How Training Data Impacts Performance in Learning-based Control

Armin Lederer, Alexandre Capone, Jonas Umlauft, and Sandra Hirche

Abstract—When first principle models cannot be derived due to the complexity of the real system, data-driven methods allow us to build models from system observations. As these models are employed in learning-based control, the quality of the data plays a crucial role for the performance of the resulting control law. Nevertheless, there hardly exist measures for assessing training data sets, and the impact of the spatial distribution of the data on the closed-loop system properties is largely unknown. This paper derives — based on Gaussian process models — an analytical relationship between the density of the training data and the control performance. We formulate a quality measure for the data set, which we refer to as ρ-gap, and derive the ultimate bound for the tracking error under consideration of the model uncertainty. We show how the ρ-gap can be applied to a feedback linearizing control law and provide numerical illustrations for our approach.

Index Terms—Machine learning, information theory and control, uncertain systems, Lyapunov methods, nonlinear systems identification.

I. INTRODUCTION

MODEL-based control requires an accurate mathematical description of the plant that is to be controlled. Classical system identification methods postulate parametric models using prior assumptions, and tune their parameters based on observations to achieve high model accuracy [1]. However, these methods are prone to yielding poor models if a wrong parametric structure is assumed, e.g., if insufficient prior knowledge on the system structure is available. This is often the case for highly complex systems, e.g., settings where humans are part of the control loop. In order to overcome these shortcomings, learning-based control employs non-parametric data-driven models, which only require little prior system knowledge in comparison to classical parametric models [2]. Such modeling techniques strongly rely on (potentially noisy) observations, making a formal analysis of the resulting control performance difficult. Hence, these techniques remain too unreliable for safety-critical applications [3].

To overcome this drawback, recent research has focused on the stability of learning-based control approaches using Gaussian process (GP) models [4]. GPs can capture model uncertainty, which allows us to derive probabilistic model error bounds [5]. The applications of GP-based methods range from safe controller optimization for quadrotors [6],

over computed torque control in robotics [7], to feedback linearization for aircraft systems [8].

Despite the widespread use of GPs in control, there exist only a few tools to assess the quality of the training set. So far, most measures used to quantify data quality are based on information-theoretical measures, e.g., information gain [9], [10]. These techniques assess data in global terms, without taking into account locally varying requirements on the data due to the control structure and the task. Since the relationship between data distributions and control performance is largely unknown, random sampling-based approaches have recently been employed to estimate the effect of data on learning-based control systems [11]. However, sampling-based approaches are computationally expensive and provide no direct insight into the interrelation between training data and control error. Therefore, deriving an analytical measure is crucial to improve the understanding of this relationship. Furthermore, it is essential to enhance the efficiency of exploration in active learning, training data selection to implement machine learning with limited computational budget, and cautious control design.

The main contribution of this paper is a novel measure, called ρ-gap, to assess training data sets from a control theoretical perspective. Based on the model uncertainty of a GP model, we investigate the uncertainty-dependent Lyapunov stability conditions for a control-affine closed-loop system. This analysis allows insights on how data should be collected, which is becoming particularly useful in exploration tasks where high data-efficiency is required. As an example, we derive a novel uncertainty-dependent ultimate bound of the tracking error for a feedback linearizing control law and show how the density of the training data affects this bound.

The paper is structured as follows: Section II defines the problem setting, after which GP regression and the required model error bounds are introduced in Section III. The proposed quality measure of data is presented in Section IV, including the definition of a control law and the derivation of the ultimate tracking error bound. The results are numerically illustrated in Section V, followed by the conclusion in Section VI.

II. PROBLEM STATEMENT

We consider a single-input system in the canonical form[1]

$$\dot{x}_1 = x_2, \quad \dot{x}_2 = x_3, \quad \ldots \quad \dot{x}_{d_x} = f(\boldsymbol{x}) + g(\boldsymbol{x})u \quad (1)$$

with state $\boldsymbol{x} = [x_1 \cdots x_{d_x}]^T \in \mathbb{X}$ in a compact set $\mathbb{X} \subset \mathbb{R}^{d_x}$, input $u \in \mathbb{U} = \mathbb{R}$, and unknown functions $f \colon \mathbb{X} \to \mathbb{R}$ and $g \colon \mathbb{X} \to \mathbb{R}$. Note that we restrict the following analysis to single-input systems due to notational convenience, but our results straightforwardly extend to multi-input control-affine systems. We assume that prior models $\hat{f} \colon \mathbb{X} \to \mathbb{R}$ and $\hat{g} \colon \mathbb{X} \to \mathbb{R}$ of the unknown functions are given, and make the following assumptions on the unknown functions and the available data:

Assumption 1: A data set

$$\mathbb{D} = \left\{ \boldsymbol{z}^{(n)} \coloneqq \begin{bmatrix} \boldsymbol{x}^{(n)} \\ u^{(n)} \end{bmatrix}, y^{(n)} = \Delta\left(\boldsymbol{z}^{(n)}\right) + \omega^{(n)} \right\}_{n=1}^{N} \quad (2)$$

is available, which contains N pairs of noiseless measurements of the state $\boldsymbol{x}^{(n)}$ and measurements of

$$\Delta(\boldsymbol{z}) = \Delta(\boldsymbol{x}, u) = f(\boldsymbol{x}) - \hat{f}(\boldsymbol{x}) + (g(\boldsymbol{x}) - \hat{g}(\boldsymbol{x}))u \quad (3)$$

perturbed by Gaussian noise $\omega^{(n)} \sim \mathcal{N}(0, s_{\text{on}}^2)$, where $\boldsymbol{z} \in \mathbb{X} \times \mathbb{U} \subset \mathbb{R}^{d_z}$, and $d_z = d_x + 1$.

Assumption 2: The unknown functions $f(\cdot)$ and $g(\cdot)$ admit Lipschitz constants L_f and L_g, respectively.

Assumption 3: The sign of $g(\cdot)$ is known and constant.

