An Uncertainty-Based Control Lyapunov Approach for Control-Affine Systems Modeled by Gaussian Process — **Source link** ↗

Jonas Umlauft, Lukas Pohler, Sandra Hirche

Institutions: Technische Universität München

Topics: Control-Lyapunov function, Gaussian process, Affine transformation, Control system and Dynamical systems theory

Related papers:

- Gaussian Processes for Machine Learning

- Information-Theoretic Regret Bounds for Gaussian Process Optimization in the Bandit Setting

- Stable Gaussian process based tracking control of Euler–Lagrange systems

- Feedback linearization using Gaussian processes

- Safe and robust learning control with Gaussian processes

An Uncertainty-based Control Lyapunov Approach for Control-affine Systems Modeled by Gaussian Process

Jonas Umlauft[1], Lukas Pöhler, Sandra Hirche[1]

Abstract—**Data-driven approaches in control allow for identification of highly complex dynamical systems with minimal prior knowledge. However, properly incorporating model uncertainty in the design of a stabilizing control law remains challenging. Therefore, this article proposes a control Lyapunov function framework which semiglobally asymptotically stabilizes a partially unknown fully actuated control affine system with high probability. We propose an uncertainty-based control Lyapunov function which utilizes the model fidelity estimate of a Gaussian process model to drive the system in areas near training data with low uncertainty. We show that this behavior maximizes the probability that the system is stabilized in the presence of power constraints using equivalence to dynamic programming. A simulation on a nonlinear system is provided.**

Index Terms—**Lyapunov methods, Machine learning, Uncertain systems, Robust control, Nonlinear systems identification.**

I. INTRODUCTION

CONTROL engineering is increasingly applied in domains where precise model descriptions cannot be derived analytically. In parallel, improved sensor, storage and processing technology make data-driven approaches feasible for modeling complex systems with high precision. Therefore, data-driven control has emerged to complement powerful machine learning algorithms with the mathematical analysis of control theory.

Traditional system identification techniques are based on parametric estimation paradigms, which suffer from finite model parameters and bias towards a particular systems class. Nonparametric models are a suitable alternative as they adapt their model complexity to the given data using minimal prior knowledge, see [1] for an overview in the linear case.

For nonlinear systems, Gaussian processes (GPs) are applied as a nonparametric, probabilistic model [2] due to following properties: First, an implicit bias-variance trade-off avoids over- or underfitting. Second, prior knowledge is directly incorporated using Bayesian principles. And, most importantly, the model quantifies its ignorance due to missing data [3]. Many application utilize GPs in a model predictive control scheme [4], [5] or reinforcement learning [6]. Inverse models for robotic manipulators are also obtained using semi-parametric methods, see [7]. However, none of these perform a stability analysis of the controlled system. This is first considered in an adaptive control scheme in [8], followed by a

computed torque approach in [9] and a feedback linearization control in [10]. These approaches focus on ultimate boundedness, but do not ensure asymptotic convergence. In addition, most applications do not take advantage of the inferred model fidelity since it is either ignored [11] or used for other control tasks, e.g., gain tuning [12]. However, areas of the state space, where the model is highly uncertain, bare a risk for any model-based control law to cause undesired behavior of the system.

The contribution of this work is a novel approach to stabilize nonlinear systems with partially unknown dynamics taking model confidence explicitly into account. We utilize techniques from dynamic programming to derive an uncertainty-based control Lyapunov function. It utilizes the model fidelity knowledge of the GP to optimally avoid uncertain regions in the state space. Based on this control Lyapunov function, the control law achieves semiglobal asymptotic stability with high probability on the partially unknown system. We show that this probability is maximized through our choice of the control Lyapunov function under input constraints.

The article is structured as follows: After defining the problem setting in Sec. II, Sec. III reviews Gaussian process state space models (GP-SSM). Section IV presents the proposed control law and the main results. A numerical illustration is shown in Sec. V followed by a conclusion in Sec. VI.

II. PROBLEM FORMULATION

Consider a fully actuated control affine system[1]

$$\dot{\boldsymbol{x}} = \boldsymbol{f}(\boldsymbol{x}) + \boldsymbol{G}(\boldsymbol{x})\boldsymbol{u}, \quad \boldsymbol{x}(0) = \boldsymbol{x}_0, \qquad (1)$$

with state $\boldsymbol{x} \in \mathcal{X} \subset \mathbb{R}^n$, compact $\mathcal{X}$ and input $\boldsymbol{u} \in \mathcal{U} \subseteq \mathbb{R}^n$. The goal is to asymptotically stabilize the system at an arbitrary point $\boldsymbol{x}_* \in \mathcal{X}$ under the following assumptions:

Assumption 1: The function $\boldsymbol{f} : \mathcal{X} \to \mathbb{R}^n$ is unknown but the function value $\boldsymbol{f}(\boldsymbol{x}_*)$ is known.
This assumption is mild, because it only requires a single (noise free) measurement of $\boldsymbol{f}$ as prior knowledge. Without loss of generality, we take $\boldsymbol{x}_*$ as an equilibrium point in the origin, thus $\boldsymbol{f}(\boldsymbol{0}) = \boldsymbol{0}$.

We also impose a restriction on the "complexity" of the function $\boldsymbol{f}$ as measured under the reproducing kernel

[1] Member of the Chair of Information-oriented Control, Department of Electrical and Computer Engineering, Technical University of Munich, D-80333 Munich, Germany [jonas.umlauft, lukas.poehler, hirche]@tum.de

[1]**Notation:** Lower/upper case bold symbols denote vectors/matrices, $\mathbb{R}_{+,0}/\mathbb{R}_+$ all real positive numbers with/without zero and $\mathbb{E}[\cdot]/\mathbb{V}[\cdot]$ the expected value/variance of a random variable, respectively. $\boldsymbol{I}_n$ denotes the $n \times n$ identity matrix, $\|\cdot\|$ the Euclidean norm and $\nabla_{\boldsymbol{x}} f$ the gradient of a function f with respect to $\boldsymbol{x}$.

Hilbert space (RKHS) norm. A RKHS is a complete subspace of the L_2, for which the inner product $\langle \cdot, \cdot \rangle_k$ fulfills the reproducing property $\langle f, k(\boldsymbol{x}, \cdot) \rangle_k = f(\boldsymbol{x})$. The induced norm $\|f\|_k = \sqrt{\langle f, f \rangle_k}$ is a measure for the smoothness of f.

