

A Control Lyapunov Perspective on Episodic Learning via Projection to State Stability

Andrew J. Taylor¹, Victor D. Dorobantu¹, Meera Krishnamoorthy,
Hoang M. Le, Yisong Yue, and Aaron D. Ames

Abstract—The goal of this paper is to understand the impact of learning on control synthesis from a Lyapunov function perspective. In particular, rather than consider uncertainties in the full system dynamics, we employ Control Lyapunov Functions (CLFs) as low-dimensional projections. To understand and characterize the uncertainty that these projected dynamics introduce in the system, we introduce a new notion: Projection to State Stability (PSS). PSS can be viewed as a variant of Input to State Stability defined on projected dynamics, and enables characterizing robustness of a CLF with respect to the data used to learn system uncertainties. We use PSS to bound uncertainty in affine control, and demonstrate that a practical episodic learning approach can use PSS to characterize uncertainty in the CLF for robust control synthesis.

I. INTRODUCTION

Properly characterizing uncertainty is a key aspect of robust control [35]. With the increasing use of learning for dynamics modelling and control synthesis [5], [10], [8], [11], [31], [24], it is correspondingly important to develop new tools to reason about the interplay between learning and robust control.

In this paper, we focus on the interplay between learning and robustness for control synthesis using Control Lyapunov Functions (CLFs) [4], [18]. The use of CLFs has seen multiple applications in recent years [19], [14], [23], and one of their primary benefits is to enable control objectives to be represented in a low-dimensional form that can be utilized with optimization-based real-time controllers [3]. This low-dimensional representation is also appealing from a learning perspective, as learning is typically more tractable in lower-dimensional spaces [32], [34], [31].

The practical implementation of CLFs remains challenging due to gap between models and hardware. In many cases, extensive tuning upon deployment is necessary [19], and even with this tuning the system is often not able to track a desired state or trajectory perfectly. Other approaches, such as those based on adaptive control [17], reduce error in the estimated CLF time derivative, but are restricted to learning over specific classes of model uncertainty.

We thus build upon ideas in robust control in order to guarantee performance in the presence of model mis-

This work was supported in part by funding and gifts from DARPA, Intel, PIMCO, and Google.

¹Both authors contributed equally.

All authors are with the Department of Computing and Mathematical Sciences, California Institute of Technology, Pasadena, CA 91125, USA
ajtaylor@caltech.edu, vdoroban@caltech.edu,
mkrishna@caltech.edu, hmle@caltech.edu,
yyue@caltech.edu, ames@caltech.edu

specification. The idea of robust CLFs is not new (cf. [12], [13]), but existing analyses focus on the full-dimensional state dynamics, which can be burdensome for learning.

In this paper, we make two main contributions. First, we propose a novel characterization called *Projection to State Stability* (PSS), which is a variant of the well-studied *Input to State Stability* (ISS) property [25], [29], [28], [33], [27], but defined on projected dynamics rather than the original state dynamics. Like ISS, PSS provides a tool to characterize tracking error in terms of the magnitude of the disturbance or uncertainty. Unlike ISS, PSS can characterize dynamic uncertainty directly in the derivative of a CLF, thus allowing a low dimensional representation of the uncertainty. In our second contribution, we demonstrate the practicality of PSS by incorporating it into an episodic learning algorithm.

Our paper is organized as follows. Section II reviews CLFs and ISS. Section III defines Projection to State Stability (PSS), and how PSS enables constructing bounds on the state of a system that depend on a projected disturbance. Section IV defines a broad class of model uncertainty for affine control systems, evaluates how this uncertainty impacts the Lyapunov derivative, and demonstrates how to restrict this uncertainty with data to determine if a system is PSS. Section V discusses how episodic learning can be used to improve PSS guarantees in practice, and presents simulation results with an uncertain inverted pendulum model.

II. PRELIMINARIES

This section provides a review of Control Lyapunov Functions (CLFs) and Input to State Stability (ISS). These tools will be used in Section III to define Projection to State Stability. This section concludes with a brief discussion of how these definitions must be modified to hold over a restriction of the domain.

Consider a state space $\mathcal{X} \subseteq \mathbb{R}^n$ and a control input space $\mathcal{U} \subseteq \mathbb{R}^m$. Assume that \mathcal{X} is path-connected and that $\mathbf{0} \in \mathcal{X}$. Consider a system governed by:

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

for state $\mathbf{x} \in \mathcal{X}$ and its derivative $\dot{\mathbf{x}}$, control input $\mathbf{u} \in \mathcal{U}$, and dynamics $\mathbf{f} : \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}^n$. The dynamics have an affine dependence on control input if for all $\mathbf{x} \in \mathcal{X}$, the map $\mathbf{u} \mapsto \mathbf{f}(\mathbf{x}, \mathbf{u})$ is affine. In this paper we assume \mathbf{f} is locally Lipschitz continuous and that $\mathbf{f}(\mathbf{0}, \mathbf{u}_0) = \mathbf{0}$ for some $\mathbf{u}_0 \in \mathcal{U}$.

The following definitions, taken from [16], are useful in analyzing stability of (1).

Definition 1 (Class \mathcal{K} Function). A continuous function $\alpha : [0, a) \rightarrow \mathbb{R}_+$, with $a > 0$, is *class \mathcal{K}* , denoted $\alpha \in \mathcal{K}$, if it is monotonically (strictly) increasing and satisfies $\alpha(0) = 0$. If the domain of α is all of \mathbb{R}_+ and $\lim_{r \rightarrow \infty} \alpha(r) = \infty$, then α is termed *radially unbounded* and *class \mathcal{K}_∞* .

Definition 2 (Class \mathcal{KL} Function). A continuous function $\beta : [0, a) \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$, with $a > 0$, is *class \mathcal{KL}* , denoted $\beta \in \mathcal{KL}$, if the function $r \mapsto \beta(r, s) \in \mathcal{K}$ for all $s \in \mathbb{R}_+$, and the function $s \mapsto \beta(r, s)$ is monotonically non-increasing with $\beta(r, s) \rightarrow 0$ as $s \rightarrow \infty$ for all $r \in [0, a)$.

We note that the strictly increasing nature of Class \mathcal{K} (\mathcal{K}_∞) functions permits an inverse Class \mathcal{K} (\mathcal{K}_∞) function $\alpha^{-1} : [0, \alpha(a)) \rightarrow \mathbb{R}_+$. We also note that the composition of Class \mathcal{K} (\mathcal{K}_∞) functions is itself a Class \mathcal{K} (\mathcal{K}_∞) function. Given these definitions, we define Control Lyapunov Functions (CLFs) as in [4], [18].

Definition 3 (Control Lyapunov Function). A continuously differentiable function $V : \mathcal{X} \rightarrow \mathbb{R}_+$ is a *CLF* for (1) on \mathcal{X} if there exist $\underline{\alpha}, \bar{\alpha}, \alpha \in \mathcal{K}_\infty$ such that:

$$\begin{aligned} \underline{\alpha}(\|\mathbf{x}\|) &\leq V(\mathbf{x}) \leq \bar{\alpha}(\|\mathbf{x}\|) \\ \inf_{\mathbf{u} \in \mathcal{U}} \dot{V}(\mathbf{x}, \mathbf{u}) &\leq -\alpha(\|\mathbf{x}\|), \end{aligned} \quad (2)$$

for all $\mathbf{x} \in \mathcal{X}$.

Under the assumption of affine dependence of \mathbf{f} on the control input, the existence of a CLF for (1) implies the existence of a continuous (away from $\mathbf{x} = \mathbf{0}$) state-feedback controller $\mathbf{k} : \mathcal{X} \rightarrow \mathcal{U}$, that renders $\mathbf{0}$ globally asymptotically stable [4]. In particular, for all $\mathbf{x} \in \mathcal{X}$, $\mathbf{k}(\mathbf{x})$ satisfies $\dot{V}(\mathbf{x}, \mathbf{k}(\mathbf{x})) \leq -\alpha(\|\mathbf{x}\|)$. Continuity at $\mathbf{x} = \mathbf{0}$ can be achieved if V satisfies the continuous control property:

Definition 4 (Continuous Control Property [4], [12]). Let $\mathbf{f}(\mathbf{0}, \mathbf{u}_0) = \mathbf{0}$ for some $\mathbf{u}_0 \in \mathcal{U}$. A CLF V is said to satisfy the *continuous control property* if for every $\epsilon > 0$, there exists $\delta > 0$ such that when $\|\mathbf{x}\| < \delta$, there exists a \mathbf{u}_1 , with $\|\mathbf{u}_1 - \mathbf{u}_0\| < \epsilon$ such that $\dot{V}(\mathbf{x}, \mathbf{u}_1) < 0$. If $\mathbf{u}_0 = \mathbf{0}$, V is said to satisfy the *small control property* [26].

