

The Impact of Data on the Stability of Learning-Based Control – Extended Version

Armin Lederer
Alexandre Capone
Thomas Beckers
Jonas Umlauf
Sandra Hirche

Chair of Information-oriented Control
Department of Electrical and Computer Engineering
Technical University of Munich
D-80333 Munich, Germany

ARMIN.LEDERER@TUM.DE
ALEXANDRE.CAPONE@TUM.DE
T.BECKERS@TUM.DE
JONAS.UMLAUFT@TUM.DE
HIRCHE@TUM.DE

Abstract

Despite the existence of formal guarantees for learning-based control approaches, the relationship between data and control performance is still poorly understood. In this paper, we propose a Lyapunov-based measure for quantifying the impact of data on the certifiable control performance. By modeling unknown system dynamics through Gaussian processes, we can determine the interrelation between model uncertainty and satisfaction of stability conditions. This allows us to directly assess the impact of data on the provable stationary control performance, and thereby the value of the data for the closed-loop system performance. Our approach is applicable to a wide variety of unknown nonlinear systems that are to be controlled by a generic learning-based control law, and the results obtained in numerical simulations indicate the efficacy of the proposed measure.

Keywords: data-driven control, Gaussian processes, data-efficient learning, safe learning-based control

1. Introduction

Learning-based control is rapidly becoming an attractive alternative to traditional control approaches, particularly in settings with limited knowledge or prohibitive system complexity (Deisenroth et al., 2015; Chua et al., 2018). This has prompted a vast amount of research into the theoretical properties of control approaches based on nonparametric and probabilistic models obtained from supervised machine learning, yielding techniques that guarantee either safety or performance requirements (Aswani et al., 2013; Berkenkamp et al., 2016; Beckers et al., 2019; Fisac et al., 2019; Capone and Hirche, 2019; Gahlawat et al., 2020; Lederer et al., 2020).

While theoretical guarantees for control performance can be obtained in various settings, the direct relationship between collected data and control performance in learning-based control with nonparametric, probabilistic models is still poorly understood. In experimental design, the value of data is often quantified using information theoretical quantities, such as mutual information or entropy (Pukelsheim, 2006). Although these quantities have been used extensively to guide exploration and control strategies (Hennig and Schuler, 2012; Alpcan and Shames, 2015; Koller et al.,

This paper is the extended version of Lederer et al. (2021a). The official publication can be found at <http://proceedings.mlr.press/v144/lederer21a/lederer21a.pdf>.

2018; Capone et al., 2020), they do not provide direct insight into the impact of data on the provable control performance. In recent years, a handful of efforts has been carried out towards understanding the influence of collected data on control theoretic properties. Lederer et al. (2021b) have proposed a Lyapunov-based measure for quantifying the value of data points with respect to a specific control task based on Gaussian process priors. In a similar vein, Capone et al. (2021) have developed an algorithm to identify the most useful data points for successfully performing multiple control tasks. However, the technique presented in Lederer et al. (2021b) is only applicable to a restricted class of systems, and scalability is a challenge for the method of Capone et al. (2021).

In this work, we present a Lyapunov-based measure for quantifying the value of training data on the certifiable stationary performance of learning-based control of nonlinear systems. Based on Gaussian process models of unknown system dynamics, the model uncertainty is quantified and a necessary condition on the uncertainty is derived for ensuring stability of unknown closed-loop systems. This condition is transformed into a required training data density, thereby providing a measure for the impact of data on the provable stationary control performance. By considering a markedly richer class of systems, this work generalizes the results from Lederer et al. (2021b).

The remainder of this paper is structured as follows. In Section 2, we formally state the considered problem. In Section 3, we briefly discuss how Gaussian processes are employed to model the unknown system, and present some preliminary results. Afterwards, in Section 4, we derive the proposed information measure. We then discuss strategies for data selection, in Section 5, after which we present some experimental results, in Section 6. Finally, some concluding remarks are provided in Section 7.

2. Problem Statement

We consider a dynamical system

$$\dot{\mathbf{x}} = \mathbf{g}(\mathbf{z}) = \mathbf{A}\mathbf{f}(\mathbf{z}), \quad (1)$$

where $\mathbf{z} = [\mathbf{x}^T \ \mathbf{u}^T]^T \in \mathbb{R}^{d_z}$, $d_z = d_x + d_u$, is the concatenation of the state $\mathbf{x} \in \mathbb{X}$ and the input $\mathbf{u} \in \mathbb{U}$ for compact sets $\mathbb{X} \subset \mathbb{R}^{d_x}$ and $\mathbb{U} \subset \mathbb{R}^{d_u}$. We assume to know the matrix $\mathbf{A} \in \mathbb{R}^{d_x \times d_f}$, whereas the function $\mathbf{f} : \mathbb{R}^{d_x + d_u} \rightarrow \mathbb{R}^{d_f}$ is unknown. This system structure is very flexible and allows general multi-dimensional nonlinear systems ($\mathbf{A} = \mathbf{I}_{d_x}$, $d_f = d_x$), as well as correlation in the outputs, which can be found, e.g., in Euler-Lagrange systems due to the symmetry of the mass matrix (Cheng and Huang, 2015). We assume to have an approximate model $\hat{\mathbf{f}} : \mathbb{R}^{d_x + d_u} \rightarrow \mathbb{R}^{d_f}$ of the unknown function $\mathbf{f}(\cdot, \cdot)$, which is often available in practice, as well as measurement data. This yields the following formal assumption, which is discussed in detail in Umlauf and Hirche (2020).

Assumption 1 *A data set containing N measurement pairs*

$$\mathbb{D}_N = \left\{ \mathbf{z}^{(n)} := \begin{bmatrix} \mathbf{x}^{(n)} \\ \mathbf{u}^{(n)} \end{bmatrix}, \mathbf{y}^{(n)} = \mathbf{g}(\mathbf{z}^{(n)}) + \boldsymbol{\epsilon}^{(n)} \right\}_{n=1}^N \quad (2)$$

is available, where $\boldsymbol{\epsilon}^{(n)} \sim \mathcal{N}(0, \boldsymbol{\Sigma}_{\text{on}})$ is i.i.d. Gaussian noise with covariance matrix $\boldsymbol{\Sigma}_{\text{on}}$.

Additionally, we assume that the unknown functions $f_i(\cdot)$ are well behaved, as expressed in the following.

Assumption 2 *The unknown functions $f_i(\cdot)$ are Lipschitz continuous with Lipschitz constants L_{f_i} , $i = 1, \dots, d_f$.*

We assume to have a learning-based control law $\pi : \mathbb{X} \times (\mathbb{X} \times \mathbb{U} \times \mathbb{R}^{d_x})^N \times \mathbb{R}_{0,+} \rightarrow \mathbb{U}$ that causally maps a state \mathbf{x} to a control input \mathbf{u} depending on the previously observed training data \mathbb{D}_N and the current time $t \in \mathbb{R}_{0,+}$. The control law is designed to achieve a control task, e.g., stabilization with respect to a reference point, or tracking of a reference trajectory. The effectiveness of the control law with respect to this task is measured via a Lyapunov function¹ $V : \mathbb{X} \times \mathbb{R}_{0,+} \rightarrow \mathbb{R}_{0,+}$ and its temporal derivative $\dot{V}(\cdot, \cdot)$ along trajectories of the closed loop system defined through

$$\tilde{\mathbf{g}}(\mathbf{x}) = \mathbf{g} \left(\left[\begin{array}{c} \mathbf{x} \\ \pi(\mathbf{x}, \mathbf{x}^{(1)}, \mathbf{u}^{(1)}, \dots, t) \end{array} \right] \right), \quad (3)$$

which only depends on \mathbf{x} , as the control inputs are specified by the policy $\pi(\cdot, \mathbf{x}^{(1)}, \mathbf{u}^{(1)}, \dots, t)$. Since we do not know the function $\mathbf{f}(\cdot)$ but only have a training data set \mathbb{D}_N and an approximate model $\hat{\mathbf{f}}(\cdot)$, the warranted control performance, measured through the size of the region in which $\dot{V}(\mathbf{x}) \leq 0$ is not guaranteed, strongly depends on the training data. In such a setting, it is crucial to understand the interrelation between training data and control performance, e.g., for determining where additional training data should be acquired to increase control performance, or when a subset of the training data must be selected to reduce the computational complexity of learning. Therefore, we consider the problem of determining the impact of data on stability certificates for learning-based control.

