

Towards Robust Data-Driven Control Synthesis for Nonlinear Systems with Actuation Uncertainty

Andrew J. Taylor¹, Victor D. Dorobantu¹, Sarah Dean¹, Benjamin Recht, Yisong Yue, and Aaron D. Ames

Abstract—Modern nonlinear control theory seeks to endow systems with properties such as stability and safety, and has been deployed successfully across various domains. Despite this success, model uncertainty remains a significant challenge in ensuring that model-based controllers transfer to real world systems. This paper develops a data-driven approach to robust control synthesis in the presence of model uncertainty using Control Certificate Functions (CCFs), resulting in a convex optimization based controller for achieving properties like stability and safety. An important benefit of our framework is nuanced data-dependent guarantees, which in principle can yield sample-efficient data collection approaches that need not fully determine the input-to-state relationship. This work serves as a starting point for addressing important questions at the intersection of nonlinear control theory and non-parametric learning, both theoretical and in application. We demonstrate the efficiency of the proposed method with respect to input data in simulation with an inverted pendulum in multiple experimental settings.

I. INTRODUCTION

Ensuring properties such as stability and safety is of significant importance in many modern control applications, including autonomous driving, industrial robotics, and aerospace vehicles. In practice, the models used to design these controllers are imperfect, with model uncertainty arising due to unmodeled dynamics and parametric errors. In the presence of such uncertainty, controllers may fail to render systems stable or safe. As real world control systems become increasingly complex, the potential for detrimental modeling errors increases, and thus it is critical to study control synthesis in the presence of uncertainty.

In this work, we propose a control synthesis process using *control certificate functions* (CCFs) [1], [2] that incorporates a data-driven approach for capturing model uncertainty. CCFs generalize popular tools from nonlinear control for achieving stability and safety such as Control Lyapunov Functions (CLFs) [3], Control Barrier Functions (CBFs) [4], and Control Barrier-Lyapunov Functions [5]. CLFs and CBFs have been successfully deployed in the context of bipedal robotics [6], [7], adaptive cruise control [4], robotic manipulators [8], and multi-agent systems [9]. Data-driven and machine learning based approaches have shown great promise for controlling systems with an uncertain model or with no model at all [10], [11], [12], [13]. The integration

¹Authors contributed equally. A.J. Taylor, V.D. Dorobantu, Y. Yue, and A.D. Ames are with the Department of Computing and Mathematical Sciences, California Institute of Technology, Pasadena, CA 91125, USA, {ajtaylor, vdoroban, yyue, ames}@caltech.edu. S. Dean and B. Recht are with the Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, Berkeley, CA 94720, USA, {sarah, brecht}@berkeley.edu.

of techniques from nonlinear control theory for achieving stability and safety with data-driven methods has become increasingly popular [14], [15], [16], [17], [18], with many approaches relying on certificate functions for theoretical guarantees [8], [19], [20], [21], [22].

Uncertainty in the effect of actuation remains a major challenge in achieving control-theoretic guarantees with data-driven methods. Many existing approaches assume certainty in how actuation enters the dynamics [23], [24], [25], use structured controllers requiring strong characterizations of this uncertainty [15], [26], or require high coverage datasets with (nearly) complete characterizations of the input-to-state relationship [27], [28]. In practice, collecting such data can be prohibitively costly or damage the system, suggesting a need for data-driven approaches that accommodate actuation uncertainty without requiring this complete characterization.

Our contribution: The contribution of this work is a novel approach for robust data-driven control synthesis via CCFs for control-affine systems with model uncertainty, including actuation uncertainty, which is broadly applicable in many real-world settings such as robotic [29] and automotive systems [30]. In Section III we incorporate data into a convex optimization based control synthesis problem as affine inequality constraints which restrict possible model uncertainties. This enables the choice of robust control inputs over convex uncertainty sets. Rather than requiring a full characterization of how input enters the system, this approach utilizes the affine structure of CCF dynamics to choose inputs. This reduces the impact of actuation uncertainty on the evolution of the certificate function and allows for guarantees of stability and safety. In summary, our results show that good performance can be achieved when training data provides sufficiently dense coverage of the state space and targeted excitation in input directions. The proposed approach provides a unique perspective for unifying nonlinear control and non-parametric machine learning that is well positioned to study both theoretical and application oriented questions at this intersection.

II. BACKGROUND

This section provides a review of certificate-based nonlinear control synthesis and an overview of how model uncertainty impacts these synthesis methods.

A. Control Certificate Functions

Consider a nonlinear control affine system given by:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{u}, \quad (1)$$

where $\mathbf{x} \in \mathbb{R}^n$, $\mathbf{u} \in \mathbb{R}^m$, and $\mathbf{f} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $\mathbf{g} : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times m}$ are locally Lipschitz continuous on \mathbb{R}^n . Further assume that the origin is an equilibrium point of the uncontrolled system ($\mathbf{f}(\mathbf{0}) = \mathbf{0}$). In this work we assume that \mathbf{u} may be chosen unbounded as in [8], [31]. Given a locally Lipschitz continuous state-feedback controller $\mathbf{k} : \mathbb{R}^n \rightarrow \mathbb{R}^m$, the closed-loop system dynamics are:

$$\dot{\mathbf{x}} = \mathbf{f}_{\text{cl}}(\mathbf{x}) \triangleq \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{k}(\mathbf{x}). \quad (2)$$

The assumption on local Lipschitz continuity of \mathbf{f} , \mathbf{g} , and \mathbf{k} implies that \mathbf{f}_{cl} is locally Lipschitz continuous. Thus for any initial condition $\mathbf{x}_0 \triangleq \mathbf{x}(0) \in \mathbb{R}^n$ there exists a maximum time interval $I(\mathbf{x}_0) = [0, t_{\text{max}})$ such that $\mathbf{x}(t)$ is the unique solution to (2) on $I(\mathbf{x}_0)$ [32].

The qualitative behavior (such as stability or safety) of the the closed-loop system (2) can be certified via the notion of a continuously differentiable *certificate function* $\mathcal{C} : \mathbb{R}^n \rightarrow \mathbb{R}$. Given a *comparison function* $\alpha : \mathbb{R} \rightarrow \mathbb{R}$ (specific to the qualitative behavior of interest), certification is specified as an inequality on the derivative of the certificate function along solutions to the closed-loop system:

$$\dot{\mathcal{C}}(\mathbf{x}) = \nabla \mathcal{C}(\mathbf{x})^\top \mathbf{f}_{\text{cl}}(\mathbf{x}) \leq -\alpha(\mathcal{C}(\mathbf{x})). \quad (3)$$

Synthesis of controllers that satisfy (3) by design motivates the following definition of a *Control Certificate Function*:

Definition 1 (Control Certificate Function (CCF)). A continuously differentiable function $\mathcal{C} : \mathbb{R}^n \rightarrow \mathbb{R}$ is a *Control Certificate Function (CCF)* for (1) with comparison function $\alpha : \mathbb{R} \rightarrow \mathbb{R}$ if for all $\mathbf{x} \in \mathbb{R}^n$:

$$\inf_{\mathbf{u} \in \mathbb{R}^m} \underbrace{\nabla \mathcal{C}(\mathbf{x})^\top \mathbf{f}(\mathbf{x})}_{L_{\mathbf{f}}\mathcal{C}(\mathbf{x})} + \underbrace{\nabla \mathcal{C}(\mathbf{x})^\top \mathbf{g}(\mathbf{x})}_{L_{\mathbf{g}}\mathcal{C}(\mathbf{x})} \mathbf{u} \leq -\alpha(\mathcal{C}(\mathbf{x})). \quad (4)$$

where $L_{\mathbf{f}}\mathcal{C} : \mathbb{R}^n \rightarrow \mathbb{R}$ and $L_{\mathbf{g}}\mathcal{C} : \mathbb{R}^n \rightarrow \mathbb{R}^m$. The control-affine nature of the system dynamics are preserved by the CCF, such that the only component of the input that impacts the evolution of the certificate function lies in the direction of $L_{\mathbf{g}}\mathcal{C}(\mathbf{x})$. Given a CCF \mathcal{C} for (1) and a corresponding comparison function α , we define the point-wise set of all control values that satisfy the inequality in (4):

$$K_{\text{ccf}}(\mathbf{x}) \triangleq \{\mathbf{u} \in \mathbb{R}^m \mid L_{\mathbf{f}}\mathcal{C}(\mathbf{x}) + L_{\mathbf{g}}\mathcal{C}(\mathbf{x})\mathbf{u} \leq -\alpha(\mathcal{C}(\mathbf{x}))\}. \quad (5)$$

Any nominal locally Lipschitz continuous controller $\mathbf{k}_d : \mathbb{R}^n \rightarrow \mathbb{R}^m$ can be modified to take values in the set $K_{\text{ccf}}(\mathbf{x})$ via the certificate-critical CCF-QP:

$$\begin{aligned} \mathbf{k}(\mathbf{x}) &= \underset{\mathbf{u} \in \mathbb{R}^m}{\text{argmin}} \frac{1}{2} \|\mathbf{u} - \mathbf{k}_d(\mathbf{x})\|_2^2 && \text{(CCF-QP)} \\ \text{s.t.} \quad & \nabla \mathcal{C}(\mathbf{x})^\top (\mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{u}) \leq -\alpha(\mathcal{C}(\mathbf{x})). \end{aligned}$$

