Learning Control Barrier Functions from Expert Demonstrations

Alexander Robey^{*,†}, Haimin Hu^{*,†}, Lars Lindemann^a, Hanwen Zhang[†],

Dimos V. Dimarogonas^{*a*}, Stephen Tu^{*b*}, Nikolai Matni[†]

Abstract-Inspired by the success of imitation and inverse reinforcement learning in replicating expert behavior through optimal control, we propose a learning based approach to safe controller synthesis based on control barrier functions (CBFs). We consider the setting of a known nonlinear control affine dynamical system and assume that we have access to safe trajectories generated by an expert — a practical example of such a setting would be a kinematic model of a selfdriving vehicle with safe trajectories (e.g., trajectories that avoid collisions with obstacles in the environment) generated by a human driver. We then propose and analyze an optimization based approach to learning a CBF that enjoys provable safety guarantees under suitable Lipschitz smoothness assumptions on the underlying dynamical system. A strength of our approach is that it is agnostic to the parameterization used to represent the CBF, assuming only that the Lipschitz constant of such functions can be efficiently bounded. Furthermore, if the CBF parameterization is convex, then under mild assumptions, so is our learning process. We end with extensive numerical evaluations of our results on both planar and realistic examples. using both random feature and deep neural network parameterizations of the CBF. To the best of our knowledge, these are the first results that learn provably safe control barrier functions from data.

I. INTRODUCTION

Consider the following safety-critical scenarios: a selfdriving car navigating through traffic, two unmanned aerial vehicles (UAVs) avoiding collision, and a robotic manipulator in a laboratory setting that must avoid injuring researchers. Although vastly different in terms of their environments, safety-specifications, and underlying dynamics, they share several key properties: (i) their dynamics are well understood and modeled, and can be accurately identified, (ii) their dynamics are inherently *nonlinear*, and (iii) *expert demonstrations* of safe and desirable behavior are readily available or can be easily collected. Motivated by these unifying properties, this paper proposes the design of safe controllers for known nonlinear dynamical systems based on *control barrier functions* learned from expert demonstrations.

Barrier functions, which are also referred to as barrier certificates, were first proposed in [1] as a means of certifying the safety of dynamical systems with respect to semialgebraic safe sets. In that work, a sum-of-squares (SOS) programming [2] approach for synthesizing polynomial barrier functions for given polynomial systems was also described. The notion of control barrier functions (CBFs) for dynamical control systems was first introduced in [3] to guarantee the existence of a control law that renders a desired safe set forward invariant. The notion of CBFs was refined by introducing reciprocal [4] and zeroing CBFs [5], which do not require that sub-level sets of the CBF be invariant within the safe set. In particular, zeroing CBFs can be used to compute a minimally invasive "correction" to a nominal control law. Importantly, this correction maintains safety by computing the solution to a quadratic program (QP) [5].

One open problem that has not been fully addressed in prior work is how such CBFs can be synthesized for general classes of systems. This challenge is similar to that which arises when addressing stability using control Lyapunov functions (CLFs) as the analog to Lyapunov functions [6]. Notably, control Lyapunov functions are a subset of control barrier functions (see [5] and [7]). Analytic and SOS based approaches to synthesizing CBFs and CLFs are summarized in [8] and have appeared in [9], [10]. These approaches, however, are known to be limited in scope and scalability.

Related work on learning and CBFs: Methods using barrier and control barrier functions to ensure safety and guide exploration during episodic supervised learning of uncertain linear dynamics include [11]–[14]. These approaches typically assume that a valid (control) barrier function is provided, and should be viewed as complementary to our results. In [15], an imitation learning based approach is used to to train a deep neural network (DNN) to replicate a CBF based controller. While the authors of [15] present empirical validation of their results, no theoretical guarantees of correctness are provided. The authors of [16] jointly learn a control Lyapunov function, a CBF, and a policy function, and then verify their validity post-hoc using Lipschitz arguments. In [17], tools from statistical learning theory, are used to learn Lyapunov functions from data for systems with unknown dynamics. Most similar in spirit to our paper are the results in [18] and [19]. In [18], the authors parameterize a CBF by a support vector machine, and use a supervised learning approach to characterize regions of the state-space as safe or unsafe based on collected data. While conceptually appealing, we note that their training procedure does not ensure a priori that there exist control actions such that the learned safe set can be made forward invariant,¹ and hence cannot guarantee safe execution of the system. In [19], a

^{*}A. Robey and H. Hu contributed equally. [†]Authors are with the Dept. of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA. ^{*a*}Authors are with the Division of Decision and Control Systems, KTH Royal Institute of Technology, Stockholm, Sweden. ^{*b*}S. Tu is with Google Brain Robotics, New York, NY.

L. Lindemann & D. V. Dimarogonas were supported in part by the Knut and Alice Wallenberg Foundation (KAW) and the SSF COIN project.

¹In particular, they do not ensure that the derivative condition $\langle \nabla h(x), f(x,u) \rangle + \alpha(h(x)) \geq 0$, holds for the learned CBF h(x) at the observed data points, with f(x, u) the system dynamics, and α an extended class \mathcal{K} function – see Section II for more details.

method is proposed which incrementally learns a *linear* CBF by clustering expert demonstrations into linear subspaces and fitting low dimensional representations. While both papers [18], [19] empirically validate their methods, neither provide proofs of correctness of the learned CBF.