Assumption 2: The function $\boldsymbol{f}(\boldsymbol{x})$ has a bounded RKHS norm under a known kernel k, $\|f_j\|_k < \infty$ for $j = 1, \ldots, n$. This assumption limits irregularities (like discontinuities) of the function [3]. The norm is computed using the eigendecomposition of the kernel using Mercer's Theorem. For all nondegenerated kernels (see Sec. III-A), the RKHS is dense in the space of continuous functions on a compact domain $\mathcal{X}$ [13]. Thus, the considered class of systems is very flexible by uniformly approximating any continuous function [14].

In addition, we assume the availability of a training set.

Assumption 3: The state and a noisy version of its derivative are measured over a finite time horizon, thus N data pairs are given in the training set

$$\mathcal{D} = \left\{ \left(\boldsymbol{x}^{(i)}, \boldsymbol{y}^{(i)} \right) \right\}_{i=1}^{N}, \quad \boldsymbol{y}^{(i)} = \boldsymbol{f}\left(\boldsymbol{x}^{(i)} \right) + \boldsymbol{\omega}_i, \quad (2)$$

where $\boldsymbol{\omega}_i$ are i.i.d. samples $\boldsymbol{\omega}_i \sim \mathcal{N}(0, \sigma_{\mathrm{on}}^2 \boldsymbol{I}_n)$, $\sigma_{\mathrm{on}}^2 \in \mathbb{R}_+$. This only requires, that the uncontrolled system $\dot{x} = \boldsymbol{f}(\boldsymbol{x})$ has a finite escape time, which is a very mild condition. The availability of measurements is essential for any data-driven design and nowadays well justified with improved sensor technology. The derivative measurement is even allowed to be noisy, i.e. from a practical perspective discrete time state measurements would suffice to compute the derivative approximately. We also assume the full actuation holds everywhere and the effect of the input on the state is perfectly known.

Assumption 4: The function $\boldsymbol{G}(\boldsymbol{x}) : \mathcal{X} \to \mathbb{R}^{n \times n}$ is known, differentiable and $\mathrm{rank}\,\boldsymbol{G}(\boldsymbol{x}) = n$, $\forall \boldsymbol{x} \in \mathcal{X}$.

This assumption holds for example for robotic manipulators, coupled multi-tank systems with inflows for each tank and many others. It is also common for control affine systems [15].

Based on these assumptions, the goal is to find a state feedback control law $\boldsymbol{u}(\boldsymbol{x})$, which asymptotically stabilizes the system (1). We derive a control Lyapunov function based on a GP representation of the dynamics. This Lyapunov function is constructed using dynamic programming to avoid areas where the uncertainty of the model is high and favor areas near training data, where model fidelity is high. The identification of the uncontrolled dynamic model $\dot{\boldsymbol{x}} = \hat{\boldsymbol{f}}(\boldsymbol{x})$ employs GP-SSMs as reviewed in the following.

III. GAUSSIAN PROCESS STATE SPACE MODELS

A. Gaussian process regression

Gaussian processes allow to model nonlinear dynamical systems and further indicate the model fidelity based on distance to training data. A GP is a stochastic process which assigns to any finite subset $\{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_M\} \subset \mathcal{X}$ in a continuous input domain a joint Gaussian distribution [3]. It is considered as a distribution over functions $\hat{f} : \mathcal{X} \to \mathbb{R}$ denoted by

$$\hat{f}(\boldsymbol{x}) \sim \mathcal{GP}(m(\boldsymbol{x}), k(\boldsymbol{x}, \boldsymbol{x}')). \quad (3)$$

The GP is fully specified by a mean function $m(\boldsymbol{x}) : \mathcal{X} \to \mathbb{R}$ and a kernel $k(\boldsymbol{x}, \boldsymbol{x}') : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$. Both depend on the hyperparameters $\boldsymbol{\psi}$ which characterize the functions over which the

GP is a distribution. For the prior mean function, it is common practice to set $m_j(\boldsymbol{x}) = 0$, for all $j = 1, \ldots, n$ if no further knowledge is available. This is also assumed in the following without loss of generality. A widely used covariance function is the nondegenerated squared exponential (SE) kernel

$$k^{\mathrm{SE}}(\boldsymbol{x}, \boldsymbol{x}') = \sigma_f^2 \exp \left(\sum_{j=1}^{n} \frac{(x_j - x_j')^2}{-2l_j^2} \right). \quad (4)$$

Its hyperparameters $\boldsymbol{\psi} = \begin{bmatrix} l_1 & \cdots & l_n & \sigma_f^2 \end{bmatrix}^{\mathsf{T}}$ are the signal variance $\sigma_f^2 \in \mathbb{R}_{+,0}$ and the lengthscales $l_j \in \mathbb{R}_+$, $j = 1, \ldots, n$. They are obtained from a likelihood maximization according to Bayesian principles. It is usually solved with gradient-based methods [3] even though it is non-convex.

Since (3) represents only functions with scalar outputs, n independent GPs are utilized to model the dynamical system (1)

$$\hat{\boldsymbol{f}}(\boldsymbol{x}) = \begin{cases} \hat{f}_1(\boldsymbol{x}) \sim \mathcal{GP}\left(0, k_1(\boldsymbol{x}, \boldsymbol{x}')\right) \\ \vdots \qquad \qquad \vdots \\ \hat{f}_n(\boldsymbol{x}) \sim \mathcal{GP}\left(0, k_n(\boldsymbol{x}, \boldsymbol{x}')\right), \end{cases} \quad (5)$$

where $\boldsymbol{\psi}_j$ with $j = 1, \ldots, n$ are corresponding hyperparameters of the kernel functions $k_j(\cdot, \cdot)$.