Before providing particular examples of certificate functions useful in control synthesis, we review the following definitions. We denote a continuous function $\alpha : [0, a) \rightarrow \mathbb{R}_+$, with $a > 0$, as *class* \mathcal{K} ($\alpha \in \mathcal{K}$) if $\alpha(0) = 0$ and α is strictly monotonically increasing. If $a = \infty$ and $\lim_{r \rightarrow \infty} \alpha(r) = \infty$, then α is *class* \mathcal{K}_∞ ($\alpha \in \mathcal{K}_\infty$). A continuous function $\alpha : (-b, a) \rightarrow \mathbb{R}$, with $a, b > 0$, is

extended class \mathcal{K} ($\alpha \in \mathcal{K}_e$) if $\alpha(0) = 0$ and α is strictly monotonically increasing. If $a, b = \infty$, $\lim_{r \rightarrow \infty} \alpha(r) = \infty$, and $\lim_{r \rightarrow -\infty} \alpha(r) = -\infty$, then α is *extended class* \mathcal{K}_∞ ($\alpha \in \mathcal{K}_{\infty,e}$). Finally, we note that $c \in \mathbb{R}$ is referred to as a *regular value* of a continuously differentiable function $h : \mathbb{R}^n \rightarrow \mathbb{R}$ if $h(\mathbf{x}) = c \implies \nabla h(\mathbf{x}) \neq \mathbf{0}_n$.

Example 1 (Stability via Control Lyapunov Functions). In the context of stabilization to the origin, a control certificate function $V : \mathbb{R}^n \rightarrow \mathbb{R}$ with a class \mathcal{K} comparison function $\alpha \in \mathcal{K}$ that satisfies:

$$\alpha_1(\|\mathbf{x}\|) \leq V(\mathbf{x}) \leq \alpha_2(\|\mathbf{x}\|), \quad (6)$$

for $\alpha_1, \alpha_2 \in \mathcal{K}$, is a *Control Lyapunov Function (CLF)* [3], [33], with stabilization to the origin achieved by controllers taking values in the point-wise set K_{ccf} given by (5) [34].

Example 2 (Safety via Control Barrier Functions). In the context of safety, defined as forward invariance [35] of a set \mathcal{S} , a control certificate function $h : \mathbb{R}^n \rightarrow \mathbb{R}$ with 0 a regular value and a comparison function $\alpha \in \mathcal{K}_{\infty,e}$ that satisfies:

$$\mathbf{x} \in \mathcal{S} \implies h(\mathbf{x}) \leq 0, \quad (7)$$

is a *Control Barrier Function (CBF)* [4], [34], with safety of the set \mathcal{S} achieved by controllers taking values in the point-wise set K_{ccf} given by (5) [36]. We adopt the opposite sign convention for h so satisfying (4) guarantees safety.

B. Model Uncertainty

In practice, uncertainty in the system dynamics (1) exists due to parametric error and unmodeled dynamics, such that the functions \mathbf{f} and \mathbf{g} are not precisely known. Control affine systems are a natural setting to study actuation uncertainty as the function \mathbf{g} can be seen as an uncertain gain multiplying the input. In this context, control synthesis is done with a nominal model that estimates the true system dynamics:

$$\hat{\mathbf{x}} = \hat{\mathbf{f}}(\mathbf{x}) + \hat{\mathbf{g}}(\mathbf{x})\mathbf{u}, \quad (8)$$

where $\hat{\mathbf{f}} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $\hat{\mathbf{g}} : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times m}$ are locally Lipschitz continuous. Adding and subtracting this expression to and from (1) implies the system evolution is described by:

$$\dot{\mathbf{x}} = \hat{\mathbf{f}}(\mathbf{x}) + \hat{\mathbf{g}}(\mathbf{x})\mathbf{u} + \underbrace{\mathbf{f}(\mathbf{x}) - \hat{\mathbf{f}}(\mathbf{x})}_{\tilde{\mathbf{f}}(\mathbf{x})} + \underbrace{(\mathbf{g}(\mathbf{x}) - \hat{\mathbf{g}}(\mathbf{x}))}_{\tilde{\mathbf{g}}(\mathbf{x})} \mathbf{u}, \quad (9)$$

where $\tilde{\mathbf{f}} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $\tilde{\mathbf{g}} : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times m}$ are the unmodeled dynamics. This uncertainty in the dynamics additionally manifests in the time derivative of a CCF for the system:

$$\begin{aligned} \dot{\mathcal{C}}(\mathbf{x}, \mathbf{u}) &= \underbrace{\nabla \mathcal{C}(\mathbf{x})^\top \hat{\mathbf{f}}(\mathbf{x})}_{L_{\hat{\mathbf{f}}}\mathcal{C}(\mathbf{x})} + \underbrace{\nabla \mathcal{C}(\mathbf{x})^\top \hat{\mathbf{g}}(\mathbf{x})}_{L_{\hat{\mathbf{g}}}\mathcal{C}(\mathbf{x})} \mathbf{u} \\ &\quad + \underbrace{\nabla \mathcal{C}(\mathbf{x})^\top \tilde{\mathbf{f}}(\mathbf{x})}_{L_{\tilde{\mathbf{f}}}\mathcal{C}(\mathbf{x})} + \underbrace{\nabla \mathcal{C}(\mathbf{x})^\top \tilde{\mathbf{g}}(\mathbf{x})}_{L_{\tilde{\mathbf{g}}}\mathcal{C}(\mathbf{x})} \mathbf{u}, \end{aligned} \quad (10)$$

where $L_{\hat{\mathbf{f}}}\mathcal{C}, L_{\tilde{\mathbf{f}}}\mathcal{C} : \mathbb{R}^n \rightarrow \mathbb{R}$, and $L_{\hat{\mathbf{g}}}\mathcal{C}, L_{\tilde{\mathbf{g}}}\mathcal{C} : \mathbb{R}^n \rightarrow \mathbb{R}^m$. The presence of uncertainty in the CCF time derivative makes it

impossible to verify whether a given control input is in the set $K_{\text{ccf}}(\mathbf{x})$ given in (5), and can lead to failure to achieve the desired qualitative behavior.

Assumption 1. The function $\mathcal{C} : \mathbb{R}^n \rightarrow \mathbb{R}$ is a valid CCF with comparison function $\alpha : \mathbb{R} \rightarrow \mathbb{R}$ for the true dynamic system (9). Mathematically this assumption appears as:

$$\inf_{\mathbf{u} \in \mathbb{R}^m} \nabla \mathcal{C}(\mathbf{x})^\top (\mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{u}) \leq -\alpha(\mathcal{C}(\mathbf{x})).$$

This assumption is structural in nature and can be met for feedback linearizable systems (such as robotic systems).

Remark 1. Though many approaches to CCF design (linearization, energy-based, numerical sums-of-squares methods) rely on knowledge of the true dynamics, it is also possible to design CCFs without explicit knowledge of the true system dynamics. For example, a valid CLF and comparison function for the true system can be designed via feedback linearization assuming only knowledge of the degree of actuation (see [37] for full details). This method also works to specify CBFs which are defined by sublevel sets of CLFs (as in our simulation results). We emphasize the difference between choosing a qualitative behavior that the system can be made to satisfy (e.g. the CCF) and actually designing the control inputs which achieve the behavior. Our work focuses on the latter: solving the problem of choosing stable/safe inputs in the presence of uncertainty.

Assumption 2. The functions $\tilde{\mathbf{f}}$ and $\tilde{\mathbf{g}}$ are globally Lipschitz continuous with known Lipschitz constants $\mathcal{L}_{\tilde{\mathbf{f}}}$ and $\mathcal{L}_{\tilde{\mathbf{g}}}$.

Remark 2. Knowledge of minimal Lipschitz constants is not necessary, but smaller Lipschitz constants are associated with improved performance of data-driven robust control methods.

III. DATA-DRIVEN ROBUST CONTROL SYNTHESIS

In this section we explore how data can be incorporated directly into an optimization based controller to robustly achieve a desired qualitative behavior specified via a CCF.

Consider a dataset consisting of N tuples of states, inputs, and corresponding state time derivatives, $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{u}_i, \dot{\mathbf{x}}_i)\}_{i=1}^N$, with $\mathbf{x}_i \in \mathbb{R}^n$, $\mathbf{u}_i \in \mathbb{R}^m$, and $\dot{\mathbf{x}}_i \in \mathbb{R}^n$ for $i = 1, \dots, N$. It may not be possible to directly measure the state time derivatives $\dot{\mathbf{x}}_i$, but they can be approximated from sequential state observations \mathbf{x}_i [37], [38]. While we do not consider noise, our construction can be modified to account for bounded noise via a constant offset to (13).

We now show how data allows us to reduce uncertainty by constraining the possible values of the functions $\tilde{\mathbf{f}}$ and $\tilde{\mathbf{g}}$ directly, without a parametric estimator. Considering the uncertain model (9) evaluated at a state and input pair $(\mathbf{x}_i, \mathbf{u}_i)$ in the dataset yields:

$$\tilde{\mathbf{F}}_i \triangleq \dot{\mathbf{x}}_i - (\tilde{\mathbf{f}}(\mathbf{x}_i) + \tilde{\mathbf{g}}(\mathbf{x}_i)\mathbf{u}_i) = \tilde{\mathbf{f}}(\mathbf{x}_i) + \tilde{\mathbf{g}}(\mathbf{x}_i)\mathbf{u}_i, \quad (11)$$

where $\tilde{\mathbf{F}}_i \in \mathbb{R}^n$ can be interpreted as the error between the true state time derivative and the nominal model (8) evaluated at the state and input pair $(\mathbf{x}_i, \mathbf{u}_i)$.