Contributions: In this paper, we propose and analyze an optimization-based approach to learning a zeroing CBF (henceforth referred to simply as a CBF) from expert trajectories for known control affine nonlinear systems. In particular, we provide precise and verifiable conditions on the expert trajectories, an additional auxiliary data-set, and the hyperparameters of the optimization problem so as to ensure that the learned CBF guarantees safe execution of the system. We further show how the underlying optimization problem can be efficiently solved when it is cast over different function spaces. In particular, we show that the problem can be solved via convex optimization when the function space lies within a (possibly infinite-dimensional) Reproducing Kernel Hilbert Space (RKHS); alternatively, when we consider the function space of Deep Neural Networks (DNNs), the problem can be solved via first-order stochastic methods such as Adam or SGD. To the best of our knowledge, these are the first such results that learn a CBF from expert demonstrations with provable safety guarantees.

Paper structure: The rest of this paper is structured as follows. In Section II, we introduce notation and formulate the general problem of learning a CBF from expert demonstrations. In Section III, we derive a set of sufficient conditions on the learned CBF and data-set that guarantee safety of the resulting closed-loop system, and we subsequently use these conditions to formulate an optimization problem for computing a function satisfying these conditions. We show in Section III-C that this optimization problem can be efficiently solved for CBFs embedded in RKHS and DNN function classes, and in Section III-D, we provide further details on the expert trajectory collection process. We present three numerical studies in Section IV: (i) a two-dimensional planar problem for which we explicitly compute and verify all of the conditions of our main theorem, showing that the conditions are indeed satisfied in practice, (ii) a two UAV collisionavoidance example where expert trajectories are generated by the closed form CBF from [20], and (iii) the same two UAV collision avoidance example, where now expert trajectories are generated by human players of a video game interface. We end with conclusions and discussions of directions for future work in Section V.

II. PRELIMINARIES AND PROBLEM FORMULATION

Let \mathbb{R} and $\mathbb{R}_{\geq 0}$ be the set of real and non-negative real numbers, respectively, and \mathbb{R}^n the set of *n*-dimensional real vectors. For $\epsilon > 0$ and $p \geq 1$, we let $\mathcal{B}_{\epsilon,p}(\bar{x}) := \{x \in \mathbb{R}^n \mid ||x - \bar{x}||_p \leq \epsilon\}$ denote the closed *p*-norm ball around $\bar{x} \in \mathbb{R}^n$. For a given set \mathcal{C} , we denote by $\mathrm{bd}(\mathcal{C})$, $\mathrm{int}(\mathcal{C})$, and \mathcal{C}^c the boundary, interior, and complement of \mathcal{C} , respectively. For two sets \mathcal{C}_1 and \mathcal{C}_2 , we denote their Minkowski sum by $\mathcal{C}_1 \oplus \mathcal{C}_2 := \{x_1 + x_2 \in \mathbb{R}^n | x_1 \in \mathcal{C}_1, x_2 \in \mathcal{C}_2\}$. A continuous function $\alpha : \mathbb{R} \to \mathbb{R}$ is an extended class \mathcal{K} function if it is strictly increasing with $\alpha(0) = 0$. The inner-product between two vectors $x, y \in \mathbb{R}^n$ is denoted by $\langle x, y \rangle$.

A. Valid Control Barrier Functions

At time $t \in \mathbb{R}_{\geq 0}$, let $x(t) \in \mathbb{R}^n$ and $u(t) \in \mathbb{R}^m$ be the state and input, respectively, of the dynamical control system described by the initial value problem

$$\dot{x}(t) = f(x(t)) + g(x(t))u(t), \ x(0) \in \mathbb{R}^n$$
 (1)

where $f : \mathbb{R}^n \to \mathbb{R}^n$ and $g : \mathbb{R}^n \to \mathbb{R}^m$ are locally Lipschitz continuous functions. Let the unique solution to (1) under a locally Lipschitz continuous control law $u : \mathbb{R}^n \to \mathbb{R}^m$ be $x : \mathcal{I} \to \mathbb{R}^n$ where $\mathcal{I} \subseteq \mathbb{R}_{\geq 0}$ is the maximum definition interval of x. Note that we do not explicitly assume forward completeness of (1) under u here, i.e., \mathcal{I} may be bounded.

Consider next a twice continuously differentiable function $h : \mathbb{R}^n \to \mathbb{R}$, and define the set

$$\mathcal{C} := \{ x \in \mathbb{R}^n \mid h(x) \ge 0 \},\tag{2}$$

which defines a set that we wish to certify as safe, i.e., that it satisfies prescribed safety specifications and can be made forward invariant through an appropriate choice of control action. We further assume that C has non-empty interior, and let D be an open set such that $D \supseteq C$. The function h(x) is said to be a *valid control barrier function* on D if there exists a locally Lipschitz continuous extended class \mathcal{K} function $\alpha : \mathbb{R} \to \mathbb{R}$ such that

$$\sup_{u \in \mathcal{U}} \langle \nabla h(x), f(x) + g(x)u \rangle \ge -\alpha(h(x))$$
(3)

holds for all $x \in \mathcal{D}$, where $\mathcal{U} \in \mathbb{R}^m$ defines constraints on the control input *u*. Consequently, we define the set of *CBF consistent inputs* induced by a valid CBF h(x) to be

$$K_{\text{CBF}}(x) := \{ u \in \mathbb{R}^m \mid \langle \nabla h(x), f(x) + g(x)u \rangle \ge -\alpha(h(x)) \}.$$

The next result follows from [5], [7].

