

Structure-Preserving Learning Using Gaussian Processes and Variational Integrators

Jan Brüdigam

JAN.BRUEDIGAM@TUM.DE

Martin Schuck

MARTIN.SCHUCK@TUM.DE

Alexandre Capone

ALEXANDRE.CAPONE@TUM.DE

Stefan Sosnowski

SOSNOWSKI@TUM.DE

Sandra Hirche

HIRCHE@TUM.DE

*Chair of Information-oriented Control
Technical University of Munich, Germany*

Editors: R. Firoozi, N. Mehr, E. Yel, R. Antonova, J. Bohg, M. Schwager, M. Kochenderfer

Abstract

Gaussian process regression is increasingly applied for learning unknown dynamical systems. In particular, the implicit quantification of the uncertainty of the learned model makes it a promising approach for safety-critical applications. When using Gaussian process regression to learn unknown systems, a commonly considered approach consists of learning the residual dynamics after applying some generic discretization technique, which might however disregard properties of the underlying physical system. Variational integrators are a less common yet promising approach to discretization, as they retain physical properties of the underlying system, such as energy conservation and satisfaction of explicit kinematic constraints. In this work, we present a novel structure-preserving learning-based modelling approach that combines a variational integrator for the nominal dynamics of a mechanical system and learning residual dynamics with Gaussian process regression. We extend our approach to systems with known kinematic constraints and provide formal bounds on the prediction uncertainty. The simulative evaluation of the proposed method shows desirable energy conservation properties in accordance with general theoretical results and demonstrates exact constraint satisfaction for constrained dynamical systems.

Keywords: Gaussian process, regression, variational integrator, dynamical system, constraint

1. Introduction

Control of complex mechanical systems such as robots typically requires good knowledge of the system dynamics to achieve high control performance. In this regard, learning-based approaches are promising, as they can accurately model additional influences on top of the nominal dynamics, such as friction or external forces. One promising learning-based method for such scenarios is Gaussian process regression, which yields good generalization from only few training samples (Rasmussen and Williams, 2006). In addition, uncertainty bounds for the learned dynamics can be given (Lederer et al., 2019; Srinivas et al., 2012), rendering the regression useful for safety-critical control applications, as shown in, for example, Umlauft et al. (2017) and Capone and Hirche (2019).

While much effort has been devoted to the theoretical properties of GP regression, the appropriate discretization of continuous-time mechanical systems in combination with GPs for simulation

and control has found little attention. Marsden and West (2001) proposed variational (symplectic) integrators as generally suitable discretization methods for mechanical systems, as they retain certain structural properties, such as energy or momentum conservation, and constraint satisfaction. As a result, they achieve higher accuracy than generic integrators such as explicit Runge-Kutta methods, which allows for larger prediction time steps—an important feature in, for example, online control. However, no attempts have been made to exploit the advantageous properties of variational integrators in combination with Gaussian process regression.

The contribution of this paper is a novel approach for physical-structure-preserving and data-driven modelling of dynamical systems. To this end, we combine a variational integrator of the nominal model dynamics and Gaussian process regression to include initially unknown system dynamics. In this way, we retain physical properties of the known system dynamics and augment them by the learned residual dynamics. This concept is also extended to learning system dynamics with known explicit kinematic constraints, for which variational integrators are particularly well suited. In addition, formal uncertainty bounds for learning unconstrained and constrained dynamical systems in combination with variational integrators are provided. The energy conservation and constraint satisfaction properties of the proposed approach are illustrated in simulations. Furthermore, the performance of our method is numerically studied for several mechanical systems including such with kinematic constraints and in different coordinates.

The rest of the paper is structured as follows: in Sec. 2 key concepts of Gaussian process regression and variational integrators are summarized, and related work is listed. Section 3 includes a derivation of the integrators for unconstrained and constrained dynamical systems in combination with Gaussian process regression, which are then evaluated and discussed in Sec. 4. Finally, we present our conclusions in Sec. 5.

2. Background

We briefly revisit the relevant concepts of Gaussian process regression and describe a first-order variational integrator. A more thorough treatment of Gaussian processes and variational integrators can be found in Rasmussen and Williams (2006) and Marsden and West (2001), respectively. The section is concluded with a review of selected related literature.

2.1. Gaussian Process Regression

A Gaussian process (GP) is employed to model an unknown function $f(\cdot)$. A GP is a collection of random variables, of which any finite subset is jointly normally distributed. As such, GPs generalize the notion of Gaussian distributions to infinite-dimensional vectors. A GP, denoted $f(\cdot) \sim \mathcal{GP}(m(\cdot), k(\cdot, \cdot))$, is specified by its mean function $m(\cdot)$ and covariance function $k(\cdot, \cdot)$.