Considering a state $\mathbf{x} \in \mathbb{R}^n$ (not necessarily present in the dataset \mathcal{D}), the second equality in (11) implies:

$$\tilde{\mathbf{f}}(\mathbf{x}) + \tilde{\mathbf{g}}(\mathbf{x})\mathbf{u}_i - \tilde{\mathbf{F}}_i = \tilde{\mathbf{f}}(\mathbf{x}) - \tilde{\mathbf{f}}(\mathbf{x}_i) + (\tilde{\mathbf{g}}(\mathbf{x}) - \tilde{\mathbf{g}}(\mathbf{x}_i))\mathbf{u}_i. \quad (12)$$

This expression provides a relationship between the possible values of the unmodeled dynamics $\tilde{\mathbf{f}}$ and $\tilde{\mathbf{g}}$ at the state \mathbf{x} and the values of the unmodeled dynamics at the data point \mathbf{x}_i . Using the local Lipschitz continuity of the unmodeled dynamics yields the following bound:

$$\begin{aligned} \left\| \tilde{\mathbf{f}}(\mathbf{x}) + \tilde{\mathbf{g}}(\mathbf{x})\mathbf{u}_i - \tilde{\mathbf{F}}_i \right\|_2 &= \left\| \tilde{\mathbf{f}}(\mathbf{x}) - \tilde{\mathbf{f}}(\mathbf{x}_i) + (\tilde{\mathbf{g}}(\mathbf{x}) - \tilde{\mathbf{g}}(\mathbf{x}_i))\mathbf{u}_i \right\|_2, \\ &\leq (\mathcal{L}_{\tilde{\mathbf{f}}} + \mathcal{L}_{\tilde{\mathbf{g}}}\|\mathbf{u}_i\|_2) \|\mathbf{x} - \mathbf{x}_i\|_2 \triangleq \epsilon_i(\mathbf{x}). \end{aligned} \quad (13)$$

where $\epsilon_i : \mathbb{R}^n \rightarrow \mathbb{R}_+$ for $i = 1, \dots, N$. We see that the bound grows with the magnitude of the Lipschitz constants $\mathcal{L}_{\tilde{\mathbf{f}}}$ and $\mathcal{L}_{\tilde{\mathbf{g}}}$ and distance of the state \mathbf{x} from the data point \mathbf{x}_i . The values of $\mathcal{L}_{\tilde{\mathbf{f}}}$ and $\mathcal{L}_{\tilde{\mathbf{g}}}$ are not explicitly data dependent, and thus the bound can be improved for a given dataset by reducing the possible model uncertainty through improved modeling. Given this construction we may define the point-wise uncertainty set:

$$\mathcal{U}_i(\mathbf{x}) \triangleq \left\{ (\mathbf{A}, \mathbf{b}) \in \mathbb{R}^{n \times m} \times \mathbb{R}^n \mid \left\| \mathbf{b} + \mathbf{A}\mathbf{u}_i - \tilde{\mathbf{F}}_i \right\|_2 \leq \epsilon_i(\mathbf{x}) \right\} \subset \mathbb{R}^{n \times m} \times \mathbb{R}^n, \quad (14)$$

noting that $(\tilde{\mathbf{g}}(\mathbf{x}), \tilde{\mathbf{f}}(\mathbf{x})) \in \mathcal{U}_i(\mathbf{x})$ and $\mathcal{U}_i(\mathbf{x})$ is closed and convex. Considering this construction over the entire dataset \mathcal{D} yields the following point-wise uncertainty set:

$$\mathcal{U}(\mathbf{x}) \triangleq \bigcap_{i=1}^N \mathcal{U}_i(\mathbf{x}) \subset \mathbb{R}^{n \times m} \times \mathbb{R}^n, \quad (15)$$

noting that $(\tilde{\mathbf{g}}(\mathbf{x}), \tilde{\mathbf{f}}(\mathbf{x})) \in \mathcal{U}(\mathbf{x})$ and $\mathcal{U}(\mathbf{x})$ is closed and convex. Therefore, $\mathcal{U}(\mathbf{x})$ consists of all possible model errors that are consistent with the observed data. This allows us to pose the following data robust control problem:

Definition 2 (Data Robust Control Certificate Function Optimization Problem).

$$\begin{aligned} \mathbf{k}_{\text{rob}}(\mathbf{x}) &= \underset{\mathbf{u} \in \mathbb{R}^m}{\text{argmin}} \frac{1}{2} \|\mathbf{u} - \mathbf{k}_d(\mathbf{x})\|_2^2 && \text{(DR-CCF-OP)} \\ \text{s.t. } &\hat{\mathcal{C}}(\mathbf{x}, \mathbf{u}) + \nabla \mathcal{C}(\mathbf{x})^\top (\mathbf{b} + \mathbf{A}\mathbf{u}) \leq -\alpha(\mathcal{C}(\mathbf{x})) \\ &\text{for all } (\mathbf{A}, \mathbf{b}) \in \mathcal{U}(\mathbf{x}). \end{aligned}$$

By construction we have that $(\tilde{\mathbf{g}}(\mathbf{x}), \tilde{\mathbf{f}}(\mathbf{x})) \in \mathcal{U}(\mathbf{x})$, implying that $\mathbf{k}_{\text{rob}}(\mathbf{x}) \in K_{\text{ccf}}(\mathbf{x})$ when the problem is feasible. Thus the closed-loop system (2) under \mathbf{k}_{rob} satisfies inequality (3). We next present one of our main results, using robust optimization [39] to yield a convex problem for synthesizing such a robust controller.

Theorem 1 (Robust Control Synthesis). *Let $\mathcal{C} : \mathbb{R}^n \rightarrow \mathbb{R}$ be a control certificate with comparison function $\alpha : \mathbb{R} \rightarrow \mathbb{R}$. The robust controller (DR-CCF-OP) is equivalently expressed as:*

$$\begin{aligned} \mathbf{k}_{\text{rob}}(\mathbf{x}) &= \underset{\substack{\mathbf{u} \in \mathbb{R}^m \\ \boldsymbol{\lambda}_i \in \mathbb{R}^n}}{\text{argmin}} \frac{1}{2} \|\mathbf{u} - \mathbf{k}_d(\mathbf{x})\|_2^2 && \text{(DR-CCF-SOCP)} \\ \text{s.t. } \widehat{\mathcal{C}}(\mathbf{x}, \mathbf{u}) - \sum_{i=1}^N \left(\boldsymbol{\lambda}_i^\top \widetilde{\mathbf{F}}_i - \|\boldsymbol{\lambda}_i\|_2 \epsilon_i(\mathbf{x}) \right) &\leq -\alpha(\mathcal{C}(\mathbf{x})), \\ \sum_{i=1}^N \boldsymbol{\lambda}_i \mathbf{u}_i^\top &= -\nabla \mathcal{C}(\mathbf{x}) \mathbf{u}^\top, \quad \sum_{i=1}^N \boldsymbol{\lambda}_i = -\nabla \mathcal{C}(\mathbf{x}). \end{aligned}$$

Proof. An input $\mathbf{u} \in \mathbb{R}^m$ is feasible if the optimal value of the optimization problem:

$$\begin{aligned} \sup_{\substack{\mathbf{A} \in \mathbb{R}^{n \times m} \\ \mathbf{b} \in \mathbb{R}^n}} \widehat{\mathcal{C}}(\mathbf{x}, \mathbf{u}) + \nabla \mathcal{C}(\mathbf{x})^\top (\mathbf{b} + \mathbf{A}\mathbf{u}) \\ \text{s.t. } \|\mathbf{b} + \mathbf{A}\mathbf{u}_i - \widetilde{\mathbf{F}}_i\|_2 \leq \epsilon_i(\mathbf{x}) \text{ for all } i \in 1, \dots, N, \end{aligned}$$

is less than or equal to $-\alpha(\mathcal{C}(\mathbf{x}))$. Each inequality constraint can be rewritten as set membership in a second-order cone. That is, $(\mathbf{b} + \mathbf{A}\mathbf{u}_i - \widetilde{\mathbf{F}}_i, \epsilon_i(\mathbf{x})) \in \mathcal{Q}_n$ for each $i \in \{1, \dots, N\}$, with $\mathcal{Q}_n \subset \mathbb{R}^{n+1}$ denoting the Lorentz cone $\mathcal{Q}_n = \{(\mathbf{y}, t) \in \mathbb{R}^n \times \mathbb{R} : \|\mathbf{y}\|_2 \leq t\}$. This yields the following dual problem:

$$\begin{aligned} \inf_{\substack{\boldsymbol{\lambda}_i \in \mathbb{R}^n \\ \nu_i \in \mathbb{R}}} \widehat{\mathcal{C}}(\mathbf{x}, \mathbf{u}) - \sum_{i=1}^N \left(\boldsymbol{\lambda}_i^\top \widetilde{\mathbf{F}}_i - \nu_i \epsilon_i(\mathbf{x}) \right) \\ \text{s.t. } \sum_{i=1}^N \boldsymbol{\lambda}_i \mathbf{u}_i^\top = -\nabla \mathcal{C}(\mathbf{x}) \mathbf{u}^\top, \quad \sum_{i=1}^N \boldsymbol{\lambda}_i = -\nabla \mathcal{C}(\mathbf{x}) \\ \|\boldsymbol{\lambda}_i\|_2 \leq \nu_i \text{ for all } i \in 1, \dots, N, \end{aligned}$$

where $(\boldsymbol{\lambda}_i, \nu_i) \in \mathcal{Q}_n$ are Lagrange multipliers corresponding to the constraints imposed by the data point $(\mathbf{x}_i, \mathbf{u}_i)$. For any $(\mathbf{x}_i, \mathbf{u}_i)$, choosing Lagrange multipliers such that $\nu_i > \|\boldsymbol{\lambda}_i\|_2$ increases the value of the dual problem compared to choosing $\nu_i = \|\boldsymbol{\lambda}_i\|_2$. We therefore simplify the problem as:

$$\begin{aligned} \inf_{\boldsymbol{\lambda}_i \in \mathbb{R}^n} \widehat{\mathcal{C}}(\mathbf{x}, \mathbf{u}) - \sum_{i=1}^N \left(\boldsymbol{\lambda}_i^\top \widetilde{\mathbf{F}}_i - \|\boldsymbol{\lambda}_i\|_2 \epsilon_i(\mathbf{x}) \right) \\ \text{s.t. } \sum_{i=1}^N \boldsymbol{\lambda}_i \mathbf{u}_i^\top = -\nabla \mathcal{C}(\mathbf{x}) \mathbf{u}^\top, \quad \sum_{i=1}^N \boldsymbol{\lambda}_i = -\nabla \mathcal{C}(\mathbf{x}). \end{aligned}$$

This optimization problem can then replace the original robust constraint in (DR-CCF-OP) to yield the final optimization problem (DR-CCF-SOCP). \square

Remark 3. Solving this optimization problem in a computationally efficient manner may require data segmentation for higher-dimensional problems [40], but a detailed consideration is outside the scope of this work, which is focused on theoretical foundations.

IV. FEASIBILITY

In this section we provide an analysis of the feasibility of the controller proposed in Section III. The feasibility of the controller (DR-CCF-SOCP) at a given state \mathbf{x} is determined by the structure of the uncertainty set $\mathcal{U}(\mathbf{x})$ defined in (15). The following lemma provides a condition on the inputs in the dataset that implies $\mathcal{U}(\mathbf{x})$ is bounded for all $\mathbf{x} \in \mathbb{R}^n$:

Lemma 1 (Bounded Uncertainty Sets). *Consider a dataset \mathcal{D} with N data points satisfying $N \geq m + 1$. If there exists a set of data points $\{(\mathbf{x}_i, \mathbf{u}_i, \dot{\mathbf{x}}_i)\}_{i=1}^{m+1} \subseteq \mathcal{D}$ such that the set of vectors:*

$$\mathcal{M} \triangleq \left\{ \begin{bmatrix} \mathbf{u}_i^\top & 1 \end{bmatrix}^\top \right\}_{i=1}^{m+1}, \quad (16)$$

are linearly independent, then the uncertainty set $\mathcal{U}(\mathbf{x})$ is bounded (and thus compact) for any $\mathbf{x} \in \mathbb{R}^n$.

Proof. Consider the set \mathcal{M} , and define:

$$\mathcal{M}_{\mathbf{u}} \triangleq \left\{ \mathbf{u}_i \in \mathbb{R}^m \mid \begin{bmatrix} \mathbf{u}_i \\ 1 \end{bmatrix} \in \mathcal{M} \right\}.$$

For arbitrary $\mathbf{x} \in \mathbb{R}^n$, let $(\mathbf{A}, \mathbf{b}) \in \mathcal{U}(\mathbf{x})$. By definition of the individual uncertainty sets $\mathcal{U}_i(\mathbf{x})$ and $\mathcal{U}(\mathbf{x})$ given in (14) and (15), respectively, we have that:

$$\|\mathbf{A}\mathbf{u}_i + \mathbf{b} - \widetilde{\mathbf{F}}_i\|_2 \leq \epsilon_i(\mathbf{x}), \quad (17)$$

for $\mathbf{u}_i \in \mathcal{M}_{\mathbf{u}}$. Defining $\mathbf{v}_i \triangleq \mathbf{A}\mathbf{u}_i + \mathbf{b} - \widetilde{\mathbf{F}}_i$, we have that:

$$\sum_{i=1}^{m+1} \|\mathbf{A}\mathbf{u}_i + \mathbf{b} - \widetilde{\mathbf{F}}_i\|_2^2 = \sum_{i=1}^{m+1} \sum_{j=1}^n |V_{ji}|^2 = \|\mathbf{V}\|_F^2,$$

where $\|\cdot\|_F$ is the Frobenius norm and $\mathbf{V} = [\mathbf{v}_1 \ \dots \ \mathbf{v}_{m+1}]$. Noting that $\|\mathbf{V}\|_F = \|\mathbf{V}^\top\|_F$, factoring and using the fact $\|\mathbf{P}\|_2 \leq \|\mathbf{P}\|_F$ for any $\mathbf{P} \in \mathbb{R}^{(m+1) \times n}$ in conjunction with (17) yields:

$$\left\| \begin{bmatrix} \mathbf{u}_1^\top & 1 \\ \vdots & \vdots \\ \mathbf{u}_{m+1}^\top & 1 \end{bmatrix} \begin{bmatrix} \mathbf{A}^\top \\ \mathbf{b}^\top \end{bmatrix} - \begin{bmatrix} \widetilde{\mathbf{F}}_1^\top \\ \vdots \\ \widetilde{\mathbf{F}}_{m+1}^\top \end{bmatrix} \right\|_2 \leq \sum_{i=1}^{m+1} \epsilon_i(\mathbf{x})^2.$$

Taking the square root of both sides and employing the reverse triangle inequality we arrive at:

$$\left\| \underbrace{\begin{bmatrix} \mathbf{u}_1^\top & 1 \\ \vdots & \vdots \\ \mathbf{u}_{m+1}^\top & 1 \end{bmatrix}}_{\mathbf{U}} \begin{bmatrix} \mathbf{A}^\top \\ \mathbf{b}^\top \end{bmatrix} \right\|_2 \leq \sqrt{\sum_{i=1}^{m+1} \epsilon_i(\mathbf{x})^2} + \left\| \begin{bmatrix} \widetilde{\mathbf{F}}_1^\top \\ \vdots \\ \widetilde{\mathbf{F}}_{m+1}^\top \end{bmatrix} \right\|_2.$$

Noting that \mathbf{U} has full rank by the linear independence of the vectors in the set \mathcal{M} , we have that:

$$\sigma_{m+1}(\mathbf{U}) \left\| \begin{bmatrix} \mathbf{A}^\top \\ \mathbf{b}^\top \end{bmatrix} \right\|_2 \leq \left\| \begin{bmatrix} \mathbf{u}_1^\top & 1 \\ \vdots & \vdots \\ \mathbf{u}_{m+1}^\top & 1 \end{bmatrix} \begin{bmatrix} \mathbf{A}^\top \\ \mathbf{b}^\top \end{bmatrix} \right\|_2,$$

where $\sigma_{m+1}(\mathbf{U}) > 0$ is the smallest singular value of \mathbf{U} . Combining the two previous inequalities shows that the uncertainty set $\mathcal{U}(\mathbf{x})$ is bounded. \square

This result shows that variety in input directions is sufficient to assert boundedness of the uncertainty set as a uniform property over the entire state space. As seen in the proof, the bound on this set may be very large if the values of $\epsilon_i(\mathbf{x})$ are large (as is the case when \mathbf{x} is far away from the data points \mathbf{x}_i associated with the inputs used to construct the bound). Alternatively, the uncertainty set will be small for a point \mathbf{x} if there is local variety in input directions within the training dataset. To understand how the size of $\mathcal{U}(\mathbf{x})$ impacts feasibility of the optimization problem, we define the following set:

$$\tilde{\mathcal{U}}_C(\mathbf{x}) \triangleq \left\{ (\mathbf{a}, b) \in \mathbb{R}^m \times \mathbb{R} \mid \exists (\mathbf{A}, \mathbf{b}) \in \mathcal{U}(\mathbf{x}) \text{ s.t.} \right. \\ \left. \mathbf{a} = (\nabla C(\mathbf{x})^\top \mathbf{A})^\top, b = \nabla C(\mathbf{x})^\top \mathbf{b} \right\}, \quad (18)$$

The set $\tilde{\mathcal{U}}_C(\mathbf{x})$ can be interpreted as the projection of the dynamics uncertainty set $\mathcal{U}(\mathbf{x})$ along the gradient of the control certificate function C , creating an $m+1$ dimensional set which represents the possible uncertainties in the Lie derivatives of C . Additionally define the set:

$$\mathcal{U}_C(\mathbf{x}) \triangleq \left\{ (L_{\tilde{\mathbf{g}}}C(\mathbf{x})^\top, L_{\tilde{\mathbf{f}}}C(\mathbf{x})) \right\} \oplus \tilde{\mathcal{U}}_C(\mathbf{x}), \quad (19)$$

where \oplus denotes a Minkowski sum. The set $\mathcal{U}_C(\mathbf{x})$ is the recentering of $\tilde{\mathcal{U}}_C(\mathbf{x})$ set around the estimate of the Lie derivatives $(L_{\tilde{\mathbf{g}}}C(\mathbf{x})^\top, L_{\tilde{\mathbf{f}}}C(\mathbf{x})) \in \mathbb{R}^{m+1}$, such that it captures the possible true Lie derivatives of C . As multiplication by $\nabla C(\mathbf{x})^\top$ is a linear transformation, $\tilde{\mathcal{U}}_C(\mathbf{x})$ and $\mathcal{U}_C(\mathbf{x})$ are convex, and if $\mathcal{U}(\mathbf{x})$ is bounded (and therefore compact), then $\tilde{\mathcal{U}}_C(\mathbf{x})$ and $\mathcal{U}_C(\mathbf{x})$ are compact.