Lemma II.1: Assume that h(x) is a valid control barrier function on \mathcal{D} and that $u : \mathcal{D} \to \mathcal{U}$ with $u(x) \in K_{\text{CBF}}(x)$ is locally Lipschitz continuous. Then it holds that $x(0) \in \mathcal{C}$ implies $x(t) \in \mathcal{C}$ for all $t \in \mathcal{I}$. If the set \mathcal{C} is compact, it additionally follows that \mathcal{C} is: 1) forward invariant, i.e., $\mathcal{I} = [0, \infty)$, and 2) asymptotically stable, which implies that x(t) approaches \mathcal{C} as $t \to \infty$ when $x(0) \in \mathcal{C}^c \cap \mathcal{D}$.

Note that $h(x) \neq 0$ when $x \in bd(\mathcal{C})$ (see [8, Remark 5]) is not required when using the Comparison Lemma instead of Nagumo's theorem to prove the above result. While the previous result provides strong guarantees of safety given a valid control barrier function, one is still left with the potentially daunting task of finding a twice continuously differentiable function h such that (i) the set \mathcal{C} defined in equation (2) captures a sufficiently large volume of "safe" states needed for the task at hand, and (ii) that it satisfies the derivative constraint (3) on an open set $\mathcal{D} \supseteq \mathcal{C}$. While safety constraints are often naturally specified on a subset of the configuration space of a system, e.g., to avoid collision, vehicles must maintain a minimum distance, ensuring that a CBF specified using such geometric intuition also satisfies constraint (3) can involve verifying complex relationships between the vector field of the system, the candidate control barrier function, and its gradient.

As described in the introduction, this challenge motivates the approach taken in this paper, wherein we propose an optimization based approach to learning a CBF from expert demonstrations for a system with known dynamics.

B. Problem Formulation

To formalize the previous discussion, we explicitly distinguish between geometric safety specifications, i.e., those that can be directly specified on (a subset) of the state-space of the system $x \in \mathbb{R}^n$, and the set C defined in equation (2) that is certified as safe by the CBF. To that end, let $S \subseteq \mathbb{R}^n$ define the aforementioned geometric safe set.

Towards the goal of learning a valid CBF, we assume that we are given a set of *expert trajectories* consisting of N_1 discretized data-points $Z_{dyn} := \{(x_i, u_i)\}_{i=1}^{N_1}$ so that $x_i \in$ int(S) (see Fig. 1 (left)). For $\epsilon > 0$, we define the set

$$\mathcal{D}' := \bigcup_{i=1}^{N_1} \mathcal{B}_{\epsilon, p}(x_i), \ \mathcal{D} := \mathcal{D}' \backslash \mathrm{bd}(\mathcal{D}') \tag{4}$$

where \mathcal{D} must satsify $\mathcal{D} \subseteq S$ to later ensure safety of the resulting control based on the learned CBF. Several comments are in order. First, note that we define \mathcal{D} based on expert trajectories for which control inputs u_i are available so that the derivative constraint (3) can be enforced during learning. Second, by construction, the x component of Z_{dyn} defines an ϵ -net over \mathcal{D} , i.e., for all $x \in \mathcal{D}$, (slightly abusing notation) there exists $x_i \in Z_{dyn}$ such that $||x_i - x||_p \leq \epsilon$. Finally, conditions on ϵ will be specified later to ensure the validity of the learned CBF.

Remark 1: We note that a conceptually similar approach, defined in terms of taking a point-wise union over previously seen safe trajectories, is used to define a safe terminal set in the Learning Model Predictive Control method of [21].

We next define the set \mathcal{N} , for $\sigma > 0$, as

$$\mathcal{N} := \{ \mathrm{bd}(\mathcal{D}) \oplus \mathcal{B}_{\sigma,p}(0) \} \setminus \mathcal{D},$$

which should be thought of as a "layer" of width σ surrounding \mathcal{D} , see Fig. 1 (center) for a graphical depiction. As will be made clear in the sequel, by enforcing that the value of the learned CBF h(x) is negative on \mathcal{N} , which can be accomplished through appropriate sampling, we ensure that the zero level set $\{x \in \mathbb{R}^n | h(x) = 0\}$ is contained within \mathcal{D} , which is a necessary condition for h(x) to be valid.