Smoothness of \mathbf{k} can be attained with assumptions on smoothness of \mathbf{f} and V [25]. We also note that $\underline{\alpha}, \bar{\alpha}, \alpha$ only need to be Class \mathcal{K} for this definition, in which case asymptotic stability is achieved locally, but we extend them to \mathcal{K}_∞ to simplify later analysis.

To accommodate disturbances or uncertainties, we consider a disturbance space $\mathcal{D} \subseteq \mathbb{R}^d$, and a modified system:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{u}, \mathbf{d}), \quad (3)$$

for disturbance $\mathbf{d} \in \mathcal{D}$ and dynamics $\mathbf{f} : \mathcal{X} \times \mathcal{U} \times \mathcal{D} \rightarrow \mathbb{R}^n$. We again assume \mathbf{f} is locally Lipschitz continuous and that $\mathbf{f}(\mathbf{0}, \mathbf{u}_0, \mathbf{0}) = \mathbf{0}$ for some $\mathbf{u}_0 \in \mathcal{U}$. The disturbance may be time-varying, state-dependent, and/or input-dependent. We assume that the disturbance is bounded for almost all times $t \geq 0$ (essentially bounded in time). This leads to the definition of ISS as in [29], [12] and ISS-CLFs as in [28]:

Definition 5 (Input to State Stability). Under a continuous state-feedback controller $\mathbf{k} : \mathcal{X} \rightarrow \mathcal{U}$, the system governed by (3) is *Input to State Stable (ISS)* if there exist $\beta \in \mathcal{KL}_\infty$ and $\gamma \in \mathcal{K}_\infty$ such that it satisfies:

$$\|\mathbf{x}(t)\| \leq \beta(\|\mathbf{x}(0)\|, t) + \gamma \left(\sup_{\tau \geq 0} \|\mathbf{d}(\tau)\| \right), \quad (4)$$

for all $t \geq 0$.

Intuitively, an ISS system is asymptotically stable to a ball with a radius that scales with the essential bound on the norm of the disturbance.

Definition 6 (Input to State Stable Control Lyapunov Function). A continuously differentiable function $V : \mathcal{X} \rightarrow \mathbb{R}_+$ is an *Input to State Stable Control Lyapunov Function (ISS-CLF)* for (3) on \mathcal{X} if there exist $\underline{\alpha}, \bar{\alpha}, \alpha, \rho \in \mathcal{K}_\infty$ such that:

$$\begin{aligned} \underline{\alpha}(\|\mathbf{x}\|) &\leq V(\mathbf{x}) \leq \bar{\alpha}(\|\mathbf{x}\|) \\ \|\mathbf{x}\| \geq \rho(\|\mathbf{d}\|) &\implies \inf_{\mathbf{u} \in \mathcal{U}} \dot{V}(\mathbf{x}, \mathbf{u}, \mathbf{d}) \leq -\alpha(\|\mathbf{x}\|), \end{aligned} \quad (5)$$

for all $\mathbf{x} \in \mathcal{X}$ and $\mathbf{d} \in \mathcal{D}$.

As with CLFs, under the assumption of affine dependence of \mathbf{f} on control inputs, an ISS-CLF for (3) satisfying the continuous control property implies the existence of a continuous state-feedback controller $\mathbf{k} : \mathcal{X} \rightarrow \mathcal{U}$ that renders the system ISS [28], [12]. Smoothness of this controller can be achieved similarly to the CLF case [25]. If the disturbance is input-dependent, it is additionally required that \mathbf{k} induces essentially bounded disturbances in time.

The condition on the Lyapunov function derivative in (2) or (5) may not be satisfied on the entire state space \mathcal{X} . In particular it may only be satisfied on a subset $\mathcal{C} \subseteq \mathcal{X}$. The system may leave \mathcal{C} during its evolution, implying the desired derivative condition may no longer be satisfiable. We therefore consider the following definition and lemma.

Definition 7 (Forward Invariance). Consider the system governed by (1). A subset $\mathcal{F} \subseteq \mathcal{X}$ can be rendered *forward invariant* if there exists a state-feedback controller $\mathbf{k} : \mathcal{X} \rightarrow \mathcal{U}$ such that the closed-loop map $\mathbf{x} \mapsto \mathbf{f}(\mathbf{x}, \mathbf{k}(\mathbf{x}))$ is Lipschitz continuous and $\mathbf{x}(0) \in \mathcal{F}$ implies $\mathbf{x}(t) \in \mathcal{F}$ for all $t \in I(\mathbf{x}(0)) \subseteq \mathbb{R}$, the maximum interval of existence.

The interval $I(\mathbf{x}(0))$ exists by Lipschitz continuity of the closed-loop map, and the solution of the differential equation is guaranteed to be unique. The definition of forward invariance applies to systems governed by (3), with disturbances appropriately restricted to subsets of \mathcal{D} if the disturbances are modeled as state-dependent and/or input-dependent. If $\mathbf{0} \in \mathcal{C}$, we may restrict Definitions 3 and 6 to a forward invariant subset $\mathcal{F} \subseteq \mathcal{C}$ with $\mathbf{0} \in \mathcal{F}$, provided such a subset exists.

Lemma 1. A sublevel set $\Omega \subseteq \mathcal{X}$ of an ISS-CLF V can be rendered forward invariant, provided $\|\mathbf{x}\| \geq \rho(\|\mathbf{d}\|)$ for all $\mathbf{x} \in \partial\Omega$ and appropriately restricted $\mathbf{d} \in \mathcal{D}$.

Proof. The condition on the Lyapunov derivative in (5) implies the existence of a state-feedback controller $\mathbf{k} :$

$\mathcal{X} \rightarrow \mathcal{U}$ satisfying $\dot{V}(\mathbf{x}, \mathbf{k}(\mathbf{x}), \mathbf{d}) < 0$ for all $\mathbf{x} \in \partial\Omega$ and appropriately restricted $\mathbf{d} \in \mathcal{D}$. Let $c = V(\mathbf{x})$ for any $\mathbf{x} \in \partial\Omega$. If $V(\mathbf{x}(0)) \in [0, c]$, then $V(\mathbf{x}(t)) \in [0, c]$ for all $t > 0$ by Nagumo's Theorem [22], [1]. Thus, if $\mathbf{x}(0) \in \Omega$, then $\mathbf{x}(t) \in \Omega$ for all $t \geq 0$. \square

III. PROJECTION TO STATE STABILITY

Input to State Stability (ISS) requires a bound on the state in terms of the norm of the disturbance as it appears in the state dynamics (see Definition 5 in Section II). This requirement does not easily permit analysis of Input to State behavior when the disturbance is more easily described by its impact in a Lyapunov function derivative. This limitation motivates Projection to State Stability (PSS), which instead provides a bound on the state in terms of a projection of the disturbance.

Definition 8 (Dynamic Projection). A continuously differentiable function $\mathbf{\Pi} : \mathcal{X} \rightarrow \mathbb{R}^k$ is a *dynamic projection* if there exist $\underline{\sigma}, \bar{\sigma} \in \mathcal{K}_\infty$ satisfying:

$$\underline{\sigma}(\|\mathbf{x}\|) \leq \|\mathbf{\Pi}(\mathbf{x})\| \leq \bar{\sigma}(\|\mathbf{x}\|), \quad (6)$$

for all $\mathbf{x} \in \mathcal{X}$.

