

Feedback Linearization based on Gaussian Processes with event-triggered Online Learning

Jonas Umlauf, *Student Member, IEEE*, Sandra Hirche, *Fellow, IEEE*

Abstract—Combining control engineering with nonparametric modeling techniques from machine learning allows for the control of systems without analytic description using data-driven models. Most existing approaches separate *learning*, i.e. the system identification based on a fixed dataset, and *control*, i.e. the execution of the model-based control law. This separation makes the performance highly sensitive to the initial selection of training data and possibly requires very large datasets. This article proposes a learning feedback linearizing control law using online closed-loop identification. The employed Gaussian process model updates its training data only if the model uncertainty becomes too large. This event-triggered online learning ensures high data efficiency and thereby reduces computational complexity, which is a major barrier for using Gaussian processes under real-time constraints. We propose safe forgetting strategies of data points to adhere to budget constraints and to further increase data-efficiency. We show asymptotic stability for the tracking error under the proposed event-triggering law and illustrate the effective identification and control in simulation.¹

Index Terms—adaptive control, machine learning, switched systems, uncertain systems, closed loop identification, data-driven control, online learning, Gaussian processes, event-based control

I. INTRODUCTION

DATA-DRIVEN control has gained a great deal of attention as costs for measuring, processing and storing data rapidly decrease and control engineering is increasingly applied in areas in which it is difficult to describe the plant using first principles. Nevertheless, a precise system description is essential for many modern model-based control algorithms, such as model predictive control and feedback linearization. Classical system identification using parametric models, such as autoregressive moving average (ARMA) or Hammerstein models [1], reaches its limits when the choice of a suitable model class is cumbersome or impossible, e.g. in systems where human behavior is part of the control loop. That is where data-driven nonparametric models have their advantages, as only minimal prior knowledge is required and they allow for higher flexibility than parametric models.

This article specifically considers Gaussian processes (GPs), which are well recognized in machine learning and control for modeling complex dynamics [2]. The Bayesian background

All authors are members of the Chair of Information-oriented Control, Department of Electrical and Computer Engineering, Technical University of Munich, D-80333 Munich, Germany, fax: +498928928340, telephone: +498928923403, [jonas.umlauf, hirche]@tum.de

The research leading to these results has received funding from the European Research Council under the European Union Seventh Framework Program (FP7/2007-2013) / ERC Starting Grant “Control based on Human Models (conhumo)” agreement no. 337654.

¹This is a preprint. For the original publication see <https://ieeexplore.ieee.org/document/8930275>

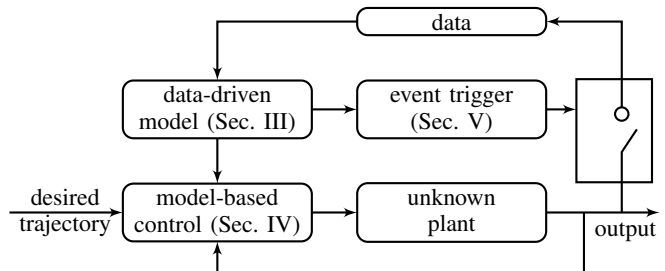


Fig. 1. Proposed concept of an online learning control law with event-triggered model updates.

allows for an implicit bias-variance trade-off [3] and, as GPs are a kernel-based method, prior knowledge (if any exists) can properly be transferred into the model [4]. The major advantage is that GP models also encode their own ignorance and therefore provide information as to whether the model is reliable for particular inputs or not.

Due to their nonparametric nature, the model complexity of a GP increases with the amount of available data. This possible unlimited expressive power is generally desired, however, it may cause difficulties from a computational point of view in the case of large training data sets. Particularly challenging are online learning schemes in which data points are accumulated over time and real-time capability is critical. This raises the question of efficient online learning strategies for nonparametric models. Time-triggered model adaptation fails to distinguish whether a new measurement or training point is necessary at the current location of the state-space or not. This calls for an event-triggered scheme, which decides upon a new measurement based on the current reliability of the model, which is expected to result in higher data-efficiency.

A. Related work

The fact that no model can initially capture all aspects of the true system motivated researchers to design robust and adaptive control methods to overcome this discrepancy [5]. The online adaptation of the control strategy or its employed model is well understood for parametric models [6], [7]. In particular for linear systems, data-driven approaches are extensively researched, see [8], and [9]. For nonlinear systems, model reference adaptive control (MRAC) is designed to effectively deal with model uncertainties or only little prior knowledge using online parameter estimation [10]. Iterative learning control (ILC) improves control performance by iteratively modulating the control signal in a repetitive task, such that experience from earlier executions is used to improve performance, see [11] and [12]. However, most existing MRAC and ILC methods

are mainly based on parametric models that suffer from limited complexity and flexibility. Model-free adaptive control (MFAC) avoids an explicit model but instead employs e.g. a dynamic linearization [13], virtual reference feedback tuning [14] or a closed-loop control parameter optimization [15]. Alternatively, a spectral analysis for nonparametric frequency-domain tuning is considered [16] or extremum seeking is employed for performance optimization [17].

Different than in classical control theory, the machine learning literature more frequently employs data-driven models with infinite expressive power for online adaptation [18]. The class of model-based reinforcement learning algorithms considers continuous model and controller updates to maximize a reward [19]. For example [20] shows a high data efficiency with Gaussian process models. These are also successfully applied in robotics [21], [22], however, most approaches lack a formal stability analysis for the system's behavior.

Very recently, several control approaches with formal guarantees for GP models have been developed that, however, keep a fixed dataset during execution of the control law [23], [24]. The work in [25] considers the control of Lagrangian systems and shows boundedness of the tracking error. The identification of a priori known stable systems with GPs is analyzed in [26], and [27] proposes an uncertainty-based control approach for which asymptotic stability is proven. However, none of these techniques updates the model while controlling the system. The work in [28] proposes a safe exploration by sequentially adding training points to the dataset, but it only stays within the region of attraction and cannot track an arbitrary trajectory in the state space.

An online learning tracking control law with time-triggered adaptation is proposed in [29]. As a result, data points are added to the training dataset irrespectively of their importance. This might compromise real-time capability, as the computational inefficiency for large datasets is a known challenge of GPs [3].

This difficulty is circumvented in [30], [31], where the unknown dynamics is estimated using high gain filters. However, these approaches suffer from the known difficulties of high gain control, i.e. a quick saturation of input signals and the amplification of noise. The latter is avoidable by combining feedback and model-based feedforward control. This idea is not just employed in this article, but was also used in [32], where a neural network identifies the dynamics without any parametric prior knowledge. Particularly [33], [34] and [35] focus on stability and performance guarantees. The work in [36] proposes a feedback linearizing control law that adapts the weights of a neural network model online. It shows boundedness of the adaptation law and the resulting controller but cannot quantify the ultimate bound because neural networks - in comparison to GPs - do not inherently provide a measure for the fidelity of the model [37]. This becomes important if the controller is applied in safety critical domains, where the tracking error must be quantified to avoid failure or damage to the system.

In summary, to date there exists no approach that adapts a nonparametric model online to guarantee the asymptotic stability of the tracking error. Thus, for universal models, which

can represent arbitrarily complex dynamics, there are missing online learning control laws to guarantee the safe behavior of the closed-loop system. Also, a data-efficient update strategy is required to keep the model computationally efficient, which is important in many real-time critical applications.

B. Contribution and structure

The main contribution of this article is an online learning feedback linearizing control law based on Gaussian processes for an initially unknown system. This control algorithm includes a closed-loop identification scheme for control affine systems exploiting compound kernels for GPs. To ensure the data-efficiency of the approach, we propose an event-triggered online learning mechanism that decides upon a model update based on its current reliability. The derivation is based on a probabilistic upper bound for the model error of a GP, and allows to provide safety guarantees in terms of the convergence properties of the closed-loop system. For noiseless training data, we show global asymptotic stability and for noisy output training data, global ultimate boundedness of the tracking error. For the case of a constraint budget for data points, we propose a forgetting strategy that maintains the convergence guarantees using a reduced number of training points.

The article is based on the preliminary work in [38], which focuses on the identification of a control affine system with GPs given a fixed dataset. In contrast, this work considers the online collection of data and updates the model while the control law is active. This allows to show asymptotic stability with a data-efficient event-triggered update rule while [38] only showed the existence of an ultimate bound.

This article is structured as follows: After formulating the considered problem formally in Sec. II, Sec. III reviews the identification of control affine systems based on GPs. In Sec. IV, the feedback linearizing tracking control law is proposed, including a convergence analysis for training data measured online at arbitrary time instances. Section V introduces an event-triggering to update the model based on its uncertainty. A numerical illustration is provided in Sec. VI followed by a conclusion in Sec. VII.

C. Notation

Lower/upper case bold symbols denote vectors/matrices, $\mathbb{R}_{+,0}/\mathbb{R}_+$ all real positive numbers with/without zero, \mathbb{N}_0/\mathbb{N} all natural numbers with/without zero, $\sigma_{\min}(\cdot), \sigma_{\max}(\cdot)$ the minimal/maximal singular value of a matrix and $\mathbb{E}[\cdot]/\mathbb{V}[\cdot]$ the expected value/variance of a random variable, respectively. \mathbf{I}_n denotes the $n \times n$ identity matrix, $\mathcal{N}(\mu, \sigma)$ a Gaussian distribution with mean μ and variance σ , $\mathbf{a}_{1:n}$ the first n elements of the vector \mathbf{a} , $\cdot \succ 0$ the positive definiteness of matrix or function and $\|\cdot\|$ the Euclidean norm if not stated otherwise.

II. PROBLEM FORMULATION

Consider a single-input system in the controllable canonical form

$$\begin{aligned} \dot{x}_1 &= x_2 \\ \dot{x}_2 &= x_3 \\ &\dots \\ \dot{x}_n &= f(\mathbf{x}) + g(\mathbf{x})u, \quad \mathbf{x}_0 = \mathbf{x}(0), \end{aligned} \quad (1)$$

with state $\mathbf{x} = [x_1 \ x_2 \ \dots \ x_n]^\top \in \mathbb{X} \subseteq \mathbb{R}^n$ and input $u \in \mathbb{U} = \mathbb{R}$; the functions $f(\cdot)$ and $g(\cdot)$ are considered unknown. The following assumptions are made.

Assumption 1: The unknown functions $f : \mathbb{X} \rightarrow \mathbb{R}$ and $g : \mathbb{X} \rightarrow \mathbb{R}$ are globally bounded and differentiable.

Differentiability is a very natural assumption, as it holds for many physical systems. The boundedness of the functions $f(\cdot), g(\cdot)$ would automatically be implied (due to the differentiability) if the set \mathbb{X} was bounded. However, we want \mathbb{X} to be possibly unbounded.

From Assumption 1, the first property is derived.

Lemma 1: Consider the system (1) under Assumption 1 with bounded and continuous $u(\mathbf{x})$. Then the solution $\mathbf{x}(t)$ does not have a finite escape time, thus $\nexists t_\infty, 0 < t_\infty < \infty$ for which

$$\lim_{t \rightarrow t_\infty} \|\mathbf{x}(t)\| = \infty. \quad (2)$$

Proof: According to [39, Theorem 3.2] the stated conditions ensure a unique solution $\mathbf{x}(t)$, for all $t > 0$ for which the finite escape time follows from the differentiability of $f(\cdot), g(\cdot)$ and the bounded control input. ■

As a stabilizing controller is not known in advance (because $f(\cdot), g(\cdot)$ are unknown), the absence of a finite escape time is important: It allows for the collection of observations of the system in any finite time interval with a “poor” controller (or also $u(\mathbf{x}) = 0$) without risking damage due to “infinite” states. Additionally, we assume the following.

Assumption 2: For system (1), $g(\mathbf{x}) > 0, \forall \mathbf{x} \in \mathbb{X}$ holds.

This ensures that the system’s relative degree is equal to the system order n for all $\mathbf{x} \in \mathbb{X}$ and the sign of $g(\cdot)$ is known. Equivalently, $g(\cdot)$ can also be taken as strictly negative, resulting in a change of sign for the control input. Assumption 2 is necessary to ensure global controllability and excludes the existence of internal dynamics. It restricts the system class, however, the focus of this work is on the online learning control, and extending it to larger system classes is part of future work.

We assume that observations are taken online while the proposed control law is active.

Assumption 3: Noiseless measurements of the state vector $\mathbf{x}^{(\kappa)} = \mathbf{x}(t_\kappa)$ and noisy measurements of the highest derivative $\mathbf{y}^{(\kappa)} = \dot{x}_n(t_\kappa) + \epsilon^{(\kappa)}$ can be taken at arbitrary time instances t_κ with $\kappa \in \mathbb{N}_0$. The observation noise $\epsilon^{(\kappa)} \sim \mathcal{N}(0, \sigma_{\text{on}}^2)$ is assumed Gaussian, independent and identically distributed. The time-varying dataset

$$\mathbb{D}_\kappa = \left\{ \mathbf{x}^{(i)}, y^{(i)} \right\}_{i=1}^{N_\kappa}, \quad (3)$$

is updated at time t_κ and remains constant until $t_{\kappa+1}$ and $N_\kappa \in \mathbb{N}_0$ denotes the current number of data points.

The exact measurement of the state is a common assumption and necessary for feedback linearization. The time derivative of the state x_n can, for practical applications, be approximated through finite differences. The approximation error is then considered as part of the measurement noise, as other additive sources of imprecision result in an overall sub-Gaussian noise distribution. Alternatively, a separate sensor for measurements of \dot{x}_n is necessary.

Throughout this article, we will refer to $\sigma_{\text{on}}^2 = 0$ as the *noiseless case* and $\sigma_{\text{on}}^2 > 0$ as the *noisy case* considering measurements of \dot{x}_n . The measurement of the state \mathbf{x} will always be assumed noise free.

Consider that N_κ is not necessarily increasing with increasing κ , as data pairs can also be discarded from the dataset if not needed anymore. However, this set \mathbb{D}_κ remains constant between two consecutive measurements, because elements are only added or removed at t_κ .