By conditioning a GP on measurement data pairs $(z_i, f(z_i))_{i=1\dots N} =: (\mathbf{Z}, \mathbf{y})$, with states $\mathbf{Z} = (z_i)_{i=1\dots N}$ and measurements $\mathbf{y} = (f(z_i))_{i=1\dots N}$, an unknown function $f(\cdot)$ can be modeled. The resulting posterior mean and variance are then given by

$$\mu(z) = \bar{m}(z) + \mathbf{k}(z, \mathbf{Z})\mathbf{K}(\mathbf{Z}, \mathbf{Z})^{-1}(\mathbf{y} - \bar{\mathbf{m}}(\mathbf{Z})), \quad (1a)$$

$$\sigma^2(z) = k(z, z) - \mathbf{k}(z, \mathbf{Z})\mathbf{K}(\mathbf{Z}, \mathbf{Z})^{-1}\mathbf{k}(z, \mathbf{Z})^T, \quad (1b)$$

where $\bar{m}(\cdot)$ is the prior mean function, $\mathbf{k}(z, \mathbf{Z}) = [k(z, z_1) \cdots k(z, z_N)]$, and the entries k_{ij} of the posterior covariance matrix \mathbf{K} are given by $k_{ij}(\mathbf{Z}, \mathbf{Z}) = k(z_i, z_j)$.

While there exist different methods for predicting multidimensional function values, in this paper we model each entry of a vector-valued function with a separate GP. This is a standard procedure in GP-based literature (Deisenroth et al., 2013). Therefore, in the following, a vector-valued GP $f(\cdot) \sim \mathcal{GP}(\mathbf{m}(\cdot), \mathbf{k}(\cdot, \cdot))$ corresponds to a vector of independent GPs.

2.2. Variational Integrators

Solving differential equations, for example the equations of motion of a dynamical system, requires numerical integration. Variational integrators, as described in Marsden and West (2001), are symplectic (“shape-preserving”) integrators for mechanical systems that are able to preserve certain physical properties of the underlying system, such as energy conservation or satisfaction of explicit constraints without drift.

Given a continuous-time mechanical system with Lagrangian $\mathcal{L}(\mathbf{x}(t), \mathbf{v}(t))$, where $\mathbf{x} \in \mathbb{R}^n$ and $\mathbf{v} \in \mathbb{R}^n$ are the generalized position and velocity of the system, respectively, as well as kinematic constraints $\mathbf{g}(\mathbf{x}(t)) = \mathbf{0} \in \mathbb{R}^c$, the action integral S of the system can be written as

$$S = \int_{t_0}^{t_N} \mathcal{L}(\mathbf{x}(t), \mathbf{v}(t)) dt + \int_{t_0}^{t_N} \boldsymbol{\lambda}(t)^\top \mathbf{g}(\mathbf{x}(t)) dt, \quad (2)$$

where $\boldsymbol{\lambda} \in \mathbb{R}^c$ is a Lagrange multiplier (constraint force) enforcing the constraints $\mathbf{g} = \mathbf{0}$.

According to the principle of least action, minimizing (2) by varying the trajectory $\mathbf{x}(t)$ yields the continuous-time differential equations of the system that could then be discretized. However, to obtain a variational integrator instead, (2) is discretized directly. For clarity, we will derive and use a first-order integrator throughout this paper, but higher-order discretizations are also possible.

A first-order variational integrator is obtained by discretizing (2) over three time steps:

$$S_d = \sum_{k=0}^1 (\mathcal{L}(\mathbf{x}_k, \mathbf{v}_k) + \boldsymbol{\lambda}_k^\top \mathbf{g}(\mathbf{x}_k)) \Delta t, \quad (3)$$

with time step Δt and where

$$\mathbf{v}_k = \frac{\mathbf{x}_{k+1} - \mathbf{x}_k}{\Delta t}. \quad (4)$$

Least action for fixed start and end points \mathbf{x}_0 and \mathbf{x}_2 , i.e., minimizing (3) with respect to the center point \mathbf{x}_1 , yields the implicit discretized equations of motion

$$\nabla_{\mathbf{x}_1} S_d =: -\mathbf{d} = \mathbf{0}. \quad (5)$$

The resulting implicit nonlinear equations of motion (variational integrator) take the form

$$\mathbf{x}_{k+1} = \mathbf{x}_k + \mathbf{v}_k \Delta t, \quad (6a)$$

$$\mathbf{d}(\mathbf{v}_{k+1}, \boldsymbol{\lambda}_{k+1}) = \mathbf{d}_0(\mathbf{v}_{k+1}) - \mathbf{G}(\mathbf{x}_{k+1})^\top \boldsymbol{\lambda}_{k+1} = \mathbf{0}, \quad (6b)$$

$$\mathbf{g}(\mathbf{x}_{k+2}) = \mathbf{g}(\mathbf{x}_{k+1}, \mathbf{v}_{k+1}) = \mathbf{0}, \quad (6c)$$

where \mathbf{d}_0 are the unconstrained dynamics, and the constraint Jacobian $\mathbf{G}(\mathbf{x}) = \frac{\partial \mathbf{g}(\mathbf{x})}{\partial \mathbf{x}}$ maps constraint forces into the dynamics. Note that the constraints are enforced for \mathbf{x}_{k+2} .

Besides the typical parameterization of mechanical systems in minimal coordinates (joint coordinates), so-called maximal coordinates can be used as well. In this case, each body of a system is

described with its six degrees of freedom and kinematic constraints represent joints connecting the bodies. The general idea for the derivation of the variational integrator remains the same, and details including the treatment of quaternion-based orientation representations are given in [Brüdigam and Manchester \(2020\)](#). Our method is also applicable to maximal-coordinate descriptions which have demonstrated numerical and control-theoretical advantages, as described in [Brüdigam and Manchester \(2021\)](#) and [Brüdigam et al. \(2021\)](#).