3. Gaussian Process Regression

For determining the impact of data on stability certificates, we employ a Gaussian process (GP) formulation, such that we can relate the control performance to the model uncertainty. We first introduce the foundations of Gaussian process regression in Section 3.1, before we explain the extension to multiple outputs, in Section 3.2. Finally, we propose an output-decoupling formulation using linear models of coregionalization and propose a novel uniform error bound, in Section 3.3.

3.1. Single Output Gaussian Processes

A Gaussian process $\mathcal{GP}(\hat{f}(\cdot, \cdot), k(\cdot, \cdot))$, uniquely defined through a prior mean function $\hat{f} : \mathbb{R}^{d_z} \rightarrow \mathbb{R}$ and a covariance function $k : \mathbb{R}^{d_z} \times \mathbb{R}^{d_z} \rightarrow \mathbb{R}_{0,+}$, is a generalization of Gaussian distributions (Rasmussen and Williams, 2006). The prior mean function $\hat{f}(\cdot)$ is often used to include approximate models in the regression, whereas the covariance function describes prior assumptions on properties such as smoothness or periodicity. A commonly used covariance function, on which we also focus in the following analysis for clarity of exposition, is the squared exponential kernel

$$k(\mathbf{z}, \mathbf{z}') = s_f^2 \exp \left(-\frac{1}{2} (\mathbf{z} - \mathbf{z}')^T \mathbf{\Lambda}^{-1} (\mathbf{z} - \mathbf{z}') \right), \quad (4)$$

where $s_f^2 \in \mathbb{R}_{0,+}$ and $\mathbf{\Lambda} \in \mathbb{R}^{d_z \times d_z}$, $\mathbf{\Lambda} \succ \mathbf{0}$, denote the signal variance and length scales, respectively.

In order to perform regression with the GP, we consider a scalar system (1) with $d_f = 1$. Then, the joint prior distribution of training targets $\mathbf{t} = [y^{(1)} \ \dots \ y^{(N)}]^T$ and the unknown function value $f(\mathbf{z})$ for input \mathbf{z} is given by

$$\begin{bmatrix} \mathbf{t} \\ f(\mathbf{z}) \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} \hat{f}(\mathbf{Z}) \\ \hat{f}(\mathbf{z}) \end{bmatrix}, \begin{bmatrix} k(\mathbf{Z}, \mathbf{Z}) + \sigma_{\text{on}}^2 \mathbf{I}_N & k(\mathbf{Z}, \mathbf{z}) \\ k^T(\mathbf{Z}, \mathbf{z}) & k(\mathbf{z}, \mathbf{z}) \end{bmatrix} \right), \quad (5)$$

1. A Lyapunov function $V : \mathbb{X} \times \mathbb{R}_{0,+} \rightarrow \mathbb{R}_{0,+}$ is positive definite, i.e., $V(\mathbf{x}, t) \geq 0$ with equality if and only if $\mathbf{x} = 0$.

where we use the abbreviations $\hat{f}(\mathbf{Z}) \in \mathbb{R}^N$, $k(\mathbf{Z}, \mathbf{Z}) \in \mathbb{R}^{N \times N}$ and $k(\mathbf{Z}, \mathbf{z}) \in \mathbb{R}^N$ with elements defined as $\hat{f}_n(\mathbf{Z}) = \hat{f}(z^{(n)})$, $k_{n,n'}(\mathbf{Z}, \mathbf{Z}) = k(z^{(n)}, z^{(n')})$ and $k_n(\mathbf{Z}, \mathbf{z}) = k(z^{(n)}, z)$, $n, n' = 1, \dots, N$, respectively. By conditioning the GP on the training data, we obtain the posterior distribution

$$f(\mathbf{z}) | y^{(1)}, \dots, y^{(N)}, z^{(1)}, \dots, z^{(N)}, \mathbf{z} \sim \mathcal{N}(\mu(\mathbf{z}), \sigma^2(\mathbf{z})) \quad (6)$$

with posterior mean and variance

$$\mu(\mathbf{z}) = \hat{f}(\mathbf{z}) + k^T(\mathbf{Z}, \mathbf{z}) (k(\mathbf{Z}, \mathbf{Z}) + \sigma_{\text{on}}^2 \mathbf{I}_N)^{-1} (\mathbf{t} - \hat{f}(\mathbf{Z})) \quad (7)$$

$$\sigma^2(\mathbf{z}) = k(\mathbf{z}, \mathbf{z}) - k^T(\mathbf{Z}, \mathbf{z}) (k(\mathbf{Z}, \mathbf{Z}) + \sigma_{\text{on}}^2 \mathbf{I}_N)^{-1} k(\mathbf{Z}, \mathbf{z}). \quad (8)$$

3.2. Multiple-Output Gaussian Process Regression

In order to apply Gaussian processes to multiple-output regression problems, we can proceed analogously to the single output case. For illustrative purposes, we assume for now that $\mathbf{A} = \mathbf{I}_{d_x}$ and $d_f = d_x$ in (1), such that we have noisy measurements of the functions $f_i(\cdot)$, $i = 1, \dots, d_f$ in the data set \mathbb{D}_N . We start again with the prior GP distribution

$$\mathbf{f}(\cdot) \sim \mathcal{GP}(\hat{\mathbf{f}}(\cdot), \mathbf{K}(\cdot, \cdot)), \quad (9)$$

where we have to consider a vector-valued prior mean function $\hat{\mathbf{f}} : \mathbb{R}^{d_z} \rightarrow \mathbb{R}^{d_f}$ and a matrix kernel function $\mathbf{K} : \mathbb{R}^{d_z} \times \mathbb{R}^{d_z} \rightarrow \mathbb{R}_{0,+}^{d_f \times d_f}$, in which each element $k_{m,m'} : \mathbb{R}^{d_z} \times \mathbb{R}^{d_z} \rightarrow \mathbb{R}_{0,+}$ is a kernel.

