

A Control Barrier Perspective on Episodic Learning via Projection-to-State Safety

Andrew J. Taylor[✉], Graduate Student Member, IEEE, Andrew Singletary[✉],
 Yisong Yue[✉], and Aaron D. Ames[✉], Senior Member, IEEE

Abstract—In this letter we seek to quantify the ability of learning to improve safety guarantees endowed by Control Barrier Functions (CBFs). In particular, we investigate how model uncertainty in the time derivative of a CBF can be reduced via learning, and how this leads to stronger statements on the safe behavior of a system. To this end, we build upon the idea of Input-to-State Safety (ISSf) to define Projection-to-State Safety (PSSf), which characterizes degradation in safety in terms of a projected disturbance. This enables the direct quantification of both how learning can improve safety guarantees, and how bounds on learning error translate to bounds on degradation in safety. We demonstrate that a practical episodic learning approach can use PSSf to reduce uncertainty and improve safety guarantees in simulation and experimentally.

Index Terms—Machine learning, Lyapunov methods, uncertain systems.

I. INTRODUCTION

ENSURING safety is of significant importance in the design of many modern control systems, from autonomous driving to industrial robotics. In practice, the models used in the control design process are imperfect, with model uncertainty arising due to parametric error and unmodeled dynamics. This uncertainty can cause the controller to render the system unsafe. As such, it is necessary to quantify how the desired safety properties degrade with uncertainty.

Control Barrier Functions (CBFs) have become increasingly popular [2], [15], [22] as a tool for synthesizing controllers that provide safety via *set invariance* [5]. Safety guarantees endowed by a controller synthesized via CBFs rely on an accurate model of a system’s dynamics, and may degrade

Manuscript received March 17, 2020; revised June 8, 2020; accepted June 29, 2020. Date of publication July 13, 2020; date of current version July 28, 2020. This work was supported by Defense Advanced Research Projects Agency under Award HR00111890035 and Award NNN12AA01C. Recommended by Senior Editor G. Cherubini. (Corresponding author: Andrew J. Taylor.)

Andrew J. Taylor and Yisong Yue are with the Department of Computing and Mathematical Sciences, California Institute of Technology, Pasadena, CA 91125 USA (e-mail: ajtaylor@caltech.edu; yyue@caltech.edu).

Andrew Singletary and Aaron D. Ames are with the Department of Mechanical and Civil Engineering, California Institute of Technology, Pasadena, CA 91125 USA (e-mail: asinglet@caltech.edu; ames@caltech.edu).

Digital Object Identifier 10.1109/LCSYS.2020.3009082

in the presence of model uncertainty. The recently proposed definition of *Input-to-State Safety* (ISSf) provides a tool for quantifying the impact on safety guarantees of such uncertainty or disturbances in the dynamics [13] by describing changes in the set kept invariant.

Due to its flexibility, it is increasingly popular to incorporate learning into safe controller synthesis [7], [8], [16], [23]. Many of these approaches seek to provide statistical guarantees on the safety via assumptions made on learning performance. In practice however, limitations on learning performance arise due to factors such as covariate shift [6], [14], limitations on model capacity, and optimization error. Thus, it is critical to understand the relationship between learning error and what safety guarantees can be ensured.

In this letter, we study how introducing learning models into safe controller synthesis done via CBFs can improve safety guarantees, and what safety guarantees can be made in the presence of learning error. In particular, we consider the episodic learning approach proposed in [21], with learning done directly on the time derivative of a CBF. We integrate this approach with Input-to-State Safety to show how learning can lead to improved safety guarantees and provide a relationship between learning error and degradation of safety guarantees.

The methods in this letter play a distinct role in the context of safe control in the presence of uncertainty. Unlike previous work using learning with Control Lyapunov Functions (CLFs) for stability [19], [20], CBFs are less conservative as they endow safety. This allows them to augment other potentially unsafe controller designs pursuing specific objectives and render them safe [10]. Additionally, these learning methods are complementary to existing methods for CBF synthesis [18], allowing them to quantify the impact of uncertainty on the product of these methods. Adaptive methods for safe control are restricted to structured, parametric uncertainty [4], but can be integrated with the results of this letter to understand how unstructured uncertainty impacts the safety of the system.

We make two main contributions in this letter. First, inspired by the idea of Projection-to-State Stability proposed in [19], we formulate general definitions of projections and projection compatible functions. Care must be taken to ensure these definitions preserve important topological properties for safety such as safe set membership. These definitions not only capture the definitions established in [19] as a special case, but

allow us to define the notion of *Projection-to-State Safety* (PSSf), which is a variant of the Input-to-State Safety property. Like ISSf, PSSf provides a tool for characterizing the degradation of safety in the presence of disturbances. Unlike ISSf, PSSf considers disturbances in a projected environment, allowing stronger guarantees on safe behavior. Second, we demonstrate the utility of PSSf by characterizing how data-driven learning models can improve safety guarantees, and how learning error leads to degradation in safety guarantees.

II. PRELIMINARIES

This section provides a review of Control Barrier Functions (CBFs) and Input-to-State Safe Control Barrier Functions (ISSf-CBFs). These tools will be used in Section III to define the notion of Projection-to-State Safety.

Consider the 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 . Given a 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 $\mathbf{x}_0 := \mathbf{x}(0) \in \mathbb{R}^n$ there exists a time interval $I(\mathbf{x}_0) = [0, t_{\max}]$ such that $\mathbf{x}(t)$ is the unique solution to (2) on $I(\mathbf{x}_0)$ [17]. In the case that \mathbf{f}_{cl} is forward complete, $t_{\max} = \infty$.

A continuous function $\alpha : [0, a) \rightarrow \mathbb{R}_+$, with $a > 0$, is said to belong to *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 said to belong to *class* \mathcal{K}_∞ ($\alpha \in \mathcal{K}_\infty$). A continuous function $\alpha : (-b, a) \rightarrow \mathbb{R}$, with $a, b > 0$, is said to belong to *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 said to belong to *extended class* \mathcal{K}_∞ ($\alpha \in \mathcal{K}_{\infty,e}$).