2.3. Related Work

Theoretical results for variational integrators are presented in [Marsden and West \(2001\)](#), and numerous extensions can be found in subsequent works, for example in [Junge et al. \(2005\)](#) and [Wenger et al. \(2017\)](#). Additionally, there exist efficient algorithms for variational integrators in minimal coordinates ([Lee et al., 2020](#); [Fan et al., 2020](#)) and maximal coordinates ([Brüdigam and Manchester, 2020](#)). With respect to learning-based approaches, variational integrators have found increasing attention in the works of [Saemundsson et al. \(2020\)](#), [Desai et al. \(2021\)](#), and [Zhong and Leonard \(2020\)](#) in combination with neural networks, but not with Gaussian process regression.

Structured learning of system dynamics with Gaussian process regression is gaining increasing attention. A common approach is to include as much information as possible about the system in the prior mean, making the residual smaller and therefore easier to learn, for example as shown in [Koller et al. \(2018\)](#) and [Capone and Hirche \(2019\)](#). Alternatively, several methods exist for constructing kernel functions of a Gaussian process such that they generate predictions obeying certain physical constraints or properties. This principle is employed in [Cheng and Huang \(2016\)](#) to generate kernels spanning a subspace that captures the Lagrangian’s projection as inverse dynamics. In a similar spirit, [Umlauf and Hirche \(2019\)](#) propose a composite kernel that captures the control-affine structure of systems where feedback linearization can be applied. However, these approaches make use of standard discretization techniques, and there exist no approaches utilizing the advantages of variational integrators. An alternative line of research considers the inclusion of constraints in GPs. [Swiler et al. \(2020\)](#) provide a survey of approaches for additionally treating constraints with Gaussian process regression. Amongst others, boundary constraints can be satisfied with so-called warping, where values are transformed to and from the constrained set with a monotone warping function, as described in [Snelson et al. \(2004\)](#). Constraints in the form of (partial) differential equations can be incorporated in Gaussian process regression as well, for example as shown in [Raissi et al. \(2017\)](#) and [Owhadi \(2015\)](#). [Geist and Trimpe \(2020\)](#) incorporate knowledge of the constraints of a continuous-time system to linearly transform a Gaussian process modeling the unconstrained accelerations, such that the accelerations adhere to constraints linear in the accelerations. A disadvantage of these methods is that they require the system at hand to satisfy strict requirements or are formulated in continuous time which might not apply in practice. In contrast, our method is formulated directly in discrete time and can handle nonlinear kinematic constraints. [Ensinger et al. \(2021\)](#) investigate structure-preserving integration for Gaussian process regression, but they do not directly integrate the GP into the integrator.

3. Combining Variational Integrators and Gaussian Process Regression

Our goal is to predict the next state $\mathbf{z}_{k+1} = [\mathbf{x}_{k+1}^T \ \mathbf{v}_{k+1}^T]^T$ of a dynamical system given the current state $\mathbf{z}_k = [\mathbf{x}_k^T \ \mathbf{v}_k^T]^T$, the nominal dynamics model (6), and training data $(\mathbf{Z}, \mathbf{y}) = (\mathbf{z}_i, \mathbf{v}_{i+1})_{i=1\dots N}$ obtained from trajectories with random initial configurations. In the following, we first present an

approach for predicting the next state for unconstrained and constrained systems. Afterwards, formal uncertainty bounds are stated, which is nontrivial for constrained systems due to the projected Gaussian process prediction.

3.1. One-Step Prediction

The next position \mathbf{x}_{k+1} is directly obtained from the variational integrator (cf. (6a)) without requiring any dynamics information:

$$\mathbf{x}_{k+1} = \mathbf{x}_k + \mathbf{v}_k \Delta t. \tag{7}$$

Therefore, the GP is only required for computing the next velocity \mathbf{v}_{k+1} . Hence, we use a posterior GP mean to model the one-step predictive model, i.e.,

$$\mathbf{v}_{k+1}(\mathbf{z}_k) := \mu(\mathbf{z}_k), \tag{8}$$

where $\mu(\mathbf{z}_k)$ is computed as in (1a).

3.1.1. UNCONSTRAINED SYSTEMS

After training the GP with the training data, a posterior mean function for the unconstrained velocity $\mathbf{v}_{k+1}(\mathbf{z}_k)$ is obtained and the integrator for the unconstrained dynamics is constructed as

$$\mathbf{x}_{k+1}(\mathbf{z}_k) = \mathbf{x}_k + \mathbf{v}_k \Delta t, \tag{9a}$$

$$\mathbf{0} = \mathbf{d}_0(\bar{\mathbf{v}}_{k+1}), \tag{9b}$$

$$\mathbf{v}_{k+1}(\mathbf{z}_k) = \bar{\mathbf{v}}_{k+1} + \mathbf{k}(\mathbf{z}_k, \mathbf{Z})K(\mathbf{Z}, \mathbf{Z})^{-1}(\mathbf{y} - \bar{\mathbf{v}}_{k+1}(\mathbf{Z})), \tag{9c}$$

where the prior mean $\bar{\mathbf{v}}_{k+1}$ is obtained by solving the (implicit) nominal dynamics (9b) resulting from the variational integrator (6), for example with Newton’s method.