While Assumptions 1 and 2 guarantee the existence of training data, and ensure that the unknown functions are well-behaved, Assumption 3 guarantees global controllability, and thereby the existence of a stabilizing control law. We consider the task of tracking a bounded reference trajectory

$$\boldsymbol{x}_d = \begin{bmatrix} x_d & \dot{x}_d & \cdots & \frac{\mathrm{d}^{d_x-1}x_d}{\mathrm{d}t^{d_x-1}} \end{bmatrix}^T, \quad (4)$$

which can be seen as the generalization of set point regularization, to which our results carry over. We employ a control law $\pi \colon \mathbb{X} \to \mathbb{U}$, whose goal is to stabilize the tracking error $\boldsymbol{e} = \boldsymbol{x} - \boldsymbol{x}_d$ with the dynamics

$$\dot{e}_1 = e_2, \quad \ldots \quad \dot{e}_{d_x} = f(\boldsymbol{x}) + g(\boldsymbol{x})\pi(\boldsymbol{x}) - \frac{\mathrm{d}^{d_x}x_d}{\mathrm{d}t^{d_x}}. \quad (5)$$

The control law $\pi(\cdot)$ is based on a model learned by GP regression with a composite kernel

$$k(\boldsymbol{z}, \boldsymbol{z}') = k_f(\boldsymbol{x}, \boldsymbol{x}') + uk_g(\boldsymbol{x}, \boldsymbol{x}')u', \quad (6)$$

and covariance functions $k_f, k_g \colon \mathbb{R}^{d_x} \times \mathbb{R}^{d_x} \to \mathbb{R}$ as suggested in [12]. Since this kernel reflects the structure of (1), it allows us to recover separate models for $f(\cdot)$ and $g(\cdot)$ [13]. While other kernel functions can also be used for a learning-based control approach [8], they generally do not allow a separation of additive model components. Since this separation is beneficial for the interpretability and intuitiveness of the relationship between training data and control performance, we focus on composite kernels in the following. The individual covariance functions represent our prior knowledge about the unknown functions.

Assumption 4: Prior knowledge of the function is expressed through prior GPs, i.e., $f(\cdot) \sim \mathcal{GP}(\hat{f}(\boldsymbol{x}), k_f(\boldsymbol{x}, \boldsymbol{x}'))$ and $g(\cdot) \sim \mathcal{GP}(\hat{g}(\boldsymbol{x}), k_g(\boldsymbol{x}, \boldsymbol{x}'))$.

This assumption imposes a probability distribution on the function space, which is shaped by the prior mean

---

[1]Notation: Lower/upper case bold symbols denote vectors/matrices, $\mathbb{R}_{+,0}/\mathbb{R}_+$ all real positive numbers with/without zero, respectively. $\boldsymbol{I}_n$ denotes the $n \times n$ identity matrix and $\|\cdot\|$ the Euclidean norm.

functions $\hat{f}(\cdot), \hat{g}(\cdot)$ and the kernel functions $k_f(\cdot, \cdot), k_g(\cdot, \cdot)$. Thereby, it implicitly requires that the covariance kernels and prior mean functions are chosen suitably, i.e., $f(\cdot), g(\cdot)$ must be expressible in terms of those functions [5].

While the stability of control laws has been investigated under these assumptions [8], [12], the derived ultimate bounds do not directly depend on the training data. Hence, the impact of training data on the performance of the learning-based controller is unknown. We address this issue by developing a flexible measure of the quality of training data with respect to the control performance. In order to illustrate the flexibility of the proposed quality measure, we derive a novel uncertainty-dependent ultimate tracking error bound for feedback linearizing control of systems with both $f(\cdot)$ and $g(\cdot)$ unknown, and apply our quality measure to this problem.

III. GAUSSIAN PROCESS REGRESSION

A Gaussian process defines a distribution $\mathcal{GP}(\mu_0(\boldsymbol{z}), k(\boldsymbol{z}, \boldsymbol{z}'))$ over functions $h \colon \mathbb{R}^{d_z} \to \mathbb{R}$ with prior mean $\mu_0 \colon \mathbb{R}^{d_z} \to \mathbb{R}$ and covariance $k \colon \mathbb{R}^{d_z} \times \mathbb{R}^{d_z} \to \mathbb{R}$, such that any finite number of evaluation points $\{\boldsymbol{z}_1, \ldots, \boldsymbol{z}_m\}$, $m \in \mathbb{N}$ is assigned a Gaussian distribution [14]. The prior mean function incorporates known parametric models into the regression, while the kernel $k(\cdot, \cdot)$ encodes information about the structure of $h(\cdot)$. Due to the definition of the modeling error $h(\cdot) = \Delta(\cdot)$ in (3), we employ a prior mean function $\mu_0(\boldsymbol{z}) = \hat{f}(\boldsymbol{x}) + \hat{g}(\boldsymbol{x})u$ and the composite kernel $k(\cdot, \cdot)$ defined in (6). For the components of the composite kernel $k(\cdot, \cdot)$, we use squared exponential kernels $k_f(\cdot, \cdot)$ and $k_g(\cdot, \cdot)$, defined as

$$k_{f,g}(\boldsymbol{x}, \boldsymbol{x}') = s_{f,g}^2 \exp\left(-\frac{\|\boldsymbol{x} - \boldsymbol{x}'\|^2}{2l_{f,g}^2}\right), \quad (7)$$

where $s_f^2, s_g^2 \in \mathbb{R}_{+,0}$ and $l_f, l_g \in \mathbb{R}_+$ denote the signal variances and length scales, respectively. Using the covariance function $k(\cdot, \cdot)$, the elements of the data covariance matrix $\boldsymbol{K}$ and the kernel vector $\boldsymbol{k}(\boldsymbol{z})$ at a test point $\boldsymbol{z}$ are given by $K_{nn'} = k(\boldsymbol{z}^{(n)}, \boldsymbol{z}^{(n')})$ and $k_n(\boldsymbol{z}) = k(\boldsymbol{z}^{(n)}, \boldsymbol{z})$, respectively. Based on these definitions, the probability of $h(\boldsymbol{z})$ conditioned on the training data $\mathbb{D}$ as well as the test point $\boldsymbol{z}$ is Gaussian with mean and variance

$$\mu(\boldsymbol{z}) = \mu_0(\boldsymbol{z}) + \boldsymbol{k}^T(\boldsymbol{z})\left(\boldsymbol{K} + s_{\text{on}}^2\boldsymbol{I}_N\right)^{-1}\boldsymbol{y}, \quad (8)$$

$$\sigma^2(\boldsymbol{z}) = k(\boldsymbol{z}, \boldsymbol{z}) - \boldsymbol{k}^T(\boldsymbol{z})(\boldsymbol{K} + s_{\text{on}}^2\boldsymbol{I}_N)^{-1}\boldsymbol{k}(\boldsymbol{z}), \quad (9)$$