The goal is to design an online learning feedback linearizing control law - based on dataset \mathbb{D}_κ - of the form

$$u_\kappa(\mathbf{x}) = \frac{1}{\hat{g}_\kappa(\mathbf{x})} \left(-\hat{f}_\kappa(\mathbf{x}) + \nu \right), \quad \kappa \in \mathbb{N}_0, \quad (4)$$

where $\nu \in \mathbb{R}$ is the input to the resulting approximately linearized system and the functions $\hat{f}_\kappa : \mathbb{X} \rightarrow \mathbb{R}, \hat{g}_\kappa : \mathbb{X} \rightarrow \mathbb{R}$ are the approximations for the unknown functions $f(\mathbf{x}), g(\mathbf{x})$. The control law (4) is switching, because the model $\hat{f}_\kappa(\mathbf{x}), \hat{g}_\kappa(\mathbf{x})$ is updated with every change of the dataset \mathbb{D}_κ at time t_κ . We would like to emphasize that measurements are not taken at a constant time interval, and updates are therefore not performed periodically. Instead, the updates will be performed when needed, i.e. triggered by an event (introduced in Sec. V) and thus t_κ for $\kappa \in \mathbb{N}_0$ are not equidistant. By definition, the κ -th update occurs at t_κ and the control law u_κ is then applied until the next event at $t_{\kappa+1}$, more formally written as

$$u(\mathbf{x}) = u_\kappa(\mathbf{x}), \quad t \in [t_\kappa, t_{\kappa+1}). \quad (5)$$

III. GAUSSIAN PROCESS LEARNING FOR CONTROL AFFINE SYSTEMS

For the closed-loop online identification of $f(\cdot)$ and $g(\cdot)$ we consider Gaussian process regression, which then provides the approximations $\hat{f}_\kappa(\cdot)$ and $\hat{g}_\kappa(\cdot)$. We will first introduce GP regression in general (Sec. III-A) before presenting our tailored solution for control affine closed-loop systems in Sec. III-B.

A. Gaussian process regression

Consider a function $f_{\text{true}} : \mathbb{X} \rightarrow \mathbb{R}$ for which noisy measurements of the image at the locations $\mathbf{x}^{(i)} \in \mathbb{X}$ are available, thus

$$y_f^{(i)} = f_{\text{true}}(\mathbf{x}^{(i)}) + \epsilon^{(i)}, \quad (6)$$

where $\epsilon^{(i)} \sim \mathcal{N}(0, \sigma_{\text{on}}^2)$ and $i = 1, \dots, N$ (where we simply write N for N_κ in this section). Modeling this function with a Gaussian process $f_{\text{GP}}(\mathbf{x})$ results in a stochastic process that assigns a Gaussian distribution to any finite subset $\{\mathbf{x}_1, \dots, \mathbf{x}_M\} \subset \mathbb{X}$ in a continuous domain. The GP is also

often considered as a distribution over functions [3], denoted by

$$f_{\mathcal{GP}}(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')), \quad (7)$$

and is fully specified by a mean $m(\mathbf{x}) : \mathbb{X} \rightarrow \mathbb{R}$ and covariance $k(\mathbf{x}, \mathbf{x}') : \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}$ function. The mean function includes prior knowledge of the function f_{true} if there is any. Otherwise, it is commonly set to zero. The covariance function, also called kernel function, determines the properties of $f_{\mathcal{GP}}(\mathbf{x})$, like the smoothness and signal variance. Mean and kernel function are described by the hyperparameters ψ .

Using Bayesian techniques, the likelihood function

$$\psi^* = \arg \max_{\psi} \log p(\mathbf{y}_f | \mathbf{X}, \psi), \quad (8)$$

$$\log p(\mathbf{y}_f | \mathbf{X}, \psi) = \frac{1}{2} (\mathbf{y}_f^T \mathbf{K}^{-1} \mathbf{y}_f - \log \det \mathbf{K} - N \log(2\pi)),$$

is maximized to obtain the optimal hyperparameters for a given set of observations. As notation, we use

$$\mathbf{X} = [\mathbf{x}^{(1)} \ \dots \ \mathbf{x}^{(N)}] \in \mathbb{R}^{n \times N}, \quad (9)$$

$$\mathbf{y}_f = [y_f^{(1)} \ \dots \ y_f^{(N)}]^T \in \mathbb{R}^N, \quad (10)$$

to denote the input/output data, respectively, and

$$\mathbf{K} = \begin{bmatrix} k(\mathbf{x}^{(1)}, \mathbf{x}^{(1)}) & \dots & k(\mathbf{x}^{(1)}, \mathbf{x}^{(N)}) \\ \vdots & \ddots & \vdots \\ k(\mathbf{x}^{(N)}, \mathbf{x}^{(1)}) & \dots & k(\mathbf{x}^{(N)}, \mathbf{x}^{(N)}) \end{bmatrix} \in \mathbb{R}^{N \times N} \quad (11)$$

concatenates the kernel evaluations of pairs of input data. Although the optimization (8) is generally non-convex, it is usually performed with conjugated gradient-based methods [3]. Each local minimum can be considered as a different interpretation of data and we discuss the effect of suboptimal identification in Sec. III-C.

In a regression task, GPs employ the joint Gaussian distribution of training data \mathbf{X}, \mathbf{y}_f and a test input \mathbf{x}^*

$$\begin{bmatrix} f_{\mathcal{GP}}(\mathbf{x}^*) \\ \mathbf{y}_f \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} m(\mathbf{x}^*) \\ \mathbf{m}^{\mathbf{X}} \end{bmatrix}, \begin{bmatrix} k^* & \mathbf{k}^T \\ \mathbf{k} & \mathbf{K} + \sigma_{\text{on}}^2 \mathbf{I}_N \end{bmatrix} \right), \quad (12)$$

where

$$\mathbf{m}^{\mathbf{X}} = [m(\mathbf{x}^{(1)}) \ \dots \ m(\mathbf{x}^{(N)})]^T, \quad (13)$$

to find the posterior mean and variance function

$$\mu(\mathbf{x}^*) := \mathbb{E}[f_{\mathcal{GP}}(\mathbf{x}^*) | \mathbf{X}, \mathbf{y}_f] \quad (14)$$

$$= m(\mathbf{x}^*) + \mathbf{k}^T (\mathbf{K} + \sigma_{\text{on}}^2 \mathbf{I}_N)^{-1} (\mathbf{y}_f - \mathbf{m}^{\mathbf{X}}),$$

$$\sigma(\mathbf{x}^*) := \mathbb{V}[f_{\mathcal{GP}}(\mathbf{x}^*) | \mathbf{X}, \mathbf{y}_f] \quad (15)$$

$$= k^* - \mathbf{k}^T (\mathbf{K} + \sigma_{\text{on}}^2 \mathbf{I}_N)^{-1} \mathbf{k},$$

through conditioning, where

$$k^* = k(\mathbf{x}^*, \mathbf{x}^*),$$

$$\mathbf{k} = [k(\mathbf{x}^{(1)}, \mathbf{x}^*) \ \dots \ k(\mathbf{x}^{(N)}, \mathbf{x}^*)]^T \in \mathbb{R}^N. \quad (16)$$

However, considering the defined problem in Sec. II, the classical GP regression framework cannot be directly applied because closed-loop measurements do not provide data points for $f(\cdot)$ and $g(\cdot)$ separately. Therefore, the following section explains how it is augmented using the given prior knowledge.

B. Closed-loop identification with prior knowledge

First, we transfer the knowledge on the positivity of the function $g(\mathbf{x})$ from Assumption 2 into the model $\hat{g}(\mathbf{x})$. It is crucial to utilize this knowledge for the model to ensure the feedback linearizing control (4) results in well-behaved control signals. Using a GP model for $\hat{g}(\mathbf{x})$, this can be ensured using a proper prior mean function.

Lemma 2: Consider the posterior mean function (14) with a bounded and differentiable kernel $k(\cdot, \cdot)$ and a dataset $(\mathbf{X}, \mathbf{y}_f)$ for which $\mathbf{x}^{(i)} \neq \mathbf{x}^{(i')}$ and $y_f^{(i)} > 0$, hold $\forall i, i' = 1, \dots, N, i \neq i'$. Then, there exists a differentiable prior mean function $m(\mathbf{x})$ such that

$$\mu(\mathbf{x}) > 0, \quad \forall \mathbf{x} \in \mathbb{X}. \quad (17)$$

Proof: Consider a prior mean function for which $0 < m(\mathbf{x}^{(i)}) < \infty$, $\forall i = 1, \dots, N$ holds, then, a differentiable $m(\mathbf{x})$ can $\forall \mathbf{x} \in \mathbb{X} \setminus \{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)}\}$ always be chosen larger than the constant $\mathbf{k}^T (\mathbf{K} + \sigma_{\text{on}}^2 \mathbf{I}_N)^{-1} (\mathbf{y}_f - \mathbf{m}^{\mathbf{X}})$, because the latter is bounded. For $\mathbf{x} \in \{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)}\}$ a choice $m(\mathbf{x}^{(i)}) = y_f^{(i)}$ (which complies with the first condition) ensures that $\mu(\mathbf{x})$ is strictly positive. ■

Remark 1: Since $g(\cdot)$ is strictly positive by Assumption 2, the condition $y_f^{(i)} > 0$ follows naturally. In case the Gaussian noise results in negative measurements $y^{(i)}$, it can be corrected using $\max(y^{(i)}, \eta)$, with an arbitrarily small $\eta > 0$. Alternatively, strictly positive noise distributions, e.g. a Gamma distribution, can also be combined with Gaussian process regression [3].

In practice, it is often sufficient to set $m(\mathbf{x})$ to a positive constant. To verify that $\mu(\mathbf{x}) > 0$ holds, the techniques in [28] can be utilized. The suitable prior mean function according to Lemma 2 will be denoted by $m_g(\mathbf{x})$.

Second, the major difficulty of closed-loop identification is to differentiate the effect of the control input and the unforced dynamics. For the control affine structure, this means that individual measurements of the functions $f(\cdot)$ and $g(\cdot)$ from (1) are not provided. Thus, functions $f(\cdot), g(\cdot)$ must be identified by only observing their sum exploiting the control affine structure. We propose utilizing a compound kernel as reviewed in Appendix A based on [4]. More specifically, we use the composite kernel

$$k(\mathbf{x}, \mathbf{x}') = k_f(\mathbf{x}, \mathbf{x}') + u(\mathbf{x}) k_g(\mathbf{x}, \mathbf{x}') u(\mathbf{x}'), \quad (18)$$

which replicates the structure of a control affine system: the first summand $k_f(\cdot, \cdot)$ represents the unknown unforced dynamics $f(\cdot)$; the second summand $u(\cdot) k_g(\cdot, \cdot) u(\cdot)$ the product of the unknown scaling of the control $g(\cdot)$ and the known state feedback control term $u(\cdot)$. As no further knowledge regarding $f(\cdot), g(\cdot)$ is given, we employ two squared exponential (SE) kernels with automatic relevance determination

$$k_f(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \exp \left(\sum_{j=1}^n \frac{(x_j - x'_j)^2}{-2l_{j,f}^2} \right), \quad (19)$$

$$k_g(\mathbf{x}, \mathbf{x}') = \sigma_g^2 \exp \left(\sum_{j=1}^n \frac{(x_j - x'_j)^2}{-2l_{j,g}^2} \right), \quad (20)$$

where the hyperparameters are the lengthscales $l_{j,f}, l_{j,g} \in \mathbb{R}_+$, $j = 1, \dots, n$ and the signal variances $\sigma_f, \sigma_g \in \mathbb{R}_+^0$. For notational convenience, they are concatenated in the vector

$$\boldsymbol{\psi}_{fg} = [l_{1,f} \ l_{1,g} \ \cdots \ l_{n,f} \ l_{n,g} \ \sigma_f^2 \ \sigma_g^2]^\top. \quad (21)$$

The SE kernel is *universal* and therefore allows for arbitrarily exact modeling of any continuous function according to [40].

Remark 2: GP models with structured kernels, like (18), must not be confused with parametric models, which have a predetermined structure and use a fixed number of parameters. In contrast, a GP with a structured kernel has potentially infinitely many parameters for each part of its structure. So the kernel encodes the knowledge that the unknown function, for example, is comprised of a sum, but each summand has unlimited flexibility.

We denote $\mathbf{U} = \text{diag}(u_1(\mathbf{x}^{(1)}), \dots, u_N(\mathbf{x}^{(N)})) \in \mathbb{R}^{N \times N}$, where u_i denotes the control law that was active at the time at which the pair $\{\mathbf{x}^{(i)}, y^{(i)}\}$ was collected for $i = 1, \dots, N$. Furthermore, $\mathbf{m}_g^{\mathbf{X}}, \mathbf{y}$ are analogously defined to (13), (10), respectively. Then

$$\mathbf{K}_{fg} = \mathbf{K}_f + \mathbf{U}^\top \mathbf{K}_g \mathbf{U} + \sigma_{\text{on}}^2 \mathbf{I}_n, \quad (22)$$

and $\mathbf{k}_f, \mathbf{k}_g, \mathbf{K}_f, \mathbf{K}_g$ are defined analogously to (16) and (11) using $k_f(\mathbf{x}, \mathbf{x}'), k_g(\mathbf{x}, \mathbf{x}')$. This notation allows for the formulation of the estimates $\hat{f}(\mathbf{x}), \hat{g}(\mathbf{x})$.

