Proceedings of Machine Learning Research vol 120:1–10, 2020 2nd Annual Conference on Learning for Dynamics and Control

Learning Constrained Dynamics with Gauss' Principle adhering Gaussian Processes

A. René Geist [†] GEIST@IS.MPG.DE
Sebastian Trimpe [†] TRIMPE@IS.MPG.DE
[†] *Intelligent Control Systems Group, Max Planck Institute for Intelligent Systems*

Editors: A. Bayen, A. Jadbabaie, G. J. Pappas, P. Parrilo, B. Recht, C. Tomlin, M. Zeilinger

Abstract

The identification of the constrained dynamics of mechanical systems is often challenging. Learning methods promise to ease an analytical analysis, but require considerable amounts of data for training. We propose to combine insights from analytical mechanics with Gaussian process regression to improve the model's data efficiency and constraint integrity. The result is a Gaussian process model that incorporates a priori constraint knowledge such that its predictions adhere to Gauss' principle of least constraint. In return, predictions of the system's acceleration naturally respect potentially non-ideal (non-)holonomic equality constraints. As corollary results, our model enables to infer the acceleration of the unconstrained system from data of the constrained system and enables knowledge transfer between differing constraint configurations.

Keywords: Constrained Lagrangian systems, Gauss' Principle, Gaussian Processes, Nonlinear system identification, Structured learning, Transfer learning

1. Introduction

The acquisition of accurate models of dynamical systems is essential for a multitude of engineering applications. If the function generating the data is unknown or too complex to be modelled from first principles, non-parametric learning models aim at inferring a function solely from the data. Gaussian Processes (GPs) are non-parametric and have been commonly used to learn dynamics. They have demonstrated their versatility in approximating continuous nonlinear functions on a plethora of real-world problems (Rasmussen and Williams, 2006), including modeling of dynamical systems (Nguyen-Tuong and Peters, 2011; Kocijan et al., 2005; Frigola et al., 2013; Mattos et al., 2016; Doerr et al., 2017; Eleftheriadis et al., 2017; Doerr et al., 2018). In this context, GPs are often preferred over alternative methods, since they provide a measure of the uncertainty about function estimates in the form of the posterior variance. Additionally, they allow for the incorporation of various model assumptions through the covariance function (kernel). However, data on real-world systems is often scarce and contains only partial information, which is why training purely data-driven models to sufficient prediction accuracy is challenging. Further, predictions made by standard GP models may violate critical system constraints compromising the predictions' integrity.

A promising approach to improve the data efficiency and constraint integrity of a GP model involves the incorporation of a priori available *structural knowledge* in the design of the covariance function. Turning to mechanical systems, the identification of fundamental structural relationships has been studied intensively by scholars for the last two centuries under the subject of analytical mechanics. In this work, we leverage the structure inherent in holonomic and non-holonomic constraint equations, which also can be non-ideal. For example, a pendulum's rod or a rolling wheel enforce a holonomic or non-holonomic constraint, respectively. A constraint is referred to as being *non-ideal* if it induces forces onto the system that produce virtual work (*e.g.*, damping and friction), and it is called *ideal* if no virtual work is produced (*e.g.*, pendulum rod). The constrained dynamics of such systems are described by the Udwadia-Kalaba equation (UKE) (Udwadia and Kalaba, 1992, 2007). The UKE is a direct result of *Gauss' principle* of least constraint (Gauß, 1829), which states that a system's constrained acceleration can be cast as the solution of a least-squares problem.

In this work, we propose a GP model that leverages mechanical constraints as prior knowledge for learning dynamics of mechanical systems. Specifically, a GP is transformed by constraint equations to satisfy Gauss' Principle. The resulting *Gauss' Principle adhering Gaussian Process* (GP^2) performs inference in a physically substantiated sub-space of the acceleration space. In return, the data-efficiency and physical integrity are improved compared to the untransformed GP.

2. Problem Formulation

The acceleration of constrained mechanical rigid-body systems is described as

$$\ddot{q} = h(q, \dot{q}, t) = M^{-1}(q, t)F(q, \dot{q}, t), \quad h : \mathbb{R}^D \to \mathbb{R}^n \tag{1}$$

with state dimension n, $D = 2n + 1$, symmetric positive definite matrix $M(q, t)$, and vector $F(q, \dot{q}, t)$. The variables q, $\dot{q}$, $\ddot{q}$ correspond to the system's positions, velocities, and accelerations, respectively. Although they depend on time t, we generally omit the time dependence for these and other variables if this is clear from the context. Inhere, the system is solely subject to constraint forces arising from sufficiently smooth *(non-)holonomic constraints*, $c_i(q, \dot{q}, t, \theta_p) = 0$, with $i = 1, 2, ..., m$, whose (second) time-derivatives are linear in $\ddot{q}$, yielding

$$A(q, \dot{q}, t)\ddot{q} = b(q, \dot{q}, t), \tag{2}$$

with functions $A : \mathbb{R}^D \to \mathbb{R}^{m\times n}$, $b : \mathbb{R}^D \to \mathbb{R}^m$, and $m < n$. We refer to (2) as the *constraining equation*. Most applications of analytical mechanics deal with such constraints (Udwadia and Kalaba, 2007, p. 80). Note that the constraints can be *non-ideal*, but naturally must be consistent. Further, it is assumed that the parametric functions $\{A, b, M\}$ are known, *e.g.*, from a preceding mechanical analysis of the "unconstrained" rigid-body dynamics (see Section 4), while the system's parameters $\theta_p = [p_1, ..., p_r]$ are potentially unknown.