With integrator (9), unmodeled and potentially non-conservative dynamics, such as friction, can be described by the learned regression model, and the desirable properties of the variational integrator, such as energy conservation, are retained for the nominal and correctly learned dynamics. So, given an underlying Lagrangian system and assuming for all \mathbf{z}_k the velocity $\mathbf{v}_{k+1}(\mathbf{z}_k)$ obtained from the integrator (9) matches the velocity obtained from the variational integrator (6) directly derived for the underlying system, i.e., both produce the same next state, then the integrators (9) and (6) are equivalent variational (symplectic) integrators with corresponding theoretical properties of variational integrators. Deviations of the prediction (9c) from the system’s true velocity no longer allow for a direct statement regarding the equivalence of the integrator (9) and variational integrator (6). Nonetheless, we obtained satisfying results in practice, for example regarding conservation of energy, as demonstrated in Sec. 4.

3.1.2. CONSTRAINED SYSTEMS

Now, systems with kinematic constraints $\mathbf{g}(\mathbf{x}) = \mathbf{0}$ are considered. We assume to know the constraints correctly in advance, i.e., the kinematic constraints need not be learned. Such constraints can be obtained from measurements or are prior knowledge.

The general approach for kinematically constrained systems remains the same as before, but now an additional projection of the predicted state onto the constraint manifold is required to satisfy the constraints. As before, a velocity $\mathbf{v}_{u,k+1}$ is calculated with (9c) from the modeled dynamics

and the Gaussian process regression. However, this velocity will not satisfy $\mathbf{g}(\mathbf{x}_{k+1}, \mathbf{v}_{u,k+1}) = \mathbf{0}$ in general, since the constraints are not explicitly incorporated in the Gaussian process regression. Note that the regression model was obtained from data from the actual constrained system and therefore the regression prediction will still be close to the constraint-satisfying velocity.

The constraint-fulfilling velocity \mathbf{v}_{k+1} is calculated with a nonlinear constrained least-squares optimization, resulting in the following integrator for constrained dynamics:

$$\mathbf{x}_{k+1}(\mathbf{z}_k) = \mathbf{x}_k + \mathbf{v}_k \Delta t, \quad (10a)$$

$$\mathbf{0} = \mathbf{d}_0(\bar{\mathbf{v}}_{k+1}) - \mathbf{G}(\mathbf{x}_{k+1})^T \boldsymbol{\lambda}_{k+1}, \quad (10b)$$

$$\mathbf{0} = \mathbf{g}(\mathbf{x}_{k+1}, \bar{\mathbf{v}}_{k+1}), \quad (10c)$$

$$\mathbf{v}_{u,k+1}(\mathbf{z}_k) = \bar{\mathbf{v}}_{k+1} + \mathbf{k}(\mathbf{z}_k, \mathbf{Z}) \mathbf{K}(\mathbf{Z}, \mathbf{Z})^{-1} (\mathbf{y} - \bar{\mathbf{v}}_{k+1}(\mathbf{Z})), \quad (10d)$$

$$\mathbf{v}_{k+1} = \underset{\mathbf{v}_{k+1}}{\operatorname{argmin}} \|\mathbf{v}_{k+1} - \mathbf{v}_{u,k+1}\|^2, \quad \text{s.t. } \mathbf{g}(\mathbf{x}_{k+1}, \mathbf{v}_{k+1}) = \mathbf{0}. \quad (10e)$$

The integrator (10) consists of the variational integrator of the nominal model (10a) - (10c), the prediction from the GP regression (10d), and the projection of the predicted velocity onto the constraint manifold (10e). Due to the projection, we recover exact constraint satisfaction also for the learned model. The optimization problem (10e) can be solved with a variety of approaches, for example with numerical methods described in [Nocedal and Wright \(2006\)](#).

As before, if the predicted velocity $\mathbf{v}_{u,k+1}$ matches the velocity obtained from the variational integrator (6) for the underlying system, we obtain equivalent variational integrators. Otherwise, no direct statement regarding the equivalence can be made. Due to the projection (10e), technically the resulting velocity \mathbf{v}_{k+1} no longer follows a jointly Gaussian distribution and therefore is not a GP anymore. However, we recover a probabilistic error bound for the output after projection in the next section and, following [Swiler et al. \(2020\)](#), we refer to the overall regression as GP regression.

3.2. Prediction Error Bound

One of the major advantages of GP regression is the availability of an error bound for the posterior mean, which comes in the form of the posterior variance multiplied by a scalar ([Chowdhury and Gopalan, 2017](#); [Lederer et al., 2019](#); [Capone et al., 2021](#); [Srinivas et al., 2012](#)).

In general, the projection (10e) is a nonlinear optimization problem which is not monotone and, hence, its inverse does not necessarily exist. Therefore, warping approaches maintaining a measure of the variance, such as those described by [Snelson et al. \(2004\)](#), cannot be applied to obtain uncertainty bounds on the velocity \mathbf{v}_{k+1} . However, a formal upper bound for the uncertainty can be given under the following assumption.