respectively, where the training outputs $y^{(n)}$ are concatenated in the target vector $\boldsymbol{y} = \begin{bmatrix} y^{(1)} & \cdots & y^{(N)} \end{bmatrix}^T$. Due to the definition of the composite kernel (6), we can express the kernel vector as $\boldsymbol{k}(\boldsymbol{z}) = \boldsymbol{k}_f(\boldsymbol{x}) + \boldsymbol{U}\boldsymbol{k}_g(\boldsymbol{x})u$, where $\boldsymbol{U} = \text{diag}([u^{(1)} \cdots u^{(N)}])$ and the elements of $\boldsymbol{k}_f(\cdot)$ and $\boldsymbol{k}_g(\cdot)$ are defined as $k_{f,n}(\boldsymbol{x}) = k_f(\boldsymbol{x}^{(n)}, \boldsymbol{x})$ and $k_{g,n}(\boldsymbol{x}) = k_g(\boldsymbol{x}^{(n)}, \boldsymbol{x})$, respectively. By exploiting this structure, it is possible to recover the posterior GPs for $f(\cdot)$ and $g(\cdot)$ from the regression, and derive probabilistic uniform error bounds that depend on the posterior standard deviation, as shown in the following lemma.

Lemma 1: Consider a GP with composite kernel given by (6), a training data set $\mathbb{D}$ and functions $f(\cdot), g(\cdot), \hat{f}(\cdot)$

and $\hat{g}(\cdot)$ satisfying Assumptions 1, 2 and 4. Then, for any $\delta \in (0,1)$ and $\tau \in \mathbb{R}_+$, it holds that

$$P\Big(|f(\boldsymbol{x})-\mu_f(\boldsymbol{x})| \leq \sqrt{\beta(\tau)}\sigma_f(\boldsymbol{x})+\gamma_f(\tau), \ \forall \boldsymbol{x}\in\mathbb{X}\Big) \geq 1-\delta \quad (10)$$

$$P\Big(|g(\boldsymbol{x})-\mu_g(\boldsymbol{x})| \leq \sqrt{\beta(\tau)}\sigma_g(\boldsymbol{x})+\gamma_g(\tau), \ \forall \boldsymbol{x}\in\mathbb{X}\Big) \geq 1-\delta, \quad (11)$$

with mean and variance components

$$\mu_f(\boldsymbol{x}) = \hat{f}(\boldsymbol{x})+\boldsymbol{k}_f^T(\boldsymbol{x})(\boldsymbol{K}+s_{\text{on}}^2\boldsymbol{I}_N)^{-1}\boldsymbol{y} \quad (12)$$

$$\mu_g(\boldsymbol{x}) = \hat{g}(\boldsymbol{x})+\boldsymbol{k}_g^T(\boldsymbol{x})\boldsymbol{U}(\boldsymbol{K}+s_{\text{on}}^2\boldsymbol{I}_N)^{-1}\boldsymbol{y} \quad (13)$$

$$\sigma_f^2(\boldsymbol{x}) = k_f(\boldsymbol{x},\boldsymbol{x})-\boldsymbol{k}_f^T(\boldsymbol{x})(\boldsymbol{K}+s_{\text{on}}^2\boldsymbol{I}_N)^{-1}\boldsymbol{k}_f(\boldsymbol{x}) \quad (14)$$

$$\sigma_g^2(\boldsymbol{x}) = k_g(\boldsymbol{x},\boldsymbol{x})-\boldsymbol{k}_g^T(\boldsymbol{x})\boldsymbol{U}(\boldsymbol{K}+s_{\text{on}}^2\boldsymbol{I}_N)^{-1}\boldsymbol{U}\boldsymbol{k}_g(\boldsymbol{x}), \quad (15)$$

and parameters

$$\beta(\tau) = 2d_x \log\left(1+\frac{r_0}{\tau}\right) - 2\log(\delta) \quad (16)$$

$$\gamma_f(\tau) = (L_{\mu_f} + L_f)\tau + \sqrt{\beta(\tau)L_{\sigma_f^2}\tau} \quad (17)$$

$$\gamma_g(\tau) = (L_{\mu_g} + L_g)\tau + \sqrt{\beta(\tau)L_{\sigma_g^2}\tau}. \quad (18)$$

Here, L_{μ_f}, L_{μ_g}, $L_{\sigma_f^2}$ and $L_{\sigma_g^2}$ are the Lipschitz constants of the mean and variance components, respectively, and $r_0 = \max_{\boldsymbol{x},\boldsymbol{x}'\in\mathbb{X}} \|\boldsymbol{x} - \boldsymbol{x}'\|$ denotes the maximum diameter of $\mathbb{X}$.

Proof: It has been shown in [13] that the independent components $f(\cdot)$ and $g(\cdot)u$ can be inferred by

$$f(\boldsymbol{x})|\mathbb{D} \sim \mathcal{N}(\mu_f(\boldsymbol{x}), \sigma_f^2(\boldsymbol{x})) \quad (19)$$

$$g(\boldsymbol{x})u|\mathbb{D} \sim \mathcal{N}(\mu_g(\boldsymbol{x})u, \sigma_g^2(\boldsymbol{x})u^2) \quad (20)$$

due to the structure of the kernel. Since $g(\boldsymbol{x})u$ depends linearly on u, we can extract

$$g(\boldsymbol{x})|\mathbb{D} \sim \mathcal{N}(\mu_g(\boldsymbol{x}), \sigma_g^2(\boldsymbol{x})). \quad (21)$$

Based on these identities, it is straightforward to adapt [5, Theorem 3.1] to obtain the uniform error bounds (10), (11). ∎

It is well known from scattered data approximation [15] that training data which covers $\mathbb{X}$ well typically leads to small posterior variances, thereby implying that the learned model has a high accuracy. Lemma 1 also exhibits this behavior, even though it additionally depends on the constants $\gamma_f(\tau)$ and $\gamma_g(\tau)$. Since these constants can be made arbitrarily small by reducing the value of τ, their effect is usually negligible. In fact, bounds (10) and (11) can be shown to converge to 0 under weak assumptions [5].