For the above-defined system, data $\mathcal{D} = \{x_k, y_k\}_{k=1}^N$ is available, consisting of input points $x_k = [q_k, \dot{q}_k, t_k]^\mathsf{T}$ and observations $y_k = h(x_k) + \epsilon_k$ where $y_k \in \mathbb{R}^n$ with ϵ_k denoting zero-mean Gaussian noise. In some cases, x_k includes a vector of control forces $u_k \in \mathbb{R}^{n_u}$ explicitly. While more general settings such as hidden states can be addressed in GP dynamics learning (Doerr et al., 2018), we here focus on the standard GP regression setting with noiseless inputs and noisy targets.

The main objective of this work is to learn h via a GP $\hat{h} \sim \mathcal{GP}(\mu_{\hat{h}}, K_{\hat{h}})$ that incorporates the structural knowledge $\{A, b, M\}$ in the GP's mean $\mu_{\hat{h}}(x|\theta_p)$ and covariance function $K_{\hat{h}}(x, x'|\theta_p)$ such that its posterior mean approximates (1) while satisfying (2) exactly.

3. Related Work

We categorize related literature into the following aspects: *(i)* leveraging functional relations arising from analytical mechanics to structure learning; *(ii)* constructing GP kernels using semi-parametric models; and *(iii)* using linear transformations to provide structure to GP models.

The incorporation of functional relationships arising from the field of analytical mechanics starts gaining attention. For example, (Ledezma and Haddadin, 2018) use the recursive nature of the Newton-Euler equations of open-chain rigid-body configurations to derive a parametric regression model. Recently, (Lutter et al., 2019) detailed the direct incorporation of the Lagrange equations describing the inverse dynamics of holonomic ideal constrained systems into neural networks. Further, (Greydanus et al., 2019) combined Hamiltonian mechanics with neural networks to predict the forward dynamics of conservative mechanical systems. Our work is linked to (Cheng and Huang, 2015), in which the operators underlying the Lagrange equations are used to derive kernels which capture inverse dynamics as the Lagrangian's projection. In comparison, our model builds on projection operations underlying mechanical constraint equations to enable structured learning of forward dynamics on possibly non-holonomic and non-ideally constrained systems.

If a parametric function is cast as a (Bayesian) linear regression problem, a degenerate GP covariance function can be derived to model the function non-parametrically. This relationship is used in (Nguyen-Tuong and Peters, 2010) to derive a structured GP kernel for learning the inverse dynamics of open-chain robot arms. In our work, we do not assume that the dynamics are linear in the system parameters, but instead, leverage that (2) is linear with respect to $\ddot{q}$.

Because GPs are closed under linear operations, and linear operators commonly occur in physical equations, such operators play a particularly important role in structured learning with GPs. Linear operations commonly used in literature are, *e.g.*, differentiation (Solak et al., 2003), integration, and wrapping of the GP inputs into a nonlinear function (Calandra et al., 2016). Furthermore, one particular application of GP regression revolves around efficiently learning the solution of differential equations by transforming GPs through convolution operators (Alvarez et al., 2009; Särkkä, 2011). (Jidling et al., 2017; Lange-Hegermann, 2018) discuss how to construct a GP kernel such that its realizations $f(x)$ fulfill a constrained equation of the form $A_x f(x) = 0$, where A_x is a linear operator. While our work also transforms a GP such that its predictions lie in a linear operator's nullspace, we emphasize that in classical mechanics such operators enable the construction of a GP whose predictions satisfy nonlinear constraints in a physically meaningful manner. In comparison, the work in (Agrell, 2019) considers modeling a GP to fulfill an inequality constraint of the form $A_x f(x) \leq b(x)$ by conditioning its posterior on carefully selected virtual observations of $b(x)$. In contrast to (Jidling et al., 2017; Lange-Hegermann, 2018) and this work, (Agrell, 2019) ensures constraint satisfaction with a certain probability at selected sample points, while we ensure constraint satisfaction over the whole input range.

None of the above works addresses the problem of constrained modeling of rigid-body mechanical systems through a tailored GP satisfying an affine equality constraint as stated in Section 2.

4. Constrained Dynamics with Gauss' Principle

In this section, we present a description of constrained mechanical systems yields an explicit form for the constraint forces acting on a system as a function of the state x and constraining equation (2) as detailed in (Udwadia and Kalaba, 2002) and (Geist and Trimpe, 2020, Sec. S3 & S4). This allows us in the following section to construct a GP that respects the constraints underlying (2).