Given a dataset $\mathcal{D}$ in (2), the GP is employed for regression. For a test input $\boldsymbol{x}^*$, the j-th component of the inferred output $\boldsymbol{y}^*$ is jointly Gaussian distributed with the training data

$$\begin{bmatrix} y_j^* \\ \boldsymbol{y}_j \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} 0 \\ \boldsymbol{0} \end{bmatrix}, \begin{bmatrix} k_j^* & \boldsymbol{k}_j^{\mathsf{T}} \\ \boldsymbol{k}_j & \boldsymbol{K}_j + \sigma_{\mathrm{on}}^2 \boldsymbol{I}_N \end{bmatrix} \right), \quad (6)$$

where $k_j^* = k_j(\boldsymbol{x}^*, \boldsymbol{x}^*) \in \mathbb{R}$, $\boldsymbol{y}_j = \begin{bmatrix} y_j^{(1)} & \cdots & y_j^{(N)} \end{bmatrix}^{\mathsf{T}} \in \mathbb{R}^N$,

$$\boldsymbol{k}_j = \begin{bmatrix} k_j\left(\boldsymbol{x}^{(1)}, \boldsymbol{x}^* \right) & \cdots & k_j\left(\boldsymbol{x}^{(N)}, \boldsymbol{x}^* \right) \end{bmatrix}^{\mathsf{T}} \in \mathbb{R}^N,$$

and

$$\boldsymbol{K}_j = \begin{bmatrix} k_j\left(\boldsymbol{x}^{(1)}, \boldsymbol{x}^{(1)} \right) & \cdots & k_j\left(\boldsymbol{x}^{(1)}, \boldsymbol{x}^{(N)} \right) \\ \vdots & \ddots & \vdots \\ k_j\left(\boldsymbol{x}^{(N)}, \boldsymbol{x}^{(1)} \right) & \cdots & k_j\left(\boldsymbol{x}^{(N)}, \boldsymbol{x}^{(N)} \right) \end{bmatrix} \in \mathbb{R}^{N \times N}. \quad (7)$$

Conditioning on test input $\boldsymbol{x}^*$ and training data $\mathcal{D}$ yields

$$\mathbb{E}[y_j^* | \mathcal{D}, \boldsymbol{x}^*, \boldsymbol{\psi}_j] = \boldsymbol{k}_j^{\mathsf{T}}(\boldsymbol{K}_j + \sigma_{\mathrm{on}}^2 \boldsymbol{I})^{-1} \boldsymbol{y}_j, \quad (8)$$

$$\mathbb{V}[y_j^* | \mathcal{D}, \boldsymbol{x}^*, \boldsymbol{\psi}_j] = k_j^* - \boldsymbol{k}_j^{\mathsf{T}}(\boldsymbol{K}_j + \sigma_{\mathrm{on}}^2 \boldsymbol{I})^{-1} \boldsymbol{k}_j. \quad (9)$$

B. Including known equilibrium to GP-SSMs

To make use of the known equilibrium point $\boldsymbol{f}(\boldsymbol{0}) = \boldsymbol{0}$ in Assumption 1, we introduce a minor modification to the GP-SSM by extending (6) as follows

$$\begin{bmatrix} y_j^* \\ \boldsymbol{y}_j \\ 0 \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} 0 \\ \boldsymbol{0} \\ 0 \end{bmatrix} \begin{bmatrix} k_j^* & \boldsymbol{k}_j^{\mathsf{T}} & k_j\left(\boldsymbol{x}^*, \boldsymbol{0} \right) \\ \boldsymbol{k}_j & \boldsymbol{K}_j + \sigma_{\mathrm{on}}^2 \boldsymbol{I}_N & \boldsymbol{k}_j^{\boldsymbol{0}} \\ k_j\left(\boldsymbol{0}, \boldsymbol{x}^* \right) & \boldsymbol{k}_j^{\boldsymbol{0}\mathsf{T}} & k_j\left(\boldsymbol{0}, \boldsymbol{0} \right) \end{bmatrix} \right),$$

where $\boldsymbol{k}_j^{\boldsymbol{0}} = \begin{bmatrix} k_j\left(\boldsymbol{x}^{(1)}, \boldsymbol{0} \right) & \cdots & k_j\left(\boldsymbol{x}^{(N)}, \boldsymbol{0} \right) \end{bmatrix}^{\mathsf{T}} \in \mathbb{R}^N$.

Note that no observation noise is added for the training point $(\boldsymbol{0}, 0)$ as this is known a priori without uncertainty and not measured. For mean and variance predictions, we define

$$\mu_j(\boldsymbol{x}^*) := \mathbb{E}[y_j^* | \mathcal{D}, \boldsymbol{x}^*, (\boldsymbol{0}, 0), \boldsymbol{\psi}_j], \quad j = 1, \ldots, n \quad (10)$$

$$\sigma_j^2(\boldsymbol{x}^*) := \mathbb{V}[y_j^* | \mathcal{D}, \boldsymbol{x}^*, (\boldsymbol{0}, 0), \boldsymbol{\psi}_j], \quad j = 1, \ldots, n \quad (11)$$

equivalent to (8) and (9) and their concatenations

$$\boldsymbol{\mu}(\boldsymbol{x}^*) := \begin{bmatrix} \mu_1(\boldsymbol{x}^*) & \cdots & \mu_n(\boldsymbol{x}^*) \end{bmatrix}^\mathsf{T}, \tag{12}$$

$$\boldsymbol{\sigma}^2(\boldsymbol{x}^*) := \begin{bmatrix} \sigma_1^2(\boldsymbol{x}^*) & \cdots & \sigma_n^2(\boldsymbol{x}^*) \end{bmatrix}^\mathsf{T}. \tag{13}$$

Due to the continuity of the SE kernel, the bounds

$$\bar{\sigma}^2 := \max_{\boldsymbol{x}^* \in \mathcal{X}} \|\boldsymbol{\sigma}^2(\boldsymbol{x}^*)\|, \qquad \bar{\mu} := \max_{\boldsymbol{x}^* \in \mathcal{X}} \|\boldsymbol{\mu}(\boldsymbol{x}^*)\|, \tag{14}$$

exist. For the variance function is the following concluded.

Lemma 1: The posterior variance function $\boldsymbol{\sigma}^2 : \mathcal{X} \to \mathbb{R}^n_{+,0}$ defined in (13) from a Gaussian process with SE kernel (4) is component-wise positive definite.