The notion of safety that we consider is formalized by specifying a *safe set* in the state space that the system must remain in to be considered safe. In particular, consider a set $\mathcal{C} \subset \mathbb{R}^n$ defined as the 0-superlevel set of a continuously differentiable function $h : \mathbb{R}^n \rightarrow \mathbb{R}$, yielding:

$$\mathcal{C} \triangleq \{\mathbf{x} \in \mathbb{R}^n : h(\mathbf{x}) \geq 0\}, \quad (3)$$

$$\partial\mathcal{C} \triangleq \{\mathbf{x} \in \mathbb{R}^n : h(\mathbf{x}) = 0\}, \quad (4)$$

$$\text{Int}(\mathcal{C}) \triangleq \{\mathbf{x} \in \mathbb{R}^n : h(\mathbf{x}) > 0\}. \quad (5)$$

We assume that \mathcal{C} is nonempty and has no isolated points, that is, $\text{Int}(\mathcal{C}) \neq \emptyset$ and $\text{Int}(\mathcal{C}) = \mathcal{C}$. We refer to \mathcal{C} as the *safe set*. This construction motivates the following definitions of forward invariant and safety:

Definition 1 (Forward Invariant & Safety): A set $\mathcal{C} \subset \mathbb{R}^n$ is *forward invariant* if for every $\mathbf{x}_0 \in \mathcal{C}$, the solution $\mathbf{x}(t)$ to (2) satisfies $\mathbf{x}(t) \in \mathcal{C}$ for all $t \in I(\mathbf{x}_0)$. The system (2) is *safe* on the set \mathcal{C} if the set \mathcal{C} is forward invariant.

Certifying the safety of the closed-loop system (2) with respect to a set \mathcal{C} may be impossible if the controller \mathbf{k}

was not chosen to enforce the safety of \mathcal{C} . Control Barrier Functions can serve as a synthesis tool for attaining the forward invariance, and thus the safety of a set:

Definition 2 (Control Barrier Function (CBF), [3]): Let $\mathcal{C} \subset \mathbb{R}^n$ be the 0-superlevel set of a continuously differentiable function $h : \mathbb{R}^n \rightarrow \mathbb{R}$ with 0 a regular value. The function h is a *Control Barrier Function* (CBF) for (1) on \mathcal{C} if there exists $\alpha \in \mathcal{K}_{\infty,e}$ such that for all $\mathbf{x} \in \mathbb{R}^n$:

$$\sup_{\mathbf{u} \in \mathbb{R}^m} \dot{h}(\mathbf{x}, \mathbf{u}) \triangleq \frac{\partial h}{\partial \mathbf{x}}(\mathbf{x})(\mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{u}) > -\alpha(h(\mathbf{x})). \quad (6)$$

We note that this definition can be relaxed such that the inequality only holds for all $\mathbf{x} \in E$ where E is an open set satisfying $\mathcal{C} \subset E \subset \mathbb{R}^n$. Given a CBF h for (1) and a corresponding $\alpha \in \mathcal{K}_{\infty,e}$, we can consider the point-wise set of all control values that satisfy (6):

$$K_{\text{cbf}}(\mathbf{x}) \triangleq \{\mathbf{u} \in \mathbb{R}^m \mid \dot{h}(\mathbf{x}, \mathbf{u}) \geq -\alpha(h(\mathbf{x}))\}.$$

One of the main results in [1], [24] relates controllers taking values in $K_{\text{cbf}}(\mathbf{x})$ to the safety of (1) on \mathcal{C} .

Theorem 1: Given a set $\mathcal{C} \subset \mathbb{R}^n$ defined as the 0-superlevel set of a continuously differentiable function $h : \mathbb{R}^n \rightarrow \mathbb{R}$, if h is a CBF for (1) on \mathcal{C} , then any Lipschitz continuous controller $\mathbf{k} : \mathbb{R}^n \rightarrow \mathbb{R}^m$, such that $\mathbf{k}(\mathbf{x}) \in K_{\text{cbf}}(\mathbf{x})$ for all $\mathbf{x} \in \mathbb{R}^n$, renders the system (1) safe with respect to the set \mathcal{C} .

Assuming $\frac{\partial h}{\partial \mathbf{x}}$ and α are Lipschitz continuous, the following optimization-based controller is Lipschitz continuous and enforces $\mathbf{k}(\mathbf{x}) \in K_{\text{cbf}}(\mathbf{x})$ for all $\mathbf{x} \in \mathbb{R}^n$ [1], [11]:

$$\begin{aligned} \mathbf{k}(\mathbf{x}) = \operatorname{argmin}_{\mathbf{u} \in \mathbb{R}^m} \quad & \frac{1}{2} \|\mathbf{u}\|_2^2 \\ \text{s.t.} \quad & \frac{\partial h}{\partial \mathbf{x}}(\mathbf{x})(\mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{u}) \geq -\alpha(h(\mathbf{x})) \end{aligned} \quad (\text{CBF-QP})$$

To accommodate disturbances or model uncertainties, we consider a disturbance space $\mathcal{D} \subset \mathbb{R}^n$ and a disturbed system:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{u} + \mathbf{d}. \quad (7)$$

with $\mathbf{d} \in \mathcal{D}$. The disturbance may be time-varying, state and/or input dependent. Viewing \mathbf{d} as a signal, $\mathbf{d}(t)$, we assume it is essentially bounded in time, and define $\|\mathbf{d}\|_\infty \triangleq \text{ess sup}_{t \geq 0} \|\mathbf{d}(t)\|$. A Lipschitz continuous state-feedback controller \mathbf{k} yields the closed-loop dynamics:

$$\dot{\mathbf{x}} = \mathbf{f}_{\text{cl}}(\mathbf{x}, \mathbf{d}) \triangleq \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{k}(\mathbf{x}) + \mathbf{d}. \quad (8)$$

In the presence of disturbances, a controller \mathbf{k} synthesized to render the set \mathcal{C} safe for the undisturbed dynamics (2) may fail to render \mathcal{C} safe for the disturbed dynamics (8). To quantify how safety degrades, we consider *input-to-state safety* [13].