We consider mechanical systems where a force F_a acting on rigid bodies results in the *unconstrained acceleration*

$$a = M(q,t)^{-1} F_a(x). \tag{3}$$

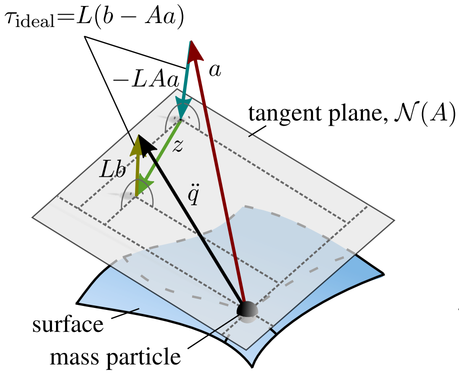

Figure 1: Mass particle sliding on a surface subject to Gauss' principle.

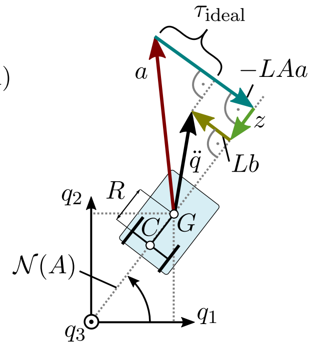

Figure 2: Unicycle subject to Gauss' Principle.

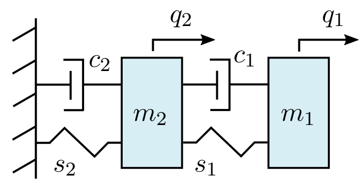

Figure 3: Duffing oscillator.

The term "unconstrained" refers to $a(x)$ describing potentially constrained Lagrangian dynamics onto which an (additional) constraining equation (2) has not yet been applied. If an (additional) constraint acts on the system, its movement changes to the *constrained acceleration* $\ddot{q} = a + \tau$. We refer to the term $\tau = \tau_{\text{ideal}} + z$ as the *constraining acceleration*, where $z(x)$ denotes the non-ideal part of $\tau(x)$ (*e.g.*, damping and friction). Udwadia et al. (1997) refers to a constraining acceleration as non-ideal if it lies in the nullspace of $A(x)$, writing $\mathcal{N}(A)$, and as ideal if it lies in the range space of $A(x)$. Gauß (1829) observed that the *ideal constraining acceleration* $\tau_{\text{ideal}}(x)$ minimizes the functional $G(x) = \tau_{\text{ideal}}^{\mathsf{T}} M \tau_{\text{ideal}}$. This fundamental principle underlying the constrained motion of rigid-body systems is referred to as *Gauss' principle of least constraint*. The minimizing solution of this least-squares problem is given by the UKE – omitting the dependencies on x for clarity – as

$$\ddot{q} = \underbrace{M^{-1}A^{\mathsf{T}}(AM^{-1}A^{\mathsf{T}})^{+}}_{L(x,\theta_p)} b + \underbrace{(I - M^{-1}A^{\mathsf{T}}(AM^{-1}A^{\mathsf{T}})^{+}A)}_{T(x,\theta_p)} \underbrace{(a+z)}_{\bar{a}(x)}, \tag{4}$$

where $(\cdot)^{+}$ denotes the Moore-Penrose pseudo (MP) inverse and I is an identity matrix. The UKE satisfies (2) by construction and corresponds to (1), for which we shall build a GP model. Next, we introduce three examples, which we use throughout for illustration. Further details can be found in (Geist and Trimpe, 2020, Sec. S6) and (Udwadia and Kalaba, 2007, p. 120, p. 213).

Ex. 1 (Particle on surface) *Consider a particle as illustrated in Fig. 1 with the known mass m sliding along a surface. While the dynamics of the particle are unknown, we want to leverage the surface geometry. The unconstrained acceleration of the particle amounts to $a = M^{-1}[u_1, u_2, u_3 - mg]^{\mathsf{T}}$ with $M^{-1} = \operatorname{diag}(1/m)$, $g = 9.81\,\frac{\mathrm{m}}{\mathrm{s}^2}$, and control forces u_i. The mass slides on the surface $q_3 = p_1 q_1^2 + p_2 q_2^2 + p_3 q_1 + p_4 \cos(p_5 q_1)$, with the states and constraint parameters being denoted as $\{q_i, \dot{q}_i\}$ and $\theta_p = [p_1, ..., p_5]$. The second time-derivative of the constraint yields (2) as*

$$\underbrace{\left[2p_1q_1 + p_3 - p_4p_5\sin(p_5q_1), \quad 2p_2q_2, \quad -1\right]}_{A(x,\theta_p)} \ddot{q} = \underbrace{\left[-2p_1\dot{q}_1^2 - 2p_2\dot{q}_2^2 + p_4p_5^2\dot{q}_1^2\cos(p_5q_1)\right]}_{b(x,\theta_p)}, \tag{5}$$

In addition to τ_{ideal} resulting from (5), a velocity quadratic damping force $F_z = Mz$ decelerates the mass non-ideally such that $z_i = -a_0(v^2/|v|)\dot{q}_i$, with the translatory velocity $v(x)$ and damping coefficient a_0. One obtains the system's constrained dynamics by inserting (5) as well as $a(x)$ and $z(x)$ into (4). While θ_p can be readily measured and $\{A, b, M\}$ are obtained from a brief mechanical analysis, modeling F_a and F_z pose a considerable challenge for a plethora of mechanical systems.