Let $\mathcal{Y} = \text{range}(\mathbf{\Pi})$, and let $\mathbf{y} = \mathbf{\Pi}(\mathbf{x})$ for all $\mathbf{x} \in \mathcal{X}$. Consider the system governed by (3). The associated projected system is governed by the dynamics:

$$\dot{\mathbf{y}} = \mathbf{D}_{\mathbf{\Pi}}(\mathbf{x})\mathbf{f}(\mathbf{x}, \mathbf{u}, \mathbf{0}) + \underbrace{\mathbf{D}_{\mathbf{\Pi}}(\mathbf{x})(\mathbf{f}(\mathbf{x}, \mathbf{u}, \mathbf{d}) - \mathbf{f}(\mathbf{x}, \mathbf{u}, \mathbf{0}))}_{\boldsymbol{\delta}}, \quad (7)$$

where $\mathbf{D}_{\mathbf{\Pi}} : \mathcal{X} \rightarrow \mathbb{R}^{k \times n}$ denotes the Jacobian of $\mathbf{\Pi}$, and $\boldsymbol{\delta}$ is implicitly a function of \mathbf{x} , \mathbf{u} , and \mathbf{d} . For the following results, we assume $\boldsymbol{\delta}$ is essentially bounded in time.

We are now ready to state our main definition. The key difference between PSS and ISS (Definition 5) is the use of $\boldsymbol{\delta}$ (7) rather than the native disturbance \mathbf{d} .

Definition 9 (Projection to State Stability). Under a continuous state-feedback controller $\mathbf{k} : \mathcal{X} \rightarrow \mathcal{U}$, a system is *Projection to State Stable* (PSS) with respect to the dynamic projection $\mathbf{\Pi}$ if there exist $\beta \in \mathcal{K}_{\mathcal{L}\infty}$ and $\gamma \in \mathcal{K}_\infty$ such that the solution to (3) satisfies:

$$\|\mathbf{x}(t)\| \leq \beta(\|\mathbf{x}(0)\|, t) + \gamma \left(\sup_{\tau \geq 0} \|\boldsymbol{\delta}(\tau)\| \right), \quad (8)$$

for all $t \geq 0$, with $\boldsymbol{\delta}$ as defined in (7).

Remark 1. If $\mathbf{\Pi}$ is an inclusion map with $k = n$, and the system can be specified as:

$$\mathbf{f}(\mathbf{x}, \mathbf{u}, \mathbf{d}) = \mathbf{f}(\mathbf{x}, \mathbf{u}, \mathbf{0}) + \mathbf{d}, \quad (9)$$

then PSS is equivalent to ISS.

Theorem 1. *Let the dynamics in (3) have an affine dependence on the control input. This corresponding system can be rendered PSS with respect to the dynamic projection $\mathbf{\Pi}$ if the system governed by (7) has an ISS-CLF satisfying the continuous control property.*

Proof. By the assumption that (3) has an affine dependence on the control input, the system governed by (7) also has an affine dependence on the control input. Therefore, a choice of continuous state-feedback controller exists such that the system governed by (7) is ISS with $\boldsymbol{\delta}$ viewed a disturbance. This implies that there exist $\beta \in \mathcal{K}_{\mathcal{L}\infty}$ and $\gamma \in \mathcal{K}_\infty$ such that:

$$\|\mathbf{\Pi}(\mathbf{x}(t))\| \leq \beta(\|\mathbf{\Pi}(\mathbf{x}(0))\|, t) + \gamma \left(\sup_{\tau \geq 0} \|\boldsymbol{\delta}(\tau)\| \right), \quad (10)$$

for all $t \geq 0$. Since $\mathbf{\Pi}$ satisfies (6) we have:

$$\|\mathbf{x}(t)\| \leq \underline{\sigma}^{-1} \left(\beta(\bar{\sigma}(\|\mathbf{x}(0)\|), t) + \gamma \left(\sup_{\tau \geq 0} \|\boldsymbol{\delta}(\tau)\| \right) \right). \quad (11)$$

Finally, define $\beta' \in \mathcal{K}_{\mathcal{L}\infty}$ and $\gamma' \in \mathcal{K}_\infty$ as:

$$\beta'(r, s) = \underline{\sigma}^{-1}(2\beta(\bar{\sigma}(r), s)) \quad (12)$$

$$\gamma'(r) = \underline{\sigma}^{-1}(2\gamma(r)). \quad (13)$$

From the weak form of the triangle inequality presented in [25], [15], it follows that:

$$\|\mathbf{x}(t)\| \leq \beta'(\|\mathbf{x}(0)\|, t) + \gamma' \left(\sup_{\tau \geq 0} \|\boldsymbol{\delta}(\tau)\| \right). \quad (14)$$

\square

We next show that a CLF V for the undisturbed dynamics of a system can be viewed as a projection, thus yielding a ISS-CLF that certifies PSS with respect to V .

Corollary 1. *Suppose $V : \mathcal{X} \rightarrow \mathbb{R}_+$ is a CLF satisfying the continuous control property on \mathcal{X} for the system $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{u}, \mathbf{0})$ and that \mathbf{f} is affine in the control input. Then the disturbed system governed by (3) is PSS with respect to the projection V .*

Proof. With the projection V we have that:

$$\delta = \nabla V(\mathbf{x})^\top (\mathbf{f}(\mathbf{x}, \mathbf{u}, \mathbf{d}) - \mathbf{f}(\mathbf{x}, \mathbf{u}, \mathbf{0})). \quad (15)$$

where $\nabla V : \mathcal{X} \rightarrow \mathbb{R}^n$ is the gradient of the Lyapunov function. The projected system is governed by:

$$\dot{V}(\mathbf{x}, \mathbf{u}, \delta) = \nabla V(\mathbf{x})^\top \mathbf{f}(\mathbf{x}, \mathbf{u}, \mathbf{0}) + \delta, \quad (16)$$

Since V is a CLF satisfying the continuous control property, there exists a continuous state-feedback controller $\mathbf{k} : \mathcal{X} \rightarrow \mathcal{U}$ satisfying:

$$\dot{V}(\mathbf{x}, \mathbf{k}(\mathbf{x}), 0) \leq -(1 - \zeta)\alpha(\|\mathbf{x}\|), \quad (17)$$

for all $\mathbf{x} \in \mathcal{X}$ and $\zeta \in (0, 1)$. Fix ζ and let $\alpha_p, \alpha_q \in \mathcal{K}_\infty$ satisfy $\alpha_p + \alpha_q = (1 - \zeta)\alpha$. Then:

$$\begin{aligned} \dot{V}(\mathbf{x}, \mathbf{k}(\mathbf{x}), \delta) &\leq -(1 - \zeta)\alpha(\|\mathbf{x}\|) + \delta \\ &\leq -\alpha_p(\|\mathbf{x}\|) - \alpha_q(\|\mathbf{x}\|) + |\delta|. \end{aligned} \quad (18)$$

Therefore:

$$\|\mathbf{x}\| \geq \alpha_q^{-1}(|\delta|) \implies \dot{V}(\mathbf{x}, \mathbf{k}(\mathbf{x}), \delta) \leq -\alpha_p(\|\mathbf{x}\|). \quad (19)$$

Since $\bar{\alpha}$ is monotonic, the left-hand side inequality of (19) is equivalent to $\bar{\alpha}(\|\mathbf{x}\|) \geq \bar{\alpha} \circ \alpha_q^{-1}(|\delta|)$, which is satisfied

whenever $V(\mathbf{x}) \geq \bar{\alpha} \circ \alpha_q^{-1}(|\delta|)$ since V is a CLF. By further weakening the implication in (19), we obtain:

$$V(\mathbf{x}) \geq \bar{\alpha} \circ \alpha_q^{-1}(|\delta|) \implies \inf_{\mathbf{u} \in \mathcal{U}} \dot{V}(\mathbf{x}, \mathbf{u}, \delta) \leq -\alpha_p \circ \bar{\alpha}^{-1}(V(\mathbf{x})), \quad (20)$$

noting that $\bar{\alpha} \circ \alpha_q^{-1}$ and $\alpha_p \circ \bar{\alpha}^{-1}$ are class \mathcal{K}_∞ . It follows that the identity map on \mathbb{R}_+ can be used as an ISS-CLF satisfying the continuous control property for the projected dynamics in (16). Therefore, the system (3) is PSS with respect to the projection V by Theorem 1. \square

IV. UNCERTAINTY MODELING & ANALYSIS

In this section we consider a structured form of uncertainty present in affine control systems. We analyze the impact of this uncertainty on a Lyapunov function derivative, and on the PSS behavior of the system.