Definition 3 (Input-to-State Safety (ISSf)): The closed-loop system (8) is *input-to-state safe* (ISSf) on a set $\mathcal{C} \subset \mathbb{R}^n$ with respect to disturbances \mathbf{d} if there exists $\bar{d} > 0$ and $\gamma \in \mathcal{K}_\infty$ such that the set $\mathcal{C}_\mathbf{d} \supset \mathcal{C}$ defined as:

$$\mathcal{C}_\mathbf{d} \triangleq \{\mathbf{x} \in \mathbb{R}^n : h(\mathbf{x}) + \gamma(\|\mathbf{d}\|_\infty) \geq 0\}, \quad (9)$$

$$\partial\mathcal{C}_\mathbf{d} \triangleq \{\mathbf{x} \in \mathbb{R}^n : h(\mathbf{x}) + \gamma(\|\mathbf{d}\|_\infty) = 0\}, \quad (10)$$

$$\text{Int}(\mathcal{C}_\mathbf{d}) \triangleq \{\mathbf{x} \in \mathbb{R}^n : h(\mathbf{x}) + \gamma(\|\mathbf{d}\|_\infty) > 0\}, \quad (11)$$

is forward invariant for all \mathbf{d} satisfying $\|\mathbf{d}\|_\infty \leq \bar{d}$.

We refer to \mathcal{C} as an *input-to-state safe* set (ISSf set) if such a set \mathcal{C}_d exists. This definition implies that though the set \mathcal{C} may not be safe, a larger set \mathcal{C}_d , depending on \mathbf{d} , is safe. \mathcal{C} can be certified as an ISSf set for the closed-loop system (8) with the following definition:

Definition 4 (Input-to-State Safe Barrier Function (ISSf-BF)): Let $\mathcal{C} \subset \mathbb{R}^n$ be the 0-superlevel set of a continuously differentiable function $h : \mathbb{R}^n \rightarrow \mathbb{R}$ with 0 a regular value. The function $h : \mathbb{R}^n \rightarrow \mathbb{R}$ is an *Input-to-State Safe Barrier Function* (ISSf-BF) for (8) on \mathcal{C} if there exists $\bar{d} > 0$, $\alpha \in \mathcal{K}_{\infty,e}$, and $\iota \in \mathcal{K}_\infty$ such that:

$$\frac{\partial h}{\partial \mathbf{x}}(\mathbf{x})(\mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{k}(\mathbf{x}) + \mathbf{d}) \geq -\alpha(h(\mathbf{x})) - \iota(\|\mathbf{d}\|), \quad (12)$$

for all $\mathbf{x} \in \mathbb{R}^n$ and $\mathbf{d} \in \mathbb{R}^n$ such that $\|\mathbf{d}\| \leq \bar{d}$.

As shown in [13], the existence of an ISSf-BF for (8) on \mathcal{C} implies \mathcal{C} is an ISSf set. Similarly to the undisturbed case, we can introduce the notion of a Control Barrier Function for synthesizing controllers that ensure input-to-state safety.

Definition 5 (ISSf Control Barrier Function (ISSf-CBF)): Let $\mathcal{C} \subset \mathbb{R}^n$ be the 0-superlevel set of a continuously differentiable function $h : \mathbb{R}^n \rightarrow \mathbb{R}$ with 0 a regular value. The function h is an *Input-to-State Safe Control Barrier Function* (ISSf-CBF) for (7) on \mathcal{C} if there exists $\bar{d} > 0$, $\alpha \in \mathcal{K}_{\infty,e}$, and $\iota \in \mathcal{K}_\infty$ such that:

$$\begin{aligned} \sup_{\mathbf{u} \in \mathbb{R}^m} \dot{h}(\mathbf{x}, \mathbf{u}, \mathbf{d}) &\triangleq \frac{\partial h}{\partial \mathbf{x}}(\mathbf{x})(\mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{u} + \mathbf{d}) \\ &> -\alpha(h(\mathbf{x})) - \iota(\|\mathbf{d}\|), \end{aligned} \quad (13)$$

for all $\mathbf{x} \in \mathbb{R}^n$ and $\mathbf{d} \in \mathbb{R}^n$ satisfying $\|\mathbf{d}\| \leq \bar{d}$.

We note that this definition is a more general definition of an ISSf-CBF compared to [13], where disturbances enter the system with the inputs. We define the pointwise set:

$$K_{\text{issf}}(\mathbf{x}) \triangleq \{\mathbf{u} \in \mathbb{R}^m \mid \dot{h}(\mathbf{x}, \mathbf{u}, \mathbf{d}) \geq -\alpha(h(\mathbf{x})) - \iota(\|\mathbf{d}\|)\},$$

noting that for a fixed input the inequality must hold for all $\mathbf{d} \in \mathbb{R}^n$ satisfying $\|\mathbf{d}\| \leq \bar{d}$. Given this result, we have the following theorem.

Theorem 2: Given a set $\mathcal{C} \subset \mathbb{R}^n$ defined as the 0-superlevel set of a continuously differentiable function $h : \mathbb{R}^n \rightarrow \mathbb{R}$, if h is an ISSf-CBF for (7) on \mathcal{C} , then any Lipschitz continuous controller $\mathbf{k} : \mathbb{R}^n \rightarrow \mathbb{R}^m$, such that $\mathbf{k}(\mathbf{x}) \in K_{\text{issf}}(\mathbf{x})$ for all $\mathbf{x} \in \mathbb{R}^n$, renders the set \mathcal{C} ISSf for (8).

This theorem follows from the fact that under the controller \mathbf{k} , h serves an ISSf-BF for (8) on \mathcal{C} .

III. PROJECTION-TO-STATE SAFETY

Input-to-State Safety describes how the safe set \mathcal{C} changes in terms of the disturbance as it appears in the state dynamics (see Definition 3 in Section II). This description does not permit analysis of how safety degrades when the disturbance is more easily characterized in a Barrier Function derivative. This limitation motivates Projection-to-State Safety (PSSf), which characterizes safety in terms of a projected disturbance.