Ex. 2 (Unicycle) *The unicycle as depicted in Fig. 2 commonly describes the motion of wheeled robots (Siciliano et al., 2010, p. 478). Here, a non-holonomic constraint* $\dot{q}_2 = \dot{q}_1 \tan(q_3)$ *only allows for instantaneous translation along the line C-G. Further,* u_1 *and* u_2 *denote control inputs in direction of C-G and around* q_3*, respectively. The system is decelerated in driving direction by* F_z*.* F_z *is induced by velocity quadratic damping as detailed in Ex. 1.*

Ex. 3 (Controlled Duffing's Oscillator) *The* Duffing's oscillator *as depicted in Fig. 3 models the behavior of two masses that are subject to cubic spring forces and liner damping. An external control force imposes the constraint* $q_2 = q_1 + p_1 \exp(-p_2 t) \sin(p_3 t)$*. The uncontrolled system is defined as the unconstrained system such that* F_a *origins from spring and damping forces.*

5. A Gaussian Process Model for Learning Constrained Dynamics

In this section, a constrained GP is derived from the UKE formulation of constrained mechanics in Section 4. To this end, we leverage that GPs are closed under linear transformations and that the operators underlying the UKE are projections. Thus, we obtain a transformed GP model for learning constrained dynamics whose mean and samples fulfill Gauss' principle.

5.1. Gaussian Process Regression

We consider learning of h in (1) with GP regression. Specifically, we approximate the true dynamics h with a multi-output GP $\hat{h}$ (Alvarez et al., 2012), $\hat{h}(x) \sim \mathcal{GP}(\mu(x), K(x, x'))$ where $\mu : \mathbb{R}^D \to \mathbb{R}^n$ denotes the prior function mean, $\mu(x) = \mathbb{E}[h(x)]$, and $K(x, x') : \mathbb{R}^{D \times D} \to \mathbb{R}^{n \times n}$ the prior covariance, $K(x, x') \triangleq K_{x,x'} = \mathbb{E}[(h(x) - \mu(x))(h(x') - \mu(x'))^\intercal]$. Given a GP model, predictions of $\hat{h}(x_*)$ are made – at a point x_* using observations y at inputs X – by computation of the covariance matrices $K(x_*, X)$ and $K(X, X)$ and conditioning the GP via

$$\mu_{x_*|X,y} = \mu_{x_*} + K_{x_*,X}(K_{X,X} + \sigma_y^2 I)^{-1}(y - \mu_X), \tag{6}$$

$$K_{x_*|X,y} = K_{x_*,x_*} - K_{x_*,X}(K_{X,X} + \sigma_y^2 I)^{-1} K_{x_*,X}^\intercal. \tag{7}$$

5.2. Gauss Principle adhering Gaussian Processes

The UKE (4) disentangles the acceleration $\ddot{q}(x)$ into a term that results from the unconstrained acceleration (being transformed), $T(x)a(x)$, one that is caused from the ideal part of the constraints (2), $L(x)b(x)$, and one caused by the non-ideal part of the constraints, $T(x)z(x)$. Hence, it provides the structure to build the sought GP model. The second term is known from the structural knowledge $A(x, \theta_p)$, $b(x, \theta_p)$, and $M(x, \theta_p)$, while on the other terms, we will place a GP prior. More specifically, we consider two cases of prior structural knowledge: *(i)* knowing the true $\theta_p = \theta_p^*$, or *(ii)* knowing only the functional form of $A(x, \theta_p)$, $b(x, \theta_p)$, $M(x, \theta_p)$, but not the true parameters θ_p^*. While we omit the dependencies on θ_p to ease the notation in the following, our work addresses both cases. As for case *(ii)*, θ_p will be treated as additional hyperparameters of the model.

We propose to model the unconstrained acceleration and non-ideal constraining acceleration jointly by placing a GP prior on $\bar{a}(x)$ such that $\hat{\bar{a}} \sim \mathcal{GP}(\mu_{\bar{a}}(x), K_{\bar{a}}(x, x'))$. As GPs are closed under linear transformations, inserting $\hat{\bar{a}}$ into (4) results in a GP modeling $\ddot{q}(x)$ as

$$\hat{h} \sim \mathcal{GP}\big(L(x)b(x) + T(x)\mu_{\bar{a}}(x),\ T(x)K_{\bar{a}}(x, x')T(x')^\intercal\big). \tag{8}$$

To the above model, we refer to as a *Gauss' Principle adhering Gaussian Process* (GP^2). By construction, the GP^2's predictions (mean and samples) satisfy (2). We further note the following favorable properties of the proposed model.