A. Uncertain Affine Systems

We consider affine control systems of the form:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{u}, \quad (21)$$

with drift dynamics $\mathbf{f} : \mathcal{X} \rightarrow \mathbb{R}^n$ and actuation matrix $\mathbf{g} : \mathcal{X} \rightarrow \mathbb{R}^{n \times m}$. If \mathbf{f} and \mathbf{g} are unknown, we may consider a nominal model of the system:

$$\dot{\mathbf{x}} = \hat{\mathbf{f}}(\mathbf{x}) + \hat{\mathbf{g}}(\mathbf{x})\mathbf{u}, \quad (22)$$

where $\hat{\mathbf{f}} : \mathcal{X} \rightarrow \mathbb{R}^n$ and $\hat{\mathbf{g}} : \mathcal{X} \rightarrow \mathbb{R}^{n \times m}$ are nominal estimates of \mathbf{f} and \mathbf{g} , respectively. In this case, (21) can be expressed as:

$$\dot{\mathbf{x}} = \hat{\mathbf{f}}(\mathbf{x}) + \hat{\mathbf{g}}(\mathbf{x})\mathbf{u} + \underbrace{(\mathbf{g}(\mathbf{x}) - \hat{\mathbf{g}}(\mathbf{x}))\mathbf{u}}_{\mathbf{A}(\mathbf{x})} + \underbrace{\mathbf{f}(\mathbf{x}) - \hat{\mathbf{f}}(\mathbf{x})}_{\mathbf{b}(\mathbf{x})}, \quad (23)$$

obtaining a representation of the dynamics as in (9). Note that the disturbance $\mathbf{d} = \mathbf{A}(\mathbf{x})\mathbf{u} + \mathbf{b}(\mathbf{x})$ is explicitly characterized as time-invariant, state-dependent, and input-dependent, with potentially unknown $\mathbf{A}(\mathbf{x})$ and $\mathbf{b}(\mathbf{x})$ for all $\mathbf{x} \in \mathcal{X}$.

As discussed in [2], [31], CLFs may be constructively formed for affine systems under proper assumptions regarding relative degree and unbounded control. Furthermore, if the true system satisfies the relative degree properties of the estimated model, then the CLF found for the estimated system can be used for the true system.

Assume \mathbf{f} , \mathbf{g} , $\hat{\mathbf{f}}$, and $\hat{\mathbf{g}}$ are Lipschitz continuous (implying \mathbf{A} and \mathbf{b} are Lipschitz continuous). Additionally, assume there exist $\mathbf{u}_0, \hat{\mathbf{u}}_0 \in \mathcal{U}$ such that $\mathbf{f}(\mathbf{0}) + \mathbf{g}(\mathbf{0})\mathbf{u}_0 = \mathbf{0}$ and $\hat{\mathbf{f}}(\mathbf{0}) + \hat{\mathbf{g}}(\mathbf{0})\hat{\mathbf{u}}_0 = \mathbf{0}$. Let V be a CLF candidate for (22). The time derivative of V is given by:

$$\dot{V}(\mathbf{x}, \mathbf{u}, \mathbf{d}) = \underbrace{(\hat{\mathbf{f}}(\mathbf{x}) + \hat{\mathbf{g}}(\mathbf{x})\mathbf{u})^\top \nabla V(\mathbf{x})}_{\hat{V}(\mathbf{x}, \mathbf{u})} + \underbrace{(\mathbf{A}(\mathbf{x})^\top \nabla V(\mathbf{x}))^\top \mathbf{u}}_{\mathbf{a}(\mathbf{x})} + \underbrace{\mathbf{b}(\mathbf{x})^\top \nabla V(\mathbf{x})}_{\mathbf{b}(\mathbf{x})}, \quad (24)$$

for all $\mathbf{x} \in \mathcal{X}$ and $\mathbf{u} \in \mathcal{U}$. As proposed in [31], we may wish to reduce the estimation error $|\dot{V} - \hat{V}|$ by improving

\hat{V} with estimates of \mathbf{a} and \mathbf{b} . Given continuous estimators $\hat{\mathbf{a}} : \mathcal{X} \rightarrow \mathbb{R}^m$ and $\hat{\mathbf{b}} : \mathcal{X} \rightarrow \mathbb{R}$, (24) may be reformulated as:

$$\dot{V}(\mathbf{x}, \mathbf{u}, \mathbf{d}) = \underbrace{(\hat{\mathbf{f}}(\mathbf{x}) + \hat{\mathbf{g}}(\mathbf{x})\mathbf{u})^\top \nabla V(\mathbf{x}) + \hat{\mathbf{a}}(\mathbf{x})^\top \mathbf{u} + \hat{\mathbf{b}}(\mathbf{x})}_{\hat{V}(\mathbf{x}, \mathbf{u})} + \underbrace{(\mathbf{A}(\mathbf{x})^\top \nabla V(\mathbf{x}) - \hat{\mathbf{a}}(\mathbf{x}))^\top \mathbf{u}}_{\mathbf{a}(\mathbf{x})} + \underbrace{\mathbf{b}(\mathbf{x})^\top \nabla V(\mathbf{x}) - \hat{\mathbf{b}}(\mathbf{x})}_{\mathbf{b}(\mathbf{x})}, \quad (25)$$

for all $\mathbf{x} \in \mathcal{X}$ and $\mathbf{u} \in \mathcal{U}$.

Both formulations decompose \dot{V} into an estimated component, \hat{V} , and a residual component. In (24) the residual terms \mathbf{a} and \mathbf{b} capture the effect of the unmodeled dynamics on the Lyapunov function derivative. In (25) the residual terms reflect the error in estimating this effect. Additionally, viewing V as a projection results in $\delta = \mathbf{a}(\mathbf{x})^\top \mathbf{u} + \mathbf{b}(\mathbf{x})$.

B. Projection to State Stability via Uncertainty Functions

If knowledge (e.g., from learning) on what values \mathbf{a} and \mathbf{b} can assume is available, the impact on the Lyapunov derivative can be constrained in a manner permitting PSS analysis of a system. Therefore, we define a function characterizing the possible uncertainties at a given state.

Definition 10 (Uncertainty Function). Let $\mathcal{P}(\mathbb{R}^m \times \mathbb{R})$ denote the set of all subsets of $\mathbb{R}^m \times \mathbb{R}$. An *uncertainty function* for (24) or (25) is a function $\Delta : \mathcal{X} \rightarrow \mathcal{P}(\mathbb{R}^m \times \mathbb{R})$ with $\Delta(\mathbf{x})$ bounded and satisfying $(\mathbf{a}(\mathbf{x}), \mathbf{b}(\mathbf{x})) \in \Delta(\mathbf{x})$ for all $\mathbf{x} \in \mathcal{X}$.

For a given $\mathbf{x} \in \mathcal{X}$, we refer to $\Delta(\mathbf{x})$ as an uncertainty set. Suppose there exists a valid uncertainty function Δ for (24) or (25). Then V satisfies:

$$\dot{V}(\mathbf{x}, \mathbf{u}, \delta) \leq \hat{V}(\mathbf{x}, \mathbf{u}) + \sup_{(\mathbf{a}, \mathbf{b}) \in \Delta(\mathbf{x})} (\mathbf{a}^\top \mathbf{u} + \mathbf{b}), \quad (26)$$

for all $\mathbf{x} \in \mathcal{X}$ and $\mathbf{u} \in \mathcal{U}$. From this point forward we limit our attention to a subset of the state space and make a critical assumption on the estimate \hat{V} for a CLF V :

Assumption 1. Let V be a CLF satisfying the continuous control property for the system governed by (22) on a subset $\mathcal{C} \subseteq \mathcal{X}$ with $\mathbf{0} \in \mathcal{C}$. We assume that:

$$\inf_{\mathbf{u} \in \mathcal{U}} \hat{V}(\mathbf{x}, \mathbf{u}) \leq -\alpha(\|\mathbf{x}\|). \quad (27)$$

for all $\mathbf{x} \in \mathcal{C}$. If \hat{V} is specified as in (24), then this assumption is satisfied by definition. If \hat{V} is specified as in (25), then this assumption states that the addition of the estimators $\hat{\mathbf{a}}$ and $\hat{\mathbf{b}}$ does not make it impossible to choose a control input such that (27) is satisfied.