We refer to a continuously differentiable function $\Pi : \mathbb{R}^n \rightarrow \mathbb{R}^k$ as a *projection*, and denote $\mathbf{y} = \Pi(\mathbf{x})$. Considering the

system governed by (7), the associated projected system is governed by the dynamics:

$$\dot{\mathbf{y}} = \mathbf{D}_\Pi(\mathbf{x})(\mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{u}) + \mathbf{D}_\Pi(\mathbf{x})\mathbf{d}, \quad (14)$$

where $\mathbf{D}_\Pi : \mathbb{R}^n \rightarrow \mathbb{R}^{k \times n}$ denotes the Jacobian of Π . As will be seen when quantifying the impact of model uncertainty in Section IV, if the disturbance can be partially characterized in terms of the state and input, we may rewrite the projected dynamics as:

$$\dot{\mathbf{y}} = \mathbf{f}_y(\mathbf{x}) + \mathbf{g}_y(\mathbf{x})\mathbf{u} + \boldsymbol{\delta}, \quad (15)$$

where $\mathbf{f}_y : \mathbb{R}^n \rightarrow \mathbb{R}^k$ and $\mathbf{g}_y : \mathbb{R}^n \rightarrow \mathbb{R}^{k \times m}$ are Lipschitz continuous on \mathbb{R}^n , and $\boldsymbol{\delta} \in \mathbb{R}^k$ is referred to as the *projected disturbance*. We note it is not explicitly necessary that the relationships $\mathbf{f}_y(\mathbf{x}) = \mathbf{D}_\Pi(\mathbf{x})\mathbf{f}(\mathbf{x})$, $\mathbf{g}_y(\mathbf{x}) = \mathbf{D}_\Pi(\mathbf{x})\mathbf{g}(\mathbf{x})$, and $\boldsymbol{\delta} = \mathbf{D}_\Pi(\mathbf{x})\mathbf{d}$ hold, but are one possible relationship between (14) and (15). For the following results, we assume that $\boldsymbol{\delta}$ is essentially bounded in time and define $\|\boldsymbol{\delta}\|_\infty \triangleq \text{ess sup}_{t \geq 0} \|\boldsymbol{\delta}(t)\|$. We relate behavior of the projected system to the original system via the following definition.

Definition 6 (Projection-to-State Safety): The closed-loop system (8) is *projection-to-state safe* (PSSf) on \mathcal{C} with respect to the projection Π and projected disturbances $\boldsymbol{\delta}$ if there exists $\bar{\delta} > 0$ and $\gamma \in \mathcal{K}_\infty$ such that the set $\mathcal{C}_\delta \supset \mathcal{C}$,

$$\mathcal{C}_\delta \triangleq \{\mathbf{x} \in \mathbb{R}^n : h(\mathbf{x}) + \gamma(\|\boldsymbol{\delta}\|_\infty) \geq 0\}, \quad (16)$$

$$\partial\mathcal{C}_\delta \triangleq \{\mathbf{x} \in \mathbb{R}^n : h(\mathbf{x}) + \gamma(\|\boldsymbol{\delta}\|_\infty) = 0\}, \quad (17)$$

$$\text{Int}(\mathcal{C}_\delta) \triangleq \{\mathbf{x} \in \mathbb{R}^n : h(\mathbf{x}) + \gamma(\|\boldsymbol{\delta}\|_\infty) > 0\}, \quad (18)$$

is forward invariant for all $\boldsymbol{\delta}$ satisfying $\|\boldsymbol{\delta}\|_\infty \leq \bar{\delta}$.

In contrast to the definition of ISSf which enlarges the safe set in terms of the disturbance \mathbf{d} , PSSf quantifies how the safe set enlarges in terms of the projected disturbance $\boldsymbol{\delta}$. To utilize safety guarantees implied by ISSf-CBFs for analyzing PSSf behavior, we require the following definition:

Definition 7 (Compatible Projection): A function $h_\Pi : \mathbb{R}^k \rightarrow \mathbb{R}$ is said to be a *compatible projection* for the function $h : \mathbb{R}^n \rightarrow \mathbb{R}$ with respect to the projection $\Pi : \mathbb{R}^n \rightarrow \mathbb{R}^k$ if there exists $\underline{\sigma}, \bar{\sigma} \in \mathcal{K}_{\infty,e}$ such that for all $\mathbf{x} \in \mathbb{R}^n$:

$$\underline{\sigma}(h(\mathbf{x})) \leq h_\Pi(\Pi(\mathbf{x})) \leq \bar{\sigma}(h(\mathbf{x})). \quad (19)$$

Remark 1: If h and h_Π are norms on \mathbb{R}^n and \mathbb{R}^k , respectively, then Π reduces to a *dynamic projection* as introduced in [19]. Whereas dynamic projections preserve the topological notion of a point between the state and projected spaces, compatible projections can preserve topological structures such as sets.

In the context of safety, if a set $\mathcal{C} \subset \mathbb{R}^n$ is defined via a continuously differentiable function h as in (3)-(5), a compatible projection h_Π for the function h with respect to Π defines a corresponding set $\mathcal{C}_\Pi \subset \mathbb{R}^k$:

$$\mathcal{C}_\Pi \triangleq \{\mathbf{y} \in \mathbb{R}^k : h_\Pi(\mathbf{y}) \geq 0\}, \quad (20)$$

$$\partial\mathcal{C}_\Pi \triangleq \{\mathbf{y} \in \mathbb{R}^k : h_\Pi(\mathbf{y}) = 0\}, \quad (21)$$

$$\text{Int}(\mathcal{C}_\Pi) \triangleq \{\mathbf{y} \in \mathbb{R}^k : h_\Pi(\mathbf{y}) > 0\}. \quad (22)$$

The inequalities in (19) preserve the notion of what states are considered safe between the state space and projected space, such that $\mathbf{x} \in \mathcal{C} \implies \Pi(\mathbf{x}) \in \mathcal{C}_\Pi$. The following theorem allows us to extend ISSf properties of the projected system on \mathcal{C}_Π to PSSf properties of the original system on \mathcal{C} .