Flexible model. If no prior knowledge about the unconstrained dynamics is available $\mu_{\bar{a}}(x)$ can be set to be the null vector. Alternatively, if a prior model for the unconstrained dynamics is known, this knowledge can be incorporated in form of the mean function $\mu_{\bar{a}}(x)$ in (8). This in return implies that $K_{\bar{a}}(x, x')$ solely models the non-ideal part of the constraint forces plus the residuals of the parametric model. For example, in Ex. 2, the functional structure of $a(x, \theta_p)$ is oftentimes known a priori. In this case, one can set $\mu_{\bar{a}}(x) = a(x, \theta_p)$. As we show in the experimental section, the parameters describing $\mu_{\bar{a}}(x)$ can then be estimated alongside the parameters of $K_{\bar{a}}(x, x')$.

Inferring the unconstrained acceleration alongside. As the constrained GP results from a linear transformation of $\bar{a}$(x), the joint distribution is obtained as

$$\begin{bmatrix} \hat{\bar{a}} \\ \hat{h} \end{bmatrix} \sim \mathcal{GP}\left(\begin{bmatrix} \mu_{\bar{a}}(x) \\ \mu_{\hat{h}}(x) \end{bmatrix}, \begin{bmatrix} K_{\bar{a}}(x, x') & K_{\bar{a}}(x) T(x')^\mathsf{T} \\ T(x) K_{\bar{a}}(x) & T(x) K_{\bar{a}}(x, x') T(x')^\mathsf{T} \end{bmatrix} \right). \tag{9}$$

That is, with (6) one can directly infer $\bar{a}(x)$ from data of the constrained system, as well as condition the constrained acceleration $\hat{h}(x)$ on prior knowledge of $\bar{a}(x)$.

Knowledge transfer between constraint configurations. In many systems, altering the constraint configuration $\{A(x, \theta_p), b(x, \theta_p)\}$ to a different known configuration $\{A'(x, \theta_p'), b'(x, \theta_p')\}$ does not change $\{\bar{a}, M\}$. For example, imagine taking a mass particle (Ex. 1) from one shape of surface to a different one with the same tribological properties. In this case, it is possible to transfer knowledge in form of $\mathcal{D}$ from one system to the different system using the joint distribution

$$\begin{bmatrix} \hat{h} \\ \hat{h}' \end{bmatrix} \sim \mathcal{GP}\left(\begin{bmatrix} \mu_{\hat{h}}(x|\theta_p) \\ \mu_{\hat{h}'}(x|\theta_p') \end{bmatrix}, \begin{bmatrix} T(x|\theta_p) K_{\bar{a}}(x, x') T(x'|\theta_p)^\mathsf{T} & T(x|\theta_p) K_{\bar{a}}(x, x') T'(x'|\theta_p')^\mathsf{T} \\ T'(x|\theta_p') K_{\bar{a}}(x, x') T(x'|\theta_p)^\mathsf{T} & T'(x|\theta_p') K_{\bar{a}}(x, x') T'(x'|\theta_p')^\mathsf{T} \end{bmatrix} \right). \tag{10}$$

6. Experimental Results

In this section, the properties of the proposed GP^2 model are analyzed on the benchmark systems detailed in Ex. 1 to 3. We compare the GP^2 to standard (multi-output) GPs.[1]

The system parameters are detailed in the supplementary material (Geist and Trimpe, 2020, Sec. S6). The input training data consists of randomly sampled observations lying inside the constrained state space. The training data was generated using the analytic ODE (1). The prediction points originate from an equidistant discretization of the constrained state-space. The training data is normalized to have zero mean and standard deviation of one.

As a first baseline for model comparison, the individual $\ddot{q}_i$ are modelled independently as $\ddot{q}_i \sim \mathcal{GP}(0, k_{\text{SE}}(x, x'))$, with squared exponential (SE) covariance function $k_{\text{SE}}(x, x')$. Further, we compare to a standard multi-output GP model, the (GPy, 2012) implementation of the LMC (Alvarez et al., 2012) with matrix $B_i = W_i W_i^\mathsf{T} + I_n \kappa$, $W_i \in \mathbb{R}^{n \times r}$, and $\kappa > 0$, reading $K(x, x') = B_1 k_{\text{SE}}(x, x') + B_2 k_{\text{bias}}(x, x') + B_3 k_{\text{linear}}(x, x')$. Inhere, k_{linear} and k_{bias} denote a linear and bias covariance function respectively (Rasmussen and Williams, 2006). The model $K(x, x') = B_1 k_{\text{SE}}(x, x')$ is referred to as ICM. The hyperparameters are optimized by maximum likelihood estimation via

1. The simulation code is available on: `https://github.com/AndReGeist/gp_squared`

Table 1: Comparison of the normalized GPs' predicted mean RMSE, and maximum constraint error for 10 runs. For the RMSE, the mean, min. (subscript) and max. (superscript) values are shown.