IV. QUALITY ASSESSMENT OF TRAINING DATA FOR LEARNING-BASED CONTROL

A. Lyapunov-Based Quality Assessment

Although GPs are frequently used in control design, the relationship between training data and the resulting performance of a control law $u = \pi(\boldsymbol{x})$ has barely been analyzed. Therefore, there is little insight on where training samples should be placed to achieve the highest improvement in control performance. In the sequel, we measure the control performance using a Lyapunov function $V : \mathbb{R}^{d_x} \to \mathbb{R}_{+,0}$. Therefore, we investigate the time derivative of the Lyapunov function for systems defined in (1), which is given by

$$\dot{V}(\boldsymbol{e}) = \sum_{i=1}^{d_x-1}\frac{\partial V(\boldsymbol{e})}{\partial e_i}e_{i+1}+\frac{\partial V(\boldsymbol{e})}{\partial e_{d_x}}\left(f(\boldsymbol{x})+g(\boldsymbol{x})\pi(\boldsymbol{x})-\frac{\mathrm{d}^{d_x}x_d}{\mathrm{d}t^{d_x}}\right). \quad (22)$$

By employing a Gaussian process model, as presented in Section III, and exploiting the uniform error bound for GPs from Lemma 1, we can bound this derivative by

$$\dot{V}(\boldsymbol{e}) \leq \dot{V}_{\text{nom}}(\boldsymbol{e}) + \dot{V}_{\sigma_f}(\boldsymbol{e}) + \dot{V}_{\sigma_g}(\boldsymbol{e}), \quad (23)$$

where the nominal component of the Lyapunov derivative is computed based on the GP mean function as

$$\dot{V}_{\text{nom}}(\boldsymbol{e}) = \sum_{i=1}^{d_x-1}\frac{\partial V(\boldsymbol{e})}{\partial e_i}e_{i+1}+\frac{\partial V(\boldsymbol{e})}{\partial e_{d_x}}\left(\mu_f(\boldsymbol{x})+\mu_g(\boldsymbol{x})\pi(\boldsymbol{x})-\frac{\mathrm{d}^{d_x}x_d}{\mathrm{d}t^{d_x}}\right). \quad (24)$$

The uncertain component of $\dot{V}(\cdot)$ is separated into components for the uncertainty about $f(\cdot)$ and $g(\cdot)$, given by

$$\dot{V}_{\sigma_f}(\boldsymbol{e}) = \left|\frac{\partial V(\boldsymbol{e})}{\partial e_{d_x}}\right|\left(\sqrt{\beta(\tau)}\sigma_f(\boldsymbol{x}) + \gamma_f(\tau)\right) \quad (25)$$

$$\dot{V}_{\sigma_g}(\boldsymbol{e}) = \left|\frac{\partial V(\boldsymbol{e})}{\partial e_{d_x}}\right|\left(\sqrt{\beta(\tau)}\sigma_g(\boldsymbol{x}) + \gamma_g(\tau)\right)|\pi(\boldsymbol{x})|. \quad (26)$$

The nominal component of the Lyapunov function derivative does not depend on the uncertainty. Hence, it does not provide insight into the relationship between training data and control performance. In contrast, $\dot{V}_{\sigma_f}(\cdot)$ and $\dot{V}_{\sigma_g}(\cdot)$ directly depend on the GP posterior standard deviations, and thereby on the training data density. In order to measure this density in a flexible way, we introduce the M-fill distance, inspired by classical concepts from scattered data approximation [15].

Definition 1: The M-fill distance $\phi_{\underline{u},\bar{u},M}(\boldsymbol{x})$ at a point $\boldsymbol{x} \in \mathbb{X}$ is defined as the minimum radius φ of a ball with center $\boldsymbol{x}$, such that the ball contains $M \in \mathbb{N}$ training samples $\boldsymbol{z}^{(n)}$ with control inputs $\underline{u} \leq |u^{(n)}| \leq \bar{u}$ for some $\underline{u}, \bar{u} \in \mathbb{R}_{+,0}$, i.e.,

$$\phi_{\underline{u},\bar{u},M}(\boldsymbol{x}) = \min_{\varphi\in\mathbb{R}_{+,0}} \varphi \quad (27a)$$

$$\text{s.t. } \left|\left\{\boldsymbol{z}^{(n)}\in\mathbb{D}: \|\boldsymbol{x}-\boldsymbol{x}^{(n)}\| \leq \varphi \wedge \underline{u} \leq |u^{(n)}| \leq \bar{u}\right\}\right| \geq M. \quad (27b)$$

The M-fill distance $\phi_{\underline{u},\bar{u},M}(\boldsymbol{x})$ measures the distance to the M closest training samples, where the parameter M is used to adapt $\phi_{\underline{u},\bar{u},M}(\boldsymbol{x})$ to the total number of training samples N. Intuitively, one should choose $M \ll N$, such that only training points in the proximity of $\boldsymbol{x}$ are relevant for $\phi_{\underline{u},\bar{u},M}(\boldsymbol{x})$, thereby making it a local measure of the data density. A low M-fill distance implies a high data density and indeed, upper bounds for $\phi_{\underline{u},\bar{u},M}(\boldsymbol{x})$ can be derived to guarantee a desired control behavior, which is expressed through upper bounds $\xi_f(\cdot)$ and $\xi_g(\cdot)$ for $\dot{V}_{\sigma_f}(\cdot)$ and $\dot{V}_{\sigma_g}(\cdot)$, respectively.

Theorem 1: If the M-fill distance $\phi_{0,\bar{u}_f,M}(\cdot)$ satisfies

$$\phi_{0,\bar{u}_f,M}^2(\boldsymbol{x}) \leq \bar{\phi}_f^2(\boldsymbol{x}) + \theta_f \quad (28)$$

for all $\boldsymbol{x} \in \mathbb{X}$, where

$$\bar{\phi}_f^2(\boldsymbol{x}) = -l_f^2\log\left(1-\frac{\left(\xi_f(\boldsymbol{e}) - \gamma_f(\tau)\left|\frac{\partial V(\boldsymbol{e})}{\partial e_{d_x}}\right|\right)^2}{\beta(\tau)s_f^2\left|\frac{\partial V(\boldsymbol{e})}{\partial e_{d_x}}\right|^2}\right) \quad (29)$$

$$\theta_f = -l_f^2\log\left(\frac{Ms_f^2+Ms_g^2\bar{u}_f^2+s_{\text{on}}^2}{Ms_f^2}\right) \quad (30)$$

for any $\bar{u}_f \in \mathbb{R}_{+,0}$, $M \in \mathbb{N}$, and $\xi_f(\boldsymbol{e}) > \gamma_f(\tau)|\partial V(\boldsymbol{e})/\partial e_{d_x}|$,

then, $\dot{V}_{\sigma_f}(e) \leq \xi_f(e), \forall x \in \mathbb{X}$.