Lemma 3: The GP posterior mean prediction for the functions $f(\mathbf{x}), g(\mathbf{x})$, based on the training data \mathbb{D}_κ in (3) for the compound kernel (18) are given by

$$\hat{f}(\mathbf{x}) := \mu_f(\mathbf{x}) = \mathbf{k}_f^\top \mathbf{K}_{fg}^{-1} (\mathbf{y} - \mathbf{U} \mathbf{m}_g^{\mathbf{X}}), \quad (23)$$

$$\hat{g}(\mathbf{x}) := \mu_g(\mathbf{x}) = m_g(\mathbf{x}) + \mathbf{k}_g^\top \mathbf{U} \mathbf{K}_{fg}^{-1} (\mathbf{y} - \mathbf{U} \mathbf{m}_g^{\mathbf{X}}), \quad (24)$$

where the prior mean function for $\hat{f}(\mathbf{x})$ is set to zero, $m_f(\mathbf{x}) = 0$, and for $\hat{g}(\mathbf{x})$, $m_g(\mathbf{x})$ is chosen according to Lemma 2.

Proof: For an input \mathbf{x} and the compound kernel (18), the joint distribution is given by

$$\begin{bmatrix} f(\mathbf{x}) \\ g(\mathbf{x}) \\ \mathbf{y} \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} 0 \\ m_g(\mathbf{x}) \\ \mathbf{U} \mathbf{m}_g^{\mathbf{X}} \end{bmatrix}, \begin{bmatrix} k_f^* & 0 & \mathbf{k}_f^\top \\ 0 & k_g^* & \mathbf{k}_g^\top \mathbf{U}^\top \\ \mathbf{k}_f & \mathbf{U} \mathbf{k}_g & \mathbf{K}_{fg} \end{bmatrix} \right), \quad (25)$$

similarly to (12). According to [4], the posterior mean functions (23) and (24) correspond equivalently to (14). ■

For these estimates, it can be shown that all prior knowledge is properly transferred into the model.

Proposition 1: Consider a control affine system (1) under Assumptions 1-3 and the compound kernel (18). Then, the estimates $\hat{f}(\mathbf{x})$ and $\hat{g}(\mathbf{x})$ in Lemma 3 are bounded and infinitely differentiable, and there exists a prior mean function $m_g(\mathbf{x})$ and a hyperparameter vector $\boldsymbol{\psi}_{fg}$ such that $\hat{g}(\mathbf{x}) > 0$ holds $\forall \mathbf{x} \in \mathbb{X}$.

Proof: The SE kernel inherits its properties differentiability and boundedness to all functions represented by the GP [3], thus also to the posterior mean functions, which are used as estimates. The strict positivity of $\hat{g}(\mathbf{x})$ follows from the fact that σ_g^2 can be made arbitrarily small, such that there always

exists a positive function m_g , such that $m_g(\mathbf{x})$ dominates the term $\mathbf{k}_g^\top \mathbf{U} \mathbf{K}_{fg}^{-1} (\mathbf{y} - \mathbf{U} \mathbf{m}_g^{\mathbf{X}})$ in (24). ■

Remark 3: The only properties of the SE kernel that are used for the derivations and proofs are its differentiability and its boundedness. Thus, the conclusions can directly be extended to other kernel functions fulfilling these properties. For the sake of focus, in this article we will consider the SE kernel only.

C. Discussion

The most obvious challenge of the closed-loop identification is that no unique yet infinitely many solutions exist for two differentiable functions to add up to the same values. Thus, only observing the sum in (1) is not promising for learning the unique correct individual functions $f(\cdot), g(\cdot)$ because it is an under-determined problem. The estimates in (23) and (24) are just one of many solutions, determined by the choice of hyperparameters, which suit the training data. Nevertheless, the optimization (8) interprets the observed data to match the kernel structure, which is shown to be successful in the simulation in Sec. VI-A. For the case in which the results are not satisfactory, we provide an extension in Appendix B to address this challenge. It merges data points of the closed-loop system with measurements from the temporary open-loop system. It thereby uses Lemma 1, which allows to safely turn off the control signal ($u = 0$) for a finite time period. Nevertheless, we want to highlight that the formal guarantees provided in the following section (Theorem 1) hold independently whether this extension is utilized or not.

Furthermore, additional knowledge like the periodicity or dependence of $f(\cdot)$ or $g(\cdot)$ on only a subset of the state variables can also be transferred into the kernel to facilitate the identification by using a periodic kernel or setting the lengthscales $l_{j,f} = l_{j',g} = \infty$ for the states j, j' of which they are independent, respectively. The latter simplifies the optimization of hyperparameters as the search space is reduced. A systematic way of constructing a more evolved kernel function (including more prior knowledge) is discussed in [4].

Considering the computational load, the inverse of \mathbf{K}_{fg} is most critical, as the number of operations increases cubically with the number of data points, thus $\mathcal{O}(N^3)$. However, adding further data points is necessary to ensure the model is precise at the current position in the state space where the most recent data points are taken from measurements. Compared to previous approaches, e.g. [38], where $\mathbf{K}_{fg}^{-1} \mathbf{y}$ is constant and can thereby be precomputed offline, here it must be recomputed with every update of the model as data points are added one at a time. This difficulty can be addressed using a rank-1 update of the inverse with the Sherman–Morrison formula [41] resulting in only $\mathcal{O}(N^2)$ operations. However, this quadratic computational complexity might still be very time consuming, and therefore lead to the data-efficient event-triggered model updates introduced in Sec. V.

Generally, Gaussian processes turn out to be very effective for the adaptive model control law: They properly transfer all prior assumptions consistently into the model (Proposition 1) and allow for the identification in a closed loop. The

nonparametric nature allows for unlimited model flexibility and the complexity increases as more data is available in a data-driven fashion. This is a crucial advantage compared to classical system identification methods, particularly for highly nonlinear systems.

IV. FEEDBACK LINEARIZING CONTROL LAW

In this section, the feedback linearizing online learning control law is proposed and the resulting closed-loop behavior is analyzed. After showing the ultimate boundedness for the most general case, we make further specific assumptions to provide stronger stability results. Here, further properties of the Gaussian process modeling technique are exploited: As the model error of the GP can be bounded and quantified, the ultimate bound of the tracking error can also be quantified.

Classical model reference adaptive control modifies the model parameters continuously over time, which is not possible here due to the nonparametric nature of the GP model. Thus, particular attention must be drawn to the resulting switching character of the control law, which stems from the time-varying dataset \mathbb{D}_κ of the Gaussian process introduced in Assumption 3.

We are interested in tracking a desired trajectories for the state x_1 , given by $x_d(t)$, with the following property.

Assumption 4: The desired trajectory $x_d(t)$ is bounded and at least $n - 1$ times differentiable, thus

$$\mathbf{x}_d(t) = \begin{bmatrix} x_d & \dot{x}_d & \cdots & \frac{d^{n-1}x_d}{dt^{n-1}} \end{bmatrix}^\top \quad (26)$$

is continuous and $\frac{d^n x_d}{dt^n}$ is bounded.²

For notational convenience, we define the tracking error

$$\mathbf{e} = \mathbf{x} - \mathbf{x}_d. \quad (27)$$

A. Control law

Consider the filtered scalar state $r \in \mathbb{R}$, defined as

$$r = [\boldsymbol{\lambda}^\top \quad 1] \mathbf{e}, \quad (28)$$

where $\boldsymbol{\lambda} = [\lambda_1 \ \lambda_2 \cdots \lambda_{n-1}]^\top \in \mathbb{R}^{n-1}$ is a coefficient vector such that for $s \in \mathbb{C}$ the polynomial $s^{n-1} + \lambda_{n-1}s^{n-2} + \cdots + \lambda_1$ is Hurwitz. Under this condition, the error converges exponentially $e \rightarrow \mathbf{0}$ as $r \rightarrow 0$ [36]. The dynamics of the filtered state is

$$\dot{r} = f(\mathbf{x}) + g(\mathbf{x})u(\mathbf{x}) + \rho, \quad (29)$$

where

$$\rho = \boldsymbol{\lambda}^\top \mathbf{e}_{2:n} - \frac{d^n x_d}{dt^n}, \quad (30)$$

with $\mathbf{e}_{2:n} = [e_2 \cdots e_n]^\top \in \mathbb{R}^{n-1}$. For the control law $u(\mathbf{x})$, we propose

$$u_\kappa(\mathbf{x}) = \frac{1}{\hat{g}_\kappa(\mathbf{x})} \left(-\hat{f}_\kappa(\mathbf{x}) - k_c r - \rho \right), \quad (31)$$

according to (4) where $\nu = -k_c r - \rho$ with $k_c \in \mathbb{R}_+$ is used. The subscript $\kappa \in \mathbb{N}_0$ indicates the κ -th time interval $t \in [t_\kappa \ t_{\kappa+1})$ for which u_κ is applied according to (5).

²The t dependency of \mathbf{x}_d, \mathbf{x} and the \mathbf{x} dependencies of $u, f, \hat{f}, g, \hat{g}$ are partially omitted for notational convenience.

The estimates $\hat{g}_\kappa(\cdot), \hat{f}_\kappa(\cdot)$ are based on N_κ training points in the time-varying dataset \mathbb{D}_κ introduced in Assumption 3. The control scheme is visualized in Fig. 2 and the adaptation procedure is provided in Algorithm 1.

Note that even though GPs themselves are probabilistic models, the control law is deterministic because it only employs the posterior mean functions as a model estimate $\hat{g}_\kappa(\cdot), \hat{f}_\kappa(\cdot)$.

Algorithm 1 Online learning for feedback linearization control

- 1: initialize $\kappa = 0, \mathbb{D}_0 = \{\}, \hat{f}_0 = 0, \hat{g}_0 = m_g(\cdot)$
 - 2: **while** simulation time not exceeded **do**
 - 3: **while** $t < t_{\kappa+1}$ **do**
 - 4: run controller u_κ in (31)
 - 5: **end while**
 - 6: set $\kappa \leftarrow \kappa + 1$
 - 7: measure $\mathbf{x}^{(\kappa)} = \mathbf{x}(t_\kappa)$ and $y^{(\kappa)} = \dot{x}_n(t_\kappa) + \epsilon^{(\kappa)}$
 - 8: add training point $\mathbb{D}_\kappa = \mathbb{D}_{\kappa-1} \cup \{(\mathbf{x}^{(\kappa)}, y^{(\kappa)})\}$
 - 9: update the estimates $\hat{f}_\kappa(\cdot), \hat{g}_\kappa(\cdot)$ in (23), (24)
 - 10: **end while**
-

B. Convergence analysis

An offline version of the control law (31), with constant dataset \mathbb{D} and estimates $\hat{f}(\mathbf{x}), \hat{g}(\mathbf{x})$ was introduced previously and was shown to be globally uniformly ultimately bounded [38, Proposition 1].

But, with the time-varying dataset and model, Algorithm 1 describes a switching control law. This switching results in a hybrid system, where some states change continuously in time but the system dynamics (or other states) change at discrete time instances. Here, the resulting closed-loop system is subject to (in general) arbitrary switching. Its convergence behavior is analyzed here based on the principle of a common Lyapunov function. It states that a Lyapunov function, which is independent of the switching signal must decrease over time along the system's trajectories. This is shown in the following.

Theorem 1: Consider the system (1) and a desired trajectory $\mathbf{x}_d(t)$ under Assumptions 1-4. Further consider the control law (31), where $f(\cdot), g(\cdot)$ are modeled by GP mean functions $\hat{f}_\kappa(\cdot), \hat{g}_\kappa(\cdot)$ in (23) and (24), respectively. The GP model is updated at arbitrary switching times t_κ according to Algorithm 1. Then, there exists a $k_c^* > 0$ such that for every $k_c \geq k_c^*$ the tracking error $\|\mathbf{e}\|$ is globally uniformly ultimately bounded.

Proof: Consider the common Lyapunov function candidate

$$V_\kappa(\mathbf{x}) = r^2/2, \quad \forall \kappa \in \mathbb{N}_0, \quad (32)$$

with time derivative

$$\begin{aligned} \dot{V}_\kappa(\mathbf{x}) &= r\dot{r} = r(f + gu_\kappa + \rho) \\ &= r \left(f + \frac{g}{\hat{g}_\kappa} (-\hat{f}_\kappa - k_c r - \rho) + \rho \right) \\ &= r \left(f - \bar{g}_\kappa \hat{f}_\kappa \right) - k_c \bar{g}_\kappa r^2 + (1 - \bar{g}_\kappa) r \rho, \end{aligned} \quad (33)$$

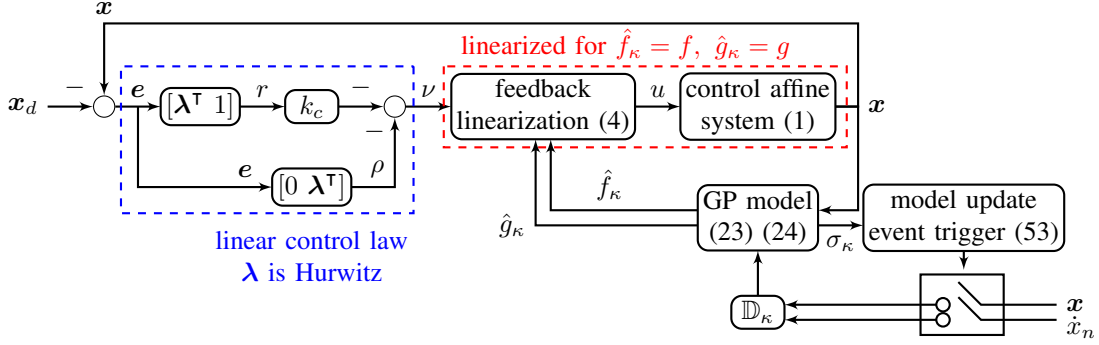


Fig. 2. The online learning feedback linearizing control scheme including the event trigger proposed in Sec. V, which controls the switching time $t_{\kappa+1}$.