	RMSE			max. constraint error		
	Surface	Unicycle	Duffing	Surface	Unicycle	Duffing
Analy. ODE	—	—	—	$2{\cdot}10^{-15}$	$6{\cdot}10^{-17}$	$2{\cdot}10^{-14}$
SE	$.235\,^{.272}_{.190}$	$0.27\,^{0.40}_{0.20}$	$.011\,^{.033}_{.004}$	2.6	0.22	6.4
ICM	$.244\,^{.277}_{.223}$	$0.21\,^{0.27}_{0.15}$	$\mathbf{.003}\,^{\mathbf{.005}}_{.002}$	1.5	0.22	0.5
LMC	$.194\,^{.233}_{.159}$	$0.21\,^{0.28}_{0.15}$	$\mathbf{.003}\,^{.006}_{\mathbf{.001}}$	1.8	0.26	0.023
GP^2, $\theta_p{=}\theta_p^*$, $\mu_{\bar{a}}{=}0$	$.058\,^{.066}_{.045}$	$0.10\,^{0.20}_{0.06}$	$.007\,^{.019}_{.001}$	$4{\cdot}10^{-12}$	$2{\cdot}10^{-12}$	$\mathbf{1{\cdot}10^{-8}}$
GP^2, $\theta_p{=}\theta_p^*$, $\mu_{\bar{a}}{\neq}0$	$\mathbf{.023}\,^{\mathbf{.032}}_{\mathbf{.018}}$	$\mathbf{0.08}\,^{\mathbf{0.15}}_{\mathbf{0.05}}$	$.005\,^{.029}_{.002}$	$\mathbf{1{\cdot}10^{-13}}$	$\mathbf{6{\cdot}10^{-13}}$	$3{\cdot}10^{-8}$
GP^2, est. θ_p, $\mu_{\bar{a}}{=}0$	$.065\,^{.071}_{.056}$	$0.13\,^{0.30}_{\mathbf{0.05}}$	$.009\,^{.028}_{.003}$	0.12	$9{\cdot}10^{-13}$	0.013
GP^2, est. θ_p, $\mu_{\bar{a}}{\neq}0$	$.027\,^{.037}_{.020}$	$0.12\,^{0.20}_{0.06}$	$.020\,^{.077}_{.004}$	0.09	$9{\cdot}10^{-13}$	0.006

L-BFGS-b (Zhu et al., 1997). For the GP^2, we model $\bar{a} \sim \mathcal{GP}(0, k_{\text{SE}}(x, x'))$, see also (Geist and Trimpe, 2020, Sec. S2), with the same optimization settings as for the other models. For $\mu_{\bar{a}} \neq 0$, the prior mean of $\bar{a}$ is set to $\mu_{\bar{a}} = a(x, \theta_p)$ for Ex. 1 and 2, while for Ex. 3 $\mu_{\bar{a}}$ models the linear part of the acceleration induced by dampers and springs. Here, the spring and damping parameters are added to θ_p and estimated alongside the other parameters.

Optimization and prediction In each of 10 optimization runs, 100 observations are sampled while the optimization is restarted 30 times for the benchmark GPs and five times for the GP^2 model. The prediction results after optimization are depicted in Table 1. For the mechanically constrained systems of Ex. 1 and 2, the GP^2 shows improved prediction accuracy. The performance can be further increased via the incorporation of additional structural knowledge in form of $\mu_{\bar{a}}$ ($\mu_{\bar{a}} \neq 0$ in Table 1). If θ_p is estimated (est. θ_p) the constraint error increases. For Ex. 1 and 3, θ_p was estimated accurately, whereas the unicycle's parameters (I_c, R) converged to their correct ratio. In the case of Ex. 3, a controller constrains the two masses to oscillate synchronously over time. Unlike the unconstrained dynamics that show nonlinear oscillatory behavior, the constrained system dynamics move similar to a single linearly damped oscillator. In this scenario, the GP^2 compares less favorable to the other models as it is learning on the more complex unconstrained dynamics. For all examples, the GP^2 demonstrates a considerable improvement with regards to constraint satisfaction.

Extrapolation and transfer For illustration of the prediction characteristics, we assume the constraint parameters as given and estimated the remaining GP's hyperparameters on 200 observations. Figure 4a illustrates that the GP^2 model *extrapolates* the prediction result for Ex. 1 ($v = 0$). Yet, extrapolation requires $\bar{a}(X) = \bar{a}(x^*)$. This is not the case for velocity input dimensions when damping plays a predominant role. For a different surface $q_3 = 0.1q_1 - 0.15q_2 - 0.1\cos(3q_1)$ resulting in $\hat{h}'(x)$ and with $\bar{a}(x) = \bar{a}'(x)$, (10) enables the transfer of the knowledge inherent in $y(x)$ to $\hat{h}'(x)$. Figure 4b (top) illustrates how the GP^2's samples of $\ddot{q}_1(x)$ at $q_3 = 90°$ and $q_3 = -90°$ are forced to zero as the unicycle can only translate in driving direction. Figure Fig. 4b (bottom), shows the posterior distribution $\bar{a}|y$ using (9). For Ex. 1 with $v = 0$ and $u = 0$, $\bar{a}_3$ is simply $g = 9.81\frac{m}{s^2}$. For Ex. 2 with $u = 0$, $\bar{a}$ solely consists of a damping force that increases with the translatory velocity.