Proof: In order to prove this lemma, we have to bound the posterior standard deviation $\sigma_f(x)$. Following the approach introduced in [16], this is achieved by considering only M training samples $z^{(n)}$ within distance $\phi_{0,\bar{u}_f,M}(x)$ to x such that the posterior variance is bounded by[2]

$$\sigma_f^2(x) \leq s_f^2 - \frac{\|k_{f,M}(x)\|^2}{\lambda_{\max}(K_M) + s_{\text{on}}^2}, \qquad (31)$$

where $k_{f,M}(x)$ and K_M denote the covariance vector and matrix based on these M samples and $\lambda_{\max}(K_M)$ denotes the maximum eigenvalue. Application of the Gershgorin theorem allows us to bound the maximum eigenvalue by

$$\lambda_{\max}(K_M) \leq M(s_f^2 + \bar{u}_f^2 s_g^2). \qquad (32)$$

Since we have $\|k_{f,M}(x)\|^2 \geq M s_f^4 \exp(-\phi_{0,\bar{u}_f,M}^2(x)/l_f^2)$, we obtain the posterior variance bound

$$\sigma_f^2(x) \leq s_f^2 - \frac{s_f^4 \exp\left(-\frac{\phi_{0,\bar{u}_f,M}^2(x)}{l_f^2}\right)}{s_f^2 + \bar{u}_f^2 s_g^2 + \frac{s_{\text{on}}^2}{M}}. \qquad (33)$$

Substituting this expression into (25) and solving for $\phi_{0,\bar{u}_f,M}^2(x)$ yields the desired result. ∎

Corollary 1: If the M-fill distance $\phi_{\underline{u}_g,\bar{u}_g,M}(\cdot)$ satisfies

$$\phi_{\underline{u}_g,\bar{u}_g,M}^2(x) \leq \bar{\phi}_g^2(x) + \theta_g, \qquad (34)$$

for all $x \in \mathbb{X}$, where

$$\bar{\phi}_g^2(x) = -l_g^2 \log\left(1 - \frac{\left(\xi_g(e) - \gamma_g(\tau)|\pi(x)|\left|\frac{\partial V(e)}{\partial e_{d_x}}\right|\right)^2}{\beta(\tau)s_g^2|\pi(x)|^2\left|\frac{\partial V(e)}{\partial e_{d_x}}\right|^2}\right) \qquad (35)$$

$$\theta_g = -l_g^2 \log\left(\frac{Ms_f^2 + Ms_g^2\bar{u}_g^2 + s_{\text{on}}^2}{Ms_g^2\underline{u}_g^2}\right) \qquad (36)$$

for any $\underline{u}_g, \bar{u}_g \in \mathbb{R}_{+,0}$, $\underline{u}_g \leq \bar{u}_g$, $M \in \mathbb{N}$, and $\xi_g(e) > \gamma_g(\tau)|\pi(x)||\partial V(e)/\partial e_{d_x}|$, then, $\dot{V}_{\sigma_g}(e) \leq \xi_g(e), \forall x \in \mathbb{X}$.

Proof: We can bound the posterior standard deviation $\sigma_f(x)$ analogously to Theorem 1 since

$$\sigma_g^2(x) \leq s_g^2 - \frac{\|k_{g,M}(x)U\|^2}{\lambda_{\max}(K_M) + s_{\text{on}}^2}. \qquad (37)$$

The remainder of this proof follows directly from a straightforward adaptation of the proof of Theorem 1. ∎

Theorem 1 and Corollary 1 allow us to directly investigate if $\dot{V}_f(\cdot)$ and $\dot{V}_g(\cdot)$ satisfy a desired behavior by measuring the M-fill distance. Therefore, they provide helpful insight on how the training data should be distributed. The quality of the data for learning the decoupling between $f(\cdot)$ and $g(\cdot)$ is measured by θ_f and θ_g. It is straightforward to see that θ_g is close to zero if $\bar{u}_g \approx \underline{u}_g$ is large. This is intuitive since the weight of $g(\cdot)$ in the function $h(z) = f(x) + g(x)u$ grows linearly with u. Thus, large control inputs are beneficial for the identification of $g(\cdot)$. In contrast, Theorem 1 shows that small control inputs $u \approx 0$ are advantageous for learning $f(\cdot)$ since the control dependency of $h(\cdot)$ disappears. Finally,

---

[2]If the dependency on x arises from the restriction of the considered training samples, we do not state this explicitly for notational simplicity.

large noise variance requires higher data density to keep the term s_{on}^2/M small in both bounds.

In contrast to θ_f and θ_g, the functions $\bar{\phi}_f^2(x)$ and $\bar{\phi}_g^2(x)$ express the dependency on the performance specification. Since $\gamma_f(\tau)$ and $\gamma_g(\tau)$ are usually negligible, the quotients in (29) and (35) can be approximated by the ratio between the squared performance specifications $\xi_f^2(\cdot)$, $\xi_g^2(\cdot)$ and the prior bounds $\beta(\tau)s_f^2|\partial V(e)/\partial e_{d_x}|^2$, $\beta(\tau)s_g^2|\pi(x)|^2|\partial V(e)/\partial e_{d_x}|^2$, respectively. Hence, small performance specifications, e.g., $\xi_f(\cdot) \ll \beta(\tau)s_f^2|\partial V(e)/\partial e_{d_x}|^2$, cause small values of $\bar{\phi}_f(\cdot)$, $\bar{\phi}_g(\cdot)$, which in turn indicate the necessity of high data density. Since the performance specifications define the allowed increase of $\dot{V}_{\sigma_f}(\cdot)$, $\dot{V}_{\sigma_g}(\cdot)$, it is natural to define them as a fraction of the nominal derivative, e.g.,

$$\xi_f(\cdot) = \chi_f \dot{V}_{\text{nom}}(\cdot), \quad \xi_g(\cdot) = \chi_g \dot{V}_{\text{nom}}(\cdot), \quad \chi_f, \chi_g \in \mathbb{R}_+, \quad (38)$$

such that stability is guaranteed for $\chi_f + \chi_g < 1$. Due to this intuitive interpretation of the M-fill distance, we propose to use it as the basis for measures of the quality of the training data distribution.