Theorem 3: Let $\mathcal{C} \subset \mathbb{R}^n$ be the 0-superlevel set of a continuously differentiable function $h : \mathbb{R}^n \rightarrow \mathbb{R}$ with 0 a regular value. The disturbed system (7) can be rendered PSSf on \mathcal{C} with respect to the projection Π and projected disturbances δ if there exists a compatible projection h_Π for h with respect to Π and Lipschitz continuous controller $\mathbf{k} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ such that h_Π is an ISSf-CBF for the projected dynamics (15) on \mathcal{C}_Π and $\mathbf{k}(\mathbf{x}) \in K_{\text{issf}}(\mathbf{x})$ with:

$$K_{\text{issf}}(\mathbf{x}) \triangleq \left\{ \mathbf{u} \in \mathbb{R}^m \mid \begin{array}{l} \dot{h}_\Pi(\Pi(\mathbf{x}), \mathbf{u}) \geq \\ -\alpha(h_\Pi(\Pi(\mathbf{x}))) - \iota(\|\delta\|) \end{array} \right\},$$

Proof: As h_Π is an ISSf-CBF for (15) on \mathcal{C}_Π and the state-feedback controller satisfies $\mathbf{k}(\mathbf{x}) \in K_{\text{issf}}(\mathbf{x})$, Theorem 2 implies that the controller \mathbf{k} renders the set \mathcal{C}_Π input-to-state safe for all δ satisfying $\|\delta\|_\infty \leq \bar{\delta}$. In particular, there exists $\gamma \in \mathcal{K}_\infty$ such that the set:

$$\mathcal{C}_{\Pi, \delta} \triangleq \left\{ \mathbf{y} \in \mathbb{R}^k \mid h_\Pi(\mathbf{y}) + \gamma(\|\delta\|_\infty) \geq 0 \right\}, \quad (23)$$

is safe. Let $\mathbf{x}_0 \in \mathbb{R}^n$ be such that $\mathbf{y}_0 = \Pi(\mathbf{x}_0) \in \mathcal{C}_{\Pi, \delta}$. With $\mathbf{x}(0) = \mathbf{x}_0$ (implying $\mathbf{y}(0) = \mathbf{y}_0$), safety of $\mathcal{C}_{\Pi, \delta}$ implies:

$$h_\Pi(\Pi(\mathbf{x}(t))) + \gamma(\|\delta\|_\infty) \geq 0, \quad (24)$$

for $t \in I(\mathbf{x}_0)$. As h_Π is a compatible projection for h with respect to Π , we have:

$$\bar{\sigma}(h(\mathbf{x}(t))) + \gamma(\|\delta\|_\infty) \geq 0, \quad (25)$$

Multiplying by $\frac{1}{2}$ and using that $\bar{\sigma} \in \mathcal{K}_{\infty, e}$, it follows that:

$$\bar{\sigma}^{-1} \left(\frac{1}{2} \bar{\sigma}(h(\mathbf{x}(t))) + \frac{1}{2} \gamma(\|\delta\|_\infty) \right) \geq 0, \quad (26)$$

The triangle inequality for class \mathcal{K} functions [12] implies:

$$h(\mathbf{x}(t)) + \underbrace{\bar{\sigma}^{-1}(\gamma(\|\delta\|_\infty))}_{\gamma'(\|\delta\|_\infty)} \geq 0, \quad (27)$$

for all $t \in I(\mathbf{x}_0)$, implying the set \mathcal{C}_δ defined as in (16)-(18) using γ' is forward invariant, and hence safe. Thus the closed-loop system (7) is PSSf on \mathcal{C} with respect to Π and corresponding projected disturbances δ . ■

Corollary 1: Let $\mathcal{C} \subset \mathbb{R}^n$ be the 0-superlevel set of a continuously differentiable function $h : \mathbb{R}^n \rightarrow \mathbb{R}$ with 0 a regular value. Viewing h as a projection such that $y = h(\mathbf{x})$, let the projected dynamics be given by:

$$\dot{y} = f_y(\mathbf{x}) + \mathbf{g}_y(\mathbf{x})\mathbf{u} + \delta \quad (28)$$

with projected disturbances $\delta \in \mathbb{R}$. If there exists a Lipschitz continuous feedback controller $\mathbf{k} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ such that:

$$f_y(\mathbf{x}) + \mathbf{g}_y(\mathbf{x})\mathbf{k}(\mathbf{x}) \geq -\alpha(y), \quad (29)$$

and there exists $\bar{\delta} > 0$ satisfying $|\delta| < \bar{\delta}$, then the disturbed system (7) can be rendered PSSf on \mathcal{C} with respect to the projection h and projected disturbances δ .

Proof: We first note that the identity map $I : \mathbb{R} \rightarrow \mathbb{R}$ is a compatible projection for h :

$$h(\mathbf{x}) \leq I(h(\mathbf{x})) \leq h(\mathbf{x}) \quad (30)$$

with $\underline{\sigma}(r) = \bar{\sigma}(r) = r$. The inequality in (29) implies the identity map is an ISSf-CBF for the projected dynamics (28):

$$\sup_{\mathbf{u} \in \mathbb{R}^m} \dot{I}(\mathbf{x}, \mathbf{u}, \delta) \geq \dot{I}(\mathbf{x}, \mathbf{k}(\mathbf{x}), \delta) \geq -\alpha(I(y)) - |\delta|, \quad (31)$$

for all $\mathbf{x} \in \mathbb{R}^n$ and $\delta \in \mathbb{R}$ satisfying $|\delta| \leq \bar{\delta}$. Therefore the system (7) can be rendered PSSf on \mathcal{C} with respect to the projection h and projected disturbances δ by Theorem 3. ■

IV. INTEGRATION WITH LEARNING

In this section we consider a structured form of uncertainty in affine control systems. We discuss the impact of this uncertainty on the PSSf behavior of the system, and demonstrate how learning can be used to mitigate said impact on safety.