While the above definition of a CBF is specified over all of \mathbb{R}^n , e.g., the definition of C in equation (2) considers all $x \in \mathbb{R}^n$ such that $h(x) \ge 0$, we make a minor modification to this definition in order to restrict the domain of interest to $\mathcal{N} \cup \mathcal{D}$, i.e., we will certify that h(x) is a valid *local* CBF over \mathcal{D} with respect to the set

$$\mathcal{C} := \{ x \in \mathcal{N} \cup \mathcal{D} \mid h(x) \ge 0 \}.$$
(5)

This restriction is natural, as we are learning a CBF h(x)from data sampled only over $\mathcal{N} \cup \mathcal{D}$, and we will show that the inclusion $\mathcal{C} \subset \mathcal{D} \subseteq S$ holds. It then follows that if h(x)is shown to satisfy the derivative constraint (3) for all $x \in \mathcal{D}$, then both C, as defined in (5), and the compact set D can be made forward invariant by some $u \in U$.

III. AN OPTIMIZATION BASED APPROACH

In this section, we define and analyze an optimization based approach to generating valid local control barrier functions from expert demonstrations.

Let \mathcal{H} be a normed function space of twice continuously differentiable functions $h : \mathbb{R}^n \to \mathbb{R}$ for which local Lipschitz bounds $L_h(x) := \sup_{x_1, x_2 \in \mathcal{B}_{\epsilon,p}(x)} \frac{|h(x_1) - h(x_2)|}{||x_1 - x_2||_p}$ can be efficiently estimated. Commonly used examples of such spaces include infinite dimensional Reproducing Kernel Hilbert Spaces (RKHS) such as those defined by Random Fourier (RF) features [22], and more recently, DNNs [23]. We defer a discussion of results specific to these two classes of CBFs to the end of this section, and focus now on a general method applicable to these, and other, spaces \mathcal{H} .

Recall the definition of Z_{dyn} and define the set $X_{safe} = \{x_i : (x_i, u_i) \in Z_{dyn}\}$. We also assume that points $X_{\mathcal{N}} = \{x_i\}_{i=1}^{N_2}$ are sampled from the set \mathcal{N} such that $X_{\mathcal{N}}$ forms an $\bar{\epsilon}$ -net of \mathcal{N} – conditions on $\bar{\epsilon}$ will be specified in the sequel. We emphasize that no associated inputs u_i are needed for the samples $X_{\mathcal{N}} \subset \mathcal{N}$, as these points are not generated by the expert, and can instead be obtained by simple computational methods such as gridding or uniform sampling.

We begin by deriving a set of sufficient conditions in terms of constraints on the learned CBF h(x), as well as conditions on the data-sets X_{safe} and X_N , that ensure that h(x) is a valid local CBF on \mathcal{D} . We then use these constraints to formulate an optimization problem that can be efficiently solved for the aforementioned function classes \mathcal{H} .

A. Guaranteeing $C \subset D \subseteq S$

We begin with the simple and intuitive requirement that the learned CBF h(x) satisfy

$$h(x_i) \ge \gamma_{\text{safe}} \text{ for all } x_i \in X_{\text{safe}},$$
 (6)

for a yet to be specified parameter $\gamma_{\text{safe}} > 0$. This in particular ensures that the set C over which $h(x) \ge 0$, as defined in equation (5), has non-empty interior.

We now derive conditions under which the learned CBF satisfies h(x) < 0 for all $x \in \mathcal{N}$, which in turn ensures that $\mathcal{C} \subset \mathcal{D} \subseteq \mathcal{S}$ due to constraint (6).

Proposition III.1: Let h(x) be Lipschitz continuous with local constant $L_h(x)$. Let $\gamma_{\text{unsafe}} > 0$, and X_N be an $\bar{\epsilon}$ -net of N with $\bar{\epsilon} < \gamma_{\text{unsafe}}/L_h(x_i)$ for all $x_i \in X_N$. Then, if

$$h(x_i) \le -\gamma_{\text{unsafe}} \text{ for all } x_i \in X_{\mathcal{N}},$$
 (7)

it holds that h(x) < 0 for all $x \in \mathcal{N}$.

Proof: By equation (7), we have that $h(x_i) \leq -\gamma_{\text{unsafe}}$ for each $x_i \in X_N$. We then have, for any $x \in N$, that there exists a point $x_i \in X_N$ satisfying $||x - x_i||_p \leq \bar{\epsilon} < \gamma_{\text{unsafe}}/L(x_i)$, from which the following chain of inequalities follows immediately

$$h(x) = h(x) - h(x_i) + h(x_i) \le |h(x) - h(x_i)| - \gamma_{\text{unsafe}}$$

$$\le L_h(x_i) ||x - x_i||_p - \gamma_{\text{unsafe}} \le L_h(x_i) \bar{\epsilon} - \gamma_{\text{unsafe}} < 0,$$



Fig. 1: Problem Setup (left): The safe set S (red box) and the set of expert trajectories (black lines). Desired result (center): The set D (orange ring) is the union of ϵ balls around the expert trajectories. The set N (black striped rings), defined around D, ensures that the learned safe set C (green ring), which is defined via the learned valid control barrier function h(x), is such that $C \subset D \subseteq S$. Control Barrier filter (right): "Artiticial" unsafe samples are no longer introduced in the center of the safe set (denoted by the blue set \mathcal{F}).

where the first inequality follows from the assumption that $h(x_i) \leq -\gamma_{\text{unsafe}}$ for all $x_i \in X_N$, the second by the local Lipschitz assumption on h(x), the third by the assumption that X_N forms an $\bar{\epsilon}$ -net of N, and the final inequality by the condition on $\bar{\epsilon}$ of the proposition.