Definition 2: The ρ_f- and ρ_g-gaps are defined as

$$\rho_f(x) = \phi_{0,\bar{u}_f,M}^2(x) - \bar{\phi}_f^2(x) \qquad (39)$$

$$\rho_g(x) = \phi_{\underline{u}_g,\bar{u}_g,M}^2(x) - \bar{\phi}_g^2(x). \qquad (40)$$

The ρ-gaps measure the difference between required squared M-fill distances $\bar{\phi}_f^2(\cdot)$, $\bar{\phi}_g^2(\cdot)$, which express the dependency of the data density on the desired bounds $\xi_f(\cdot)$, $\xi_g(\cdot)$ for $\dot{V}_{\sigma_f}(\cdot)$, $\dot{V}_{\sigma_g}(\cdot)$, and the actual squared M-fill distances $\phi_{0,\bar{u}_f,M}^2(\cdot)$, $\phi_{\underline{u}_g,\bar{u}_g,M}^2(\cdot)$, which are independent of the control problem. Therefore, searching for points $x \in \mathbb{X}$ with high ρ_f- and ρ_g-gaps leads to regions where the distance between training samples is too large to satisfy the bounds $\xi_f(\cdot), \xi_g(\cdot)$. This can be exploited, e.g., in active learning and training data selection, to find points with high ρ-gap and add them to the training set in order to reduce the violation of the performance specification. Note that θ_f and θ_g are not included in the definitions of $\rho_f(\cdot)$ and $\rho_g(\cdot)$, respectively, since they are independent of the state, and thereby do not provide useful information about the required spatial distribution of the data.

B. Ultimate Error Bound for Feedback Linearization

In order to demonstrate the intuitive relationship between the tracking error and the proposed ρ-gaps, we extend existing stability results for feedback linearization with GP models. As suggested by [12], we define the filtered state $r = [\lambda^T \ 1]e$ with Hurwitz coefficients $\lambda \in \mathbb{R}^{d_x-1}$. Based on the GP model from Section III, we define the approximately linearizing control law

$$\pi(x) = \mu_g^{-1}(x)(\nu - \mu_f(x)), \qquad (41)$$

where $\nu \in \mathbb{R}$ is the input to the linearized system, and $\mu_f(\cdot)$ and $\mu_g(\cdot)$ are defined in (12) and (13), respectively. The existence of the reciprocal value of $\mu_g(\cdot)$ can be ensured by a suitable choice of hyperparameters and prior mean $\hat{g}(\cdot)$ as shown in [12, Proposition 1]. In order to achieve asymptotic

tracking behavior in case of perfect linearization, we define the input to the linearized system as

$$\nu = -k_c r - \boldsymbol{\lambda}^T \boldsymbol{e}_{2:d_x} + \frac{\mathrm{d}^{d_x} x_d}{\mathrm{d} t^{d_x}}, \qquad (42)$$

where $k_c \in \mathbb{R}_+$ denotes the control gain and $\boldsymbol{e}_{2:d_x} = [e_2 \ \cdots \ e_{d_x}]^T$. Following the approach in Section IV-A, we can determine the ultimately bounded set for the system in (1) controlled by (41) as shown in the following theorem.

Theorem 2: Consider a system (1), a prior model $\hat{f}(\cdot)$, $\hat{g}(\cdot)$, training data $\mathbb{D}$, and let Assumptions 1-4 hold. If

$$\alpha(\boldsymbol{x}) = \frac{\sqrt{\beta(\tau)}\sigma_g(\boldsymbol{x}) + \gamma_g(\tau)}{\mu_g(\boldsymbol{x})} < \eta \quad \forall \boldsymbol{x} \in \mathbb{X} \qquad (43)$$

with

$$\eta = \min\left\{ \frac{k_c \lambda_2}{k_c \lambda_2 + \lambda_1}, \cdots, \frac{k_c \lambda_{d_x-1}}{k_c \lambda_{d_x-1} + \lambda_{d_x-2}}, \frac{k_c}{k_c + \lambda_{d_x-1}} \right\}, \qquad (44)$$

then there exists a control gain k_c, such that the tracking error $\boldsymbol{e}$ obtained with the feedback linearizing controller (41), (42) converges, with probability of at least $1 - 2\delta$, to the ultimately bounded set

$$\mathbb{B} = \left\{ \sqrt{\boldsymbol{e}^T \boldsymbol{\Lambda} \boldsymbol{e}} \leq \frac{\sqrt{\beta(\tau)}\sigma_f(\boldsymbol{x}) + \gamma_f(\tau) + \alpha(\boldsymbol{x})c(\boldsymbol{x})}{\tilde{k}_c(\boldsymbol{x})} \right\}, \qquad (45)$$

where

$$\tilde{k}_c(\boldsymbol{x}) = k_c\left(1 - \frac{\alpha(\boldsymbol{x})}{\eta}\right), \qquad c(\boldsymbol{x}) = \left| \frac{\mathrm{d}^{d_x} x_d}{\mathrm{d} t^{d_x}} - \mu_f(\boldsymbol{x}) \right|,$$
$$\boldsymbol{\Lambda} = \mathrm{diag}([\lambda_1^2 \ \cdots \ \lambda_{d_x-1}^2 \ 1]). \qquad (46)$$

Proof: It has been shown in [12] that there exists a gain k_c, such that the closed-loop system is ultimately bounded to the compact set $\mathbb{X}$. Hence, we can restrict our analysis to this set. Consider the Lyapunov function $V(r) = \frac{1}{2}r^2$, which allows us to analyze the stability of the tracking error since $\boldsymbol{\lambda}$ is Hurwitz [17]. The derivative of this Lyapunov function is decomposed into

$$\dot{V}(r) \leq -k_c r^2 + \dot{V}_{\sigma_f}(r) + \dot{V}_{\sigma_g}(r), \qquad (47)$$

where

$$\dot{V}_{\sigma_f}(r) = |r|(\sqrt{\beta(\tau)}\sigma_f(\boldsymbol{x}) + \gamma_f(\tau)) \qquad (48)$$
$$\dot{V}_{\sigma_g}(r) = |r|(\sqrt{\beta(\tau)}\sigma_g(\boldsymbol{x}) + \gamma_g(\tau))|\pi(\boldsymbol{x})|. \qquad (49)$$

We can separate the bound of the control input into feedback and feedforward components

$$|\mu_g(\boldsymbol{x})||\pi(\boldsymbol{x}))| \leq |k_c r + \boldsymbol{\lambda}^T \boldsymbol{e}_{2:d_x}| + \left| \frac{\mathrm{d}^{d_x} x_d}{\mathrm{d} t^{d_x}} - \mu_f(\boldsymbol{x}) \right|. \qquad (50)$$