where $\bar{g}_\kappa := \frac{g(\mathbf{x})}{\hat{g}_\kappa(\mathbf{x})}$ is positive and bounded $\forall \kappa$ and $\forall \mathbf{x} \in \mathbb{X}$ from Proposition 1 and Assumptions 1 and 2. As a consequence $(f - \bar{g}_\kappa \hat{f}_\kappa)$ is bounded and there exists a constant $\mathbf{a} \in \mathbb{R}^n$ such that

$$\|r(f - \bar{g}_\kappa \hat{f}_\kappa)\| \leq \|\mathbf{a}^\top \mathbf{e}\| \quad \forall \mathbf{e}, \kappa \quad (34)$$

holds, because r only grows linearly in \mathbf{e} . For similar reasons, we can find constants $\mathbf{c} \in \mathbb{R}^n$ and $\mathbf{B}, \mathbf{C} \in \mathbb{R}^{n \times n}$, with $\mathbf{B}, \mathbf{C} \succ 0$ and

$$\|\bar{g}_\kappa r^2\| \geq \mathbf{e}^\top \mathbf{B} \mathbf{e}, \quad \forall \mathbf{e}, \kappa \quad (35)$$

$$\|(1 - \bar{g}_\kappa) r \rho\| \leq \mathbf{e}^\top \mathbf{C} \mathbf{e} + \mathbf{c}^\top \mathbf{e}, \quad \forall \mathbf{e}, \kappa \quad (36)$$

which exist since $f, \hat{f}_\kappa, \bar{g}_\kappa$ are bounded. Therefore,

$$\begin{aligned} \dot{V}_\kappa(\mathbf{x}) &\leq \|\mathbf{a}\| \|\mathbf{e}\| - k_c \sigma_{\min}(\mathbf{B}) \|\mathbf{e}\|^2 + \sigma_{\max}(\mathbf{C}) \|\mathbf{e}\|^2 + \|\mathbf{c}\| \|\mathbf{e}\| \\ &= \|\mathbf{e}\| (\|\mathbf{a}\| + \|\mathbf{c}\|) + \|\mathbf{e}\|^2 (\sigma_{\max}(\mathbf{C}) - k_c \sigma_{\min}(\mathbf{B})) \end{aligned} \quad (37)$$

holds for all κ and there exists a $k_c^* > 0$ such that

$$\sigma_{\max}(\mathbf{C}) - k_c^* \sigma_{\min}(\mathbf{B}) < 0. \quad (38)$$

As a result, for every $k_c \geq k_c^*$, the Lyapunov function decreases

$$\dot{V}_\kappa(\mathbf{x}) < 0, \quad \forall \mathbf{x} \in \mathbb{X} \setminus \mathbb{B}, \quad \forall \kappa \quad (39)$$

outside of the set

$$\mathbb{B} = \left\{ \forall \mathbf{x} \in \mathbb{X} \left\| \|\mathbf{e}\| \leq \frac{\|\mathbf{a}\| + \|\mathbf{c}\|}{k_c \sigma_{\min}(\mathbf{A}) - \sigma_{\max}(\mathbf{B})} \right\} \right\}, \quad (40)$$

which forms a tube in \mathbf{x} coordinates around the desired trajectory and a ball in \mathbf{e} coordinates. Thus, we have found a common radially unbounded Lyapunov function $V(\mathbf{x})$, which decreases $\forall \kappa$ outside of the ball \mathbb{B} . According to [42, Theorem 2.1], this allows for the conclusion that for arbitrary switching sequences the tracking error converges to the ball \mathbb{B} . Since \mathbb{B} is independent of the initial state, global uniform ultimate boundedness (GUUB) holds. ■

Thus, we have shown that the tracking error is bounded by the proposed online learning control scheme for a large enough gain k_c with an arbitrary switching sequence. However, without further knowledge, a value for the critical gain k_c^* cannot be computed. We therefore make the following simplifying assumption

Assumption 5: The function $g(\mathbf{x})$ is known, thus $\hat{g}(\mathbf{x}) = g(\mathbf{x})$ and noisy training data observations of \dot{x}_n

$$y_f^{(i)} = f(\mathbf{x}^{(i)}) + \epsilon^{(i)} = \dot{x}_n^{(i)} - g(\mathbf{x}^{(i)}) u_\kappa(\mathbf{x}^{(i)}) + \epsilon^{(i)} \quad (41)$$

with $\epsilon^{(i)} \sim \mathcal{N}(0, \sigma_{\text{on}}^2)$, $i = 1, \dots, N_\kappa$ are available.

This assumption holds for many real-world systems, e.g. most Lagrangian systems, which is a considerably large class [43]. It is also a quite common assumption when working with control affine systems [44].

Analogously to (23), the unknown function is now estimated by

$$\hat{f}(\mathbf{x}) := \mu_f(\mathbf{x}) = \mathbf{k}_f^\top (\mathbf{K}_f + \sigma_{\text{on}}^2 \mathbf{I}_{N_\kappa})^{-1} \mathbf{y}_f, \quad (42)$$

where $\mathbf{y}_f = [y_f^{(1)} \dots y_f^{(N_\kappa)}]^\top$ and $\mathbf{k}_f, \mathbf{K}_f$ are computed according to (16) and (11) for the SE kernel. Under this assumption, Theorem 1 can be relaxed with respect to the choice of the gain k_c .

Corollary 1: Consider the system (1) and a desired trajectory $\mathbf{x}_d(t)$ under Assumptions 1-5. Further consider the control law (31), where $f(\cdot)$ is modeled by a GP mean function $\hat{f}_\kappa(\cdot)$ in (42), which is adapted at arbitrary switching times t_κ according to Algorithm 1. Then, the tracking error $\|\mathbf{e}\|$ of the closed-loop switching system is globally uniformly ultimately bounded for any $k_c > 0$.

Proof: Using the Lyapunov function (32), we obtain for $\bar{g}_\kappa = 1$

$$\dot{V}_\kappa(\mathbf{x}) = r\dot{r} = r(f + g u_\kappa + \rho) \quad (43)$$

$$= r(f - \hat{f}_\kappa) - k_c r^2, \quad (44)$$

which leaves us with the condition $(f - \hat{f}_\kappa) < k_c r$ for negative definiteness of $\dot{V}_\kappa(\mathbf{x})$. Thus, independent of the gain k_c , a ball exists outside of which the Lyapunov function decreases $\forall \kappa$, which leads to the GUUB $\forall k_c > 0$. ■

For completeness, we also formalize the result for dataset \mathbb{D}_κ which remains constant after t_κ , thus no further measurements are taken and no data points are deleted from the set for $t > t_\kappa$.

Corollary 2: Consider the system (1) and a desired trajectory $\mathbf{x}_d(t)$ under Assumptions 1-5. Further consider the control

law (31), where $f(\cdot)$ is modeled by a GP mean function $\hat{f}_\kappa(\cdot)$ in (42) with a fixed dataset \mathbb{D}_κ . Then, the tracking error $\|e\|$ of the closed-loop system is globally uniformly ultimately bounded for any $k_c > 0$.

Proof: The proof is straightforward as it follows from the proof for Corollary 1. ■

C. Quantifying the ultimate bound

Theorem 1 and Corollaries 1 and 2 show that an ultimate bound for the tracking error e exists, however, its size is unknown. To quantify the ultimate bound \mathbb{B} , an upper bound for the model estimate, defined as

$$\Delta f_\kappa(\mathbf{x}) = |f(\mathbf{x}) - \hat{f}_\kappa(\mathbf{x})|, \quad \forall \kappa, \quad (45)$$

is derived in this section using the variance function $\sigma_\kappa(\mathbf{x}): \mathbb{X} \rightarrow \mathbb{R}_+^0$ of the GP as defined in (15). Since the GP is a probabilistic model in nature, we cannot expect any deterministic statements regarding the error of the estimate Δf_κ . However, according to [45], it is possible to make high probability statements regarding the maximum distance from the true function $f(\mathbf{x})$ to the mean function $\mu(\mathbf{x})$ on a compact set. As known from the no-free lunch theorems [46], this generalization cannot be expected without any prior knowledge about $f(\cdot)$. Since we do not want to make any parametric assumptions that limit the complexity of $f(\cdot)$, we restrict its reproducing kernel Hilbert space (RKHS) norm as follows.

Assumption 6: The function $f(\mathbf{x})$ has a bounded reproducing kernel Hilbert space (RKHS) norm with respect to a squared exponential kernel $k(\cdot, \cdot)$, with known hyperparameters denoted by $\|f(\mathbf{x})\|_k^2 \leq B_f$.

With this additional assumption, a high probability statement regarding the precision of the mean function estimate is possible according to [45].

Lemma 4: Suppose Assumption 6 holds, then

$$\Pr \left\{ |\mu_\kappa(\mathbf{x}) - f(\mathbf{x})| \leq \beta_\kappa \sigma_\kappa(\mathbf{x}), \forall \mathbf{x} \in \tilde{\mathbb{X}}, N_\kappa \in \mathbb{N}_0 \right\} \geq 1 - \delta, \quad (46)$$

holds on a compact set $\tilde{\mathbb{X}} \subset \mathbb{R}^n$, where $\delta \in (0, 1)$, $\beta_\kappa = \sqrt{2B_f + 300\gamma_\kappa \log^3((\kappa + 1)/\delta)}$ and γ_κ is the maximum mutual information that can be obtained about $f(\cdot)$ from $\kappa + 1$ noisy samples $\mathbf{x}_{(1)}, \dots, \mathbf{x}_{(\kappa+1)}$ and $\mu_\kappa(\cdot)$ and $\sigma_\kappa(\cdot)$ are the posterior mean and variance function of a GP for N_κ data points as defined in (14) and (15), respectively.

Proof: This is a direct consequence from [45, Theorem 6]. ■

Remark 4: Consider that (46) takes all $N_\kappa \in \mathbb{N}_0$ into account at once. This means the probability δ holds not just for a single $N_\kappa \in \mathbb{N}_0$ but for all $N_\kappa \in \mathbb{N}_0$. This becomes clear when rewriting (46) as

$$\Pr \left\{ \bigcap_{N_\kappa=0}^{\infty} |\mu_\kappa(\mathbf{x}) - f(\mathbf{x})| \leq \beta_\kappa \sigma_\kappa(\mathbf{x}), \forall \mathbf{x} \in \tilde{\mathbb{X}} \right\} \geq 1 - \delta. \quad (47)$$

The model error bound in Lemma 4 only holds on a compact set $\tilde{\mathbb{X}}$. Nevertheless, we have already shown in Theorem 1

and Corollary 1 that the tracking error converges to a compact set \mathbb{B} . Thus, we set $\tilde{\mathbb{X}} = \mathbb{B}$, which leads to the following result.

Theorem 2: Consider the system (1) and a desired trajectory $\mathbf{x}_d(t)$ under Assumptions 1-6. Further consider the control law (31), where $f(\cdot)$ is modeled by a GP mean function $\hat{f}_\kappa(\cdot)$ in (42), which is adapted at arbitrary switching times t_κ according to Algorithm 1. Then, with probability $1 - \delta$, $\delta \in (0, 1)$, the tracking error $\|e\|$ is uniformly ultimately bounded for any $k_c > 0$ with the ultimate bound

$$\mathbb{B}_\kappa = \left\{ \forall \mathbf{x} \in \mathbb{X} \left\| \|e\| \leq \frac{\beta_\kappa \bar{\sigma}_\kappa}{k_c \|\lambda^\top \mathbf{1}\|} \right\}, \quad \forall \mathbf{x}_0 \in \mathbb{X}, \quad (48)$$

where $\bar{\sigma}_\kappa := \max_{\mathbf{x} \in \tilde{\mathbb{X}}} \sigma_\kappa(\mathbf{x})$ and β_κ is defined in Lemma 4.

Proof: Using the common Lyapunov candidate (32), its time derivative (44) is given in the κ -th time step for the case $g - \hat{g} = 0$ (Assumption 5) by

$$\dot{V}_\kappa(\mathbf{x}) \leq r \Delta f_\kappa(\mathbf{x}) - k_c r^2. \quad (49)$$

As Theorem 1 guarantees convergence to $\mathbb{B} = \tilde{\mathbb{X}}$, the model error must only be more closely considered in this compact set. From Lemma 4 it can be concluded that

$$\Pr \left\{ \Delta f_\kappa(\mathbf{x}) \leq \beta_\kappa \bar{\sigma}_\kappa, \forall \mathbf{x} \in \tilde{\mathbb{X}}, \kappa \in \mathbb{N}_0 \right\} \geq 1 - \delta \quad (50)$$

$$\Rightarrow \Pr \left\{ \dot{V}_\kappa(\mathbf{x}) < 0, \forall \mathbf{x} \in \tilde{\mathbb{X}} \setminus \mathbb{B}_\kappa, \kappa \in \mathbb{N}_0 \right\} \geq 1 - \delta, \quad (51)$$

which shows the convergence of r to a ball with radius $\frac{\beta_\kappa \bar{\sigma}_\kappa}{k_c}$ and the error e is ultimately bounded by \mathbb{B}_κ with a probability larger than $1 - \delta$. The attributes hold uniformly and globally from the fact that V_κ is a common, time-independent and radially unbounded Lyapunov function [42]. ■

Remark 5: Theorem 1 focuses on the existence of an ultimate bound and with \mathbb{B} in (40) therefore provides the maximum bound across all time intervals, which can be seen in (35), where B, C, c must be suitable for all κ . In contrast, Theorem 2 is more specific here and with \mathbb{B}_κ provides a quantitative bound for each time interval κ individually. Note that the tracking error e will not necessarily converge to the ball \mathbb{B}_κ by the end of the κ -th time step $t_{\kappa+1}$ because it might take infinite time. It is considered a ball that is reached by the tracking error if the control law stops adapting after the κ -th update (compare Corollary 2).

Remark 6: In contrast to Theorem 1, Theorem 2 is a stability statement that only holds with a specified probability. The reason for this lies in the uncertainty about the plant itself, but neither the plant nor any part of the controller is stochastic. Therefore, a stability analysis from deterministic control theory is applicable here, but the convergence result does not hold for all plants that fulfill the specified assumptions. A small fraction of all plants exists (specified by δ) that do not converge to the specified ultimate bound. But, if the plant does not belong to this fraction, the result will always hold and there is no stochastic stability analysis necessary. Note that the fraction δ for which the result does not hold can be made arbitrarily small.