Assumption 1 *The projection (10e) corresponds to a Lipschitz continuous function with Lipschitz constant L .*

Assumption 1 is violated if a predicted velocity lies close to multiple possible constraint-satisfying velocities, in which case the argmin value is discontinuous. Note that the same issue would occur in the nominal variational integrator if the initial guess for solving the implicit dynamics is far off. That is to say, for reasonable predictions of the GP, Assumption 1 will be fulfilled, just as the nominal variational integrator finds the correct velocity for a reasonable initial guess. In practice the Lipschitz constant L can be estimated by sampling the gradient of the function (10e) at several points in the interval of interest.

Theorem 1 *Let Assumption 1 hold. For all \mathbf{z} , let $\mathbf{v}_u(\mathbf{z}) \sim \mathcal{GP}(\bar{\mathbf{v}}_u(\mathbf{z}), \sigma_u^2(\mathbf{z}))$ denote a sample from a GP with posterior mean $\bar{\mathbf{v}}_u(\mathbf{z})$ and variance $\sigma_u^2(\mathbf{z})$, and consider a least-squares projection with nonlinear constraints $\mathbf{g} = \mathbf{0}$ of the form*

$$\mathbf{v}(\mathbf{v}_u) = \underset{\boldsymbol{\nu}}{\operatorname{argmin}} \|\boldsymbol{\nu} - \mathbf{v}_u\|^2, \text{ s.t. } \mathbf{g}(\boldsymbol{\nu}) = \mathbf{0}.$$

Then for every $\delta \in (0, 1)$, there exists a $\gamma > 0$, such that

$$\|\mathbf{v}(\mathbf{v}_u) - \mathbf{v}(\bar{\mathbf{v}}_u)\| < \sum_{i=1}^{d_v} \gamma \sigma_{u,i}(\mathbf{z})$$

holds for all states \mathbf{z} with a probability at least $1 - \delta$, where d_v is the dimension of \mathbf{v}_u and $\sigma_{u,i}(\cdot)$ is the posterior variance of the i -th entry of the Gaussian process used to model \mathbf{v}_u .

Proof Due to the Lipschitz continuity of the least-squares projection $\mathbf{v}(\cdot)$, we have $\|\mathbf{v}(\mathbf{v}_u) - \mathbf{v}(\bar{\mathbf{v}}_u)\| \leq L\|\mathbf{v}_u - \bar{\mathbf{v}}_u\|$. By applying the union bound to Lederer et al. (2019), we obtain that, for every $\delta \in (0, 1)$, there exist $\beta_1, \dots, \beta_{d_v} > 0$, such that

$$\|\mathbf{v}_u(\mathbf{z}) - \bar{\mathbf{v}}_u(\mathbf{z})\| \leq \sum_{i=1}^{d_v} |\mathbf{v}_{u,i}(\mathbf{z}) - \bar{\mathbf{v}}_{u,i}(\mathbf{z})| \leq \sum_{i=1}^{d_v} \sqrt{\beta_i} \sigma_{u,i}(\mathbf{z})$$

holds for all states \mathbf{z} with probability at least $1 - \delta$. The result then follows by setting $\gamma = L \max_i \sqrt{\beta_i}$. \blacksquare

Theorem 1 allows us to obtain an error bound for the prediction error that holds with high probability. In addition, as the position \mathbf{x}_{k+1} is a linear function of the velocity, the bound can be applied straightforwardly for the entire state prediction.

4. Evaluation and Discussion

We test our methods in several simulations to illustrate the theoretical results and evaluate the integrator’s performance. Besides demonstrating the structure-preserving properties, we also investigate errors for multi-step predictions that are relevant in, for example, model-predictive control.

For each system, training samples are drawn from 200 simulated trajectories starting with uniformly distributed random initial configurations and zero velocity. Each trajectory runs for 2 seconds with a time step of $100\mu\text{s}$. Test samples are drawn from another 100 simulated trajectories with the same setup. Each training sample consists of the full state input \mathbf{z}_k and the velocity targets \mathbf{v}_{k+1} . We use an automatic relevance determination squared-exponential kernel of suitable dimensions for all experiments. Optimal hyperparameters are determined from 100 runs of maximizing the marginal likelihood with varying initial guesses for the hyperparameters.

4.1. Energy conservation

The desirable structure-preserving properties of the combination of variational integrators and Gaussian process regressions are demonstrated on a single and double pendulum with link masses $m = 1\text{kg}$ and lengths $l = 1\text{m}$ in maximal coordinates. For the simulations, we train a Gaussian process




Figure 1: Structure-preserving properties of the integration method even without prior dynamics model. a) Energy error for an explicit Euler integrator (black), integrator (10) with 10 training samples (blue), and integrator (10) with 20 training samples (red). b) Constraint drift for an explicit Euler integrator (black), and integrator (10) with projection (blue).

on a single recorded trajectory of two seconds starting at $\theta = \frac{\pi}{2}$ and $(\theta_1, \theta_2) = (\frac{\pi}{2}, 0)$, respectively, from which the training samples are drawn. The resulting integrator is constructed without any prior dynamics knowledge according to (10). For the single pendulum, the total energy error of the integrator (10) is compared to that of an explicit Euler integration of the real conservative system dynamics on the trained trajectory. Both integrators have a time step of $\Delta t = 10\text{ms}$. The double pendulum is used for comparing the norm of the constraint drift for the explicit Euler method, and the integrator (10) with projection onto the constraints. The results are displayed in Fig. 1.