We note that as stated, the constraints (6) and (7), as well as the condition $\bar{\epsilon} < \gamma_{\text{unsafe}}/L_h(x_i)$ of Proposition III.1, may be incompatible, leading to infeasibility of an optimization problem built around them. This incompatibility arises from the fact that we are simultaneously asking for the value of h(x) to vary from γ_{safe} to γ_{unsafe} over a short distance $\bar{\epsilon}$ while having a low Lipschitz constant. In particular, as posed, the constraints require that $|h(x_s) - h(x_u)| \ge \gamma_{\text{safe}} + \gamma_{\text{unsafe}}$ for $x_s \in X_{\text{safe}}$ and $x_u \in X_{\mathcal{N}}$ safe and unsafe samples, respectively, but the samplings requirements imply that $||x_s|$ $x_u \|_2 \leq \overline{\epsilon} + \epsilon$ for at least some pair (x_s, x_u) , which in turn implies that $L(x_u) \gtrsim |h(x_s) - h(x_u)| / ||x_s - x_u||_2 \gtrsim (\gamma_{\text{safe}} +$ $\gamma_{\text{unsafe}})/(\bar{\epsilon}+\epsilon)$. Thus, if γ_{safe} and γ_{unsafe} are chosen to be too large, we may exceed the required bound of $\gamma_{\rm unsafe}/\bar{\epsilon},$ and set over which $h(x) \ge 0$ may be undesirably small (i.e., the volume of C would be too small).

We address this issue as follows: for fixed γ_{safe} , γ_{unsafe} , and $L_h := \sup_{x_i \in X_N} L_h(x_i)$, constraint (6) is relaxed to

$$h(x_i) \ge \gamma_{\text{safe}}, \quad x_i \in X_{\text{safe}},$$
 (8)

where now

$$\bar{X}_{\text{safe}} = \{ x_i \in X_{\text{safe}} \mid \inf_{x \in X_{\mathcal{N}}} \|x - x_i\|_p \ge \frac{\gamma_{\text{unsafe}} + \gamma_{\text{safe}}}{L_h} \}$$
(9)

corresponds to an inner subset of expert trajectory samples. Intuitively, this introduces a buffer region across which h(x) can vary in value from γ_{safe} to $-\gamma_{\text{unsafe}}$ without having an excessively large Lipschitz constant. A near identical argument as that used to prove Proposition III.1 can now be used to guarantee that the set C defined in equation (5) contains the set $\overline{D} = \bigcup_{x_i \in \overline{X}_{\text{safe}}} \mathcal{B}_{\epsilon,p}(x_i)$, defined as the union of ϵ -balls around the points in $\overline{X}_{\text{safe}}$, and thus, $\overline{D} \subseteq C$ can be seen as a "minimum volume" guarantee on the set C.

Corollary III.2: Let h(x) be Lipschitz continuous with local constant $L_h(x)$. Let $\gamma_{\text{safe}} > 0$, and X_{safe} be an ϵ net of \mathcal{D} with $\epsilon \leq \gamma_{\text{safe}}/L_h(x_i)$ for all $x_i \in \bar{X}_{\text{safe}}$. Then, if constraint (8) is satisfied, it holds that $h(x) \ge 0$ for all $x \in \overline{\mathcal{D}}$.

B. Guaranteeing Valid Local Control Barrier Functions

The conditions in the previous subsection guarantee that the level-sets of the learned CBF satisfy the desired properties. We now derive conditions that ensure that the derivative constraint (3) is also satisfied by the learned CBF.

Because we assume that the CBF functions h(x) are twice continuously differentiable over a compact domain $\mathcal{N} \cup \mathcal{D}$, we immediately have that $\nabla h(x)$ is Lipschitz continuous with local Lipschitz constant $L_{\nabla}(x)$.

Note that to verify that a CBF h(x) satisfying the constraints of the previous section is valid, it suffices to show that there exists a single control input $u \in \mathcal{U}$ such that the derivative constraint (3) holds. Our approach is to use the control inputs provided by the expert demonstrations. We discuss the consequences of this choice further in §III-D.

To that end, note that for a fixed u_i , the function $q(x) := \langle \nabla h(x), f(x) + g(x)u_i \rangle + \alpha(h(x))$ is Lipschitz continuous, with Lipschitz constant denoted by $L_q(x)$, as ∇h , f and g are all assumed to be Lipschitz continuous. Following a similar argument as in the previous subsection, we then have the following result guaranteeing that the learned CBF satisfies the derivative constraint (3) for all $x \in \mathcal{D}$.

Proposition III.3: Suppose q(x) is Lipschitz continuous with constant $L_q(x)$. Let $\gamma_{dyn} > 0$, and X_{safe} be an ϵ -net of \mathcal{D} with $\epsilon \leq \gamma_{dyn}/L_q(x_i)$ for all $x_i \in X_{safe}$. Then if

$$\langle \nabla h(x_i), f(x_i) + g(x_i)u_i \rangle \ge -\alpha(h(x_i)) + \gamma_{dyn}$$
 (10)

for all $x_i \in X_{\text{safe}}$, it holds that $q(x) \ge 0$ for all $x \in \mathcal{D}$.