If the nominal and true system satisfy the same relative degree property, then this assumption amounts to the addition of estimates $\hat{\mathbf{a}}$ and $\hat{\mathbf{b}}$ not violating the relative degree property.

Assumption 2. Let \mathbf{A} and \mathbf{b} be defined as in (23), and let \mathcal{C} be defined as in Assumption 1. We assume \mathbf{A} and \mathbf{b} are bounded on \mathcal{C} .

If \mathcal{C} is compact, this assumption is automatically satisfied since \mathbf{A} and \mathbf{b} are assumed continuous. Under Assumption 1, the set of admissible control inputs $\mathcal{U}(\mathbf{x})$:

$$\mathcal{U}(\mathbf{x}) = \{\mathbf{u} \in \mathcal{U} : \hat{V}(\mathbf{x}, \mathbf{u}) \leq -\alpha(\|\mathbf{x}\|)\}, \quad (28)$$

is non-empty, for all $\mathbf{x} \in \mathcal{C}$. Then the CLF V satisfies:

$$\begin{aligned} \underline{\alpha}(\|\mathbf{x}\|) \leq V(\mathbf{x}) \leq \bar{\alpha}(\|\mathbf{x}\|) \\ \inf_{\mathbf{u} \in \mathcal{U}(\mathbf{x})} \dot{V}(\mathbf{x}, \mathbf{u}, \delta) - \sup_{(\mathbf{a}, b) \in \Delta(\mathbf{x})} (\mathbf{a}^\top \mathbf{u} + b) \leq -\alpha(\|\mathbf{x}\|), \end{aligned} \quad (29)$$

for all $\mathbf{x} \in \mathcal{C}$. We now develop sufficient conditions on the uncertainty function that certifies (23) as PSS with respect to the CLF V (with V interpreted as a projection).

Theorem 2 (*Sufficient Conditions for PSS in Affine Control Systems*). *Consider the system in (23), and a CLF V for (22) with estimated time-derivative \hat{V} as defined in (24) or (25), satisfying Assumption 1. Let Δ be an uncertainty function and let $\mathbf{k} : \mathcal{X} \rightarrow \mathcal{U}$ be a continuous state-feedback controller satisfying $\mathbf{k}(\mathbf{x}) \in \mathcal{U}(\mathbf{x})$ for all $\mathbf{x} \in \mathcal{C}$, with $\mathcal{U}(\mathbf{x})$ and α related as in (28). Suppose there exists $\alpha_p, \alpha_q \in \mathcal{K}_\infty$ with $\alpha_p + \alpha_q = \alpha$ and a sublevel set $\Omega \subseteq \mathcal{C}$ of V satisfying:*

$$\|\mathbf{x}\| \geq \sup_{(\mathbf{a}, b) \in \Delta(\mathbf{x})} \alpha_q^{-1}(\mathbf{a}^\top \mathbf{k}(\mathbf{x}) + b), \quad (30)$$

for all $\mathbf{x} \in \partial\Omega$. Then the system governed by (23) is PSS with respect to the projection V on Ω .

Proof. First, note that:

$$\begin{aligned} \dot{V}(\mathbf{x}, \mathbf{k}(\mathbf{x}), \delta) - \sup_{(\mathbf{a}, b) \in \Delta(\mathbf{x})} (\mathbf{a}^\top \mathbf{k}(\mathbf{x}) + b) &\leq -\alpha(\|\mathbf{x}\|) \\ &= -\alpha_p(\|\mathbf{x}\|) - \alpha_q(\|\mathbf{x}\|), \end{aligned} \quad (31)$$

for all $\mathbf{x} \in \mathcal{C}$. Since (30) holds for all $\mathbf{x} \in \partial\Omega$ and α_q is monotonically increasing, we have:

$$\alpha_q(\|\mathbf{x}\|) \geq \sup_{(\mathbf{a}, b) \in \Delta(\mathbf{x})} (\mathbf{a}^\top \mathbf{k}(\mathbf{x}) + b), \quad (32)$$

for all $\mathbf{x} \in \partial\Omega$. It follows that:

$$\dot{V}(\mathbf{x}, \mathbf{k}(\mathbf{x}), \delta) \leq -\alpha_p(\|\mathbf{x}\|), \quad (33)$$

for all $\mathbf{x} \in \partial\Omega$. This means Ω is forward invariant, with a proof similar to that of Lemma 1. Since Ω is compact and δ is a continuous function of the state in closed-loop, we have that δ is essentially bounded. Thus, since V is a CLF for (22), Corollary 1 can be restricted to Ω ; that is, the system is PSS with respect to the projection V on Ω . \square

We may want to study a particular set of interest \mathcal{E} over which the impact of the uncertainty can be bounded. For $r > 0$, let B_r be the open ball around $\mathbf{0}$ of radius r .

Corollary 2. *Suppose there is a set \mathcal{E} and $\mu \geq 0$ satisfying:*

$$\sup_{(\mathbf{a}, b) \in \Delta(\mathbf{x})} (\mathbf{a}^\top \mathbf{k}(\mathbf{x}) + b) \leq \mu, \quad (34)$$

for all $\mathbf{x} \in \mathcal{E}$. If there exists a sublevel set Ω of V such that:

$$B_{\alpha_q^{-1}(\mu)} \subseteq \Omega \subseteq \mathcal{C} \cap \mathcal{E}, \quad (35)$$

then the system is PSS with respect to the (CLF) projection V on Ω , and the smallest sublevel set of V containing $B_{\alpha_q^{-1}(\mu)}$ is asymptotically stable.

Proof. First, note that:

$$\|\mathbf{x}\| \geq \alpha_q^{-1}(\mu) \geq \sup_{(\mathbf{a}, b) \in \Delta(\mathbf{x})} \alpha_q^{-1}(\mathbf{a}^\top \mathbf{k}(\mathbf{x}) + b), \quad (36)$$

for all $\mathbf{x} \in \partial\Omega$, and the system is PSS on Ω by Theorem 2. Asymptotic stability is achieved as:

$$\|\mathbf{x}\| \geq \alpha_q^{-1}(\mu) \geq \implies \dot{V}(\mathbf{x}, \mathbf{k}(\mathbf{x}), \delta) \leq -\alpha_p(\|\mathbf{x}\|). \quad (37)$$

\square

Improving the uncertainty set (e.g., reducing uncertainty using learning) directly leads to larger sets for a given bound, or tighter bounds on a given set. We state this formally in the next result.

Corollary 3 (*Uncertainty Function Improvement*). *Consider uncertainty functions Δ and Δ' , as well as \mathcal{E} and μ as defined in Corollary 2.*

- Fix $\mu > 0$ and let \mathcal{E}_μ be defined as:

$$\mathcal{E}_\mu = \{\mathbf{x} \in \mathcal{X} : \sup_{(\mathbf{a}, b) \in \Delta(\mathbf{x})} (\mathbf{a}^\top \mathbf{k}(\mathbf{x}) + b) \leq \mu\}. \quad (38)$$

- Fix $\mathcal{E} \subseteq \mathcal{X}$ and let $\mu_\mathcal{E}$ be defined as:

$$\mu_\mathcal{E} = \sup_{\mathbf{x} \in \mathcal{E}} \sup_{(\mathbf{a}, b) \in \Delta(\mathbf{x})} (\mathbf{a}^\top \mathbf{k}(\mathbf{x}) + b). \quad (39)$$

Suppose $\Delta'(\mathbf{x}) \subseteq \Delta(\mathbf{x})$ for all $\mathbf{x} \in \mathcal{X}$. Then the associated set \mathcal{E}'_μ and scalar $\mu'_\mathcal{E}$ satisfy $\mathcal{E}_\mu \subseteq \mathcal{E}'_\mu$ and $\mu'_\mathcal{E} \leq \mu_\mathcal{E}$.