The energy error for the explicit Euler integration in Fig. 1 a) increases over time as is typical for explicit Runge-Kutta methods. In contrast, the learned dynamics models conserve the energy of the system even for longer time periods. It is observed that for fewer training samples, a certain amount of energy error is reached before conservation occurs. One possible explanation is that a trained area of the state space has to be reached for energy conservation. The small oscillations in the energy error are typical for variational integrators (see Marsden and West (2001)). In Fig. 1 b), the constraint drift of the explicit Euler method for the double pendulum becomes clearly visible, which occurs because constraints are formulated on an acceleration level, and not a position level. In contrast, no drift occurs for the trained integrator (10).

4.2. Prediction Performance

The performance of the proposed integration methods is evaluated by comparing it to integrating a nominal dynamics model and calculating the deviation from the ground-truth simulation of a perturbed system. The ground-truth simulation with a symplectic Euler method uses a time step of $\Delta t = 0.1\text{ms}$, whereas the other integrators use a time step of $\Delta t = 10\text{ms}$. As test systems we selected a pendulum, a cartpole, a double pendulum, and a fourbar segment (closed kinematic chain). The systems are shown in Fig. 2. All links of the systems have a mass $m = 1\text{kg}$, length $l = 1\text{m}$, and inertia $J = \frac{1}{12}ml^2$.

For the Gaussian-process-based integrators, three coordinate descriptions are compared: minimal (joint) coordinates, a complex-number representation, i.e., $[\sin(\theta) \cos(\theta)]^T$ instead of joint angles θ , and maximal coordinates. We perform 20-step ahead predictions in all cases.