By concatenating the training targets $\mathbf{y}_i^{(n)}$ in the vector $\mathbf{t}^T = [y_1^{(1)} \dots y_1^{(N)} \ y_2^{(1)} \dots y_{d_x}^{(N)}]$ and conditioning the joint distribution of \mathbf{t} and $\mathbf{f}(\mathbf{z})$ on the training data \mathbb{D}_N , analogously to (6) we obtain a multivariate Gaussian distribution with mean and covariance matrix

$$\boldsymbol{\mu}(\mathbf{z}) = \hat{\mathbf{f}}(\mathbf{z}) + \mathbf{K}^T(\mathbf{Z}, \mathbf{z}) (\mathbf{K}(\mathbf{Z}, \mathbf{Z}) + \boldsymbol{\Sigma}_{\text{on}} \otimes \mathbf{I}_N)^{-1} (\mathbf{t} - \hat{\mathbf{f}}(\mathbf{Z})) \quad (10)$$

$$\boldsymbol{\Sigma}(\mathbf{z}) = \mathbf{K}(\mathbf{z}, \mathbf{z}) - \mathbf{K}^T(\mathbf{Z}, \mathbf{z}) (\mathbf{K}(\mathbf{Z}, \mathbf{Z}) + \boldsymbol{\Sigma}_{\text{on}} \otimes \mathbf{I}_N)^{-1} \mathbf{K}(\mathbf{Z}, \mathbf{z}), \quad (11)$$

where we extend the shorthand notation from Section 3.1 using

$$\mathbf{K}(\mathbf{Z}, \mathbf{Z}) = \begin{bmatrix} k_{1,1}(\mathbf{Z}, \mathbf{Z}) & \dots & k_{1,d_f}(\mathbf{Z}, \mathbf{Z}) \\ \vdots & \ddots & \vdots \\ k_{d_f,1}(\mathbf{Z}, \mathbf{Z}) & \dots & k_{d_f,d_f}(\mathbf{Z}, \mathbf{Z}) \end{bmatrix} \quad \mathbf{K}(\mathbf{Z}, \mathbf{z}) = \begin{bmatrix} k_{1,1}(\mathbf{Z}, \mathbf{z}) & \dots & k_{1,d_f}(\mathbf{Z}, \mathbf{z}) \\ \vdots & \ddots & \vdots \\ k_{d_f,1}(\mathbf{Z}, \mathbf{z}) & \dots & k_{d_f,d_f}(\mathbf{Z}, \mathbf{z}) \end{bmatrix} \quad (12)$$

$$\hat{\mathbf{f}}(\mathbf{Z}) = \begin{bmatrix} \hat{f}_1(\mathbf{Z}) & \dots & \hat{f}_{d_f}(\mathbf{Z}) \end{bmatrix}^T.$$

3.3. Output Decoupling through Linear Models of Coregionalization

While various positive definite kernels are known for scalar regression, positive definiteness is a major challenge in the multiple-output approach presented in Section 3.2 since it is not sufficient that each entry of $\mathbf{K}(\cdot, \cdot)$ is a covariance function. However, in the following we show that knowledge of the output correlation structure in the form of a matrix \mathbf{A} allows to define proper kernel matrix functions via scalar covariance functions $k_i(\cdot, \cdot)$. For this, we require the following assumption.

Assumption 3 *Prior knowledge about the functions $f_i(\cdot)$ is expressed through independent prior GP distributions with scalar kernels $k_i(\cdot, \cdot)$, $i = 1, \dots, d_f$, i.e.,*

$$\mathbf{f}(\cdot) \sim \mathcal{GP}(\hat{\mathbf{f}}(\cdot), \text{diag}([k_1(\cdot, \cdot) \ \dots \ k_{d_f}(\cdot, \cdot)])) \quad (13)$$

This assumption is not restrictive, since correlation in the training targets can be modeled through the matrix \mathbf{A} , and it is frequently used in the case where $\mathbf{A} = \mathbf{I}_{d_x}$ holds (Berkenkamp and Schoellig, 2015; Koller et al., 2018; Beckers et al., 2019; Hewing et al., 2020).

Due to Assumption 3, it directly follows that

$$\mathbf{g}(\cdot) \sim \mathcal{GP} \left(\mathbf{A} \hat{\mathbf{f}}(\cdot), \mathbf{A} \text{diag} \left([k_1(\cdot, \cdot) \ \cdots \ k_{d_f}(\cdot, \cdot)] \right) \mathbf{A}^T \right), \quad (14)$$

such that we can intuitively define a kernel matrix function through

$$\mathbf{K}(\mathbf{z}, \mathbf{z}') = \mathbf{A} \text{diag} \left([k_1(\mathbf{z}, \mathbf{z}') \ \cdots \ k_{d_f}(\mathbf{z}, \mathbf{z}')] \right) \mathbf{A}^T. \quad (15)$$

It is trivial to show that this kernel parameterization is a special case of a linear model of coregion-alization (Álvarez et al., 2011), such that we can immediately extend the approach in Duvenaud (2014) to recover models for the individual functions $f_i(\cdot)$, as shown in the following lemma².

Lemma 1 *Consider a nonlinear system (1) with matrix $\mathbf{A} = [\mathbf{a}_1 \ \cdots \ \mathbf{a}_{d_f}]$ composed of column vectors \mathbf{a}_i , for which a training data set \mathbb{D}_N and prior distributions that satisfy Assumptions 1 and 3, respectively, are given. Then, the posterior distributions are given by*

$$f_i(\mathbf{z}) | \mathbb{D} \sim \mathcal{N}(\mu_i(\mathbf{z}), \sigma_i(\mathbf{z})), \quad (16)$$

where

$$\mu_i(\mathbf{z}) = \hat{f}_i(\mathbf{z}) + (k_i^T(\mathbf{Z}, \mathbf{z}) \otimes \mathbf{a}_i^T) (\mathbf{K}(\mathbf{Z}, \mathbf{Z}) + \Sigma_{\text{on}} \otimes \mathbf{I}_N)^{-1} (\mathbf{t} - \hat{\mathbf{f}}(\mathbf{Z})) \quad (17)$$

$$\sigma_i^2(\mathbf{z}) = k_i(\mathbf{z}, \mathbf{z}) - (k_i^T(\mathbf{Z}, \mathbf{z}) \otimes \mathbf{a}_i^T) (\mathbf{K}(\mathbf{Z}, \mathbf{Z}) + \Sigma_{\text{on}} \otimes \mathbf{I}_N)^{-1} (k_i(\mathbf{Z}, \mathbf{z}) \otimes \mathbf{a}_i). \quad (18)$$

A crucial benefit of this decoupling of the outputs is that it allows the application of scalar analysis methods to uniformly bound the regression error on the input domain $\mathbb{X} \times \mathbb{U}$ as proposed in Lederer et al. (2019a). This is formalized in the following theorem.

Theorem 2 *Consider a nonlinear system (1), a training data set \mathbb{D}_N , and prior distributions satisfying Assumptions 1-3, respectively. For any $\delta \in (0, 1)$, $\tau \in \mathbb{R}_+$, and $i = 1, \dots, d_f$, it holds that*

$$P \left(|f_i(\mathbf{z}) - \mu_i(\mathbf{z})| \leq \sqrt{\beta(\delta, \tau)} \sigma_i(\mathbf{z}) + \gamma_i(\delta, \tau) \quad \forall \mathbf{z} \in \mathbb{X} \times \mathbb{U} \right) \geq 1 - \delta, \quad (19)$$

where

$$\beta(\delta, \tau) = 2d_x \log \left(1 + \frac{r_0}{\tau} \right) - \log(\delta), \quad \gamma_i(\delta, \tau) = (L_{\mu_i} + L_{f_i})\tau + \sqrt{\beta(\delta, \tau)} L_{\sigma_i^2} \tau. \quad (20)$$

Here, L_{μ_i} and $L_{\sigma_i^2}$ are the Lipschitz constants of the mean and variance functions, respectively, and $r_0 = \max_{\mathbf{z}, \mathbf{z}' \in \mathbb{X} \times \mathbb{U}} \|\mathbf{z} - \mathbf{z}'\|$.