We present our second main result in the form of a necessary and sufficient condition for feasibility of (DR-CCF-OP):

Theorem 2 (Feasibility of Data-Driven Robust Controller). *For a state $\mathbf{x} \in \mathbb{R}^n$, let the sets $\mathcal{U}(\mathbf{x})$ and $\mathcal{U}_C(\mathbf{x})$ be defined as in (15) and (19), respectively. Define the ray $\mathcal{R} \subset \mathbb{R}^{m+1}$ as $\mathcal{R} = \{\mathbf{0}_m\} \times (-\alpha(C(\mathbf{x})), \infty)$. Assuming that $\mathcal{U}(\mathbf{x})$ is bounded, the data-driven robust controller (DR-CCF-SOCP) is feasible if and only if:*

$$\mathcal{U}_C(\mathbf{x}) \cap \mathcal{R} = \emptyset. \quad (20)$$

Intuitively, the ray \mathcal{R} represents Lie derivative pairs that do not satisfy the certificate function condition with no actuation and cannot be modified through actuation. By Assumption 1, the Lie derivatives of the true system are not contained in \mathcal{R} , but the possible uncertainties permitted by data need not necessarily reflect this. If one possible uncertainty pair in the set $\mathcal{U}_C(\mathbf{x})$ is contained in \mathcal{R} , it is impossible to meet the certificate condition for that uncertainty pair.

Proof. The proof proceeds from the original structure of the robust control problem (DR-CCF-OP). The constraint on the input \mathbf{u}^* specified by this controller is given by:

$$L_{\tilde{\mathbf{f}}}C(\mathbf{x}) + \nabla C(\mathbf{x})^\top \mathbf{b} + (L_{\tilde{\mathbf{g}}}C(\mathbf{x}) + \nabla C(\mathbf{x})^\top \mathbf{A})\mathbf{u}^* \leq -\alpha(C(\mathbf{x})).$$

for all $(\mathbf{A}, \mathbf{b}) \in \mathcal{U}(\mathbf{x})$. Given the definitions of $\tilde{\mathcal{U}}_C(\mathbf{x})$ and $\mathcal{U}_C(\mathbf{x})$ in (18) and (19), this constraint can be expressed as:

$$q + \mathbf{p}^\top \mathbf{u}^* \leq -\alpha(C(\mathbf{x})) \quad (21)$$

for all $(\mathbf{p}, q) \in \mathcal{U}_C(\mathbf{x})$.

1) *Necessity:* Assume that $\mathcal{U}_C(\mathbf{x}) \cap \mathcal{R} \neq \emptyset$. This implies that there exists a $(\mathbf{p}, q) \in \mathcal{U}_C(\mathbf{x})$ such that $q > -\alpha(C(\mathbf{x}))$ and $\mathbf{p} = \mathbf{0}_m$. Thus for any input $\mathbf{u} \in \mathbb{R}^m$, we have that:

$$q + \mathbf{p}^\top \mathbf{u} = q > -\alpha(C(\mathbf{x})),$$

violating the constraint in (21). Thus the optimization problem is infeasible.

2) *Sufficiency:* Begin by defining the hyperplane $\mathcal{H} \subset \mathbb{R}^{m+1}$ with unit normal $\mathbf{n} = [\mathbf{0}_m^\top \ 1]^\top$ and offset $-\alpha(C(\mathbf{x}))$ and define the set:

$$\bar{\mathcal{U}}_C(\mathbf{x}) = \left\{ (\mathbf{p}, q) \in \mathcal{U}_C(\mathbf{x}) \mid \left\langle \begin{bmatrix} \mathbf{0}_m \\ 1 \end{bmatrix}, \begin{bmatrix} \mathbf{p} \\ q \end{bmatrix} \right\rangle \geq -\alpha(C(\mathbf{x})) \right\},$$

which corresponds to the Lie derivative pairs in the set $\mathcal{U}_C(\mathbf{x})$ that do not meet the certificate function condition (21) strictly under no input ($\mathbf{u}^* = \mathbf{0}_m$). Note that if $\mathbf{u}^* = \mathbf{0}_m$ satisfies (21) for all $(\mathbf{p}, q) \in \bar{\mathcal{U}}_C(\mathbf{x})$, then $\mathbf{u}^* = \mathbf{0}_m$ satisfies (21) for all $(\mathbf{p}, q) \in \mathcal{U}_C(\mathbf{x})$.

We also note that the set $\bar{\mathcal{U}}_C(\mathbf{x})$ is a closed subset of the compact and convex set $\mathcal{U}_C(\mathbf{x})$ that is also contained in a (convex) half-space defined by the hyperplane \mathcal{H} , meaning $\bar{\mathcal{U}}_C(\mathbf{x})$ is compact and convex (as it is the intersection of convex sets). Therefore, we can define:

$$q^{**} = \max_{(\mathbf{p}, q) \in \bar{\mathcal{U}}_C(\mathbf{x})} q = \max_{(\mathbf{p}, q) \in \bar{\mathcal{U}}_C(\mathbf{x})} \left\langle \begin{bmatrix} \mathbf{0}_m \\ 1 \end{bmatrix}, \begin{bmatrix} \mathbf{p} \\ q \end{bmatrix} \right\rangle,$$

which exists as the function $\left\langle \begin{bmatrix} \mathbf{0}_m^\top & 1 \end{bmatrix}^\top, \cdot \right\rangle : \bar{\mathcal{U}}_C(\mathbf{x}) \rightarrow \mathbb{R}$ is continuous on a compact domain. We consider the case that $q^{**} > -\alpha(C(\mathbf{x}))$, as otherwise $\mathbf{u}^* = \mathbf{0}_m$ satisfies (21).

Define the projection function $\Pi : \mathbb{R}^m \times \mathbb{R} \rightarrow \mathbb{R}^m$ such that for $(\mathbf{p}, q) \in \mathbb{R}^m \times \mathbb{R}$ we have:

$$\Pi((\mathbf{p}, q)) = \mathbf{p}.$$

As Π is a linear transform, the image of the compact and convex set $\bar{\mathcal{U}}_C(\mathbf{x})$ under Π , denoted as $\Pi(\bar{\mathcal{U}}_C(\mathbf{x}))$, is compact and convex and by assumption satisfies:

$$\Pi(\bar{\mathcal{U}}_C(\mathbf{x})) \cap \{\mathbf{0}_m\} = \emptyset.$$

We may use the strict separating hyperplane theorem [41] to separate the set $\Pi(\bar{\mathcal{U}}_C(\mathbf{x}))$ from $\{\mathbf{0}_m\}$ with the hyperplane \mathcal{H}_β with unit normal $\mathbf{s} \in S^{m-1}$ and offset $\beta \in \mathbb{R}_{++}$. This hyperplane can also be shifted to pass through the origin, given by \mathcal{H}_0 (with unit normal $\mathbf{s} \in S^{m-1}$ and offset 0). This results in the configuration seen in Figure 1.

These hyperplanes can be extended back into the ambient space $\mathbb{R}^m \times \mathbb{R}$ by defining hyperplanes $\mathcal{H}'_\beta \subset \mathbb{R}^m \times \mathbb{R}$ and $\mathcal{H}'_0 \subset \mathbb{R}^m \times \mathbb{R}$ with normal vector $\mathbf{s}' = [\mathbf{s}^\top \ 0]^\top$ and respective offsets β and 0. These hyperplanes serve to separate the vertical axis $\mathbf{0}_m \times \mathbb{R}$ from the cylinder defined by $\Pi(\bar{\mathcal{U}}_C(\mathbf{x})) \times \mathbb{R}$. Define the two following open halfspaces:

$$\mathcal{H}_0^+ \triangleq \left\{ (\mathbf{p}, q) \in \mathbb{R}^m \times \mathbb{R} \mid \left\langle \mathbf{s}', \begin{bmatrix} \mathbf{p} \\ q \end{bmatrix} \right\rangle = \mathbf{p}^\top \mathbf{s} > 0 \right\} \\ \mathcal{H}_0^- \triangleq \left\{ (\mathbf{p}, q) \in \mathbb{R}^m \times \mathbb{R} \mid \left\langle \mathbf{s}', \begin{bmatrix} \mathbf{p} \\ q \end{bmatrix} \right\rangle = \mathbf{p}^\top \mathbf{s} < 0 \right\}$$