Proof: According to [10], $\sigma_N^2(\boldsymbol{x}) \geq \sigma_{N+1}^2(\boldsymbol{x})$, $\forall \boldsymbol{x} \in \mathcal{X}$ holds, where $\sigma_N^2(\boldsymbol{x})$ and $\sigma_{N+1}^2(\boldsymbol{x})$ denote the variance functions for N and $N+1$ training data points, respectively. Considering only the training data point $(\boldsymbol{0}, 0)$, $\sigma^2(\boldsymbol{0}) = 0$ holds, as it can be directly seen from (9) and (11). This leads with the non-negativity to $\boldsymbol{\sigma}^2(\boldsymbol{0}) = \boldsymbol{0}$. For all $\boldsymbol{x} \in \mathcal{X} \backslash \{\boldsymbol{0}\}$, we can conclude that $\sigma^2(\boldsymbol{x}) > 0$ because of the measurement noise $\sigma_{\text{on}}^2 > 0$ and finite training data points. ∎

C. Model error

In the previous section, we have shown how to incorporate Assumption 1 in the model, while this section investigates the effects of Assumption 2. We limit the analysis here to the SE kernel defined in (4), however it can be generalized to other nondegenerated kernels, e.g. the Matern kernel (for an overview see [3, Table 4.1]). Based on the work in [16], we define the maximum information gain under a kernel k_j as

$$\gamma_j = \max_{\bar{\boldsymbol{x}}^{(1)},\dots,\bar{\boldsymbol{x}}^{(N+1)} \in \mathcal{X}} \frac{1}{2} \log |\boldsymbol{I}_N + \sigma_{\text{on}}^{-2} \bar{\boldsymbol{K}}_j(\boldsymbol{x}, \boldsymbol{x}')|, \tag{15}$$

where $|\cdot|$ denotes the determinate, $\boldsymbol{x}, \boldsymbol{x}' \in \{\bar{\boldsymbol{x}}^{(1)}, \dots, \bar{\boldsymbol{x}}^{(N+1)}\}$ and $\bar{\boldsymbol{K}}_j$ is the covariance matrix equivalently defined to (7). Intuitively, this is a measure for the topology of $\mathcal{X}$ by taking the best distribution of $N+1$ data points $\bar{\boldsymbol{x}}^{(i)}$ in $\mathcal{X}$ in terms of the information gain. This allows to upper bound the difference between the true function $\boldsymbol{f}(\boldsymbol{x})$ and the inferred mean $\boldsymbol{\mu}(\boldsymbol{x})$ with high probability.

Lemma 2: For any compact set $\mathcal{X} \subset \mathbb{R}^n$ and a probability $\delta \in (0,1)$ holds

$$\mathcal{P}\{\|\boldsymbol{\mu}(\boldsymbol{x}) - \boldsymbol{f}(\boldsymbol{x})\| \leq \|\boldsymbol{\beta}\| \|\boldsymbol{\sigma}(\boldsymbol{x})\|, \forall \boldsymbol{x} \in \mathcal{X}\} \geq (1-\delta)^n, \tag{16}$$

where $\boldsymbol{\mu}(\boldsymbol{x})$ and $\boldsymbol{\sigma}(\boldsymbol{x})$ are the mean and standard deviation posterior function of the GP in (12) and (13), respectively, and $\boldsymbol{\beta} = \begin{bmatrix} \beta_1 & \dots & \beta_n \end{bmatrix}^\mathsf{T}$ with

$$\beta_j = \sqrt{2\|f_j\|_{k_j}^2 + 300 \gamma_j \log^3\left(\frac{N+1}{\delta}\right)}, \quad \forall j = 1, \dots, n ,$$

where f_j is the j-th component of $\boldsymbol{f}$ and γ_j is defined in (15).

Proof: Similar to [12, Lemma 1], this is derived from [16, Theorem 6], which states for the scalar case

$$\mathcal{P}\{|\mu(\boldsymbol{x}) - f(\boldsymbol{x})| \leq \beta \sigma(\boldsymbol{x}), \ \forall \boldsymbol{x} \in \mathcal{X}\} \geq 1 - \delta. \tag{17}$$

Using the fact that $\boldsymbol{\omega}$ is uncorrelated, yields

$$(1-\delta)^n \leq \mathcal{P}\left\{ \bigcap_{j=1}^n |\mu_j(\boldsymbol{x}) - f_j(\boldsymbol{x})| \leq \beta_j \sigma_j(\boldsymbol{x}), \forall \boldsymbol{x} \in \mathcal{X} \right\}$$

$$\leq \mathcal{P}\left\{ \|\boldsymbol{\mu}(\boldsymbol{x}) - \boldsymbol{f}(\boldsymbol{x})\| \leq \left\| \begin{bmatrix} \beta_1 \sigma_1(\boldsymbol{x}) \\ \vdots \\ \beta_n \sigma_n(\boldsymbol{x}) \end{bmatrix} \right\|, \forall \boldsymbol{x} \in \mathcal{X} \right\}$$

and using the triangle and the Cauchy-Schwarz inequality

$$\left\| \begin{bmatrix} \beta_1 \sigma_1(\boldsymbol{x}) & \cdots & \beta_n \sigma_n(\boldsymbol{x}) \end{bmatrix}^\mathsf{T} \right\| \leq \sum_{j=1}^n \beta_j \sigma_j(\boldsymbol{x}) \leq \|\boldsymbol{\beta}\| \|\boldsymbol{\sigma}(\boldsymbol{x})\|,$$

yields the presented result. ∎

The information gain γ_j has a sublinear dependence on N for most kernels as shown in [16]. Thus, even though $\boldsymbol{\beta}$ increases, this bound concludes, that the model error $\|\boldsymbol{f}(\boldsymbol{x}) - \boldsymbol{\mu}(\boldsymbol{x})\|$ descreases with more training points. Therefore, Lemma 2 allows to make high probability statements regarding the maximum model error, which will be needed in the analysis of the proposed control law.

IV. CONTROL DESIGN & ANALYSIS

For the control, we consider the feedback linearizing form

$$\boldsymbol{u}(\boldsymbol{x}) = -\boldsymbol{G}^{-1}(\boldsymbol{x}) \left(\boldsymbol{\mu}(\boldsymbol{x}) + k_c \nabla_{\boldsymbol{x}}^\mathsf{T} V(\boldsymbol{x}) \right), \tag{18}$$

where $\boldsymbol{\mu}(\boldsymbol{x})$ is the mean of the GP-SSM defined in (12), $k_c > 0$ and $V : \mathcal{X} \to \mathbb{R}^+$ is a positive definite differentiable function employed as control Lyapunov function. In a first step (Sec. IV-A), we will generally show which conditions for $V(\boldsymbol{x})$ and k_c must hold to stabilize the system with a certain probability. In Sec. IV-B we analyze how this probability is maximized under finite control power constraints. Section IV-C finally proposes a specific Lyapunov candidate which takes the model uncertainty into account.