This theorem is a generalization of Lederer et al. (2021b, Lemma 1) and many properties directly transfer. Small error bounds can be achieved through small GP standard deviations $\sigma_i(\mathbf{z})$, which corresponds to high data densities. This resembles well-known relationships from scattered data approximation (Wendland, 2004) and Bayesian optimization (Srinivas et al., 2012). The dependence of the uniform error bound (19) on the constants $\gamma_i(\delta, \tau)$ does not affect this behavior, since they can be chosen arbitrarily small, and convergence to 0 can be shown under weak assumptions on $\sigma_i(\mathbf{z})$ (Lederer et al., 2019a). In general, the constant τ trades-off the effect of the data independent terms $\gamma_i(\delta, \tau)$ and the posterior standard deviations $\sigma_i(\mathbf{z})$ on the error bound. Therefore, τ should be chosen such that the uncertainty dependence of the bound dominates, i.e., $\sqrt{\beta(\delta, \tau)} \sigma_i(\mathbf{z}) \gg \gamma_i(\delta, \tau)$.

2. Proofs for all theoretical results can be found in the appendix.

Remark 3 *Theorem 2 admits the counterintuitive behavior that adding training samples can lead to a locally higher uniform error bound. This is due to the fact that adding data in some regions can increase the Lipschitz constants of $\mu(\cdot)$ and $\sigma(\cdot)$ (Lederer et al., 2019a), and thereby increase the uniform error bound in other regions. Note that a similar argument holds for uniform error bounds based on RKHS theory (Srinivas et al., 2012; Chowdhury and Gopalan, 2017).*

4. Control-Based Information Measures

While the uniform error bound in Theorem 2 establishes a connection between the training data distribution, represented by the posterior GP variance, and the regression performance, it is ignorant of the control task. In order to measure the importance of data for control performance, we consider the Lyapunov stability conditions (Khalil, 2002) for the closed loop system, which require a negative derivative of the Lyapunov function $V(\cdot, \cdot)$, i.e.,

$$\dot{V}(\mathbf{x}, t) = (\nabla_{\mathbf{x}} V(\mathbf{x}, t))^T \mathbf{A} \tilde{\mathbf{f}}(\mathbf{x}) + \frac{\partial}{\partial t} V(\mathbf{x}, t), \quad (21)$$

where we employ the shorthand notation $\tilde{\mathbf{f}}(\mathbf{x}) = \mathbf{f}([\mathbf{x}^T \ \boldsymbol{\pi}^T(\mathbf{x}, \mathbf{x}^{(1)}, \mathbf{u}^{(1)}, \dots, t)]^T)$, which is used analogously for the GP mean $\tilde{\boldsymbol{\mu}}(\cdot)$ and variance $\tilde{\boldsymbol{\sigma}}^2(\cdot)$. Although the function $\tilde{\mathbf{f}}(\cdot)$ is unknown, we can bound the Lyapunov function derivative based on the uniform error bound (19), which yields

$$\dot{V}(\mathbf{x}, t) \leq \dot{V}_{\text{nom}}(\mathbf{x}, t) + \dot{V}_{\boldsymbol{\sigma}}(\mathbf{x}, t), \quad (22)$$

where we decouple the bound into the nominal component based on the GP mean

$$\dot{V}_{\text{nom}}(\mathbf{x}, t) = (\nabla V(\mathbf{x}, t))^T \mathbf{A} \tilde{\boldsymbol{\mu}}(\mathbf{x}) + \frac{\partial}{\partial t} V(\mathbf{x}, t), \quad (23)$$

and an uncertain component depending on the GP standard deviation

$$\dot{V}_{\boldsymbol{\sigma}}(\mathbf{x}, t) = \left[\left| (\nabla_{\mathbf{x}} V(\mathbf{x}, t))^T \mathbf{a}_1 \right| \ \cdots \ \left| (\nabla_{\mathbf{x}} V(\mathbf{x}, t))^T \mathbf{a}_{d_f} \right| \right] \left(\sqrt{\beta(\delta, \tau)} \tilde{\boldsymbol{\sigma}}(\mathbf{x}) + \gamma(\delta, \tau) \right), \quad (24)$$

with $\gamma(\delta, \tau) = [\gamma_1(\delta, \tau) \ \cdots \ \gamma_{d_f}(\delta, \tau)]^T$. Since the uncertain derivative component $\dot{V}_{\boldsymbol{\sigma}}(\cdot, \cdot)$ is non-negative, a positive nominal derivative component $\dot{V}_{\text{nom}}(\cdot, \cdot)$ directly implies that (22) violates the Lyapunov stability conditions regardless of the GP posterior variance. Therefore, we assume $\dot{V}_{\text{nom}}(\mathbf{x}, t) < 0$ in the following, essentially requiring that the control law can stabilize the dynamical system defined by $\boldsymbol{\mu}(\cdot)$. As a result, stability of the closed-loop system depends on the magnitude of the uncertain Lyapunov function derivative $\dot{V}_{\boldsymbol{\sigma}}(\cdot, \cdot)$, which is strongly influenced by posterior GP standard deviations $\tilde{\boldsymbol{\sigma}}(\cdot)$. Although this establishes a direct relationship between the training data density and the control task, the dependency of $\tilde{\boldsymbol{\sigma}}(\cdot)$ on training samples is highly nonlinear and the computation of $\tilde{\boldsymbol{\sigma}}(\cdot)$ is computationally expensive. In order to mitigate these issues, we introduce the weighted M -fill distances, in analogy to the M -fill distance proposed in Lederer et al. (2021b).

Definition 4 *The weighted M -fill distance $\phi_i(\mathbf{x}, \mathbb{D}_N)$ for function $f_i(\cdot)$, $i = 1 \dots, d_f$, at a point \mathbf{x} is defined as the minimum radius φ of a ball with center $\tilde{\mathbf{z}} = [\mathbf{x}^T \ \boldsymbol{\pi}^T(\mathbf{x})]^T$, such that the ball contains M samples $\mathbf{z}^{(n)}$, i.e.,*

$$\tilde{\phi}_i(\mathbf{x}, \mathbb{D}_N) = \min_{\phi \in \mathbb{R}_{+,0}} \varphi \quad (25a)$$

$$\text{such that } \left| \left\{ \mathbf{z}^{(n)} \in \mathbb{D}_N : \left(\tilde{\mathbf{z}} - \mathbf{z}^{(n)} \right)^T \boldsymbol{\Lambda}_i^{-1} \left(\tilde{\mathbf{z}} - \mathbf{z}^{(n)} \right) \leq \varphi^2 \right\} \right| \geq M, \quad (25b)$$

where $|\cdot|$ denotes the cardinality of the set and we use the abbreviation $\boldsymbol{\pi}(\mathbf{x}) = \boldsymbol{\pi}(\mathbf{x}, \mathbf{x}^{(1)}, \mathbf{u}^{(1)}, \dots, t)$. The weighted M -fill distances measure the distance from a test point $\tilde{\mathbf{z}} = [\mathbf{x}^T \boldsymbol{\pi}^T(\mathbf{x})]^T$ to the M closest training samples in the Mahalonobis distance metric induced by the length scales $\boldsymbol{\Lambda}_i$ of the squared exponential kernels (4). By choosing a small number $M \ll N$, only training points in the proximity of the test point $\tilde{\mathbf{z}} = [\mathbf{x}^T \boldsymbol{\pi}^T(\mathbf{x})]^T$ are relevant for the weighted M -fill distance $\tilde{\phi}_i(\mathbf{x}, \mathbb{D}_N)$. This allows us to measure the local data density in a flexible way, where high training data densities are indicated by low values of $\tilde{\phi}_i(\mathbf{x}, \mathbb{D}_N)$. Moreover, it is possible to bound the posterior GP variances $\tilde{\sigma}^2(\mathbf{x})$ in terms of the weighted M -fill distances $\tilde{\phi}_i(\mathbf{x}, \mathbb{D}_N)$ ³. We exploit this property in the following theorem to derive conditions that guarantee that the summands of the uncertain Lyapunov derivative $\dot{V}_\sigma(\cdot, \cdot)$ are upper bounded by functions $\xi_i : \mathbb{R}^{d_z} \times \mathbb{R}_{0,+} \rightarrow \mathbb{R}_{0,+}$, $i = 1, \dots, d_f$. For suitably chosen functions $\xi_i(\cdot, \cdot)$, the satisfaction of these bounds implies stability of the closed-loop system.