We will now find a feasible input that lies anti-parallel to \mathbf{s} , i.e., we will consider inputs of the form $\mathbf{u}^* = -\gamma\mathbf{s}$ for $\gamma \geq 0$ throughout the rest of this proof. Finding a feasible input satisfying constraint (21) amounts to finding a γ^* such that $q - \gamma^*\mathbf{p}^\top\mathbf{s} \leq -\alpha(\mathcal{C}(\mathbf{x}))$ for all $(\mathbf{p}, q) \in \mathcal{U}_c(\mathbf{x})$. By definition we have that for any $(\mathbf{p}, q) \in \bar{\mathcal{U}}_c(\mathbf{x})$:

$$\left\langle \mathbf{s}', \begin{bmatrix} \mathbf{p} \\ q \end{bmatrix} \right\rangle = \mathbf{s}^\top \mathbf{p} \geq \beta.$$

Let $\gamma_0 > 0$ be defined as:

$$\gamma_0 = \frac{q^{**} + \alpha(\mathcal{C}(\mathbf{x}))}{\beta} > 0.$$

This value of γ_0 implies that for any $(\mathbf{p}, q) \in \bar{\mathcal{U}}_c(\mathbf{x})$ we have:

$$q - \gamma_0\mathbf{p}^\top\mathbf{s} \leq q^{**} - \gamma_0\beta \leq -\alpha(\mathcal{C}(\mathbf{x})),$$

Define the set:

$$\mathcal{V}_c(\mathbf{x}, \gamma) = \{(\mathbf{p}, q - \gamma\mathbf{p}^\top\mathbf{s}) \in \mathbb{R}^m \times \mathbb{R} \mid (\mathbf{p}, q) \in \bar{\mathcal{U}}_c(\mathbf{x})\},$$

and the function:

$$\psi(\gamma) = \max_{(\mathbf{p}, q) \in \mathcal{V}_c(\mathbf{x}, \gamma)} q,$$

that satisfies $\psi(0) = q^{**}$ and $\psi(\gamma_0) \leq -\alpha(\mathcal{C}(\mathbf{x}))$. We later show that ψ is continuous, but note the intermediate value theorem implies the existence of a $\gamma^* \in (0, \gamma_0]$ with:

$$\psi(\gamma^*) = \max_{(\mathbf{p}, q) \in \mathcal{V}_c(\mathbf{x}, \gamma^*)} q = -\alpha(\mathcal{C}(\mathbf{x})).$$

Fixing $\mathbf{u}^* = -\gamma^*\mathbf{s}$, we have that condition (21) is satisfied for all points $(\mathbf{p}, q) \in \mathcal{H}_0^+ \cap \bar{\mathcal{U}}_c(\mathbf{x})$. Likewise, consider a point $(\mathbf{p}, q) \in \mathcal{H}_0^+ \cap (\mathcal{U}_c(\mathbf{x}) \setminus \bar{\mathcal{U}}_c(\mathbf{x}))$, for which the condition (21) is satisfied with no input as $q < -\alpha(\mathcal{C}(\mathbf{x}))$, and note that $\mathbf{p}^\top\mathbf{s} > 0$. We have that:

$$q + \mathbf{p}^\top\mathbf{u}^* = q - \gamma^*\mathbf{p}^\top\mathbf{s} < -\alpha(\mathcal{C}(\mathbf{x})).$$

Combining these results, we have that condition (21) is satisfied for all $(\mathbf{p}, q) \in \mathcal{H}_0^+ \cap \mathcal{U}_c(\mathbf{x})$. Additionally, for any $(\mathbf{p}, q) \in \mathcal{H}_0' \cap \mathcal{U}_c(\mathbf{x})$, $\mathbf{p}^\top\mathbf{s} = 0$, as \mathbf{s}' is normal to \mathcal{H}_0' . As strict separation implies that $q < -\alpha(\mathcal{C}(\mathbf{x}))$ for any $(\mathbf{p}, q) \in \mathcal{H}_0' \cap \mathcal{U}_c(\mathbf{x})$, we have $q - \gamma^*\mathbf{p}^\top\mathbf{s} < -\alpha(\mathcal{C}(\mathbf{x}))$ and condition (21) is satisfied.

Lastly, we must consider points $(\mathbf{p}, q) \in \mathcal{H}_0^- \cap \mathcal{U}_c(\mathbf{x})$. To this end, define the set:

$$\mathcal{W}_c(\mathbf{x}, \gamma^*) = \{(\mathbf{p}, q - \gamma^*\mathbf{p}^\top\mathbf{s}) \in \mathbb{R}^m \times \mathbb{R} \mid (\mathbf{p}, q) \in \mathcal{U}_c(\mathbf{x})\}.$$

This set can be interpreted as an invertible linear transformation of the set $\mathcal{U}_c(\mathbf{x})$ as we have:

$$\begin{bmatrix} \mathbf{p}' \\ q' \end{bmatrix} = \begin{bmatrix} \mathbf{p} \\ q - \gamma^*\mathbf{p}^\top\mathbf{s} \end{bmatrix} = \begin{bmatrix} \mathbf{I}_{m \times m} & \mathbf{0}_m \\ -\gamma^*\mathbf{s}^\top & 1 \end{bmatrix} \begin{bmatrix} \mathbf{p} \\ q \end{bmatrix},$$

implying that $\mathcal{W}_c(\mathbf{x}, \gamma^*)$ is compact and convex. Similarly, define the set:

$$\bar{\mathcal{W}}_c(\mathbf{x}, \gamma^*) = \left\{ (\mathbf{p}, q) \in \mathcal{W}_c(\mathbf{x}, \gamma^*) \mid \left\langle \begin{bmatrix} \mathbf{0}_m \\ 1 \end{bmatrix}, \begin{bmatrix} \mathbf{p} \\ q \end{bmatrix} \right\rangle \geq -\alpha(\mathcal{C}(\mathbf{x})) \right\}.$$




Fig. 1. Projected view of Lie derivative uncertainty set used in proving feasibility of the (DR-CCF-OP). The green region corresponds to $\Pi(\mathcal{U}_c(\mathbf{x}))$ while the orange region highlights the set $\bar{\mathcal{U}}_c(\mathbf{x})$ lying above the hyperplane \mathcal{H} . The line \mathcal{H}_β represents the strictly separating hyperplane with normal \mathbf{s} and offset β , while \mathcal{H}_0 is the hyperplane shifted to pass through the origin.

The set $\bar{\mathcal{W}}_c(\mathbf{x}, \gamma^*)$ is convex and contains points in $\mathcal{W}_c(\mathbf{x}, \gamma^*)$ that are in or above the hyperplane \mathcal{H} , or points that meet with equality or violate (21), respectively.

By the specific choice of γ^* , there exists at least one point $\mathbf{v}^+ \in \mathcal{H}_0^+ \cap \bar{\mathcal{W}}_c(\mathbf{x}, \gamma^*)$. Furthermore, by the preceding strict separation, we have that $\mathcal{H}_0' \cap \bar{\mathcal{W}}_c(\mathbf{x}, \gamma^*) = \emptyset$. Now assume for contradiction that $\mathcal{H}_0^- \cap \bar{\mathcal{W}}_c(\mathbf{x}, \gamma^*) \neq \emptyset$, or that there exists $\mathbf{v}^- \in \mathcal{H}_0^- \cap \bar{\mathcal{W}}_c(\mathbf{x}, \gamma^*)$. As the set $\bar{\mathcal{W}}_c(\mathbf{x}, \gamma^*)$ is convex and $\langle \mathbf{s}', \mathbf{v}^+ \rangle > 0$ and $\langle \mathbf{s}', \mathbf{v}^- \rangle < 0$, there is a $\lambda^* \in (0, 1)$ satisfying

$$\langle \mathbf{s}', (1 - \lambda^*)\mathbf{v}^+ + \lambda^*\mathbf{v}^- \rangle = 0,$$

implying $(1 - \lambda^*)\mathbf{v}^+ + \lambda^*\mathbf{v}^- \in \mathcal{H}_0' \cap \bar{\mathcal{W}}_c(\mathbf{x}, \gamma^*)$. This contradicts the fact $\mathcal{H}_0' \cap \bar{\mathcal{W}}_c(\mathbf{x}, \gamma^*) = \emptyset$, implying that $\mathcal{H}_0^- \cap \bar{\mathcal{W}}_c(\mathbf{x}, \gamma^*) = \emptyset$, or that condition (21) is satisfied for all points $(\mathbf{p}, q) \in \mathcal{H}_0^- \cap \mathcal{U}_c(\mathbf{x})$. Combining all previous results, condition (21) is satisfied for all points $(\mathbf{p}, q) \in \mathcal{U}_c(\mathbf{x})$ for the input $\mathbf{u}^* = -\gamma^*\mathbf{s}$, ensuring feasibility.