Proof.

$$\sup_{(\mathbf{a}, b) \in \Delta'(\mathbf{x})} (\mathbf{a}^\top \mathbf{k}(\mathbf{x}) + b) \leq \sup_{(\mathbf{a}, b) \in \Delta(\mathbf{x})} (\mathbf{a}^\top \mathbf{k}(\mathbf{x}) + b). \quad (40)$$

\square

C. Uncertainty Function Construction

We now provide a constructive method for creating an uncertainty function from a dataset of state and control values generated by a system. Assume \mathbf{A} and \mathbf{b} are Lipschitz continuous with constants L_A and L_b , respectively. Additionally, assume that \mathbf{A} and \mathbf{b} are bounded on \mathcal{C} by constants $\|\mathbf{A}\|_\infty$ and $\|\mathbf{b}\|_\infty$, respectively. Consider a dataset $D \subseteq (\mathcal{X} \times \mathcal{U}) \times \mathbb{R}$ consisting of data-measurement pairs $((\mathbf{x}, \mathbf{u}), \dot{V}(\mathbf{x}, \mathbf{u}, \delta))$. Such measurements of \dot{V} can be obtained through numerical differentiation of computed values of V . For notational convenience, let $D_0 = \{(\mathbf{x}, \mathbf{u}) : ((\mathbf{x}, \mathbf{u}), \dot{V}(\mathbf{x}, \mathbf{u}, \delta)) \in D\}$.

Proposition 1. Given a dataset D , an uncertainty function Δ can be constructed as:

$$\begin{aligned} \Delta(\mathbf{x}) = \{(\mathbf{a}, b) \in \mathbb{R}^m \times \mathbb{R} : \pm(\mathbf{a}^\top \mathbf{u}' + b) \leq \epsilon(\mathbf{x}, \mathbf{x}', \mathbf{u}') \\ \text{for all } (\mathbf{x}', \mathbf{u}') \in D_0\}, \end{aligned} \quad (41)$$

for all $\mathbf{x} \in \mathcal{X}$, where $\epsilon : \mathcal{X} \times \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}_+$ is continuous.

Remark 2. For all $\mathbf{x} \in \mathcal{X}$, $\Delta(\mathbf{x})$ is a closed, symmetric polyhedron and is bounded given sufficiently diverse control

inputs in the dataset. In this case, $\Delta(\mathbf{x})$ is a compact, convex set. The supremum present in Theorem 2 and Corollary 2 becomes a linear program (LP) and can be efficiently solved.

Proof of Proposition 1. Define observed error as:

$$\ell(\mathbf{x}, \mathbf{u}) = \left| \dot{V}(\mathbf{x}, \mathbf{u}, \delta) - \hat{V}(\mathbf{x}, \mathbf{u}) \right|, \quad (42)$$

for all $(\mathbf{x}, \mathbf{u}) \in D_0$. Consider a test point $(\mathbf{x}, \mathbf{u}) \in \mathcal{X} \times \mathcal{U}$ and a data point $(\mathbf{x}', \mathbf{u}') \in D_0$. Note that $\ell(\mathbf{x}', \mathbf{u}')$ satisfies:

$$\begin{aligned} \ell(\mathbf{x}', \mathbf{u}') &= |\mathbf{a}(\mathbf{x}')^\top \mathbf{u}' + b(\mathbf{x}')| \\ &= |\mathbf{a}(\mathbf{x})^\top \mathbf{u}' + b(\mathbf{x}) + (\mathbf{a}(\mathbf{x}') - \mathbf{a}(\mathbf{x}))^\top \mathbf{u}' \\ &\quad + b(\mathbf{x}') - b(\mathbf{x})| \\ &\geq |\mathbf{a}(\mathbf{x})^\top \mathbf{u}' + b(\mathbf{x})| \\ &\quad - \|\mathbf{a}(\mathbf{x}') - \mathbf{a}(\mathbf{x})\|_2 \|\mathbf{u}'\|_2 - |b(\mathbf{x}') - b(\mathbf{x})|, \end{aligned} \quad (43)$$

where the inequality follows from the reverse triangle inequality, triangle inequality, and Cauchy-Schwarz inequality.

For simplicity we proceed with the construction assuming the estimated Lyapunov function derivative is specified as in (24). The resulting bound will be modified to include estimators as in (25). Note that:

$$\begin{aligned} \|\mathbf{a}(\mathbf{x}') - \mathbf{a}(\mathbf{x})\|_2 &= \|\mathbf{A}(\mathbf{x}')^\top \nabla V(\mathbf{x}') - \mathbf{A}(\mathbf{x})^\top \nabla V(\mathbf{x})\|_2 \\ &= \|(\mathbf{A}(\mathbf{x}') - \mathbf{A}(\mathbf{x}))^\top \nabla V(\mathbf{x}') \\ &\quad + \mathbf{A}(\mathbf{x})^\top (\nabla V(\mathbf{x}') - \nabla V(\mathbf{x}))\|_2 \\ &\leq L_{\mathbf{A}} \|\mathbf{x}' - \mathbf{x}\|_2 \|\nabla V(\mathbf{x}')\|_2 \\ &\quad + \|\mathbf{A}\|_\infty \|\nabla V(\mathbf{x}') - \nabla V(\mathbf{x})\|_2, \end{aligned} \quad (44)$$

where the inequality follows from the triangle inequality, submultiplicativity of matrix norms, and Lipschitz and bounded assumptions for \mathbf{A} . Since it is also true that:

$$\begin{aligned} \|\mathbf{a}(\mathbf{x}') - \mathbf{a}(\mathbf{x})\|_2 &= \|\mathbf{A}(\mathbf{x}')^\top (\nabla V(\mathbf{x}') - \nabla V(\mathbf{x})) \\ &\quad + (\mathbf{A}(\mathbf{x}') - \mathbf{A}(\mathbf{x}))^\top \nabla V(\mathbf{x})\|_2, \end{aligned} \quad (45)$$

then the following bound holds:

$$\begin{aligned} \|\mathbf{a}(\mathbf{x}') - \mathbf{a}(\mathbf{x})\|_2 &\leq L_{\mathbf{A}} \|\mathbf{x}' - \mathbf{x}\|_2 \|\nabla V(\mathbf{x})\|_2 \\ &\quad + \|\mathbf{A}\|_\infty \|\nabla V(\mathbf{x}') - \nabla V(\mathbf{x})\|_2. \end{aligned} \quad (46)$$

Let $\epsilon_L(\mathbf{x}, \mathbf{x}') = \|\mathbf{x} - \mathbf{x}'\| \min \{\|\nabla V(\mathbf{x})\|_2, \|\nabla V(\mathbf{x}')\|_2\}$, and let $\epsilon_\infty(\mathbf{x}, \mathbf{x}') = \|\nabla V(\mathbf{x}) - \nabla V(\mathbf{x}')\|_2$. Observe that ϵ_L and ϵ_∞ are continuous functions. Next, note that:

$$\|\mathbf{a}(\mathbf{x}') - \mathbf{a}(\mathbf{x})\|_2 \leq \epsilon_L(\mathbf{x}, \mathbf{x}') L_{\mathbf{A}} + \epsilon_\infty(\mathbf{x}, \mathbf{x}') \|\mathbf{A}\|_\infty. \quad (47)$$

Similarly,

$$|b(\mathbf{x}') - b(\mathbf{x})| \leq \epsilon_L(\mathbf{x}, \mathbf{x}') L_{\mathbf{b}} + \epsilon_\infty(\mathbf{x}, \mathbf{x}') \|\mathbf{b}\|_\infty. \quad (48)$$

Therefore,

$$\begin{aligned} |\mathbf{a}(\mathbf{x})^\top \mathbf{u}' + b(\mathbf{x})| &\leq \ell(\mathbf{x}', \mathbf{u}') + \epsilon_L(\mathbf{x}, \mathbf{x}') (L_{\mathbf{A}} \|\mathbf{u}'\|_2 + L_{\mathbf{b}}) \\ &\quad + \epsilon_\infty(\mathbf{x}, \mathbf{x}') (\|\mathbf{A}\|_\infty \|\mathbf{u}'\| + \|\mathbf{b}\|_\infty). \end{aligned} \quad (49)$$