V. EVENT-TRIGGERED MODEL UPDATE

The results in the previous section all hold for arbitrary switching sequences (any definition of $t_{\kappa+1}$ is possible) because we have so far not specified when new training data points are taken and the model is updated accordingly. Our goal is a data-efficient online learning scheme and thereby we only want to add training data if necessary. Thus, switching should not occur synchronously (after a specified fixed time interval) but asynchronously (whenever needed), which is known as event-triggered control.

The general idea of event-based control is to utilize a scarce resource (sensor measurements, computational power, communication channel, etc.) only when required. In contrast to time-triggered control, where the resource is used periodically (synchronously), it is thereby typically more resource-conserving [47]. In our setting, we aim to reduce the number of model updates and measurements for training data to keep the computational complexity low. The key idea of our data-efficient online learning is therefore to take only new training data into account if there is a necessity based on the current uncertainty in the model.

Previous work in [29] uses a time-triggered model, thus measurements are taken and training points are added after a specific time, thus $t_{\kappa+1} = t_\kappa + \Delta t$ with fixed interval $\Delta t > 0$. However, this causes the following difficulties: First, it is unknown whether the current estimate $\hat{f}_\kappa(\cdot)$ of the function $f(\cdot)$ is precise enough to ensure a further decrease of the Lyapunov function. From (49) it is clear that the estimates must become more precise as r gets smaller to guarantee the negative definiteness of \dot{V}_κ for $\forall \mathbf{x} \in \mathbb{X}$.

Considering that at some parts of the state space more data points are necessary to model the function $f(\cdot)$ precisely than in others shows that choosing a constant Δt properly is impossible without knowing the function $f(\cdot)$. Second, over an infinite time horizon, the time-triggered update causes infinitely many (possibly unnecessary) measurements. This is critical even for finite time, because the number of operations to update the GP model increases with $\mathcal{O}(N^3)$ (or $\mathcal{O}(N^2)$ at best, when using the Sherman–Morrison formula) [3].

In summary, for the time-triggered design, there is a trade-off between the precision and the computational complexity of the model when choosing the update rate. If more points are added to the dataset, the variance of the GP model and thereby the maximum model error decreases according to Lemma 4. However, many training points increase the time to compute the model estimate and possibly result in a loss of the real-time capability [48].

Therefore, in the interest of data efficiency and the associated computational complexity, we trigger measurements and their intake to the dataset in an event-based fashion. An intuitive idea is to add training points as soon as the error Δf_κ becomes too large, which is specified based on the Lyapunov stability condition. Generally, to guarantee stability, an event must be triggered before the temporal derivative of the Lyapunov function turns non-negative, thus

$$t_{\kappa+1} := \left\{ t > t_\kappa \mid \dot{V}(\mathbf{x}) \geq 0 \right\}. \quad (52)$$

However, since an exact computation of $\dot{V}(\cdot)$ is not possible, we have to evaluate an upper bound as presented in the following. First, we will consider noiseless measurements of the highest state derivative for the training data before addressing the case in which these measurements are corrupted by noise.

A. Asymptotic stability for noiseless measurements

We first define the noise free case formally in an assumption.

Assumption 7: Measurements of \dot{x}_n are available noise free, thus $\sigma_{\text{on}}^2 = 0$ in Assumption 3.

A well-suited indicator for the necessity to add a new training point is the variance function of the GP $\sigma_\kappa(\cdot)$ in (15) as it bounds the maximum error with high probability according to Lemma 4. Based on this intuition, we propose the following event

$$t_{\kappa+1} := \{t > t_\kappa \mid \beta_\kappa \sigma_\kappa(\mathbf{x}) \geq k_c |r|\}, \quad (53)$$

where the triggering time $t_{\kappa+1}$ is defined as the first time after t_κ when $\beta_\kappa \sigma_\kappa(\mathbf{x})$ becomes larger than or equal to $k_c |r|$.

Remark 7: In the time instance after each update $t = t_\kappa$, $\sigma_\kappa(\mathbf{x}(t_\kappa)) = 0$ generally holds, which implies $\sigma_\kappa(\mathbf{x}(t_\kappa)) \leq k_c |r(t_\kappa)|$. Since both σ_κ and r are continuous over time between two events, the event will always be triggered at the equality, thus $\beta_\kappa \sigma_\kappa(\mathbf{x}(t_{\kappa+1})) = k_c |r(t_{\kappa+1})|$.

Using the proposed event (53) in Algorithm 1 as a trigger to update the model, the following is concluded.

Theorem 3: Consider the system (1) and a desired trajectory $\mathbf{x}_d(t)$ under Assumptions 1-7. Further consider the control law (31), where $f(\cdot)$ is modeled by a GP mean function $\hat{f}_\kappa(\cdot)$ in (42) which is updated according to the event-triggering law (53) and Algorithm 1. Then, the tracking error e is globally asymptotically stable for any $k_c > 0$ and the inter-event time $\Delta t_\kappa := t_{\kappa+1} - t_\kappa$ is lower bounded by a positive constant $t_{\text{lb}} > 0$, for all $\kappa \in \mathbb{N}$ with probability $1 - \delta$.

Proof: We consider again the common Lyapunov candidate (32) and its time derivative

$$\dot{V}_\kappa(\mathbf{x}) \leq r \Delta f_\kappa(\mathbf{x}) - k_c r^2, \quad (54)$$

where Δf_κ is the model error defined in (45). With noiseless measurements, a GP mean function passes through each training point [49]. Thus, the estimate \hat{f}_κ is exact for the time step t_κ , and $V_\kappa(\mathbf{x}(t_\kappa)) = -k_c r^2$. For $t_\kappa < t < t_{\kappa+1}$ the estimation error $\Delta f_\kappa(\mathbf{x}(t))$ continuously changes and is generally larger than zero. But the term $k_c r^2$ will dominate $r \Delta f_\kappa(\mathbf{x})$ with probability $1 - \delta$ by the design of the triggering condition (53) and Lemma 4, thus

$$\Pr \left\{ \dot{V}_\kappa(\mathbf{x}) < 0, \forall \mathbf{x} \in \tilde{\mathbb{X}}, \kappa \in \mathbb{N}_0 \right\} \geq 1 - \delta. \quad (55)$$

From Theorem 1 it is known that the system reaches a compact set $\tilde{\mathbb{X}}$ for any initial condition $\mathbf{x}_0 \in \mathbb{X}$. Therefore, Lemma 4 is applicable and with the radial boundedness of V , the global asymptotic stability with probability $1 - \delta$ is shown.

To show that the inter-event time is lower bounded, we define the Lipschitz constant $L_\sigma > 0$, such that $\dot{\sigma}_\kappa \leq L_\sigma \dot{r}$,

which exists due to the differentiability of σ_κ with respect to r . Following the lines of [50]

$$\begin{aligned} \frac{d}{dt} \left| \frac{\sigma_\kappa}{r} \right| &= \frac{d}{dt} \frac{\sqrt{\sigma_\kappa^2}}{\sqrt{r^2}} = \frac{\dot{\sigma}_\kappa r - \sigma_\kappa \dot{r}}{r^2} \\ &\leq \left| \frac{\dot{\sigma}_\kappa}{r} \right| + \left| \frac{\sigma_\kappa \dot{r}}{r^2} \right| \\ &\leq \left| \frac{L_\sigma(\Delta f_\kappa - k_c r)}{r} \right| + \left| \frac{\sigma_\kappa(\Delta f_\kappa - k_c r)}{r^2} \right| \\ &\leq L_\sigma \left| \frac{\Delta f_\kappa}{r} \right| + L_\sigma k_c + \left| \frac{\Delta f_\kappa \sigma_\kappa}{r^2} \right| + k_c \left| \frac{\sigma_\kappa}{r} \right|, \end{aligned}$$

and using Lemma 4 yields

$$\Pr \left\{ \frac{d}{dt} \left| \frac{\sigma_\kappa}{r} \right| \leq L_\sigma \beta_\kappa \left| \frac{\sigma_\kappa}{r} \right| + L_\sigma k_c + \beta_\kappa \left| \frac{\sigma_\kappa}{r} \right|^2 + k_c \left| \frac{\sigma_\kappa}{r} \right|, \forall \mathbf{x} \in \tilde{\mathbb{X}}, k \in \mathbb{N}_0 \right\} \geq 1 - \delta$$

for which we define $\phi = \left| \frac{\sigma_\kappa}{r} \right|$. The differential equation

$$\dot{\phi} = \beta_\kappa \phi^2 + \phi(L_\sigma \beta_\kappa + k_c) + L_\sigma k_c, \quad (56)$$

with initial condition $\phi(t_\kappa) = 0$ (from $\sigma_\kappa(\mathbf{x}(t_\kappa)) = 0$) yields

$$\phi(t) = \frac{1}{2\beta_\kappa} \left(c_1 \tan \left(\frac{1}{2} ((t - t_\kappa)c_1 \pm c_2) \right) - L_\sigma \beta_\kappa - k_c \right), \quad (57)$$

according to [51] for the time interval $t \in [t_\kappa, t_{\kappa+1}]$ where $c_1 = \sqrt{4\beta_\kappa L_\sigma k_c - (L_\sigma \beta_\kappa + k_c)^2}$ and $c_2 = 2 \arccos \left(\frac{-c_1}{2\sqrt{\beta_\kappa L_\sigma k_c}} \right)$. By design, the event is triggered at $\phi = k_c/\beta_\kappa$, which leads to the lower bound on the inter-event time of

$$\begin{aligned} \Delta t_\kappa &\geq \left(2 \arctan \left((3k_c + L_\sigma \beta_\kappa)/c_1 \right) + c_2 \right) / c_1 \\ &\geq (\pi + c_2) / c_1 =: t_{\text{lb}}, \end{aligned}$$

where $\arctan(\xi) < \pi/2, \forall \xi > 0$ is used. ■

Alternatively, we consider the scenario in which the model error can continuously be monitored.

Assumption 8: Measurements of $\mathbf{x}, \dot{\mathbf{x}}_n$ are continuously available without effort.

To take advantage of this assumption, we propose the following event-trigger

$$t_{\kappa+1} := \{t > t_\kappa \mid \Delta f_\kappa(\mathbf{x}) \geq k_c |r|\}, \quad (58)$$

which allows to drop Assumption 6 and the probabilistic nature of Theorem 3 ($\delta = 0$) as formalized in the following.

Corollary 3: Consider the system (1) and a desired trajectory $\mathbf{x}_d(t)$ under Assumptions 1-5, 7 and 8. Further consider the control law (31), where $f(\cdot)$ is modeled by a GP mean function $\hat{f}_\kappa(\cdot)$ in (42) which is updated according to the event-triggering law (58) and Algorithm 1. Then, the tracking error e is globally asymptotically stable for any $k_c > 0$ and the inter-event time $\Delta t_\kappa := t_{\kappa+1} - t_\kappa$ is lower bounded by a positive constant $t_{\text{lb}} > 0$.

Proof: This is a direct consequence of the proof for Theorem 3. ■

We want to highlight that this requires measurements at any continuous time instance, which is generally not possible

due to nonzero update rates of digital sensors. Therefore, Corollary 3 is rather stated for completeness. Nevertheless, the algorithm remains data-efficient despite the infinity measurements in finite time, because data points are only stored if actually needed.

B. Ultimate boundedness for noisy measurements

In case of noisy measurements of $\dot{\mathbf{x}}_n$ (Assumption 7 does not hold), it is possible to find an ultimate bound to which the system converges. The difference to Theorem 2 is that we now make use of the event-triggered model update (Theorem 2 allowed arbitrary updates), which shrinks the size of the ultimate bound to a size that is proportional to the noise level. From Theorem 3 we derive the following results.

Corollary 4: Consider the system (1) and a desired trajectory $\mathbf{x}_d(t)$ under Assumptions 1-6. Further consider the control law (31), where $f(\cdot)$ is modeled by a GP mean function $\hat{f}_\kappa(\cdot)$ in (42) which is updated according to the event-triggering law

$$t_{\kappa+1} := \{t > t_\kappa \mid \beta_\kappa \sigma_\kappa(\mathbf{x}) \geq k_c |r| \cap e \notin \mathbb{B}_{\sigma_{\text{on}}}\}. \quad (59)$$

and Algorithm 1 where

$$\mathbb{B}_{\sigma_{\text{on}}} = \left\{ e \in \tilde{\mathbb{X}} \mid \|e\| \leq \frac{\sigma_{\text{on}} \beta_\kappa}{k_c \|\lambda^\top \mathbf{1}\|} \right\}. \quad (60)$$

Then, the tracking error e is GUUB to the set $\mathbb{B}_{\sigma_{\text{on}}}$ for any $k_c > 0$ and the inter-event time Δt_κ is lower bounded by a positive constant $t'_{\text{lb}} > 0$, for all $\kappa \in \mathbb{N}$ with probability $1 - \delta$.

Proof: In contrast to Theorem 3, a measurement at time t_κ does not lead to $\Delta f_\kappa(\mathbf{x}(t_\kappa)) = 0$, but we make use of the fact that the variance function of a GP (15) at any training point can be upper bounded in terms of the measurement noise. Considering the variance for a single training data point at $\mathbf{x}(t_\kappa)$ as an upper bound for the variance function (which holds according to [38]), the following is concluded

$$\sigma_\kappa(\mathbf{x}(t_\kappa)) \leq \sqrt{\sigma_f^2 - \frac{\sigma_f^4}{\sigma_f^2 + \sigma_{\text{on}}^2}} = \sqrt{\frac{\sigma_{\text{on}}^2}{1 + \sigma_{\text{on}}^2/\sigma_f^2}} < \sigma_{\text{on}} \quad (61)$$

for $\sigma_f^2 < \infty$ using $k_f(\mathbf{x}, \mathbf{x}) = \sigma_f^2$ in (15). Considering again the Lyapunov function (32) and its time derivative

$$\dot{V}_\kappa(\mathbf{x}(t_\kappa)) \leq |r|(\beta_\kappa \sigma_{\text{on}} - k_c |r|), \quad (62)$$

it is clear that inside $\mathbb{B}_{\sigma_{\text{on}}}$ the negative definiteness of \dot{V}_κ cannot be ensured. But outside of this ball it is negative definite and therefore GUUB can be shown similarly to Theorem 3.