Finally, we show that ψ is continuous. For $\gamma, \gamma' > 0$ and $(\mathbf{p}, q) \in \bar{\mathcal{U}}_c(\mathbf{x})$, we have:

$$\begin{aligned} & \min_{(\mathbf{p}', q') \in \mathcal{V}_c(\mathbf{x}, \gamma')} \left\| \begin{bmatrix} \mathbf{p} \\ q - \gamma\mathbf{p}^\top\mathbf{s} \end{bmatrix} - \begin{bmatrix} \mathbf{p}' \\ q' \end{bmatrix} \right\|_2 \\ & \leq \left\| \begin{bmatrix} \mathbf{p} \\ q - \gamma\mathbf{p}^\top\mathbf{s} \end{bmatrix} - \begin{bmatrix} \mathbf{p} \\ q - \gamma'\mathbf{p}^\top\mathbf{s} \end{bmatrix} \right\|_2 \leq \mathbf{p}^\top\mathbf{s}|\gamma - \gamma'|, \end{aligned}$$

and similarly, for $(\mathbf{p}', q') \in \bar{\mathcal{U}}_c(\mathbf{x})$, we have:

$$\min_{(\mathbf{p}, q) \in \mathcal{V}_c(\mathbf{x}, \gamma)} \left\| \begin{bmatrix} \mathbf{p} \\ q \end{bmatrix} - \begin{bmatrix} \mathbf{p}' \\ q' - \gamma'(\mathbf{p}')^\top\mathbf{s} \end{bmatrix} \right\|_2 \leq (\mathbf{p}')^\top\mathbf{s}|\gamma - \gamma'|.$$

Therefore, the Hausdorff distance between $\mathcal{V}_c(\mathbf{x}, \gamma)$ and $\mathcal{V}_c(\mathbf{x}, \gamma')$ is bounded by:

$$d_H(\mathcal{V}_c(\mathbf{x}, \gamma), \mathcal{V}_c(\mathbf{x}, \gamma')) \leq \left(\max_{(\mathbf{p}, q) \in \bar{\mathcal{U}}_c(\mathbf{x})} \mathbf{p}^\top\mathbf{s} \right) |\gamma - \gamma'|,$$

implying $\mathcal{V}_c(\mathbf{x}, \gamma)$ is Lipschitz continuous (with respect to the Hausdorff metric) as a function of γ . The support function of a nonempty, compact, and convex set $\mathcal{A} \subset \mathbb{R}^m \times \mathbb{R}$

given by $h_{\mathcal{A}} : \mathbb{R}^{m+1} \rightarrow \mathbb{R}$ is defined as:

$$h_{\mathcal{A}}(\mathbf{v}) = \max_{(\mathbf{p}, q) \in \mathcal{A}} \left\langle \mathbf{v}, \begin{bmatrix} \mathbf{p} \\ q \end{bmatrix} \right\rangle$$

Recalling $\mathbf{n} = [\mathbf{0}_m^\top \ 1]^\top$, the Hausdorff distance between two nonempty, compact, and convex sets $\mathcal{A}, \mathcal{B} \subset \mathbb{R}^m \times \mathbb{R}$ satisfies:

$$d_H(\mathcal{A}, \mathcal{B}) = \max_{\mathbf{v} \in \mathcal{S}^m} |h_{\mathcal{A}}(\mathbf{v}) - h_{\mathcal{B}}(\mathbf{v})| \geq |h_{\mathcal{A}}(\mathbf{n}) - h_{\mathcal{B}}(\mathbf{n})|.$$

Noting that ψ can be expressed in terms of the support function as a composition, with:

$$\psi(\gamma) = \max_{(\mathbf{p}, q) \in \mathcal{V}_c(\mathbf{x}, \gamma)} q = h_{\mathcal{V}_c(\mathbf{x}, \gamma)}(\mathbf{n}),$$

implies ψ is a continuous function. \square

V. SIMULATION

To demonstrate the capabilities of the proposed controller, we run simulated experiments in the setting of an inverted pendulum, described by the system model:

$$\frac{d}{dt} \begin{bmatrix} \theta \\ \dot{\theta} \end{bmatrix} = \begin{bmatrix} \dot{\theta} \\ \hat{g} \sin(\theta) \end{bmatrix} + \begin{bmatrix} 0 \\ \frac{1}{\hat{m}\hat{\ell}^2} \end{bmatrix} u \quad (22)$$

with state $\mathbf{x} = [\theta, \dot{\theta}]^\top \in \mathbb{R}^2$, gravitational acceleration estimate $\hat{g} = 10$, length estimate $\hat{\ell} = 0.63$, and mass estimate $\hat{m} = 0.63$. We assume that the unknown true system has modified inverted pendulum dynamics given by:

$$\frac{d}{dt} \begin{bmatrix} \theta \\ \dot{\theta} \end{bmatrix} = \begin{bmatrix} \dot{\theta} \\ \frac{g}{\ell} \sin(\theta) \end{bmatrix} + \begin{bmatrix} 0 \\ \frac{1 - 0.75 \exp(-\theta^2)}{m\ell^2} \end{bmatrix} u \quad (23)$$

with gravitational acceleration $g = 10$, length $\ell = 0.7$, and mass $m = 0.7$. We note the inclusion of a state-dependent input gain given by $1 - 0.75 \exp(-\theta^2)$, which attenuates the input most significantly when the pendulum is vertical.

Consider the functions $V : \mathbb{R}^2 \rightarrow \mathbb{R}_+$ and $h : \mathbb{R}^2 \rightarrow \mathbb{R}$, given by $V(\mathbf{x}) = \mathbf{x}^\top \mathbf{P} \mathbf{x}$ and $h(\mathbf{x}) = \mathbf{x}^\top \mathbf{P} \mathbf{x} - c$ with

$$\mathbf{P} = \begin{bmatrix} \sqrt{3} & 1 \\ 1 & \sqrt{3} \end{bmatrix} \quad (24)$$

and a constant $c = 0.2$. Noting that both the system model (22) and the true system (23) are feedback linearizable, V

and h satisfy the CCF condition (4) for the comparison function $\alpha(r) = \lambda_{\min}(\mathbf{Q})r / \lambda_{\max}(\mathbf{P})$, for both the model and the true system (implying Assumption 1 is met). In particular, V may serve as a CLF, and h as a CBF.

We explore data-driven stability and safety with this system for different methods of gathering training data. We compare the robust data-driven controller with an oracle controller, which knows the true model, and a nominal controller, which treats the estimated model as if it were true, both specified via the (CCF-QP). In each setting, the system model underestimates the pendulum mass and length and assumes that the input gain is independent of the state. As a result, the Lipschitz constants of the errors can be bounded by $\mathcal{L}_{\tilde{f}} = g|\ell - \hat{\ell}|/(\ell\hat{\ell})$ and $\mathcal{L}_{\tilde{g}} = 0.75\sqrt{2} \exp(-\frac{1}{2})/(m\ell^2)$.

For the stability experiment, we generate data sets by gridding the state and input spaces. We consider θ in the interval $[0, 1]$ and $\dot{\theta}$ in the interval $[-0.25, 0.25]$, with grid sizes $\epsilon_\theta = \epsilon_{\dot{\theta}} = \frac{1}{40}$. We generate a *sparse* data set by considering $u \in \{-5, -1\}$ and a *dense* data set by considering $u \in \{-5, -3, -1, 1, 3, 5\}$. We set the regularizing controller \mathbf{k}_d to be a feedback linearizing controller designed using the estimated dynamics and linear gains $k_p = 1/2$ and $k_d = \sqrt{3}/2$. The system is simulated from the initial condition $\mathbf{x} = [0.8, 0.1]$ for 10 seconds, with control inputs specified at 100 Hz. For the safety experiment, we consider a similar pair of sparse and dense data sets with θ in the interval $[0, 0.25]$ and $\dot{\theta}$ in the same interval as the stability experiment, the same grid sizes, and the same sets of control inputs. The system is simulated with no regularizing controller $\mathbf{k}_d \equiv 0$ from the initial condition $\mathbf{x} = [0.1, 0.1]$ for the same amount of time and the same control input frequency.

The results of the simulations may be seen in Figure 2. In both experiments, we see that the nominal controller fails to achieve the specified objective, while the oracle controller succeeds. Furthermore, we see that the data-driven controllers perform nearly identically for both the sparse and dense input data sets. This similarity indicates that greater variety in input directions and coverage of the input space by the data set are not needed to achieve satisfactory closed-loop behavior.

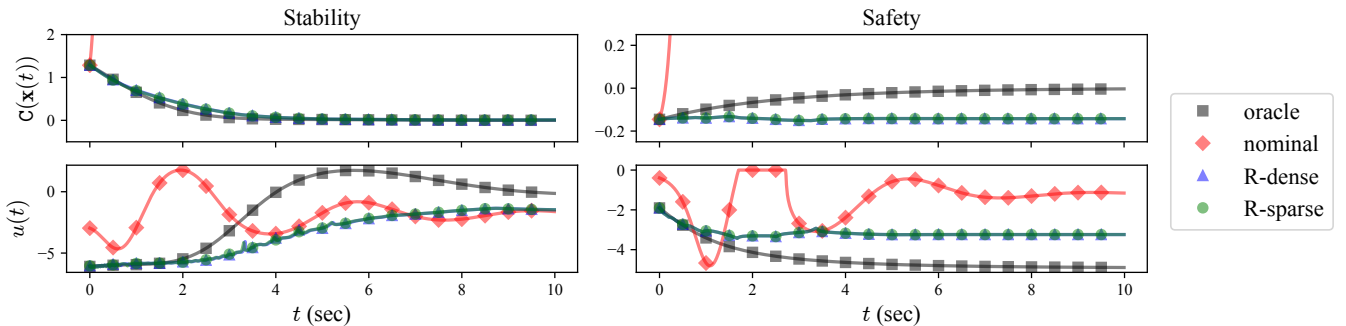


Fig. 2. The value of the certificate function $C(\mathbf{x})$ (top) and input \mathbf{u} (bottom) over time for the stability (left) and safety (right) experiments. We compare the behavior of an oracle controller which has access to the true dynamics (black square), a nominal controller which assumes that the estimated dynamics are true (red diamond), a robust controller using dense input data (blue triangle), and a robust controller using sparse input data (green circle). For stability, the desired behavior for $C(\mathbf{x})$ is to converge towards 0, while for safety it is to remain nonpositive.