Proof: Following a similar argument as the proof of Proposition III.1, we note that by equation (10), we have that $q(x_i) \ge \gamma_{\text{dyn}}$ for each $x_i \in X_{\text{safe}}$. We then have, for any $x \in \mathcal{D}$, that there exists a point $x_i \in X_{\text{safe}}$ satisfying $\|x - x_i\|_p \le \epsilon \le \gamma_{\text{dyn}}/L_q(x_i)$, from which the following chain of inequalities follows immediately

$$q(x) = q(x_i) + q(x) - q(x_i) \ge \gamma_{\text{dyn}} - |q(x) - q(x_i)|$$

$$\ge \gamma_{\text{dyn}} - L_q(x_i) ||x - x_i||_p \ge \gamma_{\text{dyn}} - L_q(x_i) \epsilon \ge 0,$$

where the first inequality follows from the assumption that $q(x_i) \ge \gamma_{\text{dyn}}$ for all $x_i \in X_{\text{safe}}$, the second by the Lipschitz assumption on q(x), the third by the assumption that X_{safe} forms an ϵ -net of X_{safe} , and the final inequality by the condition on ϵ of the proposition.

The following theorem, which follows immediately from the previous results, states a set of sufficient conditions guaranteeing that a learned CBF is locally valid on the domain $\mathcal{N} \cup \mathcal{D}$. We next use these conditions to formulate an optimization based approach to learning a CBF from data.

Theorem III.4: Let a twice continously differentiable function h(x) be a candidate CBF, and let the sets S, N, \mathcal{D} , \mathcal{C} , and $\overline{\mathcal{D}}$, and the data-sets $X_{\mathcal{N}}$, X_{safe} , and $\overline{X}_{\text{safe}}$ be defined as above. Suppose that $X_{\mathcal{N}}$ forms a $\overline{\epsilon}$ -net of \mathcal{N} satisfying the conditions of Proposition III.1, and that X_{safe} forms an ϵ -net of \mathcal{D} satisfying the conditions of Corollary III.2 & Proposition III.3. Then if h(x) satisfies constraints (7), (8), and (10), it holds that the set \mathcal{C} is non-empty, $\overline{\mathcal{D}} \subseteq \mathcal{C} \subset \mathcal{D} \subseteq S$, and the function h(x) is a valid local control barrier function on \mathcal{D} with domain $\mathcal{N} \cup \mathcal{D}$.

Control Barrier Filters: We introduce a simple and natural extension to the notion of a local CBF. Consider the exact same scenario as above, with the additional set $\mathcal{F} \subseteq \mathcal{S} \setminus \mathcal{C}$ that satisfies the following condition: for each $\zeta_0 \in \mathcal{F}$ there exists no continuous signal $\zeta : \mathbb{R}_{\geq 0} \to \mathbb{R}^n$ with $\zeta(0) := \zeta_0$ and with $\zeta(t') \notin \mathcal{S}$ for some t' > 0 and $\zeta(t'') \notin \mathcal{C}$ for all t'' > 0. This means that the set \mathcal{C} filters all trajectories starting from \mathcal{F} , i.e., each trajectory starting from \mathcal{F} has to pass through \mathcal{C} to escape \mathcal{S} and thereby renders \mathcal{F} safe (see Fig. 1(right)). This follows in the spirit of set invariance [24]. As illustrated in Fig. 1(right), this allows us to remove the perhaps counter-intuitive requirement of having to introduce "artificial" unsafe samples in a region that is clearly safe, further reducing the conservatism of the resulting controller.

C. Computing a Control Barrier Function

Using the results of the previous subsection, we propose solving the following optimization problem to learn a CBF from expert trajectories:

$$\min_{h \in \mathcal{H}} \|h\|$$
s.t. $h(x_i) \ge \gamma_{\text{safe}}, \quad \forall x_i \in \bar{X}_{\text{safe}}(L_h)$
 $h(x_i) \le -\gamma_{\text{unsafe}},$

$$\operatorname{Lip}(h(x_i), \bar{\epsilon}) \le L_h, \quad \forall x_i \in X_{\mathcal{N}},$$
 $q(x_i, u_i) := \langle \nabla h(x_i), f(x_i, u_i) \rangle + \alpha(h(x_i)) \ge \gamma_{\text{dyn}},$

$$\operatorname{Lip}(q(x_i, u_i), \epsilon) \le L_q, \quad \forall (x_i, u_i) \in Z_{\text{dyn}}.$$
(11b)

The positive constants γ_{safe} , γ_{unsafe} , γ_{dyn} , L_h and L_q are hyperparameters that are set according to the conditions of Theorem III.4 given data-sets X_{safe} and X_N defining corresponding ϵ and $\bar{\epsilon}$ -nets. Here the constraints defined in equations (11a) and (11b) assume that there exists a function $\text{Lip}(\cdot, \epsilon)$ that returns an upper bound on the Lipschitz constant of its argument in an ϵ neighborhood. We note that it may be difficult to enforce these bounds while solving the optimization problem, in which case we must resort to bootstrapping the values of L_h and L_q by iteratively solving optimization problem (11), computing the values L_h and L_q for the learned CBF h(x), verifying if the conditions of Theorem III.4 hold, and readjusting the hyperparameters accordingly if not. This is a standard approach to hyperparameter tuning, and we show in Section IV that it can indeed be successfully applied to verify the conditions of Theorem III.4 in practice.