Trajectory prediction Figure 4c illustrates trajectory predictions on Ex. 1 computed by a Runke-Kutta-45 (RK45) ODE solver. The solver uses either (1), the SE model, or the GP^2 model. While the

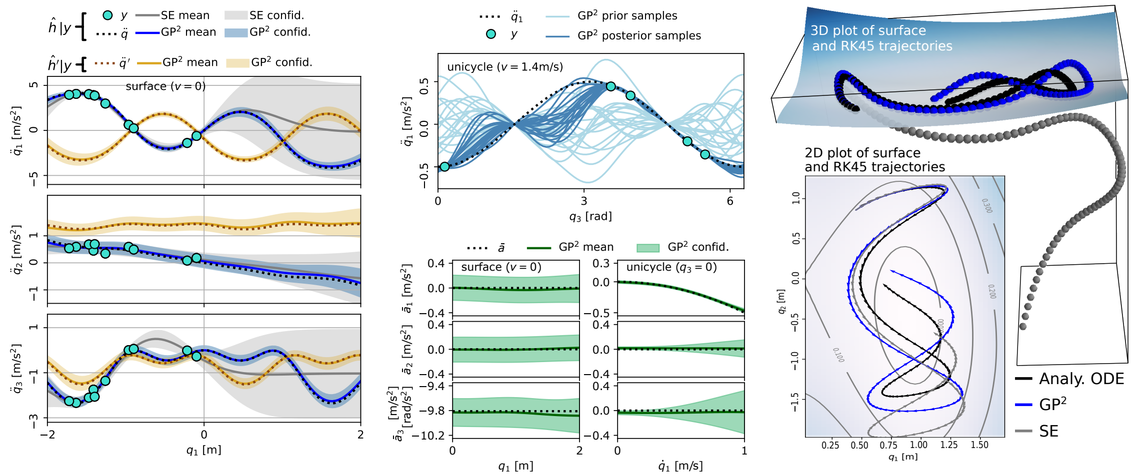

(a) (b) (c)

Figure 4: GP^2 predictions. Figure 4a: Given y from one surface (blue), predictions of $\ddot{q}$ on the same and $\ddot{q}'$ on another surface (yellow). Figure 4b: (Top) Samples of $\hat{h}$ before and after conditioning on data y; (Bottom) Prediction of $\bar{a}|y$. Figure 4c: RK45 trajectory predictions of Ex. 1.

SE trajectory prediction leaves the surface, the GP^2's prediction remains on the surface independent of the overall prediction performance. Inhere, the GP^2 predictions' Euclidean error to the surface increases in the same order of magnitude as with the analytical ODE. In the supplementary material (Geist and Trimpe, 2020, Sec. S7), we further illustrate that the trajectory predictions of the GP^2 with estimated hyper-parameters also show improved constraint integrity.

7. Concluding Remarks

We propose a new GP model for learning Lagrangian dynamics that are subject to a non-ideal equality constraint. We leverage that the constraint equation and the system's mass matrix are often straightforward to obtain from a prior mechanical analysis while the constrained dynamics – with non-ideal forces acting on the system – require considerable effort to be modeled parametrically. For typical mechanical examples, the numerical results demonstrate improved data efficiency and constraint satisfaction. While the numerical results herein treat low-dimensional examples chosen for the purpose of illustration, the method also applies when several constraints act onto the system. Investigating the model's benefits on high-dimensional and hardware experiments is subject of ongoing work, and likewise further analysis of the cases of a singular $M(x)$.

Acknowledgments

The authors thank F. Solowjow, A. von Rohr, H. Haresamudram, and C. Fiedler for the helpful discussions. We further thank the Cyber Valley Initiative and the international Max Planck research school for intelligent systems (IMPRS-IS) for supporting A. René Geist.

References

Christian Agrell. Gaussian processes with linear operator inequality constraints. *Journal of Machine Learning Research*, 20:1–36, 2019.

Mauricio Alvarez, David Luengo, and Neil D Lawrence. Latent force models. In *Artificial Intelligence and Statistics*, pages 9–16, 2009.

Mauricio A Alvarez, Lorenzo Rosasco, Neil D Lawrence, et al. Kernels for vector-valued functions: A review. *Foundations and Trends in Machine Learning*, 4(3):195–266, 2012.

Roberto Calandra, Jan Peters, Carl Edward Rasmussen, and Marc Peter Deisenroth. Manifold Gaussian processes for regression. In *Inte. Joint Conference on Neural Networks*. IEEE, 2016.

Ching-An Cheng and Han-Pang Huang. Learn the Lagrangian: A vector-valued rkhs approach to identifying lagrangian systems. *Transactions on cybernetics*, 46(12):3247–3258, 2015.