To exclude Zeno behavior, only $e \notin \mathbb{B}_{\sigma_{\text{on}}}$ must be analyzed, since inside $\mathbb{B}_{\sigma_{\text{on}}}$ no events are triggered. The lower bound on the inter-event time is derived along the lines of Theorem 3. Hence, the dynamics of $\phi(t)$ as derived in (56) are the same for the noisy case, but the initial condition $\phi(t_\kappa)$ is now unequal to zero (due to the noise). However, it can be upper bounded by

$$\phi(t_\kappa) < \sqrt{\frac{\sigma_{\text{on}}^2}{1 + \sigma_{\text{on}}^2/\sigma_f^2}} / |r| := \phi_0.$$

The solution for the zero initial condition in (57) is adapted to a nonzero initial condition ϕ_0 according to [51] by changing c_2 to

$$c'_2 = 2 \arctan \left(\frac{2\beta_\kappa \phi_0 + L_\sigma \beta_\kappa + k_c}{c_1} \right). \quad (63)$$

Accordingly, the lower bound on the inter event time is

$$\Delta t_\kappa \geq (\pi + c'_2)/c_1 =: t'_{\text{lb}},$$

which concludes the proof. \blacksquare

C. Forgetting strategies

With the event in (53) and Algorithm 1, we have proposed a strategy that adds data points to the dataset only if necessary to ensure further convergence of the system. However, this still leads to a growing computational burden for computing the GP model as the cardinality of the dataset \mathbb{D}_κ monotonically increases with time. Particularly, if the desired trajectory covers a large area in the state space or when high precision tracking is required, keeping up the real-time capability of the adaptation algorithm is challenging. A common technique to circumvent this problem is a forgetting mechanism (deleting old data points when new ones are added) if a particular budget is reached. While most other works, e.g. [29], use a heuristic for the forgetting strategy, we propose a *safe* forgetting rule, which requires to store only a single data point.

Corollary 5: Consider the system (1) and a desired trajectory $\mathbf{x}_d(t)$ under Assumptions 1-7. Further consider the control law (31), where $f(\cdot)$ is modeled by a GP mean function $\hat{f}_\kappa(\cdot)$ in (42). This estimate is updated at the event $t_{\kappa+1}$ in (53), where at each event κ all old data points are eliminated from the dataset, thus

$$\mathbb{D}_\kappa = \{\mathbf{x}(t_\kappa), \dot{\mathbf{x}}_n(t_\kappa)\}.$$

Then, with probability $1 - \delta$, the tracking error e is globally asymptotically stable for any $k_c > 0$.

Proof: This follows along the lines of the proof of Theorem 3. With the continuity of $\sigma(\mathbf{x})$, which is zero at the single training point, $\sigma(\mathbf{x}(t_\kappa)) = 0$, it follows that there exists a neighborhood of $\mathbf{x}(t_\kappa)$ for which $\sigma_\kappa(\mathbf{x}) < k_c|r|$ holds. Thus the results from Corollary 6 and Theorem 3 are applicable. \blacksquare Deleting all old data points is consequent in terms of data efficiency, but in general triggers events more frequently. Thus, by storing more than one data point, future measurements can be avoided particularly for periodic desired trajectories. For a fixed budget $\bar{N} \in \mathbb{N}$, we can also forget unnecessary points and still guarantee stability.

Corollary 6: Consider the system (1) and a desired trajectory $\mathbf{x}_d(t)$ under Assumptions 1-7. Further consider the control law (31), where $f(\cdot)$ is modeled by a GP mean function $\hat{f}_\kappa(\cdot)$ in (42). This estimate is updated at the event $t_{\kappa+1}$ in (53), where at each event κ the dataset \mathbb{D}_κ is limited to hold at most $\bar{N} \in \mathbb{N}$ data points, such that

$$\beta_\kappa^{\text{elim}} \sigma_\kappa^{\text{elim}}(\mathbf{x}) < k_c|r| \quad (64)$$

remains true, where $\beta_\kappa^{\text{elim}}$ and $\sigma_\kappa^{\text{elim}}$ denote the values after the elimination. Then, with probability $1 - \delta$, the tracking error e is globally asymptotically stable for any $k_c > 0$.

Proof: This follows along the lines of the proof of Theorem 3. From Corollary 5, it is known that there always exists a reduced dataset that fulfills (64) for $\bar{N} \geq 1$. \blacksquare Finding the reduced dataset is not a trivial combinatorial problem, but we refer to the existing literature for efficient algorithms [52]. Note that the reduced dataset must necessarily contain the most recent measurement at t_κ as otherwise the event would not have been triggered.

D. Discussion

From a control perspective, the most important advantage of GPs is the quantification of the uncertainty, i.e. an upper bound of the model error as given by Lemma 4. We note that the prerequisite for this lemma, the bounded RKHS norm in Assumption 6, is difficult to verify, however, minimal assumptions are necessary, as otherwise a generalization beyond the training data is impossible [46].

Also, the maximum mutual information γ_κ in Lemma 4 cannot be computed analytically for a general kernel, but we refer to the existing literature [45], which provides upper bounds on γ_κ for different kernels (including the squared exponential kernel). Since β_κ is not trivial to find, we would like to point out that β_κ always appears in the ratio with k_c , thus any conservatism/approximation in β_κ can generally be compensated for by the designer's choice of the control gain k_c .

Overcoming these challenges, the GP allows - based on event-triggered online learning - to design a feedback linearizing control law that asymptotically stabilizes an initial unknown system (with high probability). This is made possible by the error bounds on the model, which is the most significant advantage of a GP over alternative modeling approaches like neural networks [53].

As the model update is event-triggered, only data points that are necessary to increase the precision of the model are collected. This reduces the frequency at which measurements are taken and increases data efficiency. With Corollary 5 we have shown that only a single data point must be stored to guarantee asymptotic convergence. This is a significant advantage of the locally linearizing control law in comparison to predictive control laws or reinforcement learning algorithms, where an accurate global model is required. Accurate global models require active exploration, e.g. through exploration noise, which sacrifices control performance (known as exploration-exploitation trade-off).

VI. NUMERICAL ILLUSTRATION

To illustrate the proposed approach, we present simulations³ for the control affine system

$$\begin{aligned} \dot{x}_1 &= x_2, \\ \dot{x}_2 &= \underbrace{1 - \sin(x_1) + s(x_2)}_{=f(\mathbf{x})} + \underbrace{\left(1 + \frac{1}{2} \sin(x_2/2)\right)}_{=g(\mathbf{x})} u, \end{aligned} \quad (65)$$

³The code is available at <https://gitlab.lrz.de/ga68car/adaptFeLi4GPs>




Fig. 3. Scenario 1: The black solid line illustrates the actual, the green dashed line the desired value for the state x_1 . The system converges to the desired state over time.

where $s(x_2) = \frac{0.5}{1 + \exp(-x_2/10)}$ is the sigmoidal function. It is a modified pendulum system and fulfills Assumptions 1 and 2. To ensure Assumption 6 holds, we do not simulate directly on (65) but use a GP mean that was trained on it with a high density of training points. As we are working in simulation, Assumptions 3, 5 and 7 do hold or do not hold by design in the following two different scenarios, which we use to illustrate the proposed approach. An overview of the employed parameters is given in Table I

A. Scenario 1: Time-triggered updates

In Scenario 1 (S1), we illustrate the results from Sec. IV, which are shown to hold for an arbitrary switching sequence. Therefore, we utilize a periodic, time-triggered model update, thus $t_{\kappa+1} - t_{\kappa} = \Delta t$, $\forall \kappa$, with $\Delta t = 0.5$ and follow Algorithm 1. We consider $f(\mathbf{x})$ and $g(\mathbf{x})$ to be unknown, so Assumption 5 does not hold, but we know that $g(\mathbf{x})$ is positive, so Assumption 2 holds. As the reference trajectory

$$x_d(t) = 1 - \frac{1}{1 + \exp(-20(t - 10))} \quad (66)$$

is used, which describes a “soft” jump from $x_1 = 1$ to $x_1 = 0$ at $t = 10$ and it fulfills the required smoothness in Assumption 4. The scenario works on noisy measurements, thus Assumption 7 does not hold. As this scenario does not utilize Assumption 6, we consider the kernel’s hyperparameters to be unknown. Therefore, a hyperparameter optimization according to (8) is performed at each model update step κ . The simulation is stopped manually after $T_{\text{sim}} = 20$, which leads to $N = 40$ data points.

Figure 3 shows the desired trajectory and the corresponding tracking performance of the controller over time. Figure 4 illustrates the resulting trajectory in the state space. In Figs. 5 and 6, the true system dynamics, $f(\mathbf{x}), g(\mathbf{x})$ are compared with the approximations $\hat{f}(\mathbf{x}), \hat{g}(\mathbf{x})$ at the end of the simulation. It turns out that the hyperparameters for k_{fg} are well identified: With $l_{1,f} \ll l_{2,f}$, the estimate for $f(\mathbf{x})$ shows that $f(\mathbf{x})$ mainly depends on x_1 (and vice versa for $g(\mathbf{x})$ with $l_{1,g} \gg l_{2,g}$). As it can be seen in Figs. 5 and 6, the estimates are more precise near the training data.

It can be seen that the two stationary points of the desired trajectory, $(1, 0)$ and $(0, 0)$, are approached with high precision




Fig. 4. Scenario 1: Black crosses indicate the collected training points, the black solid line illustrates the actual, the green dashed the desired trajectory. The system approaches the desired states as more training points are collected.




Fig. 5. Scenario 1: The surface illustrates the relative error between the true function $f(\mathbf{x})$ and the model estimate $\hat{f}_{40}(\mathbf{x})$ after taking 40 training points (black crosses). The error is the lowest (in terms of absolute value) near the training data.




Fig. 6. Scenario 1: The surface illustrates the relative error between the true function $g(\mathbf{x})$ and the model estimate $\hat{g}_{40}(\mathbf{x})$ after taking 40 training points (black crosses). The error is the lowest (in terms of absolute value) near the training data.

k_c	λ	$m_g(\mathbf{x})$	$\sigma_{\text{on}}^2(\text{S1})$	$\sigma_{\text{on}}^2(\text{S2})$	\mathbf{x}_0	β	r_{min}
1	1	$= 2, \forall \mathbf{x}$	10^{-6}	10^{-16}	$[3 \ 2]^\top$	7	10^{-5}

TABLE I
SIMULATION PARAMETERS

in the steady state. However, for both a few measurements must be collected in the corresponding area of the state space and the following model updates are required to achieve this high precision. Once the steady state has been reached, the time-triggered implementation keeps adding unnecessary data points even though the model already has a high precision in this area. This is improved with the event-triggered model update as illustrated in Scenario 2.

B. Scenario 2: Event-triggered updates

In Scenario 2 (S2), the results in Sec. V are illustrated, which utilizes the event-triggered model update described by (53). For this scenario, $g(\mathbf{x})$ is assumed to be known (Assumption 5 holds). The reference trajectory

$$x_d(t) = \sin(t) \quad (67)$$

is used, which describes a circle with radius 1 in the state space. The scenario works on noise free measurements, thus Assumption 7 does hold, however, for numerical stability a minimal noise is assumed ($\sigma_{\text{on}}^2 = 10^{-16}$). The simulation is stopped manually after $T_{\text{sim}} = 100$.

As this scenario utilizes Assumption 6, we take the kernel hyperparameters to be known at $\sigma_f^2 = 5, l_{1/2,f}^2 = 5$ and do not update these at any of the triggered events. Additionally, we set β_κ constant $\forall \kappa$ and refer to the discussion in Sec. V-D and [28]. Additionally, we enforce a lower bound $r > r_{\text{min}}$ to avoid numerical difficulties.

Figure 8 shows the tracking error until $t = 30$, which initially decreases approximately exponentially until a numerical limit is reached. In the event-triggered setup, a total of 51 events are triggered until sufficient training points are collected around the desired trajectory. This is also visualized in the state space view in Fig. 7. In comparison, the time-triggered approach requires to store 200 data points and would keep adding points for longer simulations, which the event-based would not.

The time-trigger results in a higher computational burden than the event-trigger: The simulation of the latter only takes $\approx 5s$ for the MATLAB 2019a implementation on a i5-6200U CPU with 2.3GHz and uses 0.20957MB. The time-triggered approach takes $\approx 11s$ and uses 0.47924MB of memory.

Figure 8 also shows that the stability criteria in Theorem 3 are fulfilled for the event-triggered case, since $\beta_\kappa \sigma_\kappa(\mathbf{x}) \leq k_c |r|$ holds for any time. In contrast - for the time-triggered case - this condition is violated frequently, which means that negative definiteness of the common Lyapunov function - and thereby stability - cannot be shown.

VII. CONCLUSION

This article proposes an online learning feedback linearizing control law based on Gaussian process models. The closed-loop identification of the initially unknown system exploits




Fig. 7. Scenario 2: Black crosses indicate the collected training points, the black solid line illustrates the actual, the green dashed line the desired trajectory. The color map shows the variance function (15) for the GP $\sigma_{51}^2(\mathbf{x})$ after the 51st update, where yellow indicates low variance and blue high variance.

the control affine structure by utilizing a composite kernel. The model is updated in an event-triggered fashion, taking advantage of the uncertainty measure of the GP. The control law results in the global asymptotic stability of the tracking error in the noiseless case and in global uniform ultimate boundedness for noisy measurements (of the highest state derivative) with high probability. We therefore propose a safe and data-efficient online learning control approach because model updates occur only if required to ensure stability. Zeno behavior is excluded, as a lower bound on the inter-event time is derived. The proposed techniques are illustrated using simulations to support the theoretical results.

APPENDIX A EXPRESSING STRUCTURE IN KERNELS

According to [4], the kernel of the GP does not only determine the smoothness properties of the resulting functions but can also be utilized to express prior knowledge regarding the structure of the unknown function.