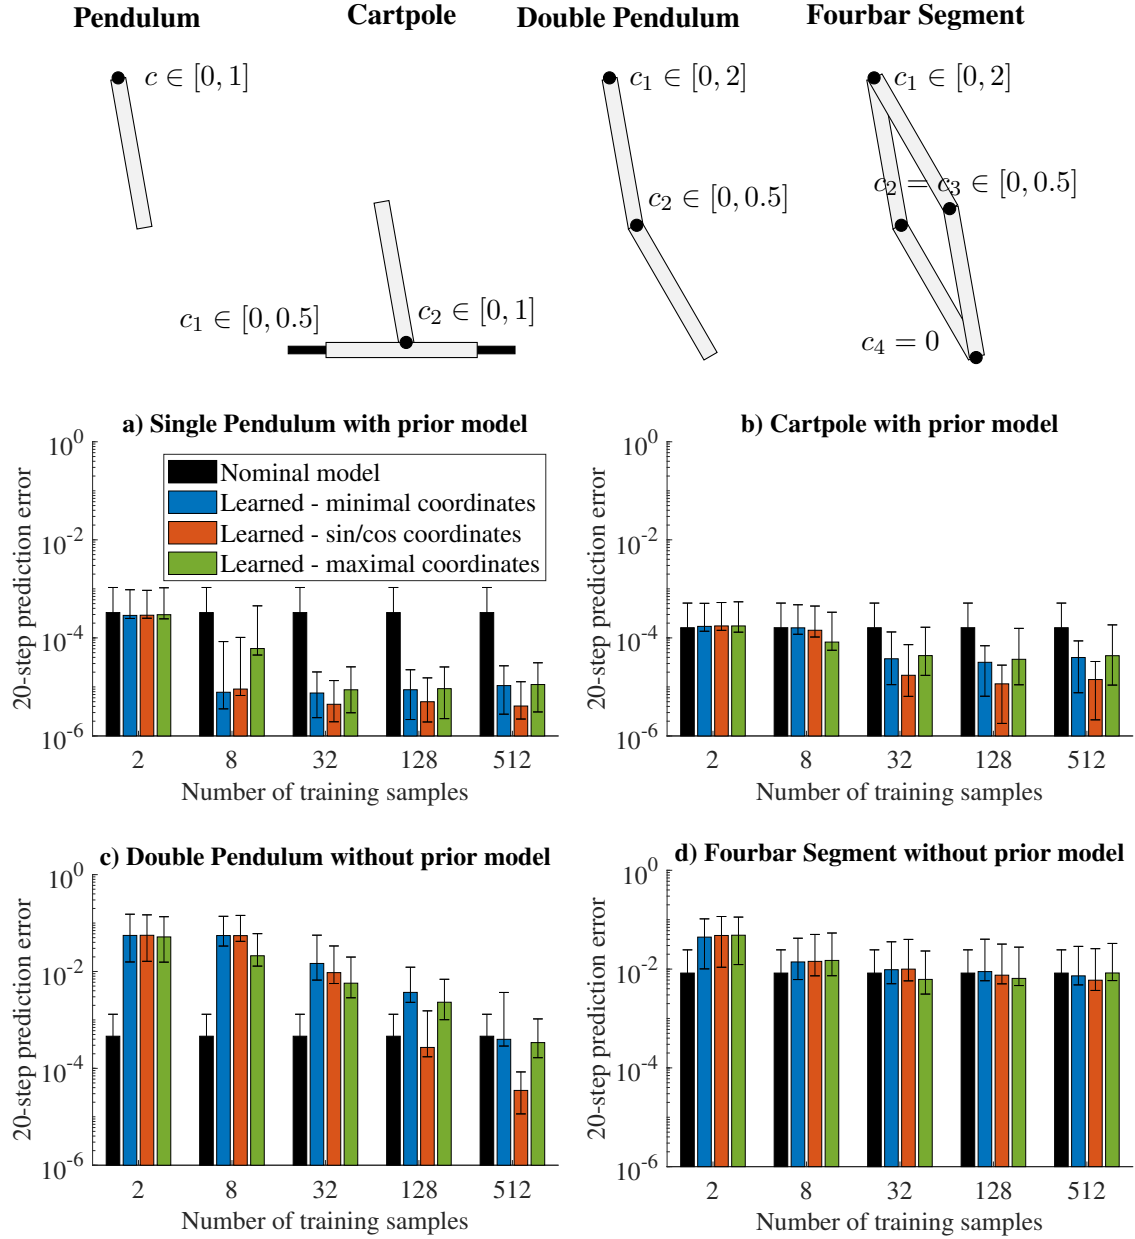


Figure 2: Median mean-square position error and 10 and 90 percentiles for 20-step predictions of perturbed systems (top, with friction coefficients) trained on varying sample numbers and integrated with different integrators: variational integrator of the nominal dynamics model without learning (black), integrator (9) trained in minimal coordinates (blue), integrator (9) trained in sin/cos coordinates (red), integrator (10) trained in maximal coordinates (green). a) Error for a single pendulum with known prior mean dynamics. b) Error for a cartpole with known prior mean dynamics, c) Error for a double pendulum without prior mean dynamics, d) Error for a fourbar segment without prior mean dynamics.

4.2.1. PREDICTION WITH PRIOR MEAN

For the evaluation with a prior mean, i.e., a nominal dynamics model, a single pendulum and cart-pole are used. The masses and inertias of the links of the nominal system models are distorted by random uniformly distributed ($[0.9, 1.1]$) perturbation factors, for example $m' = 0.9m$. Additionally, viscous friction was added for the joints (see uniformly distributed coefficients c in Fig. 2). Each observed state is disturbed with additive zero mean Gaussian noise ($\sigma = 10^{-3}$) for training. The results are shown in Fig. 2 a) and b).

As expected, the results in Figure 2 a) and b) indicate that including a prior mean model yields reasonable results for the proposed integrators despite noise and distortion. The uncertainty and error decreases (note the log scale) with additional training samples, highlighting the advantage of including Gaussian process regression for learning residual dynamics compared to just using the nominal model. Minimal and maximal coordinates achieve similar results, justifying the use of the projection for systems with explicit constraints, while the sin/cos parameterization outperforms the other two coordinate descriptions. It appears that no further improvement is achieved once a certain number of training sample are used. One possible explanation for this result could be that the optimization of hyperparameters gets stuck in local minima.

4.2.2. PREDICTION WITHOUT PRIOR MEAN

We also compare the performance of the learning-augmented integrator without a prior dynamics model, i.e., without a prior mean, on a double pendulum and a fourbar segment. The Gaussian process regression uses the constant mean from the training data as its mean function, a commonly used strategy in situations without a prior mean. The observation noise and friction as well as the perturbation factors for masses and inertias remain the same as in Sec. 4.2.1. The results are shown in Fig. 2 c) and d).

In Fig. 2, unlike the predictions with a prior mean, the errors and uncertainties reached are much higher for few training samples compared to integrating the nominal dynamics. As the number of training samples increases, the accuracy improves until it is close to the prediction accuracy of the integrators with a prior nominal model and partially surpasses the pure nominal dynamics. As before, minimal and maximal coordinates result in similar behavior. In the case of the pendulum, the sin/cos parameterization outperforms both. We only evaluate up to 512 samples due to the long computation times for optimizing hyperparameters, which would also be undesirable in an experimental implementation.

5. Conclusions

We have presented an approach for structure-preserving learning of mechanical systems by combining variational integrators and Gaussian process regression that accurately model and integrate such systems. Known kinematic constraints can be treated with this approach as well. For correct predictions with the Gaussian process regression, the developed method is a variational integrator with corresponding symplectic properties. The evaluation of the proposed methods on a variety of mechanisms in simulation shows satisfactory results both for unconstrained and constrained parameterizations. The data-efficiency of Gaussian process regression and accurate predictions of variational integrators even for large time steps make the proposed method interesting for real-time control applications.

Acknowledgments

We thank Armin Lederer, Samuel Tesfazgi, and Petar Bevanda for their help in preparing this manuscript. This work was supported by the European Union’s Horizon 2020 research and innovation programme under grant agreement no. 871295 ”SeaClear”.