The feedback component can be bounded by

$$|k_c r + \boldsymbol{\lambda}^T \boldsymbol{e}_{2:d_x}| \leq \frac{k_c}{\eta}|r|, \qquad (51)$$

while the feedforward component is a bounded state dependent function [12]. Hence, we obtain the bound

$$\dot{V}(r) \leq -\tilde{k}_c(\boldsymbol{x})r^2 + |r|\left(\sqrt{\beta(\tau)}\sigma_f(\boldsymbol{x}) + \gamma_f(\tau) + \alpha(\boldsymbol{x})c(\boldsymbol{x})\right) \qquad (52)$$

due to the definition of $\alpha(\boldsymbol{x})$ and $\tilde{k}_c(\boldsymbol{x})$. The quotient $\alpha(\boldsymbol{x})/\eta$ is smaller than one by assumption, such that $\tilde{k}_c(\boldsymbol{x}) > 0$ and

the Lyapunov function derivative becomes negative for all $r > (\sqrt{\beta(\tau)}\sigma_f(\boldsymbol{x}) + \gamma_f(\tau) + \alpha(\boldsymbol{x})c(\boldsymbol{x}))/\tilde{k}_c(\boldsymbol{x})$. ∎

While the condition on the sufficiently high control gain k_c is theoretically important to ensure the global ultimate boundedness, it is practically sufficient to analyze the conditions of Theorem 2 on a compact set $\mathbb{X}$ and choose k_c high enough to ensure that all points $\boldsymbol{x}_d + \boldsymbol{e}$ with $\boldsymbol{e} \in \mathbb{B}$ are in the interior of $\mathbb{X}$. Based on this local analysis, an ultimate bound is obtained which holds for initial values $\boldsymbol{x}(0)$ in a neighborhood of $\boldsymbol{x}_d(0)$. Moreover, the condition on $\alpha(\cdot)$ stems from the uncertainty about $g(\cdot)$, and ensures that its sign is robustly known under the posterior GP distribution. Due to this uncertainty, the effect of the feedback control on the tracking error bound is reduced, resulting in a diminished effective control gain $\tilde{k}_c(\cdot)$. Furthermore, the ultimate error bound can be made arbitrarily small by increasing the effective gain $\tilde{k}_c$. This is achieved by increasing the nominal gain k_c or reducing the uncertainty about $g(\cdot)$. In fact, if the function $g(\cdot)$ is known exactly, this effect disappears and we recover the ultimate bound which has been proposed in different forms in [5], [12].

V. NUMERICAL EVALUATION

A. Experimental Setting

We investigate the learning-based controller (41) and the corresponding ultimate bound on the nonlinear system

$$f(\boldsymbol{x}) = 1 - \sin(x_1) + \frac{1}{1 + \exp(-x_2)} \qquad (53)$$

$$g(\boldsymbol{x}) = 20\left(1 + \frac{1}{2}\sin\left(\frac{x_2}{4}\right)\right). \qquad (54)$$

This dynamical system is well-suited for illustrating the proposed ρ-gaps, as explained in the following. Both functions are slowly varying, such that the GP is capable of extrapolating to regions without data. Therefore, we do not need well-distributed data and can investigate the effect of an increasing distance to training data on the ρ-gaps. We express prior knowledge about this system using the approximate models $\hat{f}(\boldsymbol{x}) = 0$, $\hat{g}(\boldsymbol{x}) = 20$. Moreover, we define the reference trajectory $x_d(t) = 2\sin(t)$ and generate $N = 1000$ training samples by applying a high gain feedback linearizing controller based on the approximate models $\hat{f}(\cdot)$, $\hat{g}(\cdot)$ to track the reference trajectory. We add zero mean Gaussian noise with standard deviation $s_{\mathrm{on}} = 0.5$ to the observed accelerations and train a GP using log-likelihood maximization [14]. We approximate the Lipschitz constants L_μ and L_{σ^2} of the resulting GP numerically and bound the Lipschitz constants of $f(\cdot)$ and $g(\cdot)$ by twice the nominal value.

In the numerical experiment, we simulate the system for $T = 30$ starting at $\boldsymbol{x}(0) = \boldsymbol{0}$ with a control gain of $k_c = 40$ and $\lambda = 1$. The constants $\beta(\tau)$, $\gamma_f(\tau)$ and $\gamma_g(\tau)$ are computed using $\tau = 10^{-4}$, $\delta = 10^{-2}$, and the conditions for the ultimately bounded set (45) are investigated on the set $\mathbb{X} = \{\boldsymbol{x} \in \mathbb{R}^2 : \|\boldsymbol{x}\| \leq 2.5\}$. For the definition of the performance specifications, we use (38) with $\chi_f = \chi_g = 0.25$, such that their satisfaction ensures stability. We simplify the ρ-gaps analogously to the proof of Theorem 2 and measure the density of informative points for the identification of $g(\cdot)$ by
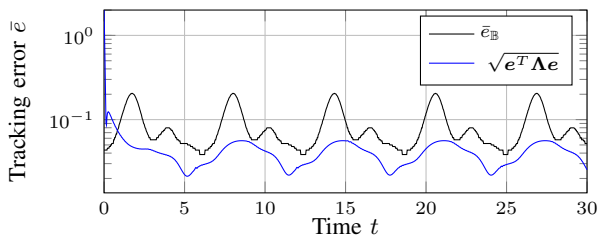
Fig. 1: Radius of the ultimately bounded set (black) and tracking error observed in simulation (blue).

defining $\bar{u}_g = \max_{n \leq N} |u^{(n)}|$. Moreover, we define $\underline{u}_g$ such that 90% of the control inputs $|u^{(n)}|$ are smaller than $\underline{u}_g$, and choose $M = 1$. Similarly, we define $\bar{u}_f$ such that 90% of the control inputs $|u^{(n)}|$ are larger than $\bar{u}_f$.

B. Results

The evolution of the observed tracking error $\sqrt{e^T \Lambda e}$ and the maximum extension of the ultimately bounded set $\bar{e}_\mathbb{B}$ are depicted in Fig. 1. It can be clearly observed that the tracking error indeed satisfies the ultimate bound for the given confidence level $\delta = 10^{-2}$ in the steady state after a brief transient period. The curves for the ultimate error bound $\bar{e}_\mathbb{B}$ and the tracking error $\sqrt{e^T \Lambda e}$ exhibit a similar behavior with minima and maxima occurring at almost identical times.