A. Conditions for asymptotic stability

Theorem 1: Consider the unknown system (1) under Assumptions 1-4 represented by a Gaussian process with SE kernel in (12), (13) and a positive definite differentiable function $V(\boldsymbol{x})$ with $\|\nabla_{\boldsymbol{x}} V(\boldsymbol{0})\| = 0$ and

$$\|\boldsymbol{\sigma}(\boldsymbol{x})\| - \|\nabla_{\boldsymbol{x}} V(\boldsymbol{x})\| \leq 0, \qquad \forall \boldsymbol{x} \in \mathcal{X}. \tag{19}$$

Further consider the control law (18) with $k_c > \|\boldsymbol{\beta}\|$ with $\boldsymbol{\beta}$ from Lemma 2. Then, the origin of the closed loop system is semiglobally asymptotically stable with probability at least $(1-\delta)^n$ for all $\boldsymbol{x}_0 \in \mathcal{X}$.

Proof: For $V(\boldsymbol{x})$ as Lyapunov candidate holds $\forall \boldsymbol{x} \in \mathcal{X}$

$$\dot{V}(\boldsymbol{x}) = \nabla_{\boldsymbol{x}} V(\boldsymbol{x}) \dot{\boldsymbol{x}} = \nabla_{\boldsymbol{x}} V(\boldsymbol{x})(\boldsymbol{f}(\boldsymbol{x}) + \boldsymbol{G}(\boldsymbol{x})\boldsymbol{u}(\boldsymbol{x}))$$

$$= \nabla_{\boldsymbol{x}} V(\boldsymbol{x}) \left(\boldsymbol{f}(\boldsymbol{x}) - \boldsymbol{\mu}(\boldsymbol{x}) - k_c \nabla_{\boldsymbol{x}}^\mathsf{T} V(\boldsymbol{x}) \right) \tag{20}$$

$$\leq \|\nabla_{\boldsymbol{x}} V(\boldsymbol{x})\| \|\boldsymbol{f}(\boldsymbol{x}) - \boldsymbol{\mu}(\boldsymbol{x})\| - k_c \|\nabla_{\boldsymbol{x}} V(\boldsymbol{x})\|^2,$$

where the inequality yields from the Cauchy-Schwarz inequality. Employing Lemma 2 allows the conclusion

$$\mathcal{P}\left\{ \forall \boldsymbol{x} \in \mathcal{X}, \dot{V}(\boldsymbol{x}) \leq \|\nabla_{\boldsymbol{x}} V(\boldsymbol{x})\| \left(\|\boldsymbol{\beta}\| \|\boldsymbol{\sigma}(\boldsymbol{x})\| - k_c \|\nabla_{\boldsymbol{x}} V(\boldsymbol{x})\| \right) \right\}$$

$$\geq (1-\delta)^n \tag{21}$$

which yields

$$\mathcal{P}\left\{\forall \boldsymbol{x} \in \mathcal{X} \setminus \{\mathbf{0}\}, \ \dot{V}(\boldsymbol{x}) < 0\right\} \geq (1-\delta)^n \qquad (22)$$

under condition (19) and $k_c > \|\boldsymbol{\beta}\|$. The strict inequality holds because $\|\nabla_{\boldsymbol{x}} V(\boldsymbol{x})\|$ is lower bounded by the positive definite function $\|\boldsymbol{\sigma}(\boldsymbol{x})\|$. Additionally, $\dot{V}(\mathbf{0}) = 0$ holds from condition $\|\nabla_{\boldsymbol{x}} V(\mathbf{0})\| = 0$. ∎

Remark 1: Here, semiglobal stability is achieved because the set $\mathcal{X}$ is bounded to apply Lemma 2 properly. However, the region of attraction, the compact set $\mathcal{X}$, can be chosen arbitrarily large as long as it is finite. For any compact $\mathcal{X}$ (no matter how large), there exists a gain k_c which results in asymptotic stability on the entire set.

B. Considering power limitations

In principle many control Lyapunov functions stabilize the system according to Theorem 1. However, it is not trivial to find a suitable candidate if we consider the practical limitation, that only finite control power $u_{\max} > 0$ is available

$$\|\boldsymbol{u}(\boldsymbol{x})\| \leq u_{\max}, \qquad \forall \boldsymbol{x} \in \mathcal{X}. \qquad (23)$$

Given this constraint, the goal for the control design is to maximize the probability that the system is asymptotically stabilized. We will refer to this probability as the *confidence* $\Delta = (1-\delta)^n$ and denote $\|\boldsymbol{\beta}\| = \beta(\Delta)$ from Lemma 2. Thus, we want to choose $V(\boldsymbol{x})$ and k_c such that the confidence Δ is maximized, the power constraint (23) is satisfied and the equilibrium is stabilized in the sense of Theorem 1. This is formalized and solved in the following:

Lemma 3: Consider the upper bound for the control input

$$\|\boldsymbol{u}(\boldsymbol{x})\| \leq (\bar{\mu} + k_c \|\nabla_{\boldsymbol{x}}^{\mathsf{T}} V(\boldsymbol{x})\|)/\underline{\lambda}_{\boldsymbol{G}} \leq u_{\max} \quad \forall \boldsymbol{x} \in \mathcal{X}, \quad (24)$$

where $\bar{\mu}$ is defined in (14) and $\underline{\lambda}_{\boldsymbol{G}}$ is the smallest singular value of $\boldsymbol{G}$, $\forall \boldsymbol{x} \in \mathcal{X}$. If the input power constraint (24) fulfills $u_{\max} > \bar{\mu}/\underline{\lambda}_{\boldsymbol{G}}$, the maximization of the confidence

$$\Delta^* = \underset{\|\nabla_{\boldsymbol{x}} V(\boldsymbol{x})\|, k_c}{\arg\max} \ \Delta \qquad (25a)$$

$$\text{s.t. } (\bar{\mu} + k_c\|\nabla_{\boldsymbol{x}} V(\boldsymbol{x})\|)/\underline{\lambda}_{\boldsymbol{G}} \leq u_{\max}, \ k_c > \beta(\Delta),$$
$$\|\nabla_{\boldsymbol{x}} V(\boldsymbol{x})\| \geq \|\boldsymbol{\sigma}(\boldsymbol{x})\|, \quad \forall \boldsymbol{x} \in \mathcal{X} \qquad (25b)$$

is solved by

$$\|\nabla_{\boldsymbol{x}} V(\boldsymbol{x})\| = \|\boldsymbol{\sigma}(\boldsymbol{x})\|, \qquad k_c = (u_{\max}\underline{\lambda}_{\boldsymbol{G}} - \bar{\mu})/\bar{\sigma}. \quad (26)$$