Theorem 5 Choose τ such that $\sqrt{\beta(\delta, \tau)}\tilde{\sigma}_i(\mathbf{x}) > \gamma_i(\delta, \tau)$ holds for all $\mathbf{x} \in \mathbb{X}$ and $\xi_i : \mathbb{R}^{d_z} \times \mathbb{R}_{0,+} \rightarrow \mathbb{R}_{0,+}$ such that

$$\dot{V}_{\sigma_{i,0}}^2(\mathbf{x}, t) = 4\beta(\delta, \tau)s_{f_i}^2 \left| (\nabla_{\mathbf{x}} V(\mathbf{x}, t))^T \mathbf{a}_i \right|^2 > \xi_i^2(\mathbf{x}, t). \quad (26)$$

If the M -fill distance $\tilde{\phi}_i(\cdot, \mathbb{D}_N)$ satisfies $\tilde{\phi}_i^2(\mathbf{x}, \mathbb{D}_N) \leq \bar{\phi}_i^2(\mathbf{x}, t) + \theta_i^2$ for all $\mathbf{x} \in \mathbb{X}$, where

$$\bar{\phi}_i^2(\mathbf{x}, t) = -\log \left(1 - \frac{\xi_i^2(\mathbf{x}, t)}{\dot{V}_{\sigma_{i,0}}^2(\mathbf{x}, t)} \right) \quad (27)$$

$$\theta_i^2 = \log \left(s_{f_i}^2 \|\mathbf{a}_i\|_2^2 \right) - \log \left(\max_{m=1, \dots, d_x} \sum_{n=1}^{d_f} |a_{m,n}| \|\mathbf{a}_n\|_1 s_{f_n}^2 + \frac{\lambda_{\max}(\Sigma_{\text{on}})}{M} \right), \quad (28)$$

then, with probability of at least $1 - \delta$, it holds for all $\mathbf{x} \in \mathbb{X}$ that

$$\dot{V}_{\sigma_i}(\mathbf{x}, t) = \left(\sqrt{\beta(\delta, \tau)}\tilde{\sigma}_i(\mathbf{x}) + \gamma_i(\delta, \tau) \right) \left| (\nabla_{\mathbf{x}} V(\mathbf{x}, t))^T \mathbf{a}_i \right| \leq \xi_i(\mathbf{x}, t). \quad (29)$$

Condition (26) is necessary to ensure the existence of the logarithm in (27), but it is not restrictive since $\xi_i^2(\cdot, \cdot)$ are upper bounds. Hence, we can simply tighten the bounds until $\xi_i^2(\cdot, \cdot)$ satisfies condition (26). The expressions (27) and (28) have different roles. The values θ_i express the difficulty of recovering the functions $f_i(\cdot)$ from the noisy measurements of $\mathbf{A}\mathbf{f}(\cdot)$, which in turn depends on the signal variances $s_{f_i}^2$ of the independent covariance functions $k_i(\cdot, \cdot)$ and the magnitude of the elements of \mathbf{A} . In contrast, $\bar{\phi}_i^2(\cdot, \cdot)$ captures the dependency on the control task. The numerator in (27) corresponds to a summand of the uncertain Lyapunov derivative component (24) without any training data as $s_{f_i}^2$ corresponds to the prior GP variance. Therefore, $\bar{\phi}_i^2(\mathbf{x}, t)$ goes to ∞ when the prior uncertain Lyapunov component converges to the bound $\xi_i^2(\mathbf{x}, t)$. This intuitively reflects the fact that no data are required when the GP prior is already sufficient to guarantee (29).

As mentioned previously, Theorem 5 can be used to analyze the stability of the closed loop system. More specifically, stability is guaranteed if the negated nominal Lyapunov derivative $\dot{V}_{\text{nom}}(\cdot, \cdot)$ is larger than the uncertain component $\dot{V}_\sigma(\cdot, \cdot)$. This can be trivially checked with Theorem 5 by defining $\xi_j(\cdot, \cdot)$ such that $\sum_{i=1}^{d_f} \xi_i(\mathbf{x}, t) \leq |\dot{V}_{\text{nom}}(\mathbf{x}, t)|$. A natural choice satisfying this condition

3. A bound for the posterior variance in terms of the M -fill distance is derived in the appendix.

together with constraint (26) is given by

$$\xi_i(\mathbf{x}, t) = \min \left\{ -\frac{\dot{V}_{\text{nom}}(\mathbf{x}, t) \|\mathbf{a}_i\|_1}{\sum_{j=1}^{d_f} \|\mathbf{a}_j\|_1}, \dot{V}_{\sigma_{i,0}}(\mathbf{x}, t) - \nu \right\}, \quad (30)$$

where $\nu \in \mathbb{R}_+$ is an arbitrarily small constant. Furthermore, convergence rates can be examined in a similar way. For example, an exponential rate of convergence is achieved by guaranteeing that a requirement similar to (30) is satisfied, where $\dot{V}_{\text{nom}}(\mathbf{x}, t)$ is replaced by $\dot{V}_{\text{nom}}(\mathbf{x}, t) + V(\mathbf{x}, t)$.

Due to the intuitive interpretation of Theorem 5, we propose to use it as the basis for a measure of the importance of training data for control. This naturally leads to the definition of the ρ -gap.

Definition 6 *The ρ -gap is defined as*

$$\rho(\mathbf{x}, t, \mathbb{D}_N) = \sum_{i=1}^{d_f} \max\{0, \phi_i^2(\mathbf{x}, \mathbb{D}_N) - \bar{\phi}_i^2(\mathbf{x}, t) - \theta_i^2\}. \quad (31)$$

Essentially, the ρ -gap measures the discrepancy between the required data density, which is expressed through $\bar{\phi}_i^2(\mathbf{x}, t) + \theta_i^2$ and depends on the desired bounds $\xi_i(\cdot, \cdot)$, the Lyapunov derivative and the signal standard deviations $s_{f_i}^2$, and the actual data density represented by the M -fill distances $\phi_i^2(\mathbf{x}, \mathbb{D}_N)$, which are independent of the control problem and only depend on the available data.

5. Data Selection Strategies

Based on the information measure proposed in Section 4, the data set can be preprocessed to contain only the most relevant information for the given control task. This becomes particularly important in scenarios where the prediction of the GP model must be performed under tight real-time constraints. Even with a precomputation of the matrix inverse in (17) (which takes $\mathcal{O}(N^3)$), the N kernel evaluations (for $k_m^T(\mathbf{Z}, \mathbf{z})$) and the corresponding multiplications ($\mathcal{O}(N)$ for the posterior mean, $\mathcal{O}(N^2)$ for the posterior variance) must still be performed online. In practice, this imposes an upper bound for the number of points that can be considered by the model. We formulate the resulting computational constraint independent of the hardware and specific real-time limit as follows.