In practice, the system dynamics (1) are not known during control design due to parametric error and unmodeled dynamics. Instead, a nominal model of the system is utilized:

$$\hat{\mathbf{x}} = \hat{\mathbf{f}}(\mathbf{x}) + \hat{\mathbf{g}}(\mathbf{x})\mathbf{u}, \quad (32)$$

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 assumed to be Lipschitz continuous on \mathbb{R}^n . By adding and subtracting (32) to (1), the dynamics of the system can be expressed as:

$$\dot{\mathbf{x}} = \hat{\mathbf{f}}(\mathbf{x}) + \hat{\mathbf{g}}(\mathbf{x})\mathbf{u} + \underbrace{\mathbf{f}(\mathbf{x}) - \hat{\mathbf{f}}(\mathbf{x})}_{\mathbf{b}(\mathbf{x})} + \underbrace{(\mathbf{g}(\mathbf{x}) - \hat{\mathbf{g}}(\mathbf{x}))\mathbf{u}}_{\mathbf{A}(\mathbf{x})}, \quad (33)$$

where the unknown disturbance $\mathbf{d} = \mathbf{b}(\mathbf{x}) + \mathbf{A}(\mathbf{x})\mathbf{u}$ explicitly depends on the state and input to the system. If the function $h : \mathbb{R}^n \rightarrow \mathbb{R}$ is a CBF for the nominal model (32) on \mathcal{C} , the uncertainty directly manifests in the time derivative of h :

$$\begin{aligned} \dot{h}(\mathbf{x}, \mathbf{u}) &= \overbrace{\frac{\partial h}{\partial \mathbf{x}}(\mathbf{x})(\hat{\mathbf{f}}(\mathbf{x}) + \hat{\mathbf{g}}(\mathbf{x})\mathbf{u})}^{\hat{h}(\mathbf{x}, \mathbf{u})} \\ &\quad + \underbrace{\frac{\partial h}{\partial \mathbf{x}}(\mathbf{x})\mathbf{b}(\mathbf{x})}_{b(\mathbf{x})} + \underbrace{\frac{\partial h}{\partial \mathbf{x}}(\mathbf{x})\mathbf{A}(\mathbf{x})\mathbf{u}}_{\mathbf{a}(\mathbf{x})^\top}. \end{aligned} \quad (34)$$

Given that h is a CBF for (32) on \mathcal{C} , let $\mathbf{k} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ be a Lipschitz continuous state-feedback controller such that:

$$\sup_{\mathbf{u} \in \mathbb{R}^m} \hat{h}(\mathbf{x}, \mathbf{u}) \geq \hat{h}(\mathbf{x}, \mathbf{k}(\mathbf{x})) \geq -\alpha(h(\mathbf{x})). \quad (35)$$

Letting the projected disturbance be defined as:

$$\delta = b(\mathbf{x}) + \mathbf{a}(\mathbf{x})^\top \mathbf{k}(\mathbf{x}), \quad (36)$$

Corollary 1 implies that if there exists a $\bar{\delta} > 0$ such that $|b(\mathbf{x}) + \mathbf{a}(\mathbf{x})^\top \mathbf{k}(\mathbf{x})| \leq \bar{\delta}$ for all $\mathbf{x} \in \mathbb{R}^n$, the uncertain system (1) can be rendered PSSf on \mathcal{C} with respect to the projection h and projected disturbances δ .

As in [21], we may wish to reduce the error between h and \hat{h} by utilizing data-driven models to estimate the functions

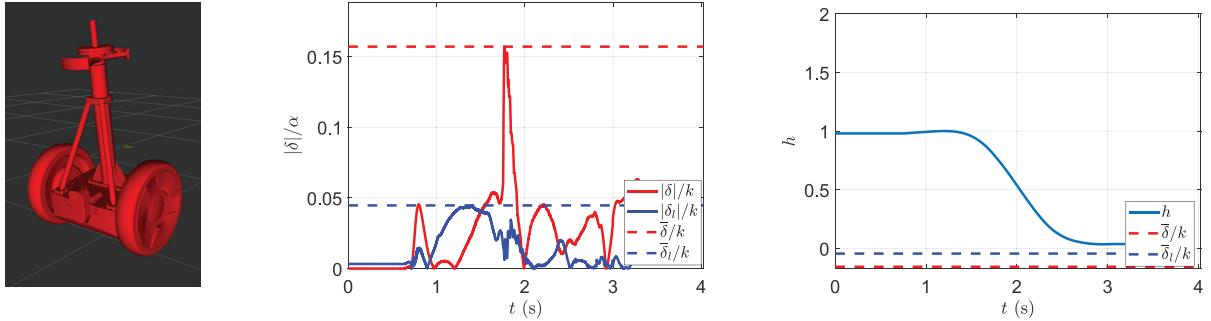


Fig. 1. Simulation results with Segway platform demonstrating improvement in PSSf behavior. **(Left)** Robotic Segway platform model used in simulation. **(Center)** Absolute value of the projected disturbance δ along the trajectory without learning models ((36), red) and with learning models ((38), blue), with learning reducing the worse case projected disturbance $(\bar{\delta}/\alpha)$. **(Right)** The value of the barrier satisfies the corresponding worst case lower bound with and without learning being used to compute δ . The worst case lower bound is raised with learning (the blue dashed line lies above the red dashed line).

b and \mathbf{a} . In particular, given Lipschitz continuous estimators $\hat{b} : \mathbb{R}^n \rightarrow \mathbb{R}$ and $\hat{\mathbf{a}} : \mathbb{R}^n \rightarrow \mathbb{R}^m$, (34) can be reformulated as:

$$\begin{aligned} \hat{h}(\mathbf{x}, \mathbf{u}) &= \underbrace{\frac{\partial h}{\partial \mathbf{x}}(\mathbf{x})(\hat{\mathbf{f}}(\mathbf{x}) + \hat{\mathbf{g}}(\mathbf{x})\mathbf{u}) + \hat{b}(\mathbf{x}) + \hat{\mathbf{a}}(\mathbf{x})^\top \mathbf{u}}_{\tilde{b}(\mathbf{x})} \\ &+ \underbrace{\frac{\partial h}{\partial \mathbf{x}}(\mathbf{x})\mathbf{b}(\mathbf{x}) - \hat{b}(\mathbf{x}) + \left(\frac{\partial h}{\partial \mathbf{x}}(\mathbf{x})\mathbf{A}(\mathbf{x}) - \hat{\mathbf{a}}(\mathbf{x})^\top \right) \mathbf{u}}_{\tilde{\mathbf{a}}(\mathbf{x})^\top} \end{aligned} \quad (37)$$

Under the assumption that the introduction of the estimators does not violate the CBF condition, such that there exists a Lipschitz state-feedback controller \mathbf{k} satisfying (35) with \hat{h} defined as in (37), we define the projected disturbance as:

$$\delta = \tilde{b}(\mathbf{x}) + \tilde{\mathbf{a}}(\mathbf{x})^\top \mathbf{k}(\mathbf{x}) \quad (38)$$

As before, if there exists $\bar{\delta} > 0$ such that $|\tilde{b}(\mathbf{x}) + \tilde{\mathbf{a}}(\mathbf{x})^\top \mathbf{k}(\mathbf{x})| \leq \bar{\delta}$ for all $\mathbf{x} \in \mathbb{R}^n$, Corollary 1 can be used to certify (1) as PSSf on \mathcal{C} with respect to the projection h and projected disturbances δ . This statement is formalized as follows.