VI. CONCLUSION

We propose a novel approach for robust data-driven control synthesis under model uncertainty and show that good performance can be achieved when training data provides sufficiently dense coverage of the state space with targeted excitation in input directions. This approach is well positioned to make progress on both theoretical and application-oriented problems at the intersection nonlinear control and non-parametric machine learning. Future work includes investigating strategies for data segmentation to enable efficient computation for closed-loop control and understanding continuity and recursive feasibility of the controller.

REFERENCES

- [1] R. Dimitrova and R. Majumdar, "Deductive control synthesis for alternating-time logics," in *International Conference on Embedded Software (EMSOFT)*. IEEE, 2014, pp. 1–10.
- [2] N. M. Boffi, S. Tu, N. Matni, J.-J. E. Slotine, and V. Sindhvani, "Learning stability certificates from data," *arXiv preprint arXiv:2008.05952*, 2020.
- [3] Z. Artstein, "Stabilization with relaxed controls," *Nonlinear Analysis: Theory, Methods & Applications*, vol. 7, no. 11, pp. 1163–1173, 1983.
- [4] A. Ames, J. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs with application to adaptive cruise control," in *Conference on Decision & Control (CDC)*. IEEE, 2014, pp. 6271–6278.
- [5] S. Prajna and A. Jadbabaie, "Safety verification of hybrid systems using barrier certificates," in *International Workshop on Hybrid Systems: Computation and Control (HSCC)*. Springer, 2004, pp. 477–492.
- [6] A. D. Ames, K. Galloway, K. Sreenath, and J. W. Grizzle, "Rapidly exponentially stabilizing control lyapunov functions and hybrid zero dynamics," *Transactions on Automatic Control*, vol. 59, no. 4, pp. 876–891, 2014.
- [7] Q. Nguyen, A. Hereid, J. W. Grizzle, A. D. Ames, and K. Sreenath, "3d dynamic walking on stepping stones with control barrier functions," in *Conference on Decision and Control (CDC)*. IEEE, 2016, pp. 827–834.
- [8] S. M. Khansari-Zadeh and A. Billard, "Learning control lyapunov function to ensure stability of dynamical system-based robot reaching motions," *Robotics and Autonomous Systems*, vol. 62, no. 6, 2014.
- [9] D. Pickem, P. Glotfelter, L. Wang, M. Mote, A. Ames, E. Feron, and M. Egerstedt, "The robotarium: A remotely accessible swarm robotics research testbed," in *International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 1699–1706.
- [10] J. Kober, J. A. Bagnell, and J. Peters, "Reinforcement learning in robotics: A survey," *The International Journal of Robotics Research*, vol. 32, no. 11, pp. 1238–1274, 2013.
- [11] G. Shi, X. Shi, M. O'Connell, R. Yu, K. Azizzadenesheli, A. Anandkumar, Y. Yue, and S. Chung, "Neural lander: Stable drone landing control using learned dynamics," in *International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 9784–9790.
- [12] R. Cheng, G. Orosz, R. M. Murray, and J. W. Burdick, "End-to-end safe reinforcement learning through barrier functions for safety-critical continuous control tasks," in *Conference on Artificial Intelligence*, vol. 33. AAAI, 2019, pp. 3387–3395.
- [13] J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, "Learning quadrupedal locomotion over challenging terrain," *Science robotics*, vol. 5, no. 47, 2020.
- [14] A. Aswani, H. Gonzalez, S. S. Sastry, and C. Tomlin, "Provably safe and robust learning-based model predictive control," *Automatica*, vol. 49, no. 5, pp. 1216–1226, 2013.
- [15] T. Beckers, D. Kulić, and S. Hirche, "Stable gaussian process based tracking control of euler-lagrange systems," *Automatica*, vol. 103, pp. 390–397, 2019.
- [16] F. Berkenkamp and A. P. Schoellig, "Safe and robust learning control with gaussian processes," in *European Control Conference (ECC)*. IEEE, 2015, pp. 2496–2501.
- [17] J. H. Gillula and C. J. Tomlin, "Guaranteed safe online learning via reachability: tracking a ground target using a quadrotor," in *International Conference on Robotics and Automation (ICRA)*. IEEE, 2012, pp. 2723–2730.
- [18] G. Qu, C. Yu, S. Low, and A. Wierman, "Combining model-based and model-free methods for nonlinear control: A provably convergent policy gradient approach," *arXiv preprint arXiv:2006.07476*, 2020.
- [19] A. Lederer, A. Capone, and S. Hirche, "Parameter optimization for learning-based control of control-affine systems," *Proceedings of Machine Learning Research (PMLR)*, vol. 120, pp. 465–475, 2020.
- [20] J. Choi, F. Castañeda, C. Tomlin, and K. Sreenath, "Reinforcement Learning for Safety-Critical Control under Model Uncertainty, using Control Lyapunov Functions and Control Barrier Functions," in *Robotics: Science and Systems*, Corvallis, Oregon, USA, July 2020.
- [21] M. Cohen and C. Belta, "Approximate optimal control for safety-critical systems with control barrier functions," *arXiv preprint arXiv:2008.04122*, 2020.
- [22] F. Castañeda, J. J. Choi, B. Zhang, C. J. Tomlin, and K. Sreenath, "Gaussian process-based min-norm stabilizing controller for control-affine systems with uncertain input effects," *arXiv preprint arXiv:2011.07183*, 2020.
- [23] J. F. Fisac, A. K. Akametalu, M. N. Zeilinger, S. Kaynama, J. Gillula, and C. J. Tomlin, "A general safety framework for learning-based control in uncertain robotic systems," *Transactions on Automatic Control*, 2018.
- [24] J. Umlauf, L. Pöhler, and S. Hirche, "An uncertainty-based control lyapunov approach for control-affine systems modeled by gaussian process," *Control Systems Letters*, vol. 2, no. 3, pp. 483–488, 2018.
- [25] L. Zheng, J. Pan, R. Yang, H. Cheng, and H. Hu, "Learning-based safety-stability-driven control for safety-critical systems under model uncertainties," *arXiv preprint arXiv:2008.03421*, 2020.
- [26] A. Lederer, A. Capone, J. Umlauf, and S. Hirche, "How training data impacts performance in learning-based control," *arXiv preprint arXiv:2005.12062*, 2020.
- [27] F. Berkenkamp, R. Moriconi, A. P. Schoellig, and A. Krause, "Safe learning of regions of attraction for uncertain, nonlinear systems with gaussian processes," in *Conference on Decision & Control (CDC)*. IEEE, 2016, pp. 4661–4666.
- [28] F. Berkenkamp, M. Turchetta, A. Schoellig, and A. Krause, "Safe model-based reinforcement learning with stability guarantees," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2017, pp. 908–918.
- [29] R. M. Murray, Z. Li, and S. S. Sastry, *A mathematical introduction to robotic manipulation*. CRC press, 1994.
- [30] P. A. Ioannou and C.-C. Chien, "Autonomous intelligent cruise control," *IEEE Transactions on Vehicular Technology*, vol. 42, no. 4, pp. 657–672, 1993.
- [31] M. Jankovic, "Robust control barrier functions for constrained stabilization of nonlinear systems," *Automatica*, vol. 96, pp. 359–367, 2018.
- [32] L. Perko, *Differential equations and dynamical systems*. Springer Science & Business Media, 2013, vol. 7.
- [33] E. D. Sontag, "A 'universal' construction of artstein's theorem on nonlinear stabilization," *Systems & Control Letters*, vol. 13, no. 2, pp. 117–123, 1989.
- [34] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs for safety critical systems," *Transactions on Automatic Control*, vol. 62, no. 8, pp. 3861–3876, 2017.
- [35] F. Blanchini, "Set invariance in control," *Automatica*, vol. 35, no. 11, pp. 1747–1767, 1999.
- [36] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: Theory and applications," in *European Control Conference (ECC)*. IEEE, 2019, pp. 3420–3431.
- [37] A. J. Taylor, V. D. Dorobantu, H. M. Le, Y. Yue, and A. D. Ames, "Episodic learning with control lyapunov functions for uncertain robotic systems," in *International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2019, pp. 6878–6884.
- [38] A. J. Taylor, A. Singletary, Y. Yue, and A. D. Ames, "Learning for safety-critical control with control barrier functions," *Proceedings of Machine Learning Research (PMLR)*, vol. 120, pp. 708–717, 2020.
- [39] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski, *Robust optimization*. Princeton University Press, 2009, vol. 28.
- [40] A. Lederer, A. J. O. Conejo, K. Maier, W. Xiao, and S. Hirche, "Real-time regression with dividing local gaussian processes," *arXiv preprint arXiv:2006.09446*, 2020.
- [41] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge University Press, 2004.