While ϵ_L and ϵ_∞ decrease as the test point approaches data points, without estimators as in (25), the observed loss

term ℓ can remain large. By including such estimators, the observed loss term may be reduced, but the bound must be modified with the following additional continuous function:

$$\epsilon_{\mathcal{H}}(\mathbf{x}, \mathbf{x}', \mathbf{u}') = |(\hat{\mathbf{a}}(\mathbf{x}) - \hat{\mathbf{a}}(\mathbf{x}'))^\top \mathbf{u}' + \hat{b}(\mathbf{x}) - \hat{b}(\mathbf{x}')|, \quad (50)$$

which accounts for potential error in the estimation at the test point. ϵ is then specified as the total upper bound. \square

Corollary 4. *The uncertainty set Δ as specified in (41) is continuous with respect to the Hausdorff metric.*

Proof. Note that the inequality constraints in (41) can be rewritten as:

$$\Delta(\mathbf{x}) = \left\{ (\mathbf{a}, b) \in \mathbb{R}^m \times \mathbb{R} : \Xi \begin{bmatrix} \mathbf{a} \\ b \end{bmatrix} \preceq \xi(\mathbf{x}) \right\}, \quad (51)$$

where $\Xi \in \mathbb{R}^{2|D| \times (m+1)}$, $\xi : \mathcal{X} \rightarrow \mathbb{R}^{2|D|}$ and \preceq denotes elementwise inequality. The function ξ is continuous since ϵ is continuous; therefore, the results established in [7] show that the point-to-set map Δ is continuous. \square

V. INTEGRATION WITH LEARNING

We now explore the practical interplay between learning and systematic improvement of PSS properties, in particular by decreasing the upper bound in (49). By decreasing this bound, the uncertainty set in (41) can be made smaller, which in turn can increase the state space region over which PSS properties can be certified and/or achieve reduced degradation (see Corollary 3). As discussed previously, using PSS rather than ISS enables lower-dimensional learning objectives and upper bounds which can be efficiently evaluated during and after learning.

Learning also offers direct ways to decrease the upper bound in (49). As discussed in Section IV, estimators can be used to reduce the observed loss in (42), which appears directly in the upper bound. We can use supervised learning to train such estimators. One complication is that, using baseline controllers, it may not be possible to collect data in regions we wish to certify PSS properties. As the distances between a point of interest and previously collected data grow, ϵ_L and ϵ_∞ can grow larger, weakening uncertainty bounds at the point of interest. By refining the baseline controller using learned models, the system may be controlled towards these regions of interest.

A. Episodic Learning Framework

We demonstrate the practicality of PSS by incorporating it into an episodic learning framework based on learning CLF time derivatives [31]. Controller improvement is achieved by alternating between executing a controller to gather data and refining estimates of residual uncertainty. As data collection and learning progresses, the size of uncertainty sets decreases, enabling stronger PSS certifications for the system.

We briefly describe the DaCLyF (Dataset Aggregation for Control Lyapunov Functions) learning approach from [31].

Algorithm 1 Dataset Aggregation for Control Lyapunov Functions (DaCLyF) [31]

Require: Lyapunov function V , Lyapunov function derivative estimate \hat{V}_0 , model classes \mathcal{H}_a and \mathcal{H}_b , loss function \mathcal{L} , set of initial conditions \mathcal{X}_0 , nominal state-feedback controller \mathbf{u}_0 , number of experiments T , sequence of trust coefficients $0 \leq w_1 \leq \dots \leq w_T \leq 1$

```

D = ∅                                ▷ Initialize dataset
for k = 1, ..., T do
  x0 ← sample( $\mathcal{X}_0$ )                ▷ Sample initial condition
  Dk ← experiment(x0, uk-1)      ▷ Execute experiment
  D ← D ∪ Dk                        ▷ Aggregate dataset
   $\hat{\mathbf{a}}, \hat{\mathbf{b}}$  ← ERM( $\mathcal{H}_a, \mathcal{H}_b, \mathcal{L}, D, \hat{V}_0$ )  ▷ Fit estimators
   $\hat{V}_k$  ←  $\hat{V}_0 + \hat{\mathbf{a}}^\top \mathbf{u} + \hat{\mathbf{b}}$     ▷ Update derivative estimator
  uk ← u0 + wk · augment(u0,  $\hat{V}_k$ ) ▷ Update controller
end for
return D,  $\hat{V}_T$ , uT

```

Let \mathcal{H}_a and \mathcal{H}_b be nonlinear estimator classes, and let \mathcal{H} be the class of estimators of the form:

$$\hat{W}(\mathbf{x}, \mathbf{u}) = \hat{V}(\mathbf{x}, \mathbf{u}) + \hat{\mathbf{a}}(\mathbf{x})^\top \mathbf{u} + \hat{\mathbf{b}}(\mathbf{x}), \quad (52)$$

for all $\mathbf{x} \in \mathcal{X}$, given estimators $\hat{\mathbf{a}} \in \mathcal{H}_a$ and $\hat{\mathbf{b}} \in \mathcal{H}_b$. Equipped with a loss function $\mathcal{L} : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}_+$ and a dataset $D = \{((\mathbf{x}_i, \mathbf{u}_i), \dot{V}_i)\}$ obtained from experiments with a baseline controller, we approximately solve the Empirical Risk Minimization (ERM) problem over the class \mathcal{H} to update the estimate of the Lyapunov function derivative. Then, the controller is updated with an augmenting controller of the form:

$$\begin{aligned} \mathbf{u}'(\mathbf{x}) = \arg \min_{\mathbf{u}' \in \mathbb{R}^m} & \frac{1}{2} \begin{bmatrix} \mathbf{u}(\mathbf{x}) \\ \mathbf{u}' \end{bmatrix}^\top \mathbf{P} \begin{bmatrix} \mathbf{u}(\mathbf{x}) \\ \mathbf{u}' \end{bmatrix} + \mathbf{q}^\top \begin{bmatrix} \mathbf{u}(\mathbf{x}) \\ \mathbf{u}' \end{bmatrix} + r \\ \text{s.t.} & \hat{V}(\mathbf{x}, \mathbf{u}(\mathbf{x}) + \mathbf{u}') \leq -\alpha(\|\mathbf{x}\|) \\ & \mathbf{u}(\mathbf{x}) + \mathbf{u}' \in \mathcal{U}, \end{aligned} \quad (53)$$

for all $\mathbf{x} \in \mathcal{C}$, with $\mathbf{P} \in \mathbb{S}_+^{2m}$, $\mathbf{q} \in \mathbb{R}^{2m}$, $r > 0$, and \hat{V} denoting the updated estimator. Here \mathbb{S}_+^{2m} denotes the set of positive semidefinite matrices of size $2m \times 2m$. The augmenting controller, weighted by a trust factor, is additively incorporated with the baseline controller. As more data is collected, the trust factor is increased. This entire procedure is outlined in Algorithm 1.

The estimator class \mathcal{H} can include a wide variety of nonlinear functions. Should \mathcal{H}_a and \mathcal{H}_b be classes of Lipschitz continuous estimators, the upper bound (50) can be weakened further using the associated Lipschitz constants to permit further analysis of the uncertainty function specified in (41). Importantly, this motivates the use of spectrally normalized deep neural networks [20], [6] for this estimation problem.

B. Simulation Results

In this section we apply Algorithm 1 to an inverted pendulum model with parametric uncertainty. The pendulum

is modeled as a massless rod with torque input at a fixed base. The true mass and the length are perturbed by up to 30% of their nominal values. The baseline controller is a linear proportional derivative (PD) controller to track angle and angle rate trajectories. The estimators are chosen from the class of two layer neural networks with 200 hidden units and ReLU nonlinearities, mapping concatenated state and Lyapunov function gradients to \mathbb{R}^m and \mathbb{R} . The trust factors are chosen in a sigmoid fashion. Naive exploratory control is introduced as in [31], with perturbations chosen uniformly at random, independently in each coordinate, and scaled by 25% of the norm of the current control input.

A comparison of the baseline controller and final augmented controller demonstrating improved tracking performance is shown in Fig. 1. A comparison of PSS bounds for the model-based QP controller and the final augmented controller is shown in Fig. 2, with observed trajectories superimposed. The QP controller is unable to keep the system in regions in which the bound is shown to be small. On the other hand, the augmented controller keeps the system close to the desired trajectory, consistently near training data. The bounds are small along the observed trajectory, in comparison.