Assumption 4 *The computational constraints allow a maximum of \bar{N} data points to be considered by the GP regression model (17).*

For the non-trivial case $N > \bar{N}$, this makes a selection of an active data set $\mathbb{D}_{\bar{N}} \subset \mathbb{D}_N$ necessary. Such a data selection has been considered for general function learning in Krause et al. (2008), and specifically for control tasks in Umlauf et al. (2020). But both employ entropy-based criteria, which only aims to optimize the precision of the model but does not consider the closed-loop control performance. Therefore, we utilize the ρ -gap as the measure for the control performance of a data set to find the optimal active data set⁴

$$\mathfrak{I}^* = \arg \min_{\mathfrak{I} \in \mathbb{P}_{\bar{N}}^{\mathbb{I}}} \max_{t \in \mathbb{T}, \mathbf{x} \in \mathbb{X}} \rho(\mathbf{x}, t, \mathfrak{I}), \quad (32)$$

where $\mathbb{T} = [0, T[$ with initial time $t_0 \in \mathbb{R}_{0,+}$, $t_0 < \infty$ and (possibly infinite) final time $T \in \mathbb{R} \cup \infty$. The optimal active data set is then given by $\mathbb{D}_{\bar{N}} = \{\mathbf{z}^{(\mathfrak{I}(i))}, \mathbf{y}^{(\mathfrak{I}(i))}\}_{i=1}^{\bar{N}}$. If the desired trajectory

4. To simplify notation, we introduce the index set $\mathbb{I} = \{1, \dots, N\}$ and all possible subsets with size \bar{N} , denoted as $\mathbb{P}_{\bar{N}}^{\mathbb{I}}$. $\mathfrak{I}(i)$ denotes the i -th entry of the index set \mathfrak{I} .

Algorithm 1: Greedy optimization for optimal subset selection

Input: $\mathbb{D}_N, \rho(\cdot, \cdot), \mathbb{T}_0, \dots, \mathbb{T}_S$
Output: $\mathbb{D}_{\bar{N}}^{\mathbb{T}_0}, \dots, \mathbb{D}_{\bar{N}}^{\mathbb{T}_S}$
for $\mathbb{T} = \mathbb{T}_0, \dots, \mathbb{T}_S$ **do**
 $\mathbb{D}_{\bar{N}}^{\mathbb{T}} \leftarrow \emptyset, \mathbb{I} = \{1, \dots, N\}$
 for $n = 0, \dots, \bar{N}$ **do**
 $i^*, t^* \leftarrow \arg \max_{i \in \mathbb{I}, t \in \mathbb{T}} \rho(\mathbf{x}^{(i)}, t, \mathbb{D}_{\bar{N}}^{\mathbb{T}})$
 $\mathbb{D}_{\bar{N}}^{\mathbb{T}} \leftarrow \mathbb{D}_{\bar{N}}^{\mathbb{T}} \cup \{\mathbf{z}^{(i^*)}, \mathbf{y}^{(i^*)}\}$
 $\mathbb{I} \leftarrow \mathbb{I} \setminus \{i^*\}$
 end
end

has a wide spread or \bar{N} is small, then the selected subset might not lead to a satisfactory control performance. For such a case, we can partition the task in $S \in \mathbb{N}$ time intervals $\mathbb{T}_0 = [0, t_1[, \mathbb{T}_1 = [t_1, t_2[, \dots, \mathbb{T}_S = [t_S, T[$ and compute the corresponding optimal subsets $\mathbb{D}_{\bar{N}}^{\mathbb{T}_0}, \dots, \mathbb{D}_{\bar{N}}^{\mathbb{T}_S}$.

Due to its mixed nature (combinatorial in \mathbb{T} , continuous in t and \mathbf{x}), the optimization problem (32) is not trivial to solve. However, the optimization can be performed offline, assuming that sufficient memory capacity is available to store all precomputed subsets. Furthermore, in the field of function learning it has been shown that greedy algorithms can show near-optimal behavior (Krause et al., 2008). Therefore, we propose the greedy data selection procedure shown in Algorithm 1.

6. Numerical Evaluation

In order to evaluate the proposed importance measure, we consider the nonlinear system

$$\dot{\mathbf{x}} = \mathbf{x} + \frac{1}{1 + \exp(-2x_1)} \begin{bmatrix} 1 \\ -1 \end{bmatrix} + 0.5 \begin{bmatrix} \sin(\pi x_2) \\ \cos(\pi x_1) \end{bmatrix} + \mathbf{u}, \quad (33)$$

which is a slight modification of the example proposed in (Umlauft et al., 2018). We assume a prior model $\hat{\mathbf{f}}(\mathbf{z}) = \mathbf{x} + \mathbf{u}$, and define the kernel matrix using

$$\mathbf{A} = \begin{bmatrix} 1 & 0 \\ -1 & 1 \end{bmatrix} \quad (34)$$

and squared exponential kernels $k_1(\mathbf{x}, \mathbf{x}'), k_2(x_1, x_1')$. This ensures that the correlation between the outputs caused by the second summand in (33) is properly modeled. We employ a control law

$$\boldsymbol{\pi}(\mathbf{x}, t) = -(\boldsymbol{\mu}(\mathbf{x}) + K(\mathbf{x} - \mathbf{x}_{\text{ref}}(t)) - \dot{\mathbf{x}}_{\text{ref}}(t)),$$

with gain $K = 15$ and references $\mathbf{x}_{\text{ref}}(t) = [c_1 \sin(t) \ c_2 \cos(t)]^T$ with randomly drawn $c_i \sim \mathcal{N}(0, 1)$. Nominal stability of the closed loop is shown using the Lyapunov function $V(\mathbf{x}, t) = (\mathbf{x} - \mathbf{x}_{\text{ref}}(t))^T (\mathbf{x} - \mathbf{x}_{\text{ref}}(t))$. The training set is generated by simulating the closed-loop system with prior mean $\boldsymbol{\mu}(\cdot) = \hat{\mathbf{f}}(\cdot)$ and sampling $N = 100$ data points during the interval $t \in [0, T]$ with $T = 10$. We divide the period of the reference trajectory into $S = 10$ equally long intervals \mathbb{T}_s and select subsets of cardinality $\bar{N} = 10$ for each interval using Algorithm 1 with $M = 1$ and $\xi_i(\cdot)$ as defined in (30).

Snapshots of the selected subsets and the resulting system trajectories are illustrated in Fig. 1. It can be clearly seen that the training samples are chosen close to the reference in the considered time intervals \mathbb{T}_s . This is because the feedback $K(\mathbf{x} - \mathbf{x}_{\text{ref}})$ ensures that the Lyapunov stability condition




Figure 1: Snapshots of the desired and actual trajectory of the dynamical system. Data selected by Algorithm 1 lies close to the future reference and ensures a high tracking accuracy.

is satisfied far away from the reference regardless of training data. Moreover, the data density grows as the distance to the reference decreases due to the vanishing effect of the feedback in its proximity.

We evaluate our technique by carrying out 100 control law roll-outs with randomly drawn trajectory parameters c_i . Furthermore, we compare the results with the performance of the full data set, as well as data selected with a greedy maximization of the mutual information with respect to a uniform grid over $[-1.5, 1.5]^2$, and maximization of the mutual information with respect to the considered trajectory interval (Umlauft et al., 2020). The results are depicted in Table 1. A reduction in the average computation time by a factor of approximately 10 for all subset selection methods can be observed, which is a straightforward consequence of the linear complexity of predictions (Rasmussen and Williams, 2006). Moreover, the steady state mean squared tracking error is the lowest for the subset selected based on the ρ -gap, and even smaller than for the full data set. The approach from (Umlauft et al., 2020) exhibits a similar performance. Although it might seem unintuitive that reducing the number of samples benefits control performance, this is not excluded by Theorem 2, as discussed in Remark 3. This underlines the importance of selecting training data for learning-based control.

7. Conclusion

We presented the ρ -gap, a measure that quantifies the value of data for a broad class of control tasks. The proposed quantity is used to identify the optimal data set for control tasks under computational constraints. Simulations demonstrate that the data subsets selected using the presented measure are highly correlated with the control task and can even be beneficial for control performance.