A. Sum of functions

Consider $f_a, f_b: \mathbb{X} \rightarrow \mathbb{R}$ which both originate from two independent GP priors

$$f_a(\mathbf{x}) \sim \mathcal{GP}(m_a(\mathbf{x}), k_a(\mathbf{x}, \mathbf{x}')), \quad (68)$$

$$f_b(\mathbf{x}) \sim \mathcal{GP}(m_b(\mathbf{x}), k_b(\mathbf{x}, \mathbf{x}')), \quad (69)$$

and add up to $f_{\text{sum}}: \mathbb{X} \rightarrow \mathbb{R}$, thus $f_{\text{sum}}(\mathbf{x}) = f_a(\mathbf{x}) + f_b(\mathbf{x})$. Then,

$$f_{\text{sum}}(\mathbf{x}) \sim \mathcal{GP}(m_a(\mathbf{x}) + m_b(\mathbf{x}), k_a(\mathbf{x}, \mathbf{x}') + k_b(\mathbf{x}, \mathbf{x}')) \quad (70)$$

is also a GP with kernel $k_a(\mathbf{x}, \mathbf{x}') + k_b(\mathbf{x}, \mathbf{x}')$.

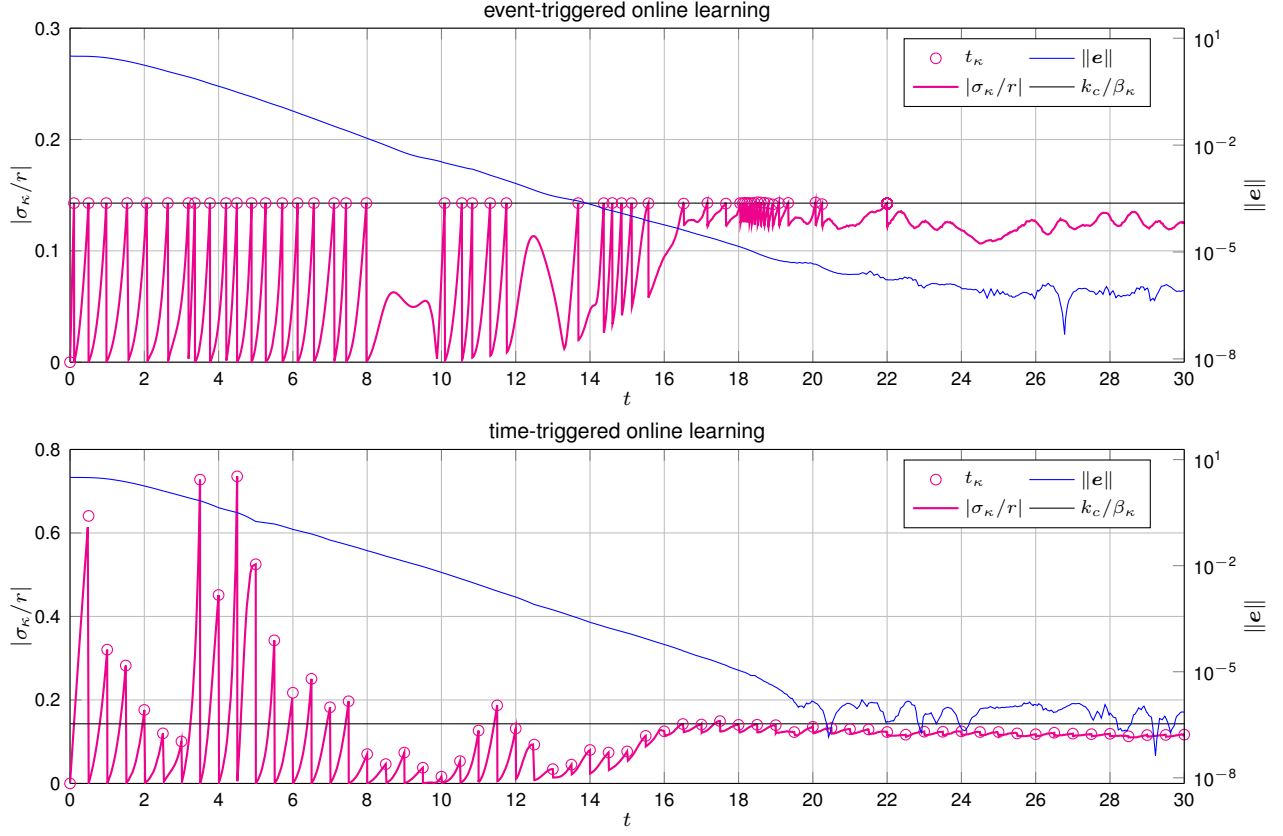


Fig. 8. Scenario 2: Comparison of the event-triggered (top) and time-triggered (bottom) online learning. For the first, events (magenta circles) are triggered when the threshold k_c/β_κ (black vertical line) is reached by σ_κ/r as proposed in (53). For the latter, events are triggered after a fixed time interval ($\Delta t = 0.5$). The blue lines show the norm of the tracking error $\|e\|$.

For regression, where noisy measurements with $\epsilon^{(i)} \sim \mathcal{N}(0, \sigma_{\text{on}}^2)$ of the sum of the two functions are available

$$y_{\text{sum}}^{(i)} = f_{\text{sum}}(\mathbf{x}^{(i)}) + \epsilon^{(i)} = f_a(\mathbf{x}^{(i)}) + f_b(\mathbf{x}^{(i)}) + \epsilon^{(i)}, \quad (71)$$

with $i = 1, \dots, N$, the joint distribution of the individual functions and the observations is given by

$$\begin{bmatrix} f_a(\mathbf{x}^*) \\ f_b(\mathbf{x}^*) \\ \mathbf{y}_{\text{sum}} \end{bmatrix} \sim \mathcal{N} \left(\mathbf{0}, \begin{bmatrix} k_a^* & \mathbf{0} & \mathbf{k}_a^\top \\ \mathbf{0} & k_b^* & \mathbf{k}_b^\top \\ \mathbf{k}_a & \mathbf{k}_b & \mathbf{K}_a + \mathbf{K}_b + \sigma_{\text{on}}^2 \mathbf{I}_N \end{bmatrix} \right), \quad (72)$$

where the prior mean functions are set to zero $m_a(\mathbf{x}) = m_b(\mathbf{x}) = 0$ for notational simplicity and $\mathbf{k}_a, \mathbf{k}_b, k_a^*, k_b^*$ are defined according to (16). By conditioning, the output of f_a and f_b are inferred for a test point \mathbf{x}^*

$$f_a(\mathbf{x}^*) | \mathbf{X}, \mathbf{y}_{\text{sum}} \sim \mathcal{N}(\mathbf{k}_a^\top \mathbf{K}_{\text{sum}}^{-1} \mathbf{y}_{\text{sum}}, k_a^* - \mathbf{k}_a^\top \mathbf{K}_{\text{sum}}^{-1} \mathbf{k}_a), \quad (73)$$

$$f_b(\mathbf{x}^*) | \mathbf{X}, \mathbf{y}_{\text{sum}} \sim \mathcal{N}(\mathbf{k}_b^\top \mathbf{K}_{\text{sum}}^{-1} \mathbf{y}_{\text{sum}}, k_b^* - \mathbf{k}_b^\top \mathbf{K}_{\text{sum}}^{-1} \mathbf{k}_b), \quad (74)$$

where $\mathbf{K}_{\text{sum}} = \mathbf{K}_a + \mathbf{K}_b + \sigma_{\text{on}}^2 \mathbf{I}_N$ with $\mathbf{K}_a, \mathbf{K}_b$ according to (11). Similarly to (8), the extended hyperparameter vector $\boldsymbol{\psi}_{\text{sum}} = [\boldsymbol{\psi}_a^\top \ \boldsymbol{\psi}_b^\top]^\top$ is obtained through the optimization

of the likelihood, where $\mathbf{K} = \mathbf{K}_{\text{sum}}$ and $\mathbf{y}_f = \mathbf{y}_{\text{sum}}$. This allows to predict a value of the individual functions f_a, f_b even though only their sum has been measured.

B. Product with known function

Consider an unknown function $f_h(\mathbf{x}): \mathbb{X} \rightarrow \mathbb{R}$, which is multiplied with the known function $h(\mathbf{x}): \mathbb{X} \rightarrow \mathbb{R}$, we can model f_h using a GP with a scaled kernel function and noisy measurements

$$y_{\text{prod}}^{(i)} = f_{\text{prod}}(\mathbf{x}^{(i)}) + \epsilon^{(i)} = f_h(\mathbf{x}^{(i)}) h(\mathbf{x}^{(i)}) + \epsilon^{(i)} \quad (75)$$

of the product with $\epsilon^{(i)} \sim \mathcal{N}(0, \sigma_{\text{on}}^2)$, $i = 1, \dots, N$. Thus, if $f_h \sim \mathcal{GP}(0, k_h(\mathbf{x}, \mathbf{x}'))$ is a GP, then $f_{\text{prod}}(\mathbf{x})$ is also a GP with kernel

$$k_{\text{prod}}(\mathbf{x}, \mathbf{x}') = h(\mathbf{x}) k_h(\mathbf{x}, \mathbf{x}') h(\mathbf{x}'), \quad (76)$$

where the prior mean is set to zero $m_h(\mathbf{x}) = 0$ for notational simplicity.

The joint distribution of the measurements and the inferred output of f_{prod} at a test input \mathbf{x}^* is given by

$$\begin{bmatrix} f_h(\mathbf{x}^*) \\ \mathbf{y}_{\text{prod}} \end{bmatrix} \sim \mathcal{N} \left(\mathbf{0}, \begin{bmatrix} k_h^* & \mathbf{k}_h^\top \mathbf{H}^\top \\ \mathbf{H} \mathbf{k}_h & \mathbf{H}^\top \mathbf{K}_h \mathbf{H} + \sigma_{\text{on}}^2 \mathbf{I}_N \end{bmatrix} \right), \quad (77)$$

where $\mathbf{H} = \text{diag}(h(\mathbf{x}^{(1)}), \dots, h(\mathbf{x}^{(N)})) \in \mathbb{R}^{N \times N}$ and $\mathbf{k}_h, k_h^*, \mathbf{K}_h$ are defined similarly to (16) and (11),

respectively. Through the conditioning on the training data and the input, the function f_h is inferred by

$$f_h(\mathbf{x}^*) | \mathbf{X}, \mathbf{y}_{\text{prod}} \sim \mathcal{N}(\mathbf{k}_h^\top \mathbf{H}^\top \mathbf{K}_{\text{prod}}^{-1} \mathbf{y}_{\text{prod}}, \mathbf{k}_h^* - \mathbf{k}_h^\top \mathbf{H}^\top \mathbf{K}_{\text{prod}}^{-1} \mathbf{H} \mathbf{k}_h), \quad (78)$$

where $\mathbf{K}_{\text{prod}} = \mathbf{H}^\top \mathbf{K}_h \mathbf{H} + \sigma_{\text{on}}^2 \mathbf{I}_N$.

Remark 8: Instead of scaling the kernel, it seems more straight forward to use $y_{\text{prod}}^{(i)}/h(\mathbf{x}^{(i)})$ as training data for a GP with an unscaled kernel. However, this would scale the observation noise undesirably, is numerically unstable and is not compatible with the summation of kernels in Appendix A-A, which we combine in our identification approach in Sec. III-B.

APPENDIX B IMPROVING IDENTIFICATION

From Lemma 2 it is known that the state remains bounded for any finite $0 < T < \infty$ without any control input. Thus, without risking damage to the system, one can set $u = 0$ for time interval T and record an open-loop training point

$$y^{(i_{\text{ol}})} = f(\mathbf{x}^{(i_{\text{ol}})}) + \epsilon^{(i_{\text{ol}})}, \quad (79)$$

which is highly beneficial, as it only measures $f(\mathbf{x})$ (with the usual noise ϵ). The GP framework allows to merge these $i_{\text{ol}} = 1, \dots, N_{\text{ol}}$ observations with the closed-loop training points in \mathbb{D}_κ to improve the prediction as follows: Consider the extension of the joint distribution (25) (where $u = 1$ is assumed in the close loop measurements and $m_g(\mathbf{x}) = 0$ for notational convenience)

$$\begin{bmatrix} f(\mathbf{x}^*) \\ g(\mathbf{x}^*) \\ \mathbf{y} \\ \mathbf{y}_{\text{ol}} \end{bmatrix} \sim \mathcal{N} \left(\mathbf{0}, \begin{bmatrix} k_f^* & 0 & \mathbf{k}_f^\top & \mathbf{k}_{f,\text{ol}}^\top \\ 0 & k_g^* & \mathbf{k}_g^\top & \mathbf{0}_{1 \times N_{\text{ol}}} \\ \mathbf{k}_f & \mathbf{k}_g & \mathbf{K}_{fg} & \mathbf{K}_{\text{ol},\text{cl}}^\top \\ \mathbf{k}_{f,\text{ol}} & \mathbf{0}_{N_{\text{ol}} \times 1} & \mathbf{K}_{\text{ol},\text{cl}} & \mathbf{K}_{\text{ol}} \end{bmatrix} \right),$$

where $\mathbf{K}_{\text{ol},\text{cl}}, \mathbf{K}_{\text{ol}}$ are the pairwise evaluation of $k(\mathbf{x}^{(i_{\text{ol}})}, \mathbf{x}^{(i)})$, $k(\mathbf{x}^{(i_{\text{ol}})}, \mathbf{x}^{(i_{\text{ol}})})$ and $\mathbf{k}_{f,\text{ol}}$ evaluates $k(\mathbf{x}^*, \mathbf{x}^{(i_{\text{ol}})})$. for all $i = 1, \dots, N$, $i_{\text{ol}}, i'_{\text{ol}} = 1, \dots, N_{\text{ol}}$. Then, the estimates are given by

$$\hat{f}(\mathbf{x}^*) = [\mathbf{k}_f^\top \ \mathbf{k}_{f,\text{ol}}^\top] \mathcal{K}^{-1} \tilde{\mathbf{y}}, \quad \hat{g}(\mathbf{x}^*) = [\mathbf{k}_g^\top \ \mathbf{0}_{1 \times N_{\text{ol}}}] \mathcal{K}^{-1} \tilde{\mathbf{y}}$$

with $\mathcal{K} = \begin{bmatrix} \mathbf{K}_{fg} & \mathbf{K}_{\text{ol},\text{cl}}^\top \\ \mathbf{K}_{\text{ol},\text{cl}} & \mathbf{K}_{\text{ol}} \end{bmatrix}$, and $\tilde{\mathbf{y}} = \begin{bmatrix} \mathbf{y} \\ \mathbf{y}_{\text{ol}} \end{bmatrix}$. (80)

We do not further investigate this extension, since it does not provide any additional formal guarantees regarding the convergence, however, in practice, a significant improvement of the identification can be expected.