Fig. 1. Comparison of tracking performance for PD controller and final augmented controller. The final augmented controller tracks the desired angle trajectory more effectively.

VI. CONCLUSION

We presented a novel low-dimensional view of stability for uncertain systems and a method of evaluating PSS behavior using experimental data. This method constructs a bound on disturbances to a CLF derivative, and can be integrated with a machine learning framework to improve PSS behavior. Finally, we validate this procedure on a simulated system.

Future work includes incorporating the upper bounds into online learning settings and developing optimal exploration strategies. Quantifying the impact of learning on PSS provides an objective for deciding how to collect data, also known as the exploration problem in learning literature [21], [5], [9], [8], [30]. In particular, reductions of the uncertainty bound may be used to formulate regret in online learning settings or reward in imitation and reinforcement learning settings. Additional future work includes extending the notion of PSS to Control Safety and Barrier Functions,




Fig. 2. Comparison of disturbance upper bounds with model based QP controller (left) and final augmented controller (right). The disturbance upper bound is computed from the maximum over uncertainty sets in (41) for both controllers, and observed trajectories are displayed with dashed lines. The augmenting controller keeps the system in regions with lower disturbance bounds while the system leaves the region around the training data under the QP controller. The maps were generated by sampling states randomly about training data points and evaluating the upper bound for each sampled state. The results were then discretized for ease of visualization. Each bin is colored by the maximum disturbance observed in the bin.

more thoroughly studying the benefits of learning low-dimensional representations of the dynamics versus the full-order dynamics, and utilizing PSS to augment controller synthesis for complex real-world robotic systems.

REFERENCES

- [1] R. Abraham, J. Marsden, and T. Ratiu. *Manifolds, tensor analysis, and applications*, volume 75. Springer S&B Media, 2012.
- [2] A. Ames, K. Galloway, K. Sreenath, and J. Grizzle. Rapidly exponentially stabilizing control lyapunov functions and hybrid zero dynamics. *IEEE Transactions on Automatic Control*, 59(4):876–891, 2014.
- [3] A. Ames and M. Powell. Towards the unification of locomotion and manipulation through control lyapunov functions and quadratic programs. In *Control of Cyber-Physical Systems*, pages 219–240. Springer, 2013.
- [4] Z. Artstein. Stabilization with relaxed controls. *Nonlinear Analysis: Theory, Methods & Applications*, 7(11):1163–1173, 1983.
- [5] A. Aswani, H. Gonzalez, S. Sastry, and C. Tomlin. Provably safe and robust learning-based model predictive control. *Automatica*, 49(5):1216–1226, 2013.
- [6] P. Bartlett, D. Foster, and M. Telgarsky. Spectrally-normalized margin bounds for neural networks. In *Advances in Neural Information Processing Systems*, pages 6240–6249, 2017.
- [7] R. Batson. Combinatorial behavior of extreme points of perturbed polyhedra. *Journal of mathematical analysis and applications*, 127(1):130–139, 1987.
- [8] F. Berkenkamp, R. Moriconi, A. Schoellig, and A. Krause. Safe learning of regions of attraction for uncertain, nonlinear systems with gaussian processes. In *2016 IEEE 55th Conference on Decision and Control (CDC)*, pages 4661–4666. IEEE, 2016.
- [9] F. Berkenkamp, A. Schoellig, and A. Krause. Safe controller optimization for quadrotors with gaussian processes. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 491–496. IEEE, 2016.
- [10] G. Chowdhary, H. Kingravi, J. How, and P. Vela. Bayesian nonparametric adaptive control using gaussian processes. *IEEE transactions on neural networks and learning systems*, 26(3):537–550, 2015.
- [11] S. Dean, H. Mania, N. Matni, B. Recht, and S. Tu. Regret bounds for robust adaptive control of the linear quadratic regulator. In *Advances in Neural Information Processing Systems*, pages 4192–4201, 2018.
- [12] R. Freeman and P. Kokotovic. Inverse optimality in robust stabilization. *SIAM journal on control and optimization*, 34(4):1365–1391, 1996.
- [13] R. Freeman and P. Kokotovic. *Robust nonlinear control design: state-space and Lyapunov techniques*. Springer S. & B. Media, 2008.
- [14] K. Galloway, K. Sreenath, A. Ames, and J. Grizzle. Torque saturation in bipedal robotic walking through control lyapunov function-based quadratic programs. *IEEE Access*, 3:323–332, 2015.
- [15] C. Kellett. A compendium of comparison function results. *Mathematics of Control, Signals, and Systems*, 26(3):339–374, 2014.
- [16] H.K. Khalil. *Nonlinear Systems - 3rd Edition*. PH, Upper Saddle River, NJ, 2002.
- [17] M. Krstić and P. Kokotović. Control lyapunov functions for adaptive nonlinear stabilization. *Systems & Control Letters*, 26(1):17–23, 1995.
- [18] Y. Lin and E. Sontag. A universal formula for stabilization with bounded controls. *Systems & Control Letters*, 16(6):393–397, 1991.
- [19] W-L. Ma, S. Kolathaya, E. Ambrose, C. Hubicki, and A. Ames. Bipedal robotic running with durus-2d: Bridging the gap between theory and experiment. In *Proc. of the 20th International Conf. on Hybrid Systems: CC*, pages 265–274. ACM, 2017.
- [20] T. Miyato, T. Kataoka, M. Koyama, and Y. Yoshida. Spectral normalization for generative adversarial networks. *arXiv preprint arXiv:1802.05957*, 2018.
- [21] T. Moldovan and P. Abbeel. Safe exploration in markov decision processes. *arXiv preprint arXiv:1205.4810*, 2012.
- [22] M. Nagumo. Über die lage der integralkurven gewöhnlicher differentialgleichungen. *Proceedings of the Physico-Mathematical Society of Japan. 3rd Series*, 24:551–559, 1942.
- [23] Q. Nguyen and K. Sreenath. Optimal robust control for bipedal robots through control lyapunov function based quadratic programs. In *Robotics: Science and Systems*, 2015.
- [24] G. Shi, X. Shi, M. O’Connell, R. Yu, K. Azizzadenesheli, A. Anandkumar, Y. Yue, and S-J. Chung. Neural lander: Stable drone landing control using learned dynamics. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2019.
- [25] E. Sontag. Smooth stabilization implies coprime factorization. *IEEE transactions on automatic control*, 34(4):435–443, 1989.
- [26] E. Sontag. A universal construction of artstein’s theorem on nonlinear stabilization. *Systems & control letters*, 13(2):117–123, 1989.
- [27] E. Sontag. Input to state stability: Basic concepts and results. In *Nonlinear and optimal control theory*, pages 163–220. Springer, 2008.
- [28] E. Sontag and Y. Wang. On characterizations of input-to-state stability with respect to compact sets. In *Nonlinear Control Systems Design 1995*, pages 203–208. Elsevier, 1995.
- [29] E. Sontag and Y. Wang. On characterizations of the input-to-state stability property. *Systems & Control Letters*, 24(5):351–359, 1995.
- [30] Y. Sui, V. Zhuang, J. Burdick, and Y. Yue. Stagewise safe bayesian optimization with gaussian processes. In *International Conference on Machine Learning (ICML)*, 2018.
- [31] A. Taylor, V. Dorobantu, H. Le, Y. Yue, and A. Ames. Episodic learning with control lyapunov functions for uncertain robotic systems. In *IROS 2019*, 2019.
- [32] V. Vapnik. An overview of statistical learning theory. *IEEE transactions on neural networks*, 10(5):988–999, 1999.
- [33] Y. Wang. A converse lyapunov theorem with applications to iss-disturbance attenuation. In *Proc. 13th IFAC World Congress*, pages 79–84, 1996.
- [34] D-X. Zhou. The covering number in learning theory. *Journal of Complexity*, 18(3):739–767, 2002.
- [35] K. Zhou, J. Doyle, K. Glover, et al. *Robust and optimal control*, volume 40. Prentice hall New Jersey, 1996.