The largest possible confidence for any arbitrarily small $\epsilon > 0$ with $(u_{\max}\underline{\lambda}_{\boldsymbol{G}} - \bar{\mu})/\bar{\sigma} > \epsilon > 0$ is

$$\Delta^* = \beta^{-1}\left((u_{\max}\underline{\lambda}_{\boldsymbol{G}} - \bar{\mu})/\bar{\sigma} - \epsilon\right), \qquad (27)$$

where β^{-1} is the inverse of $\beta(\Delta)$.

Proof: Since $\beta(\Delta)$ is monotone, the problem can be rewritten to optimize over β. Since the constraint set is then convex and the objective is linear, the maximum is attained on the boundary. As k_c must not depend on $\boldsymbol{x}$, but must hold $\forall \boldsymbol{x} \in \mathcal{X}$, the lower bound $(u_{\max}\underline{\lambda}_{\boldsymbol{G}} - \bar{\mu})/\bar{\sigma}$ is taken for k_c. The inverse β^{-1} exists since $\beta(\Delta)$ is monotone. ∎

Remark 2: Lemma 3 does not maximize the reliability under the constraint (23) but uses an upper bound for the control input as constraint, given by (24). The solution (27) is a lower bound to the actual achievable confidence, because (24) is less restrictive than (23). The condition for optimality $\|\nabla_{\boldsymbol{x}} V(\boldsymbol{x})\| = \|\boldsymbol{\sigma}(\boldsymbol{x})\|$ also holds for the true constraint (23) and will mainly be used in the following.

C. Uncertainty-based Lyapunov function

From Lemma 3, we obtain a condition for the optimal (with respect to the confidence) Lyapunov function given by

$$\|\boldsymbol{\sigma}(\boldsymbol{x})\| = \|\nabla_{\boldsymbol{x}} V(\boldsymbol{x})\|, \qquad \forall \boldsymbol{x} \in \mathcal{X} \qquad (28)$$
$$\text{with boundary condition } V(\mathbf{0}) = 0.$$

Solving this differential equation is not trivial, however it is a well known problem: A necessary condition for the optimal value function in dynamic programming takes the same form as derived from Hamilton-Jacobi-Bellman (HJB) equation. This equivalence is formalized in the following.

Theorem 2: Consider the value function $\hat{V} : \mathcal{X} \to \mathbb{R}_{+,0}$ defined as

$$\hat{V}(\boldsymbol{x}_0) = \min_{\hat{\boldsymbol{x}}(s)} \int_0^S c_{\boldsymbol{\sigma}}\left(\hat{\boldsymbol{x}}(s)\right) ds \qquad (29)$$
$$\text{s.t. } \hat{\boldsymbol{x}}(0) = \boldsymbol{x}_0, \quad \hat{\boldsymbol{x}}(S) = \mathbf{0}, \qquad (30)$$

representing the accumulated cost along an optimal path $\hat{\boldsymbol{x}}(s)$ from the current state $\boldsymbol{x}_0$ to the origin, which is parameterized by the arclength $s \in [0 \ S]$, $S \in \mathbb{R}_+$ for which holds $\|\frac{d\hat{\boldsymbol{x}}}{ds}\| = 1$. For a stage cost $c_{\boldsymbol{\sigma}} : \mathcal{X} \to \mathbb{R}_{+,0}$, given by

$$c_{\boldsymbol{\sigma}}(\boldsymbol{x}) = \|\boldsymbol{\sigma}(\boldsymbol{x})\|, \qquad (31)$$

the value function $\hat{V}$ is a control Lyapunov function $V(\boldsymbol{x})$ which solves the optimization (25).

Proof: Derived from the HJB equation, a necessary condition for the value function is

$$\|\nabla_{\boldsymbol{x}} \hat{V}(\boldsymbol{x})\| = c_{\boldsymbol{\sigma}}(\boldsymbol{x}), \qquad \forall \boldsymbol{x} \in \mathcal{X} \qquad (32)$$
$$\text{with boundary condition } \hat{V}(\mathbf{0}) = 0.$$

This is also known as the inhomogeneous Eikonal equation [17]. From the choice (31), it can be directly seen that the conditions for deriving the value function (32) and for maximizing the achievable confidence in (28) are equivalent. ∎

Remark 3: The solution to the inhomogeneous Eikonal equation is for most cost functions $c_{\boldsymbol{\sigma}}$ not differentiable. However, it is common practice to consider viscosity solutions [18], which exist for continuous, bounded and positive definite cost functions, which is the case for $c_{\boldsymbol{\sigma}}(\boldsymbol{x})$ in (31) (Lemma 1). We conclude that any solution of the Eikonal equation taken as control Lyapunov function also maximizes the confidence according to Lemma 3. The equivalence is beneficial from a practical point of view: It allows to use the value function from dynamic programming [19] as control Lyapunov function to maximize the achievable confidence. To compute it efficiently, various techniques, e.g. *fast marching methods* [20], are available. It is also interesting from a theoretical perspective as it

matches an intuitive understanding: Under a power constraint, the maximum confidence is achieved if the control minimizes the uncertainty about the model along the path. The control drives the system towards areas of the state space with training points to avoid areas with low model fidelity. The control law tracks the optimal path (in terms of lowest uncertainty) which is generated by

$$\dot{\boldsymbol{x}} = -\nabla_{\boldsymbol{x}}\hat{V}(\boldsymbol{x}). \tag{33}$$

Applying this uncertainty-based control Lyapunov function in the control law (18) allows the following conclusion.

Proposition 1: Consider the unknown system (1) under Assumptions 1-4 modeled by a GP in (12) with SE kernel (4) and the power constraint (24). Further consider the control law (18) where $k_c = (u_{\max}\lambda_{\boldsymbol{G}} - \bar{\mu})/\bar{\sigma}$ and $V(\boldsymbol{x})$ is the viscosity solution of the Eikonal equation (32) with $c_{\boldsymbol{\sigma}}(\boldsymbol{x}) = \|\boldsymbol{\sigma}(\boldsymbol{x})\|$. Then, the origin of the closed loop system is semiglobally asymptotically stable with at least probability Δ^* which is the maximum in terms of Lemma 3.