Convexity: We first note that optimization problem (11) is convex in *h* if the function α is linear in its argument, and if we exclude the bounds (11a) and (11b), and instead verify them via the bootstrapping method described above. Therefore, if \mathcal{H} is parameterized as $\mathcal{H} = \{h_{\theta}(\cdot) = \langle \phi(\cdot), \theta \rangle : \theta \in \Theta\}$ with Θ a convex set and $\phi(\cdot)$ a known but possibly nonlinear transformation, then problem (11) is convex, and can be solved efficiently using standard solvers. Note that very rich function classes such as infinite dimensional RKHS from statistical learning theory can be approximated to arbitrary accuracy as such a \mathcal{H} [22].

In the more general case when $\mathcal{H} = \{h_{\theta}(\cdot) : \theta \in \Theta\}$, such as when h is a DNN or when α is a general nonlinear function of its argument, optimization problem (11) is nonconvex. Due to the computational complexity of general nonlinear constrained programming, we propose an unconstrained relaxation of problem (11) which can be solved efficiently in practice by first order gradient based methods. Let $[x]_{+} = \max\{x, 0\}$ for $x \in \mathbb{R}$. Our unconstrained relaxation is the following optimization problem:

$$\min_{\theta \in \Theta} \|\theta\|^{2} + \lambda_{s} \sum_{x_{i} \in \bar{X}_{safe}} [\gamma_{safe} - h_{\theta}(x_{i})]_{+}$$

$$+ \lambda_{u} \sum_{x_{i} \in X_{\mathcal{N}}} [h_{\theta}(x_{i}) + \gamma_{unsafe}]_{+}$$

$$+ \lambda_{d} \sum_{(x_{i}, u_{i}) \in Z_{dyn}} [\gamma_{dyn} - (\langle \nabla h_{\theta}(x_{i}), f(x_{i}, u_{i}) \rangle$$

$$+ \alpha(h_{\theta}(x_{i}))]_{+}.$$
(12)

The positive parameters λ_s , λ_u , λ_d allow us to trade off the relative importance of each of the terms in the optimization. While equation (12) is in general a non-convex optimization problem, it can be solved efficiently in practice with stochastic first-order gradient methods such as Adam or SGD.

Lipschitz continuity of \mathcal{H} : As described earlier, because we assume that functions in \mathcal{H} are twice continuously differentiable and we restrict ourselves to a compact domain $\mathcal{N} \cup \mathcal{D}$, we immediately have that h and ∇h are both uniformly Lipschitz over $\mathcal{N} \cup \mathcal{D}$. We show here two examples of \mathcal{H} where it is computationally efficient to estimate an upper bound on the Lipschitz constants of functions $h \in \mathcal{H}$.

In the case of random Fourier features with ℓ random features, where $h(x) = \langle \phi(x), \theta \rangle$ and $\phi(x) \in \mathbb{R}^{\ell}$ is

$$\phi(x) = \sqrt{2/\ell} (\cos(\langle x, w_1 \rangle + b_1), ..., \cos(\langle x, w_\ell \rangle + b_\ell)),$$

then we can analytically compute upper bounds as follows. First, we have by the Cauchy-Schwarz inequality $|h(x_1) - h(x_2)| \leq \|\phi(x_1) - \phi(x_2)\|_2 \|\theta\|_2$. To bound $\|\phi(x_1) - \phi(x_2)\|_2$, we bound the spectral norm of the Jacobian $D\phi(x)$, which is a matrix where the *i*-th row is $-\sqrt{2/\ell} \sin(\langle x, w_i \rangle + b_i) w_i^{\mathsf{T}}$. Let $s_i = \sin(\langle x, w_i \rangle + b_i)$ and observe that $\|D\phi(x)\| = \sqrt{2/\ell} \sup_{\|v\|_2=1} (\sum_{i=1}^{\ell} s_i^2 \langle w_i, v \rangle^2)^{1/2} \leq \sqrt{2/\ell} \sup_{\|v\|_2=1} (\sum_{i=1}^{\ell} \langle w_i, v \rangle^2)^{1/2} = \sqrt{2/\ell} \|W\|$, where W is a matrix with the *i*-th row equal to w_i . While the bound $\sqrt{2/\ell} \|W\|$ can be used in computations, we can further understand order-wise scaling of the bound as follows. For random Fourier features corresponding to the popular Gaussian radial basis function kernel, $w_i \stackrel{\text{iid}}{\sim} \mathsf{N}(0, \sigma^2 I)$ where σ^2 is the (inverse) bandwidth of the Gaussian kernel. Therefore, by standard results in non-asymptotic random matrix theory [25], we have that $\|W\| \leq \sigma(\sqrt{\ell} + \sqrt{n} + \sqrt{2\log(1/\delta)})$ w.p. at least $1 - \delta$. Combining these calculations, we have that the Lipschitz constant of h can be bounded by $\sqrt{2\sigma^2}(1 + \sqrt{n/\ell} + \sqrt{(2/\ell)\log(1/\delta)}) \|\theta\|_2$ w.p. at least $1 - \delta$.