Snapshots of the state trajectory as well as the corresponding ultimate bound are depicted in Fig. 2. On the left hand side the maximum of the ultimate bound occurs at the maximum of $\rho_f(\cdot)$ along the reference trajectory. Taking a closer look at the training data, it can be seen that the training samples exhibit large control inputs $u^{(n)}$ in this region, which leads to high uncertainty in the identification of $f(\cdot)$. Therefore, there is a lack of training data, as indicated by $\rho_f(\cdot)$. In contrast, the right hand side of Fig. 2 shows the ρ_g-gap, which is minimal at the minimum of the ultimate bound $\bar{e}_\mathbb{B}$. This minimum of $\rho_g(\cdot)$ is a consequence of the choice of $g(\cdot)$, which is monotonously increasing with respect to x_2 in the considered state space $\mathbb{X}$. Therefore, model uncertainties have a slightly weaker effect in the upper half plane than in the lower half plane. Furthermore, it can be clearly observed that $\rho_g(\cdot)$ exhibits a maximum along the trajectory at $[-2 \ 0]^T$. In fact, this lack of training data with high control inputs $u^{(n)}$ causes the second local maximum in the ultimate bound at $2\pi m + 4$, $m \in \mathbb{N}$. Finally, both ρ-gaps show different behavior with increasing distance to the reference trajectory: Due to the increasing nominal derivative $\dot{V}_{\text{nom}}(\cdot)$, $\rho_f(\cdot)$ decreases strongly, while a similar growth of the control input $\pi(\boldsymbol{x})$ compensates this effect, causing a growing $\rho_g(\cdot)$.

VI. CONCLUSION

This paper introduces a quality measure for training sets in data-driven control. We establish a relationship between the distribution of the training data and the ultimate bound of the tracking error for a Gaussian process-based control law. In contrast to state-of-the-art information-theoretical quantities, our measure allows us to determine the most useful data points for control. In future work, this can be used to design exploration algorithms that collect data such that the control performance of the control law is maximized.
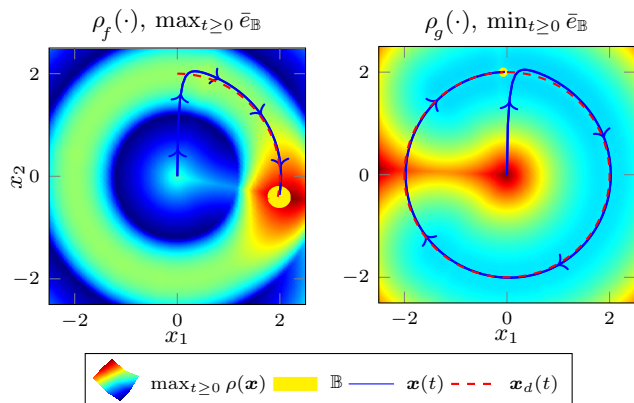

Fig. 2: Snapshots of the state trajectory (blue) and reference trajectory (red) of the system controlled with $N = 1000$ training samples obtained using a high gain controller. Large ultimate bounds are strongly related to large values of $\rho_f(\cdot)$ and $\rho_g(\cdot)$ (red background).

REFERENCES

[1] L. Ljung, *System Identification*. NJ, USA: Prentice Hall PTR, 1998.
[2] J. Kocijan, R. Murray-Smith, C. E. Rasmussen, and A. Girard, "Gaussian Process Model Based Predictive Control," in *Proceedings of the American Control Conference*, 2004, pp. 2214–2219.
[3] J. F. Fisac, A. K. Akametalu, M. N. Zeilinger, S. Kaynama, J. Gillula, and C. J. Tomlin, "A General Safety Framework for Learning-Based Control in Uncertain Robotic Systems," *IEEE Transactions on Automatic Control*, vol. 64, no. 7, pp. 2737–2752, 2019.
[4] J. Umlauft, A. Lederer, and S. Hirche, "Learning Stable Gaussian Process State Space Models," in *Proceedings of the American Control Conference*, 2017, pp. 1499–1504.
[5] A. Lederer, J. Umlauft, and S. Hirche, "Uniform Error Bounds for Gaussian Process Regression with Application to Safe Control," in *Advances in Neural Information Processing Systems*, 2019.
[6] F. Berkenkamp, A. P. Schoellig, and A. Krause, "Safe Controller Optimization for Quadrotors with Gaussian Processes," in *Proceedings of the IEEE International Conference on Robotics and Automation*, 2016, pp. 491–496.
[7] T. Beckers, D. Kulić, and S. Hirche, "Stable Gaussian Process based Tracking Control of Euler-Lagrange Systems," *Automatica*, vol. 103, no. 23, pp. 390–397, 2019.
[8] G. Chowdhary, H. A. Kingravi, J. P. How, and P. A. Vela, "Bayesian nonparametric adaptive control using Gaussian processes," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 26, no. 3, pp. 537–550, 2015.
[9] T. Koller, F. Berkenkamp, M. Turchetta, and A. Krause, "Learning-based Model Predictive Control for Safe Exploration," in *Proceedings of the IEEE Conference on Decision and Control*, 2018, pp. 6059–6066.
[10] T. Alpcan and I. Shames, "An Information-Based Learning Approach to Dual Control," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 26, no. 11, pp. 2736–2748, 2015.
[11] A. Capone, A. Lederer, J. Umlauft, and S. Hirche, "Data Selection for Multi-Task Learning under Dynamic Constraints," 2020. [Online]. Available: http://arxiv.org/abs/2005.03270
[12] J. Umlauft and S. Hirche, "Feedback Linearization based on Gaussian Processes with event-triggered Online Learning," *IEEE Transactions on Automatic Control*, 2020.
[13] D. K. Duvenaud, "Automatic Model Construction with Gaussian Processes," Ph.D. dissertation, University of Cambridge, 2014.
[14] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. Cambridge, MA: The MIT Press, 2006.
[15] H. Wendland, *Scattered Data Approximation*. Cambridge University Press, 2004.
[16] A. Lederer, J. Umlauft, and S. Hirche, "Posterior Variance Analysis of Gaussian Processes with Application to Average Learning Curves," 2019. [Online]. Available: http://arxiv.org/abs/1906.01404
[17] A. Yesildirek and F. Lewis, "Feedback Linearization Using Neural Networks," in *Proceedings of the IEEE Conference on Decision and Control*, 1995, pp. 2539–2544.