Proof: The solution to the Eikonal equation $V(\boldsymbol{x})$ is positive definite from the imposed boundary condition $V(\mathbf{0}) = 0$ and the fact that the path integral over a positive definite cost function is larger zero for any starting point outside the origin. The condition (19) is fulfilled with equality and $k_c > \beta(\Delta^*)$ holds from (27). Therefore, the stability follows directly from Theorem 1 and the maximum confidence from Lemma 3 and Theorem 2. ∎

V. NUMERICAL EVALUATION

This section discusses the implementation of the proposed approach and presents a numerical evaluation on a two dimensional system.

A. Setup and implementation

Consider the unknown unstable dynamics

$$\begin{aligned} \dot{x}_1 &= x_1 + (\cos(x_1) - 1)x_2 + u_1, \\ \dot{x}_2 &= -\mathrm{s}(x_1) + x_2 + u_2, \end{aligned} \tag{34}$$

where $\mathrm{s}(x_1) = \frac{1}{1+\exp(-2x_1)} - 0.5$ is a shifted sigmoid function. The dynamics is continuous, has therefore a finite RKHS norm under the squared exponential kernel on a compact set and complies to Assumption 2 according to [16]. To obtain the dataset $\mathcal{D}$, the uncontrolled system is initialized at $\boldsymbol{x}_0 = \begin{bmatrix} 0.3 & 0 \end{bmatrix}^{\mathsf{T}}$ and $\boldsymbol{x}_0 = \begin{bmatrix} -0.3 & 0.1 \end{bmatrix}^{\mathsf{T}}$, simulated for time $T = 2$ and observations are taken with sampling interval $\Delta t = 0.3$ and noise $\sigma_{\mathrm{on}}^2 = 0.01$ according to Assumption 3. This leads to $N = 20$ on $\mathcal{X} = [-5\ 5]^2$. The training point $(\mathbf{0}, \mathbf{0})$ is added according to Assumption 1. For training the Gaussian process, the hyperparameters are obtained through likelihood maximization using a quasi-Newton method. The maximum input power is assumed to $u_{\max} = 24$. According to Lemma 3, we estimate $k_c \approx 17$ and set $\epsilon = 2$. To obtain the confidence Δ, the following Monte Carlo approach is employed because the constants γ_j and $\|f_j\|_{k_j}$ are hard to compute and the bound in Lemma 2 becomes very conservative: By sampling 10^4 realizations of the GPs, we observe that the

condition in Lemma 2 for a $\|\boldsymbol{\beta}\| = 15$ holds in $\Delta^* = 99.87\%$ of the cases. If the likelihood optimization converged to a local minimum, this sampling-based method accounts for the suboptimality of the hyperparameters, because it directly evaluates the probability of a large model error in (16). From a theoretical perspective, this suboptimality causes a different RKHS norm $\|\boldsymbol{f}\|_k$, leading to a higher $\|\boldsymbol{\beta}\|$.

The most critical part of the implementation is the computation of the control Lyapunov function V in (29) itself. For its numerical evaluation, we discretize the state-space $\mathcal{X}$ uniformly with 10^4 points and use the Dijkstra algorithm on the resulting graph to compute the cost from the origin to every node in the graph. This computation is performed offline after the GP is trained. While the controller is running, a linear interpolation between the costs of each node in the grid is used to obtain the value function in the continuous space.[2] To apply the value function as a control Lyapunov function, its gradient is evaluated by first order finite differences (with finite difference $\epsilon_{\mathrm{fd}} = 10^{-4}$). For more sophisticated algorithms, we refer to the existing literature dealing with fast marching methods [20]. The time-continuous forward simulation is performed with a variable-step 4-order Runge-Kutta solver until the trajectories reach proximity of the origin with $\|\boldsymbol{x}\| < 0.05$.

B. Results

The training data and the resulting GP-SSM are visualized in Fig. 1. It shows the generalization by the mean function which indicates the diverging behavior. The norm of the variance, which is chosen as the cost $c_{\boldsymbol{\sigma}}(\boldsymbol{x})$, increases with increasing distance to the training data. Figure 2 shows the uncertainty-based control Lyapunov function and the resulting behavior of the controlled system under the proposed control law (18) for different initial positions. The control law first pushes the system towards training data before driving in the area of high model fidelity towards the origin.

C. Discussion

Generally, the approach allows to stabilize arbitrary complex control affine fully actuated systems under the introduced assumptions due to the model flexibility. The guarantee for asymptotic convergence is obtained with probability Δ, specified by Lemma 3. The proposed Lyapunov candidate maximizes Δ in terms of Lemma 3 under a given input power constraint. Thus, there is a trade-off in the design between the achieved confidence Δ, the number of training data N, the complexity of the system to control $\|\boldsymbol{f}\|_k$ and the maximum power $u_{\max}$. For the example system (34), we obtained a confidence of $\Delta = 99.87\%$ for $u_{\max} = 24$. However, it reduces to $\Delta = 92.4\%$ for $u_{\max} = 23$ and $\Delta = 28.45\%$ for $u_{\max} = 22$. Thus the confidence is very sensitive to the maximum power. On the downside, the GP and the proposed uncertainty-based control Lyapunov function have a high computational complexity. However, most of the heavy computation is performed

