$ [m]. A peculiarity of this experiment in comparison to the others seen before is a wide range of initial conditions. In fact, the policy must learn how to control the ball to the center starting from any position on the plate’s surface. Hence, the initial distribution considered for b_0^x and b_0^y is the uniform $\mathcal{U}(-0.15, 0.15)$ [m]. The measurements provided by the camera are affected by a significant quantization error. For instance, Fig. 10 reports the measured differences between consecutive ball positions during a trial. Consider that these quantities are the targets of the GPs in the VF model. In such a context, ball velocity estimation can be very challenging. In fact, for applying MC-PILCO4PMS on the same system, methods like finite differences and low-pass filtering were not sufficient, and it was necessary to implement a Kalman filter (online) and a Kalman smoother (offline), whose tuning was a delicate and time-consuming procedure of critical importance for the success of the algorithm. On the contrary, VF-MC-PILCO managed to solve the task by working directly with raw position measurements, without the need of applying any kind of filtering. Besides that, VF-MC-PILCO proved to be surprisingly data-efficient, being able to solve the task at the




Fig. 11: Ten different ball trajectories obtained by VF-MC-PILCO policy. Steady-state positions are marked with black crosses. The dashed circle has the same diameter as that of the ball.

second trial, after only 7.97 seconds of experience, whereas MC-PILCO4PMS solved the task after 11.33 seconds.

We tested the policy starting from ten different points in order to compare the two policies obtained by VF-MC-PILCO (Fig. 11) and MC-PILCO4PMS. The mean steady-state error, i.e. the average distance of the last ball position from the center observed in the ten tests, was 0.0134 [m], while MC-PILCO4PMS final policy obtained a slightly better result, with a mean error of 0.0099 [m]. This might be due to the difference between the two policy inputs: MC-PILCO4PMS relies on a Kalman filter, while VF-MC-PILCO works directly with raw measurements (in presence of significant noise). Nevertheless, this performance difference is quite negligible, given the dimension of the ball whose radius is 0.016 [m].

VII. CONCLUSIONS

We presented VF-MC-PILCO, a novel MBRL algorithm, specifically designed to learn from scratch how to control mechanical systems, without the need of computing any explicit velocity estimates. In our opinion, this may be a critical advantage when dealing with real systems affected by significant measurement noise, since, in this kind of scenario, the design of accurate velocity estimators can be a tedious task. The algorithm uses GPR to model the joint position changes, based on the history of past control actions and measurements. VF-MC-PILCO was tested both in simulated environments (cart-pole and UR5 robot) as well as in two real mechanical systems (Furuta pendulum and ball-and-plate rig). It proved able to solve all the tasks, with results that are in line with the performance of our previous MBRL algorithm (MC-PILCO4PMS), which instead works with a complete state representation and must perform velocity estimation.

REFERENCES

[1] Athanasios S Polydoros and Lazaros Nalpantidis. Survey of model-based reinforcement learning: Applications on robotics. *Journal of Intelligent & Robotic Systems*, 86(2):153–173, 2017.

[2] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.

[3] Christopher G Atkeson and Juan Carlos Santamaria. A comparison of direct and model-based reinforcement learning. In *Proceedings of international conference on robotics and automation*, volume 4, pages 3557–3564. IEEE, 1997.

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

[5] Marc Deisenroth and Carl E Rasmussen. Pilco: A model-based and data-efficient approach to policy search. In *Proceedings of the 28th International Conference on machine learning (ICML-11)*, pages 465–472, 2011.

[6] Marc Peter Deisenroth, Carl Edward Rasmussen, and Dieter Fox. Learning to control a low-cost manipulator using data-efficient reinforcement learning. *Robotics: Science and Systems VII*, pages 57–64, 2011.

[7] Felix Berkenkamp, Matteo Turchetta, Angela Schoellig, and Andreas Krause. Safe model-based reinforcement learning with stability guarantees. In *Advances in neural information processing systems*, pages 908–918, 2017.

[8] Konstantinos Chatzilygeroudis, Roberto Rama, Rituraj Kaushik, Dorian Goepp, Vassilis Vassiliades, and Jean-Baptiste Mouret. Black-box data-efficient policy search for robotics. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 51–58. IEEE, 2017.

[9] Kurtland Chua, Roberto Calandra, Rowan McAllister, and Sergey Levine. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. In *Advances in Neural Information Processing Systems*, pages 4754–4765, 2018.

[10] Thanard Kurutach, Ignasi Clavera, Yan Duan, Aviv Tamar, and Pieter Abbeel. Model-ensemble trust-region policy optimization. *arXiv preprint arXiv:1802.10592*, 2018.

[11] 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.

[12] D. Romeres, M. Zorzi, R. Camoriano, S. Traversaro, and A. Chiuso. Derivative-free online learning of inverse dynamics models. *IEEE Transactions on Control Systems Technology*, 28(3):816–830, 2020.

[13] A. Dalla Libera, D. Romeres, D. K. Jha, B. Yerazunis, and D. Nikovski. Model-based reinforcement learning for physical systems without velocity and acceleration measurements. *IEEE Robotics and Automation Letters*, 5(2):3548–3555, 2020.

[14] Marc Peter Deisenroth, Dieter Fox, and Carl Edward Rasmussen. Gaussian processes for data-efficient learning in robotics and control. *IEEE transactions on pattern analysis and machine intelligence*, 37(2):408–423, 2013.

[15] Diego Romeres, Devsh K Jha, Alberto DallaLibera, Bill Yerazunis, and Daniel Nikovski. Semiparametrical gaussian processes learning of forward dynamical models for navigating in a circular maze. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 3195–3202. IEEE, 2019.

[16] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

[17] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *International conference on machine learning*, pages 1278–1286. PMLR, 2014.

[18] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[19] 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. *Advances in Neural Information Processing Systems 32*, pages 8024–8035, 2019.

[20] Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033. IEEE, 2012.

[21] Benjamin Seth Cazzolato and Zebb Prime. On the dynamics of the furuta pendulum. *Journal of Control Science and Engineering*, 2011, 2011.