Theorem 4: Let $\mathcal{C} \subset \mathbb{R}^n$ be the 0-superlevel set of a continuously differentiable function $h : \mathbb{R}^n \rightarrow \mathbb{R}$ with 0 a regular value, and let $\hat{h} : \mathbb{R}^n \rightarrow \mathbb{R}$ be defined as in (34) (or as in (37)). If there exists a Lipschitz continuous state-feedback controller $\mathbf{k} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ satisfying (35), and $\bar{\delta} > 0$ such that the corresponding projected disturbance defined as in (36) (or as in (38)) satisfies $|\delta| \leq \bar{\delta}$, then (7) is PSSf on \mathcal{C} with respect to the projection h and projected disturbances δ .

We note that this theorem captures both the case without estimators (using definitions (34) and (36)) and the case with estimators (using definitions (37) and (38)). In the presence of estimators, this theorem defines a quantitative relationship between the prediction error of the estimators, $|\hat{h}(\mathbf{x}, \mathbf{k}(\mathbf{x})) - \hat{h}(\mathbf{x}, \mathbf{k}(\mathbf{x}))| = |\delta|$, and the degradation of the safety of the closed-loop system. As the prediction error is reduced (via additional training data or more complex learning models), the set kept safe more closely resembles \mathcal{C} .

V. SIMULATION & EXPERIMENTAL VALIDATION

To demonstrate the ability of learning to improve safety guarantees via Projection-to-State Safety, we deployed the

episodic learning framework with CBFs established in [21] on a robotic Segway platform, seen in Figure 1 and 2, in simulation and experimentally. The input to the system is specified as torque about the wheels of the Segway. Each wheel is given the same torque, constraining the system to the plane, and making the states of interest the position, velocity, pitch angle, and pitch angle rate. The dynamics for the system are given by the unconstrained Euler-Lagrange equations.¹

In both simulation and experiment a sequence of episodes were ran to train estimators \hat{b} and $\hat{\mathbf{a}}$, represented by 4-layer neural networks with ReLu nonlinearities. In each episode the Segway was set to track a desired trajectory in the pitch angle space without violating a barrier function on a portion of its state, using the controller (CBF-QP) and safety-critical control formulation in [10]. After the sequence of episodes, the Segway was ran once more with a learning-informed controller, and the projected disturbance δ as defined in (36) and (38) was computed. The worst case disturbance $\bar{\delta}$ was found, and a lower bound on h for that trajectory was determined using the fact $h \leq \alpha^{-1}(\bar{\delta}) \implies \dot{h} \geq 0$. In both simulation and experiment, $\alpha(r) = kr$ with $k > 0$.

In simulation, the Segway was given a bound on its position in space, constraining it to a one meter distance from its starting location. The CBF was generated through the backup controller method [9]. The value of the CBF is computed at each time-step by integrating the system forward in time under a backup control law. Sensitivity analysis along the trajectory is used to compute the gradient of the CBF. This simulation result highlights the ability of learning to reduce worst case disturbances for complex CBFs that cannot be expressed in closed-form. The simulation was done in a Robot Operating System (ROS) based C++ environment similar to [18]. The simulation environment accurately simulates the physical system by adding input delay, sensor noise, and state estimation. The code is identical to that on the robot for the state estimator, controller, and CBF, apart from the ROS functionality. Experimentally, a CBF was specified to limit the pitch angle and pitch angle rate of the Segway to an ellipse about the Segway's equilibrium state. The desired pitch angle

¹The simulation code (including an expression for the full system dynamics) can be found at https://github.com/DrewSingletary/cyberpod_sim_ros.

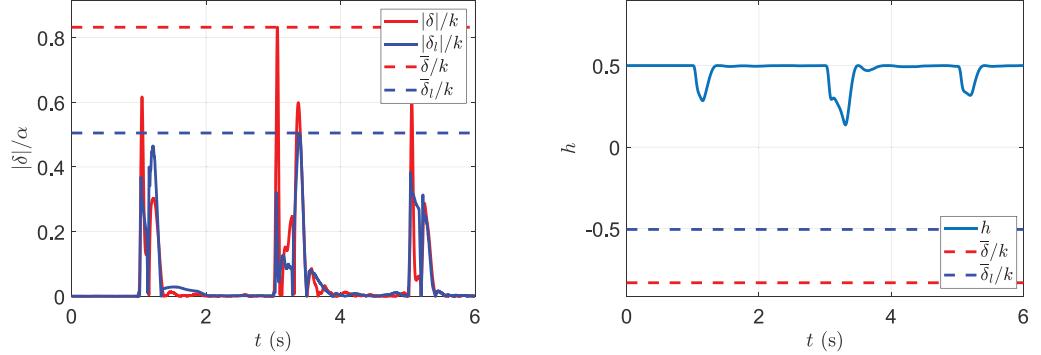


Fig. 2. Experimental results with Segway platform demonstrating improvement in PSSf behavior. **(Left)** Physical robotic Segway platform used in experimentation. **(Center)** Absolute value of the projected disturbance δ along the trajectory without learning models ((36), red) and with learning models ((38), blue), with learning reducing the worse case projected disturbance $(\bar{\delta}/\alpha)$. **(Right)** The value of the barrier satisfies the corresponding worst case lower bound with and without learning being used to compute δ . The worst case lower bound with learning (the blue dashed line lies above the red dashed line).

trajectory leads the Segway to tip quickly, leaving the safety set in the absence of the CBF and safety-critical controller.