---

[2]The linear interpolation is guaranteed positive definite if one grid point is positioned at the origin. It is differentiable almost everywhere (except for the grid points). The alternative spline interpolation is differentiable, but not guaranteed positive everywhere.
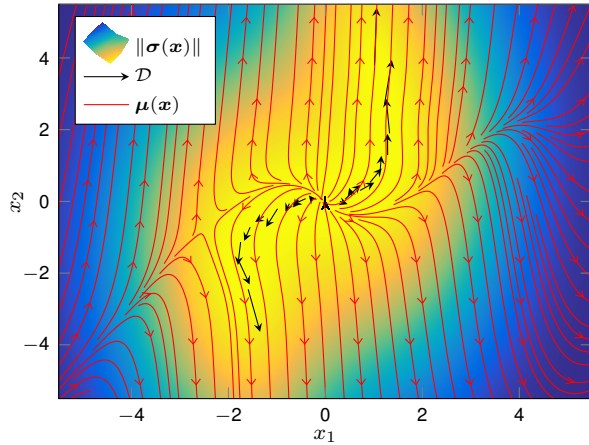
Fig. 1. The training data $\mathcal{D}$ (black arrows) obtained from the uncontrolled simulation of (34) and the mean prediction of the GP-SSM (red streamlines) defined in (12). The norm of the predicted variance (13) is shown as colormap (yellow is low, blue is high).
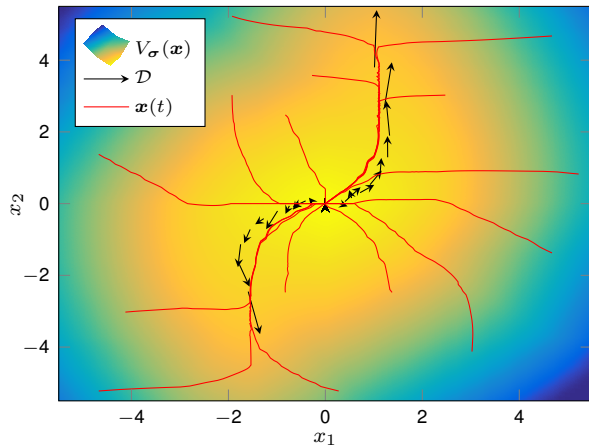

Fig. 2. The value function $V(\boldsymbol{x})$ from (29), used as uncertainty-based control Lyapunov function, is shown as colormap (yellow is low, blue is high) with training data (black arrows). Trajectories (red lines) show the closed loop behavior under the proposed control law (18) for different initial states.

offline including the collection of training data (with $\boldsymbol{u} = \boldsymbol{0}$), the training of the GP model, the Monte Carlo simulation to compute the confidence and the computation of the Lyapunov function using a gridded discretization of the state space and Dijkstra's algorithm. For the controller itself, we utilize an interpolation of the discretized Lyapunov function which is sufficiently fast in the online phase. Generally, the uncontrolled dynamics $\boldsymbol{f}(\boldsymbol{x})$ are assumed unknown, but the approach is limited to fully actuated systems with the effect of each input being perfectly known (Assumption 4).

VI. CONCLUSION

This paper introduces a control Lyapunov function approach to ensure asymptotic stability using a data-driven Gaussian process state space model. Using the specified assumptions for the fully actuated control affine dynamics, we derive general conditions on the control Lyapunov candidate which achieve asymptotic stability with high probability on the true system. To maximize this probability under control constraints, we propose an uncertainty-aware control Lyapunov function

which favors areas with low model uncertainty. We propose a numerically efficient method for its computation and illustrate it in simulation on an unstable nonlinear system.

REFERENCES

[1] G. Pillonetto, F. Dinuzzo, T. Chen, G. De Nicolao, and L. Ljung, "Kernel methods in system identification, machine learning and function estimation: A survey," *Automatica*, vol. 50, no. 3, pp. 657–682, 2014.

[2] J. Kocijan, *Modelling and Control of Dynamic Systems Using Gaussian Process Models*. Springer, 2016.

[3] C. E. Rasmussen and C. K. Williams, *Gaussian Processes for Machine Learning*. Cambridge, MA, USA: MIT Press, Jan. 2006.

[4] R. Murray-Smith, D. Sbarbaro, C. E. Rasmussen, and A. Girard, "Adaptive, cautious, predictive control with Gaussian process priors," in *IFAC Symposium on System Identification*, Aug. 2003, pp. 1195–1200.

[5] J. Kocijan, R. Murray-Smith, C. E. Rasmussen, and A. Girard, "Gaussian process model based predictive control," in *American Control Conference (ACC)*, vol. 3. IEEE, 2004, pp. 2214–2219.

[6] M. P. Deisenroth, D. Fox, and C. E. Rasmussen, "Gaussian processes for data-efficient learning in robotics and control," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 37/2,pp.408–423, 2015.

[7] D. Romeres, M. Zorzi, R. Camoriano, and A. Chiuso, "Online semi-parametric learning for inverse dynamics modeling," in *Conference on Decision and Control (CDC)*. IEEE, Dec 2016, pp. 2945–2950.

[8] G. Chowdhary, H. A. Kingravi, J. P. How, and P. A. Vela, "Bayesian nonparametric adaptive control using Gaussian processes," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 26, no. 3, pp. 537–550, Mar. 2015.

[9] T. Beckers, J. Umlauft, 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, 2017, pp. 3877 – 3884.

[10] J. Umlauft, 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.

[11] D. Nguyen-Tuong and J. Peters, "Learning robot dynamics for computed torque control using local Gaussian processes regression," in *ECSIS Symposium on Learning and Adaptive Behaviors for Robotic Systems*. IEEE, 2008, pp. 59–64.

[12] T. Beckers, J. Umlauft, 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.

[13] 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.

[14] N. Srinivas, A. Krause, S. Kakade, and M. Seeger, "Gaussian process optimization in the bandit setting: No regret and experimental design," in *International Conference on Machine Learning (ICML)*, J. Fürnkranz and T. Joachims, Eds. Haifa, Israel: Omnipress, June 2010, pp. 1015–1022. [Online]. Available: http://www.icml2010.org/papers/422.pdf

[15] 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.

[16] 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.

[17] I. M. Mitchell and S. Sastry, "Continuous path planning with multiple constraints," in *Conference on Decision and Control (CDC)*, vol. 5. IEEE, 2003, pp. 5502–5507.

[18] M. G. Crandall and P.-L. Lions, "Viscosity solutions of hamilton-jacobi equations," *Transactions of the American Mathematical Society*, vol. 277, no. 1, pp. 1–42, 1983.

[19] D. P. Bertsekas, *Dynamic programming and optimal control*. Athena scientific Belmont, MA, 1995, vol. 1, no. 2.

[20] J. A. Sethian, "Fast marching methods," *SIAM review*, vol. 41, no. 2, pp. 199–235, 1999.

[21] J. Umlauft, L. Pöhler, and S. Hirche, "An uncertainty-based control lyapunov approach for control-affine systems modeled by Gaussian process," *IEEE Control Systems Letters*, vol. 2/3, pp. 483–488, 2018.