Table 1: Tracking errors and GP prediction times resulting from different subset selection criteria. The ρ -gap significantly outperforms existing methods regarding control performance.

critierion	full data set	mutual information w.r.t. uniform grid	mutual information w.r.t. reference	ρ -gap
steady-state MSE ($\cdot 10^{-3}$)	1.15	1.32	0.38	0.16
prediction time (μs)	437	45.0	45.0	45.0

Acknowledgments

This work was supported by the European Research Council Consolidator Grant Safe data-driven control for human-centric systems (CO-MAN) under grant agreement number 864686. Armin Lederer gratefully acknowledges financial support from the German Academic Scholarship Foundation.

References

- Tansu Alpcan and Iman Shames. An Information-Based Learning Approach to Dual Control. *IEEE Transactions on Neural Networks and Learning Systems*, 26(11):2736–2748, 2015.
- Mauricio A. Álvarez, Lorenzo Rosasco, and Neil D. Lawrence. Kernels for Vector-Valued Functions: A Review. *Foundations and Trends in Machine Learning*, 4(3):195–266, 2011.
- Anil Aswani, Humberto Gonzalez, S. Shankar Sastry, and Claire Tomlin. Provably Safe and Robust Learning-Based Model Predictive Control. *Automatica*, 49(5):1216–1226, 2013.
- Thomas Beckers, Dana Kulić, and Sandra Hirche. Stable Gaussian Process based Tracking Control of Euler-Lagrange Systems. *Automatica*, 103(23):390–397, 2019.
- Felix Berkenkamp and Angela P. Schoellig. Safe and Robust Learning Control with Gaussian Processes. In *Proceedings of the European Control Conference*, pages 2496–2501, 2015.
- Felix Berkenkamp, Angela P. Schoellig, and Andreas Krause. Safe Controller Optimization for Quadrotors with Gaussian Processes. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 491–496, 2016.
- 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, Gerrit Noske, Jonas Umlauft, Thomas Beckers, Armin Lederer, and Sandra Hirche. Localized Active Learning of Gaussian Process State Space Models. In *Learning for Dynamics & Control*, pages 490–499, 2020.
- Alexandre Capone, Armin Lederer, Jonas Umlauft, and Sandra Hirche. Data Selection for Multi-Task Learning under Dynamic Constraints. *IEEE Control Systems Letters*, 5(3):959–964, 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(10):3247–3258, 2015.
- Sayak Ray Chowdhury and Aditya Gopalan. On Kernelized Multi-armed Bandits. In *Proceedings of the International Conference on Machine Learning*, pages 844–853, 2017.
- 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.
- 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, 2015.

- David Kristjanson Duvenaud. *Automatic Model Construction with Gaussian Processes*. PhD thesis, University of Cambridge, 2014.
- Jaime F. Fisac, Anayo K. Akametalu, Melanie N. Zeilinger, Shahab Kaynama, Jeremy Gillula, and Claire J. Tomlin. A General Safety Framework for Learning-Based Control in Uncertain Robotic Systems. *IEEE Transactions on Automatic Control*, 64(7):2737–2752, 2019.
- Aditya Gahlawat, Pan Zhao, Andrew Patterson, Naira Hovakimyan, and Evangelos A. Theodorou. $\mathcal{RL}_1\text{-GP}$: \mathcal{L}_1 Adaptive Control with Bayesian Learning. In *Learning for Dynamics & Control*, pages 826–837, 2020.
- Philipp Hennig and Christian J. Schuler. Entropy Search for Information-Efficient Global Optimization. *Journal of Machine Learning Research*, 13:1809–1837, 2012.
- Lukas Hewing, Elena Arcari, Lukas P. Fröhlich, and Melanie N. Zeilinger. On Simulation and Trajectory Prediction with Gaussian Process Dynamics. In *Learning for Dynamics & Control*, pages 424–434, 2020.
- Hassan K. Khalil. *Nonlinear Systems*. Prentice-Hall, Upper Saddle River, NJ, third edition, 2002.
- Torsten Koller, Felix Berkenkamp, Matteo Turchetta, and Andreas Krause. Learning-based Model Predictive Control for Safe Exploration. In *Proceedings of the IEEE Conference on Decision and Control*, pages 6059–6066, 2018.
- A Krause, Aarti Singh, and C Guestrin. Near-optimal Sensor Placements in Gaussian Processes: Theory, Efficient Algorithms and Empirical Studies. *Journal of Machine Learning Research*, 9: 235–284, 2008.
- Armin Lederer, Jonas Umlauf, and Sandra Hirche. Uniform Error Bounds for Gaussian Process Regression with Application to Safe Control. In *Advances in Neural Information Processing Systems*, pages 659–669, 2019a.
- Armin Lederer, Jonas Umlauf, and Sandra Hirche. Posterior Variance Analysis of Gaussian Processes with Application to Average Learning Curves. 2019b. URL <http://arxiv.org/abs/1906.01404>.
- Armin Lederer, Alexandre Capone, and Sandra Hirche. Parameter Optimization for Learning-based Control of Control-Affine Systems. In *Learning for Dynamics and Control*, volume 120, pages 465–475, 2020.
- Armin Lederer, Alexandre Capone, Thomas Beckers, Jonas Umlauf, and Sandra Hirche. The Impact of Data on the Stability of Learning-Based Control. In *Learning for Dynamics and Control*, volume 144, pages 623–635, 2021a. URL <http://proceedings.mlr.press/v144/lederer21a.html>.
- Armin Lederer, Alexandre Capone, Jonas Umlauf, and Sandra Hirche. How Training Data Impacts Performance in Learning-Based Control. *IEEE Control Systems Letters*, 5(3):905–910, 2021b.
- Friedrich Pukelsheim. *Optimal Design of Experiments*. Society for Industrial and Applied Mathematics, 2006.

Carl E. Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. The MIT Press, Cambridge, MA, 2006.

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.

Jonas Umlauf and Sandra Hirche. Feedback Linearization based on Gaussian Processes with event-triggered Online Learning. *IEEE Transactions on Automatic Control*, 2020.

Jonas Umlauf, Lukas Pöhler, and Sandra Hirche. An Uncertainty-Based Control Lyapunov Approach for Control-Affine Systems Modeled by Gaussian Process. *IEEE Control Systems Letters*, 2(3):483–488, 2018.

Jonas Umlauf, Thomas Beckers, Alexandre Capone, Armin Lederer, and Sandra Hirche. Smart Forgetting for Safe Online Learning with Gaussian Processes. In *Learning for Dynamics & Control*, pages 160–169, 2020.

Holger Wendland. *Scattered Data Approximation*. Cambridge University Press, 2004.