In both cases, we see that introducing learning estimators into the computation of the projected disturbance decreases the worse case disturbance ($\bar{\delta} > \bar{\delta}_l$). This leads to a greater lower bound on h , and thus a stronger guarantee on the PSSf behavior of the system. We note that the conservative nature of the lower bounds on h arise from the fact that the worst case disturbance $\bar{\delta}$ along the trajectory is used. If the worst case disturbance can be reduced (by data-aware control synthesis), stronger guarantees on safety can be made.

VI. CONCLUSION

We presented a novel method for assessing the impact of disturbances on safety in a project environment via Projection-to-State Safety, and considered how it can be utilized with learning to mitigate the impact of model uncertainty on safety. We demonstrate the ability of learning to improve the guarantees endowed by PSSf in simulation and experimentally on a Segway platform. Future work includes developing methods for quantifying the worst case projected disturbance and synthesizing controllers that reduce the projected disturbance.

REFERENCES

- [1] A. Ames, J. Grizzle, and P. Tabuada, “Control barrier function based quadratic programs with application to adaptive cruise control,” in *Proc. IEEE Conf. Decis. Control (CDC)*, 2014, pp. 6271–6278.
- [2] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, “Control barrier functions: Theory and applications,” in *Proc. IEEE Eur. Control Conf. (ECC)*, 2019, pp. 3420–3431.
- [3] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, “Control barrier function based quadratic programs for safety critical systems,” *IEEE Trans. Autom. Control*, vol. 62, no. 8, pp. 3861–3876, Aug. 2017.
- [4] E. Arabi, K. Garg, and D. Panagou, “Safety-critical adaptive control with nonlinear reference model systems,” 2019. [Online]. Available: arXiv:1909.07916.
- [5] F. Blanchini and S. Miani, *Set-Theoretic Methods in Control*. Cham, Switzerland: Springer, 2008.
- [6] X. Chen, M. Monfort, A. Liu, and B. D. Ziebart, “Robust covariate shift regression,” in *Proc. Int. Conf. Artif. Intell. Stat. (AISTATS)*, 2016, pp. 1270–1279.
- [7] 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 *Proc. Conf. Artif. Intell.*, vol. 33, 2019, pp. 3387–3395.
- [8] J. F. Fisac, A. K. Akametalu, M. N. Zeilinger, S. Kaynama, J. Gillula, and C. J. Tomlin, “A general safety framework for learning-based control in uncertain robotic systems,” *IEEE Trans. Autom. Control*, vol. 64, no. 7, pp. 2737–2752, Jul. 2019.
- [9] T. Gurriet, M. Mote, A. D. Ames, and E. Feron, “An online approach to active set invariance,” in *Proc. IEEE Conf. Decis. Control (CDC)*, 2018, pp. 3592–3599.
- [10] T. Gurriet, A. Singletary, J. Reher, L. Ciarletta, E. Feron, and A. Ames, “Towards a framework for realizable safety critical control through active set invariance,” in *Proc. Int. Conf. Cyber Phys. Syst. (ICCPs)*, 2018, pp. 98–106.
- [11] M. Jankovic, “Robust control barrier functions for constrained stabilization of nonlinear systems,” *Automatica*, vol. 96, pp. 359–367, Oct. 2018.
- [12] C. M. Kellett, “A compendium of comparison function results,” *Math. Control Signals Syst.*, vol. 26, no. 3, pp. 339–374, 2014.
- [13] S. Kolathaya and A. D. Ames, “Input-to-state safety with control barrier functions,” *Control Syst. Lett.*, vol. 3, no. 1, pp. 108–113, 2018.
- [14] A. Liu, G. Shi, S.-J. Chung, A. Anandkumar, and Y. Yue, “Robust regression for safe exploration in control,” 2020. [Online]. Available: arXiv:1906.05819.
- [15] Q. Nguyen and K. Sreenath, “Exponential control barrier functions for enforcing high relative-degree safety-critical constraints,” in *Proc. Amer. Control Conf. (ACC)*, 2016, pp. 322–328.
- [16] M. Ohnishi, L. Wang, G. Notomista, and M. Egerstedt, “Barrier-certified adaptive reinforcement learning with applications to Brushbot navigation,” *IEEE Trans. Robot.*, vol. 35, no. 5, pp. 1186–1205, Aug. 2019.
- [17] L. Perko, *Differential Equations and Dynamical Systems*, vol. 7. New York, NY, USA: Springer, 2013.
- [18] A. Singletary, P. Nilsson, T. Gurriet, and A. D. Ames, “Online active safety for robotic manipulators,” in *Proc. IEEE Int. Conf. Intell. Robots Syst. (IROS)*, 2019, pp. 173–178.
- [19] A. J. Taylor, V. D. Dorobantu, M. Krishnamoorthy, H. M. Le, Y. Yue, and A. D. Ames, “A control Lyapunov perspective on episodic learning via projection to state stability,” in *Proc. IEEE Conf. Decis. Control (CDC)*, 2019, pp. 1448–1455.
- [20] 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 *Proc. Int. Conf. Intell. Robots Syst. (IROS)*, 2019, pp. 6878–6884.
- [21] A. J. Taylor, A. Singletary, Y. Yue, and A. D. Ames, “Learning for safety-critical control with control barrier functions,” 2020. [Online]. Available: arXiv:1912.10099.
- [22] L. Wang, A. D. Ames, and M. Egerstedt, “Safety barrier certificates for collisions-free multirobot systems,” *IEEE Trans. Robot.*, vol. 33, no. 3, pp. 661–674, Jun. 2017.
- [23] L. Wang, E. A. Theodorou, and M. Egerstedt, “Safe learning of quadrotor dynamics using barrier certificates,” in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, 2018, pp. 2460–2465.
- [24] X. Xu, P. Tabuada, J. W. Grizzle, and A. D. Ames, “Robustness of control barrier functions for safety critical control,” *IFAC PapersOnLine*, vol. 48, no. 27, pp. 54–61, 2015.