REFERENCES

- [1] L. Ljung, *System Identification*. NJ, USA: Prentice Hall PTR, 1998.
- [2] J. Kocijan, *Modelling and Control of Dynamic Systems Using Gaussian Process Models*. Springer, 2016.
- [3] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. Cambridge, MA, USA: MIT Press, Jan. 2006.
- [4] D. Duvenaud, "Automatic model construction with Gaussian processes," Ph.D. dissertation, Computational and Biological Learning Laboratory, University of Cambridge, 2014.
- [5] P. A. Ioannou and J. Sun, *Robust adaptive control*. PTR Prentice-Hall Upper Saddle River, NJ, 1996, vol. 1.

- [6] M. Krstic, I. Kanellakopoulos, P. V. Kokotovic *et al.*, *Nonlinear and adaptive control design*. Wiley New York, 1995, vol. 222.
- [7] K. J. Åström and B. Wittenmark, *Adaptive control*. Courier Corporation, 2013.
- [8] A. S. Bazanella, L. Campestri, and D. Eckhard, *Data-driven controller design: the H2 approach*. Springer Science & Business Media, 2011.
- [9] G. R. Gonçalves da Silva, A. S. Bazanella, C. Lorenzini, and L. Campestri, "Data-driven lqr control design," *IEEE Control Systems Letters*, vol. 3, no. 1, pp. 180–185, Jan. 2019.
- [10] L. Campestri, D. Eckhard, A. S. Bazanella, and M. Gevers, "Data-driven model reference control design by prediction error identification," *Journal of the Franklin Institute*, vol. 354, no. 6, pp. 2628–2647, 2017, special issue on recent advances on control and diagnosis via process measurements. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S001600321630271X>
- [11] D. A. Bristow, M. Tharayil, and A. G. Alleyne, "A survey of iterative learning control," *IEEE Control Systems Magazine*, vol. 26, no. 3, pp. 96–114, Jun. 2006.
- [12] M.-B. Rădac, R.-E. Precup, E. M. Petriu, and S. Preitl, "Iterative data-driven tuning of controllers for nonlinear systems with constraints," *IEEE Transactions on Industrial Electronics*, vol. 61, no. 11, pp. 6360–6368, Nov. 2014.
- [13] Z. Hou, R. Chi, and H. Gao, "An overview of dynamic-linearization-based data-driven control and applications," *IEEE Transactions on Industrial Electronics*, vol. 64, no. 5, pp. 4076–4090, May 2017.
- [14] M. C. Campi and S. M. Savarese, "Direct nonlinear control design: the virtual reference feedback tuning (vrft) approach," *IEEE Transactions on Automatic Control*, vol. 51, no. 1, pp. 14–27, Jan. 2006.
- [15] H. Hjalmarsson, M. Gevers, S. Gunnarsson, and O. Lequin, "Iterative feedback tuning: theory and applications," *IEEE Control Systems Magazine*, vol. 18, no. 4, pp. 26–41, Aug. 1998.
- [16] L. C. Kammer, R. R. Bitmead, and P. L. Bartlett, "Direct iterative tuning via spectral analysis," *Automatica*, vol. 36, pp. 1301–1307, 2000.
- [17] N. J. Killingsworth and K. Miroslav, "Pid tuning using extremum seeking: online, model-free performance optimization," *IEEE Control Systems Magazine*, vol. 26, no. 1, pp. 70–79, Feb. 2006.
- [18] E. Theodorou, J. Buchli, and S. Schaal, "Reinforcement learning of motor skills in high dimensions: A path integral approach," in *International Conference on Robotics and Automation (ICRA)*. IEEE, May 2010, pp. 2397–2403.
- [19] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. Cambridge, MA, USA: MIT Press, 1998.
- [20] M. P. Deisenroth and C. E. Rasmussen, "PILCO: A model-based and data-efficient approach to policy search," in *International Conference on Machine Learning (ICML)*, 2011, pp. 465–472.
- [21] D. Nguyen-Tuong and J. Peters, "Model learning for robot control: a survey," *Cognitive processing*, vol. 12, no. 4, pp. 319–340, 2011.
- [22] D. Nguyen-Tuong, M. Seeger, and J. Peters, "Model learning with local Gaussian process regression," *Advanced Robotics*, vol. 23, no. 15, pp. 2015–2034, 2009.
- [23] T. Beckers, J. Umlauf, and S. Hirche, "Stable model-based control with Gaussian process regression for robot manipulators," in *World Congress of the International Federation of Automatic Control (IFAC)*, vol. 50, no. 1. Toulouse, France: Elsevier, Jul. 2017, pp. 3877–3884.
- [24] Y. Fanger, J. Umlauf, and S. Hirche, "Gaussian processes for dynamic movement primitives with application in knowledge-based cooperation," in *International Conference on Intelligent Robots and Systems (IROS)*. IEEE, Oct. 2016, pp. 3913–3919.
- [25] T. Beckers, J. Umlauf, D. Kulic, and S. Hirche, "Stable Gaussian process based tracking control of Lagrangian systems," in *Conference on Decision and Control (CDC)*. IEEE, Dec. 2017, pp. 5180–5185. [Online]. Available: <https://ieeexplore.ieee.org/document/8264427>
- [26] J. Umlauf, A. Lederer, and S. Hirche, "Learning stable Gaussian process state space models," in *American Control Conference (ACC)*, IEEE, May 2017, pp. 1499–1504. [Online]. Available: <https://ieeexplore.ieee.org/document/7963165>
- [27] J. Umlauf, L. Pöhler, and S. Hirche, "An uncertainty-based control Lyapunov approach for control-affine systems modeled by Gaussian process," *IEEE Control Systems Letters*, vol. 2, no. 3, pp. 483–488, Jul. 2018. [Online]. Available: <https://ieeexplore.ieee.org/document/8368325>
- [28] F. Berkenkamp, R. Moriconi, A. Schoellig, and A. Krause, "Safe learning of regions of attraction for uncertain, nonlinear systems with Gaussian processes," *arXiv preprint arXiv:1603.04915*, 2016.
- [29] 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, Mar. 2015.


- [30] A. Chakraborty and M. Arcak, "Robust stabilization and performance recovery of nonlinear systems with unmodeled dynamics," *IEEE Transactions on Automatic Control (TAC)*, vol. 54, no. 6, pp. 1351–1356, May 2009.
- [31] A. Chakraborty and M. Arcak, "Time-scale separation redesigns for stabilization and performance recovery of uncertain nonlinear systems," *Automatica*, vol. 45, no. 1, pp. 34–44, 2009. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0005109808003701>
- [32] M. M. Polycarpou and P. A. Ioannou, *Identification and control of nonlinear systems using neural network models: Design and stability analysis*. University of Southern Calif., 1991.
- [33] F. L. Lewis, A. Yesildirek, and K. Liu, "Multilayer neural-net robot controller with guaranteed tracking performance," *IEEE Transactions on Neural Networks*, vol. 7, no. 2, pp. 388–399, 1996.
- [34] F. L. Lewis, K. Liu, and A. Yesildirek, "Neural net robot controller with guaranteed tracking performance," *IEEE Transactions on Neural Networks*, vol. 6, no. 3, pp. 703–715, 1995.
- [35] R. M. Sanner and J.-J. E. Slotine, "Stable adaptive control and recursive identification using radial Gaussian networks," in *Conference on Decision and Control (CDC)*. IEEE, 1991, pp. 2116–2123.
- [36] A. Yesildirek and F. L. Lewis, "Feedback linearization using neural networks," *Automatica*, vol. 31, no. 11, pp. 1659–1664, 1995.
- [37] J. Umlauf, Y. Fanger, and S. Hirche, "Bayesian uncertainty modeling for programming by demonstration," in *International Conference on Robotics and Automation (ICRA)*, May 2017, pp. 6428–6434.
- [38] J. Umlauf, T. Beckers, M. Kimmel, and S. Hirche, "Feedback linearization using Gaussian processes," in *Conference on Decision and Control (CDC)*. IEEE, Dec. 2017, pp. 5249–5255. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/8264435/>
- [39] H. K. Khalil and J. W. Grizzle, *Nonlinear systems*. Prentice hall New Jersey, 1996, vol. 3.
- [40] M. W. Seeger, S. M. Kakade, and D. P. Foster, "Information consistency of nonparametric Gaussian process methods," *IEEE Transactions on Information Theory*, vol. 54, no. 5, pp. 2376–2382, May 2008.
- [41] J. Sherman and W. J. Morrison, "Adjustment of an inverse matrix corresponding to a change in one element of a given matrix," *Ann. Math. Statist.*, vol. 21, no. 1, pp. 124–127, Mar. 1950. [Online]. Available: <https://doi.org/10.1214/aoms/117729893>
- [42] D. Liberzon, *Switching in systems and control*. Springer Science & Business Media, 2012.
- [43] T. Beckers, D. Kulić, and S. Hirche, "Stable Gaussian process based tracking control of Euler-Lagrange systems," *Automatica*, vol. 23, no. 103, pp. 390–397, 2019.
- [44] J.-J. E. Slotine and J. Karl Hedrick, "Robust input-output feedback linearization," *International Journal of Control*, vol. 57, no. 5, pp. 1133–1139, 1993.
- [45] N. Srinivas, A. Krause, S. M. Kakade, and M. W. Seeger, "Information-theoretic regret bounds for Gaussian process optimization in the bandit setting," *IEEE Transactions on Information Theory*, vol. 58, no. 5, pp. 3250–3265, May 2012.
- [46] D. H. Wolpert, "The supervised learning no-free-lunch theorems," in *Soft Computing and Industry*. Springer, 2002, pp. 25–42.
- [47] W. P. M. H. Heemels, K. H. Johansson, and P. Tabuada, "An introduction to event-triggered and self-triggered control," in *Conference on Decision and Control (CDC)*, Dec. 2012, pp. 3270–3285.
- [48] D. Nguyen-Tuong, J. R. Peters, and M. Seeger, "Local Gaussian process regression for real time online model learning," in *Advances in Neural Information Processing Systems (NeurIPS)*, D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, Eds. Curran Associates, Inc., 2009, pp. 1193–1200.
- [49] J. Umlauf, T. Beckers, and S. Hirche, "A scenario-based optimal control approach for Gaussian process state space models," in *European Control Conference (ECC)*, 6 2018, pp. 1386–1392. [Online]. Available: <https://ieeexplore.ieee.org/document/8550458>
- [50] P. Tabuada, "Event-triggered real-time scheduling of stabilizing control tasks," *IEEE Transactions on Automatic Control (TAC)*, vol. 52, no. 9, pp. 1680–1685, 2007.
- [51] Wolfram—Alpha, "Solution first-order nonlinear ordinary differential equation," Online, Sep. 2018. [Online]. Available: <http://www.wolframalpha.com>
- [52] P. C. Chu and J. E. Beasley, "A genetic algorithm for the multidimensional knapsack problem," *Journal of heuristics*, vol. 4, no. 1, pp. 63–86, 1998.
- [53] F. L. Lewis, S. Jagannathan, and A. Yesildirek, *Neural network control of robot manipulators and non-linear systems*. CRC Press, 1998.



Jonas Umlauf (S'14) received the B.Sc. and M.Sc. degree in electrical engineering and information technology from the Technical University of Munich, Germany, in 2013 and 2015, respectively.

His Master's thesis was completed at the Computational and Biological Learning Group at the University of Cambridge, UK. Since May 2015, he has been a PhD student at the Chair of Information-oriented Control, Department of Electrical and Computer Engineering at the Technical University of Munich, Germany. His current research interests

include the stability of data-driven control systems and system identification based on Gaussian processes.



Sandra Hirche (M'03-SM'11-F'20) received the Dipl.Ing degree in aeronautical engineering from the Technical University of Berlin, Germany, in 2002 and the Dr. Ing. degree in electrical engineering from the Technical University of Munich, Germany, in 2005.

From 2005 to 2007 she was awarded a Postdoc scholarship from the Japanese Society for the Promotion of Science at the Fujita Laboratory, Tokyo Institute of Technology, Tokyo, Japan. From 2008 to 2012, she was an associate professor at the Technical

University of Munich. Since 2013, she has served as TUM Liesel Beckmann Distinguished Professor and has the Chair of Information-oriented Control in the Department of Electrical and Computer Engineering at the Technical University of Munich. Her main research interests include cooperative, distributed and networked control with applications in human-machine interaction, multi-robot systems, and general robotics. She has published more than 150 papers in international journals, books and refereed conferences. Dr. Hirche has served on the Editorial Boards of the IEEE Transactions on Control of Network Systems, IEEE Transactions on Control Systems Technology, and the IEEE Transactions on Haptics. She has received multiple awards such as the Rohde & Schwarz Award for her PhD thesis, the IFAC World Congress Best Poster Award in 2005 and - together with students - the 2018 Outstanding Student Paper Award of the IEEE Conference on Decision and Control as well as Best Paper Awards from IEEE Worldhaptics and the IFAC Conference of Manoeuvring and Control of Marine Craft in 2009.