Appendix A. Component-wise Uniform Error Bounds for Multiple-Output Gaussian Process Regression

Proof of Lemma 1 It can be easily checked that we can express the kernel matrix function as

$$\mathbf{K}(\mathbf{z}, \mathbf{z}') = \sum_{i=1}^{d_f} \mathbf{B}_i k_i(\mathbf{z}, \mathbf{z}'), \quad (35)$$

where

$$\mathbf{B}_i = \begin{bmatrix} a_{1,i} \\ \vdots \\ a_{d_x,i} \end{bmatrix} \underbrace{\begin{bmatrix} a_{1,i} & \cdots & a_{d_x,i} \end{bmatrix}}_{\mathbf{a}_i^T}. \quad (36)$$

Therefore, the kernel matrix is linear in the scalar kernel functions $k_m(\cdot, \cdot)$, such that the posterior of $\mathbf{a}_i f_i(\cdot)$ can be obtained as

$$\mathbf{a}_i f_i(\cdot) | \mathbb{D}_N \sim \mathcal{N}(\mu_{\mathbf{a}_i f_i}(\cdot), \Sigma_{\mathbf{a}_i f_i}(\cdot)), \quad (37)$$

where

$$\mu_{\mathbf{a}_i f_i}(\mathbf{z}) = (\mathbf{B}_i \otimes k_i^T(\mathbf{Z}, \mathbf{z})) \left(\Sigma_{\text{on}} \otimes \mathbf{I}_N + \sum_{j=1}^{d_f} \mathbf{B}_j \otimes k_j(\mathbf{Z}, \mathbf{Z}) \right)^{-1} (\mathbf{t} - \hat{\mathbf{f}}(\mathbf{Z})) \quad (38)$$

$$\Sigma_{\mathbf{a}_i f_i}(\mathbf{z}) = k_i(\mathbf{z}, \mathbf{z}) \mathbf{B}_i - (\mathbf{B}_i \otimes k_i^T(\mathbf{Z}, \mathbf{z})) \left(\Sigma_{\text{on}} \otimes \mathbf{I}_N + \sum_{j=1}^{d_f} \mathbf{B}_j \otimes k_j(\mathbf{Z}, \mathbf{Z}) \right)^{-1} (\mathbf{B}_i \otimes k_i(\mathbf{Z}, \mathbf{z})) \quad (39)$$

This follows from a trivial extension of the results in (Duvenaud, 2014) to multiple-output GPs. Due to the definition of \mathbf{B}_i , we can equivalently write

$$\mu_{\mathbf{a}_i f_i}(\mathbf{z}) = \mathbf{a}_i \left((\mathbf{a}_i^T \otimes k_i^T(\mathbf{Z}, \mathbf{z})) \left(\boldsymbol{\Sigma}_{\text{on}} \otimes \mathbf{I}_N + \sum_{j=1}^{d_f} \mathbf{a}_j \mathbf{a}_j^T \otimes k_j(\mathbf{Z}, \mathbf{Z}) \right)^{-1} (\mathbf{t} - \hat{\mathbf{f}}(\mathbf{Z})) \right) \quad (40)$$

$$\boldsymbol{\Sigma}_{\mathbf{a}_i f_i}(\mathbf{z}) = \mathbf{a}_i k_i(\mathbf{z}, \mathbf{z}) \mathbf{a}_i^T \quad (41)$$

$$- \mathbf{a}_i \left((\mathbf{a}_i^T \otimes k_i^T(\mathbf{Z}, \mathbf{z})) \left(\boldsymbol{\Sigma}_{\text{on}} \otimes \mathbf{I}_N + \sum_{j=1}^{d_f} \mathbf{a}_j \mathbf{a}_j^T \otimes k_j(\mathbf{Z}, \mathbf{Z}) \right)^{-1} (\mathbf{a}_i \otimes k_i(\mathbf{Z}, \mathbf{z})) \right) \mathbf{a}_i^T, \quad (42)$$

from which we can directly deduce the identities (17) and (18). \blacksquare

Proof of Theorem 2 The result follows from Lemma 1 and a straightforward adaption of (Lederer et al., 2019a, Theorem 3.1). \blacksquare

Appendix B. Variance Bounds and Lyapunov-Based Data Densities

Lemma 7 *The posterior variance $\tilde{\sigma}_i^2(\mathbf{x})$ defined in (18) is bounded by*

$$\tilde{\sigma}_j^2(\mathbf{x}) \leq s_{f_j}^2 - \frac{s_{f_j}^4 \exp(-\tilde{\phi}_j^2(\mathbf{x})) \sum_{i=1}^{d_x} a_{i,j}^2}{\max_{m=1, \dots, d_x} \sum_{n=1}^{d_f} \sum_{i=1}^{d_x} a_{m,n} a_{i,n} s_{f_n}^2 + \frac{\lambda_{\max}(\boldsymbol{\Sigma}_{\text{on}})}{M}}. \quad (43)$$

Proof This result is a direct extension of (Lederer et al., 2019b, Corollary 3.1) to multiple-output GPs with linear coregionalization and we pursue the proof analogously. Since the posterior variance is non-increasing, we can consider only training samples $\mathbf{z}^{(n)}$ within distance at most $\tilde{\phi}_j(\mathbf{x})$ to $[\mathbf{x}^T \quad \boldsymbol{\pi}^T(\mathbf{x})]^T$ in the posterior variance calculation. Therefore, we obtain

$$\sigma_j^2(\mathbf{x}) \leq s_{f_j}^2 - \frac{\|k_j^T(\mathbf{Z}_{\tilde{\phi}_j(\mathbf{x})}, \mathbf{z}) \otimes \mathbf{a}_j^T\|^2}{\lambda_{\max} \left(\boldsymbol{\Sigma}_{\text{on}} \otimes \mathbf{I}_N + \sum_{n=1}^{d_f} \mathbf{a}_n \mathbf{a}_n^T \otimes k_n(\mathbf{Z}_{\tilde{\phi}_j(\mathbf{x})}, \mathbf{Z}_{\tilde{\phi}_j(\mathbf{x})}) \right)}, \quad (44)$$

where $\mathbf{Z}_{\tilde{\phi}_j(\mathbf{x})}$ denotes the training samples with distance at most $\tilde{\phi}_j(\mathbf{x})$ to $[\mathbf{x}^T \quad \boldsymbol{\pi}^T(\mathbf{x})]^T$. We trivially obtain the bound

$$\|k_j(\mathbf{Z}_{\tilde{\phi}_j(\mathbf{x})}, \mathbf{z}) \otimes \mathbf{a}_j\| \geq M s_{f_j}^4 \exp(-\tilde{\phi}_j^2(\mathbf{x})) \sum_{i=1}^{d_x} a_{i,j}^2 \quad (45)$$

due to the distance restriction. Moreover, the application of Gershgorin's theorem yields

$$\lambda_{\max} \left(\boldsymbol{\Sigma}_{\text{on}} \otimes \mathbf{I}_N + \sum_{n=1}^{d_f} \mathbf{a}_n \mathbf{a}_n^T \otimes k_n(\mathbf{Z}_{\tilde{\phi}_j(\mathbf{x})}, \mathbf{Z}_{\tilde{\phi}_j(\mathbf{x})}) \right) \leq \lambda_{\max}(\boldsymbol{\Sigma}_{\text{on}}) + \max_{m=1, \dots, d_x} \sum_{n=1}^{d_f} \sum_{i=1}^{d_x} a_{m,n} a_{i,n} M s_{f_n}^2 \quad (46)$$

due to the definition of the M -fill distance $\tilde{\phi}_j(\mathbf{x})$ in Theorem 4. Substituting the bounds (45) and (46) in (44) finally yields the result. \blacksquare

Proof of Theorem 5 Since $\sqrt{\beta(\tau)}\tilde{\sigma}_j(\mathbf{x}) > \gamma_j(\tau)$ by assumption, we can simplify

$$\sqrt{\beta(\tau)}\tilde{\sigma}_j(\mathbf{x}) + \gamma_j(\tau) \leq 2\sqrt{\beta(\tau)}\tilde{\sigma}_j(\mathbf{x}). \quad (47)$$

Therefore, satisfaction of the condition

$$4 \left(\sum_{i=1}^{d_x} |a_{i,j} \frac{\partial}{\partial x_i} V(\mathbf{x}, t)| \right)^2 \beta(\tau) \tilde{\sigma}_j^2(\mathbf{x}) \leq \xi_j^2(\mathbf{x}, t) \quad (48)$$

implies the statement of Theorem 5. Hence, we can substitute (43) and solve for $\tilde{\phi}_j^2(\mathbf{x})$ in order to prove Theorem 5. ■