Andreas Doerr, Christian Daniel, Duy Nguyen-Tuong, Alonso Marco, Stefan Schaal, Marc Toussaint, and Sebastian Trimpe. Optimizing long-term predictions for model-based policy search. In *Conference on Robot Learning*, pages 227–238, 2017.

Andreas Doerr, Christian Daniel, Martin Schiegg, Nguyen-Tuong Duy, Stefan Schaal, Marc Toussaint, and Trimpe Sebastian. Probabilistic recurrent state-space models. volume 80 of *Proceedings of Machine Learning Research*, Stockholm Sweden, 10–15 Jul 2018. PMLR.

Stefanos Eleftheriadis, Tom Nicholson, Marc Deisenroth, and James Hensman. Identification of Gaussian process state space models. In *Advances in neural information processing systems*, pages 5309–5319, 2017.

Roger Frigola, Fredrik Lindsten, Thomas B Schön, and Carl Edward Rasmussen. Bayesian inference and learning in Gaussian process state-space models with particle MCMC. In *Advances in Neural Information Processing Systems*, pages 3156–3164, 2013.

Carl Friedrich Gauß. Über ein neues allgemeines Grundgesetz der Mechanik. *Journal für die reine und angewandte Mathematik*, 4:232–235, 1829.

A. René Geist and Sebastian Trimpe. Supplementary material: Learning constrained dynamics with Gauss' principle adhering Gaussian processes (including supplementary material). *arXiv preprint arXiv:2004.11238*, 2020.

GPy. A Gaussian process framework in python. `github.com/SheffieldML/GPy`, 2012.

Samuel Greydanus, Misko Dzamba, and Jason Yosinski. Hamiltonian neural networks. In *Advances in Neural Information Processing Systems*, pages 15353–15363, 2019.

Carl Jidling, Niklas Wahlström, Adrian Wills, and Thomas B Schön. Linearly constrained Gaussian processes. In *Advances in Neural Information Processing Systems*, pages 1215–1224, 2017.

Juš Kocijan, Agathe Girard, Blaž Banko, and Roderick Murray-Smith. Dynamic systems identification with Gaussian processes. *Mathematical and Computer Modelling of Dynamical Systems*, 11(4):411–424, 2005.

Markus Lange-Hegermann. Algorithmic linearly constrained Gaussian processes. In *Advances in Neural Information Processing Systems*, pages 2141–2152, 2018.

Fernando Díaz Ledezma and Sami Haddadin. Fop networks for learning humanoid body schema and dynamics. In *2018 IEEE-RAS 18th International Conference on Humanoid Robots (Humanoids)*, pages 1–9. IEEE, 2018.

M. Lutter, C. Ritter, and J. Peters. Deep lagrangian networks: Using physics as model prior for deep learning. In *7th International Conference on Learning Representations (ICLR)*, May 2019.

César Lincoln C Mattos, Andreas Damianou, Guilherme A Barreto, and Neil D Lawrence. Latent autoregressive Gaussian processes models for robust system identification. *IFAC-PapersOnLine*, 49(7):1121–1126, 2016.

Duy Nguyen-Tuong and Jan Peters. Using model knowledge for learning inverse dynamics. In *International Conference on Robotics and Automation (ICRA)*, pages 2677–2682. IEEE, 2010.

Duy Nguyen-Tuong and Jan Peters. Model learning for robot control: a survey. *Cognitive processing*, 12(4):319–340, 2011.

Carl Edward Rasmussen and Christopher KI Williams. *Gaussian processes for machine learning*. The MIT press, 2006.

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

Bruno Siciliano, Lorenzo Sciavicco, Luigi Villani, and Giuseppe Oriolo. *Robotics: modelling, planning and control*. Springer Science & Business Media, 2010.

Ercan Solak, Roderick Murray-Smith, William E Leithead, Douglas J Leith, and Carl E Rasmussen. Derivative observations in Gaussian process models of dynamic systems. In *Advances in neural information processing systems*, pages 1057–1064, 2003.

Firdaus E Udwadia and Robert Kalaba. *Analytical dynamics: a new approach*. Cambridge University Press, 2007.

Firdaus E Udwadia and Robert E Kalaba. A new perspective on constrained motion. *Proceedings of the Royal Society of London.*, 439(1906):407–410, 1992.

Firdaus E Udwadia and Robert E Kalaba. On the foundations of analytical dynamics. *International Journal of non-linear mechanics*, 37(6):1079–1090, 2002.

Firdaus E Udwadia, Robert E Kalaba, and Hee-Chang Eun. Equations of motion for constrained mechanical systems and the extended d'alembert's principle. *Quarterly of Applied Mathematics*, 55(2):321–331, 1997.

Ciyou Zhu, Richard H Byrd, Peihuang Lu, and Jorge Nocedal. Algorithm 778: L-bfgs-b: Fortran subroutines for large-scale bound-constrained optimization. *ACM Transactions on Mathematical Software (TOMS)*, 23(4):550–560, 1997.