References

- Jan Brüdigam and Zachary Manchester. Linear-Time Variational Integrators in Maximal Coordinates. In *Workshop on the Algorithmic Foundations of Robotics (WAFR)*, 2020.
- Jan Brüdigam and Zachary Manchester. Linear-Quadratic Optimal Control in Maximal Coordinates. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9775–9781. IEEE, 2021.
- Jan Brüdigam, Jana Janeva, Stefan Sosnowski, and Sandra Hirche. Linear-Time Contact and Friction Dynamics in Maximal Coordinates using Variational Integrators. *arXiv:2109.07262 [cs]*, September 2021.
- Alexandre Capone and Sandra Hirche. Backstepping for partially unknown nonlinear systems using gaussian processes. *IEEE Control Systems Letters*, 3(2):416–421, 2019.
- Alexandre Capone, Armin Lederer, and Sandra Hirche. Gaussian process uniform error bounds with unknown hyperparameters for safety-critical applications. *arXiv:2109.02606 [cs]*, 2021.
- Ching-An Cheng and Han-Pang Huang. Learn the Lagrangian: A Vector-Valued RKHS Approach to Identifying Lagrangian Systems. *IEEE Transactions on Cybernetics*, 46(12):3247–3258, December 2016.
- Sayak Ray Chowdhury and Aditya Gopalan. On kernelized multi-armed bandits. In *International Conference on Machine Learning*, pages 844–853. PMLR, 2017.
- 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.
- Shaan A. Desai, Marios Mattheakis, and Stephen J. Roberts. Variational integrator graph networks for learning energy-conserving dynamical systems. *Phys. Rev. E*, 104:035310, Sep 2021.
- Katharina Ensinger, Friedrich Solowjow, Sebastian Ziesche, Michael Tiemann, and Sebastian Trimpe. Structure-preserving gaussian process dynamics. *arXiv:2102.01606 [cs]*, 2021.
- Taasha Fan, Jarvis Schultz, and Todd Murphey. Efficient Computation of Higher-Order Variational Integrators in Robotic Simulation and Trajectory Optimization. In *Algorithmic Foundations of Robotics XIII*, pages 689–706, Cham, 2020. Springer International Publishing.
- A. René Geist and Sebastian Trimpe. Learning Constrained Dynamics with Gauss Principle adhering Gaussian Processes. In *Proceedings of the 2nd Conference on Learning for Dynamics and Control*, volume 120 of *Proceedings of Machine Learning Research (PMLR)*, pages 225–234. PMLR, June 2020.

- Oliver Junge, Jerrold Marsden, and Sina Ober-Blöbaum. Discrete Mechanics and Optimal Control. *IFAC Proceedings Volumes*, 38(1):538–543, 2005. ISSN 14746670.
- Torsten Koller, Felix Berkenkamp, Matteo Turchetta, and Andreas Krause. Learning-based model predictive control for safe exploration. In *2018 IEEE Conference on Decision and Control (CDC)*, pages 6059–6066. IEEE, 2018.
- A Lederer, J Umlauft, and S Hirche. Uniform error bounds for gaussian process regression with application to safe control. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2019.
- Jeongseok Lee, C. Karen Liu, Frank C. Park, and Siddhartha S. Srinivasa. A Linear-Time Variational Integrator for Multibody Systems. In *Algorithmic Foundations of Robotics XIII*, pages 352–367, Cham, 2020. Springer International Publishing.
- J. Marsden and M. West. Discrete Mechanics and Variational Integrators. *Acta Numerica*, 10: 357–514, 2001.
- J. Nocedal and S. Wright. *Numerical Optimization*. Springer, 2006.
- Houman Owhadi. Bayesian Numerical Homogenization. *Multiscale Modeling & Simulation*, 13(3):812–828, January 2015.
- Maziar Raissi, Paris Perdikaris, and George Em Karniadakis. Machine learning of linear differential equations using gaussian processes. *Journal of Computational Physics*, 348:683–693, 2017.
- Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian processes for machine learning*. Adaptive computation and machine learning. MIT Press, Cambridge, Mass, 2006.
- Steindor Saemundsson, Alexander Terenin, Katja Hofmann, and Marc Deisenroth. Variational integrator networks for physically structured embeddings. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pages 3078–3087. PMLR, 26–28 Aug 2020.
- Edward Snelson, Zoubin Ghahramani, and Carl Rasmussen. Warped Gaussian Processes. In S. Thrun, L. Saul, and B. Schölkopf, editors, *Advances in Neural Information Processing Systems*, volume 16. MIT Press, 2004.
- Niranjan Srinivas, Andreas Krause, Sham M. Kakade, and Matthias W. Seeger. Information-theoretic regret bounds for gaussian process optimization in the bandit setting. *IEEE Transactions on Information Theory*, 58(5):3250–3265, 2012.
- Laura P. Swiler, Mamikon Gulian, Ari L. Frankel, Cosmin Safta, and John D. Jakeman. A Survey of Constrained Gaussian Process Regression: Approaches and Implementation Challenges. *Journal of Machine Learning for Modeling and Computing*, 1(2):119–156, 2020.
- Jonas Umlauft and Sandra Hirche. Feedback linearization based on gaussian processes with event-triggered online learning. *IEEE Transactions on Automatic Control*, 65(10):4154–4169, 2019.
- Jonas Umlauft, Armin Lederer, and Sandra Hirche. Learning stable gaussian process state space models. In *2017 American Control Conference (ACC)*, pages 1499–1504, 2017.

Theresa Wenger, Sina Ober-Blöbaum, and Sigrid Leyendecker. Construction and analysis of higher order variational integrators for dynamical systems with holonomic constraints. *Advances in Computational Mathematics*, 43(5):1163–1195, October 2017.

Yaofeng Desmond Zhong and Naomi Leonard. Unsupervised learning of lagrangian dynamics from images for prediction and control. *Advances in Neural Information Processing Systems*, 33, 2020.