We now bound the Lipschitz constant of the gradient $\nabla h(x) = D\phi(x)^{\mathsf{T}}\theta$. We do this by bounding the spectral norm of the Hessian $\nabla^2 h(x) = -\sqrt{2/\ell} \sum_{i=1}^{\ell} c_i \theta_i w_i w_i^{\mathsf{T}}$, with $c_i = \cos(\langle x, w_i \rangle + b_i)$. A simple bound is $\|\nabla^2 h(x)\| \leq \sqrt{2/\ell} \|\theta\|_{\infty} \|W\|^2 \leq 3\sqrt{2} \|\theta\|_{\infty} \sigma^2 (\ell + n + 2\log(1/\delta))/\sqrt{\ell}$, where the last inequality holds w.p. at least $1 - \delta$.

When h(x) is a DNN, accurately estimating the Lipschitz constant is more involved. In general, the problem of exactly computing the Lipschitz constant of h is known to be NPhard [26]. Notably, because most commonly-used activation functions ϕ are known to be 1-Lipschitz (e.g. ReLU, tanh, sigmoid), a naive upper bound on the Lipschitz constant of his given by the product of the norms of the weight matrices; that is, $L_h \leq \prod_k ||W^k||$. However, this bound is known to be quite loose [23]. Recently, the authors of [23] proposed a semidefinite-programming based approach to efficiently compute an accurate upper bound on L_h . In particular, this approach relies on incremental quadratic constraints to represent the couplings between pairs of neurons in the neural network h. On the other hand, there are relatively few results that provide accurate upper bounds for the Lipschitz constant of the gradient of h when h is a neural network. While ongoing work looks to extend the results from [23] to compute upper bounds on $\operatorname{Lip}(\nabla h)$, to the best of our knowledge, the only general method for computing an upper bound on $\operatorname{Lip}(\nabla h)$ is through post-hoc sampling [27].

D. Data Collection

We briefly comment on how data should be collected to ensure that the conditions of Theorem III.4 are satisfied.

What should the experts do?: At a high level, our results state that if a smooth CBF can be found that satisfies the constraints (6), (7), and (10) over a sufficiently fine sampling of the state-space, then the resulting function is a valid CBF. We focus here on the derivative constraint (10), which must be verified to hold for some $u \in \mathcal{U}$, by using the expert example data (x_i, u_i) . In particular, the more transverse the vector field $f(x_i, u_i)$ is to the level sets of the learned CBF $h(x_i)$ (i.e., the more parallel it is to the inward pointing normal $\nabla h(x_i)$), the larger the inner-product term in constraint (10) without increasing the Lipschitz constant of h(x). In words, this says that the expert demonstrations should demonstrate how to move away from the unsafe set.

Constructing ϵ -nets: In order to construct a ϵ -net of a set S, a simple randomized algorithm which repeatedly uniformly samples from S works with high probability (see, for example, [25]). Hence, as long as we can efficiently sample from S (e.g. when S is a basic primitive set or has an efficient set-membership oracle), uniform sampling is a viable strategy. Alternatively, a gridding approach can be taken. We note that in either case, for a set of diameter r on the order of $O((\frac{r}{\epsilon})^d)$ samples are required. While this exponential dependence is undesirable, we observe that in practice, the expert demonstrations allow us to focus on a subset of the state-space associated with desirable behavior, significantly reducing the diameters of the sets to be sampled.

IV. NUMERICAL EXPERIMENTS

All code is publicly available at https://github. com/unstable-zeros/learning-cbfs.

1) *Planar Example:* Our first experiment is the following two dimensional planar system adapted from [28]:

$$\dot{x}_1 = -x_1 + (x_1^2 + \delta)u_1, \dot{x}_2 = -x_2 + (x_2^2 + \delta)u_2,$$
 (13)

where $\delta > 0$ is a fixed parameter guaranteeing that the system is globally feedback linearizable. We set $\delta = 1$ in our experiments. The desired safe set is $S = \{x : x_1 \leq 1, x_2 \leq 1\}$. We generate expert data for this system as follows. Because the system is feedback linearizable, given a desired trajectory $x_d(t)$, we can easily design a nominal controller which tracks $x_d(t)$. We can then construct a safe controller (w.r.t. S) by solving the CBF-QP problem [4], [5] with the CBF $h(x) = \min\{1 - x_1, 1 - x_2\}$.

We design two sets of desired trajectories. Let the unit vector $v(\theta) = [-\cos\theta \sin\theta]^{\mathsf{T}}$. The first set is defined for a fixed r > 0 as $x_d(t) = rv(t)$ from $t \in [0, 2\pi]$. We do this for $r \in \{0.2666, 0.3, 0.3333\}$, sampling 80 time equispaced points along each curve. The second set of desired trajectories are for a fixed $\theta \in [0, 2\pi]$, where we consider a trajectory that starts at $x(0) = 0.4666v(\theta)$ and ends up at $x(t_f) = 0.3666v(\theta)$, and one where $x(0) = 0.1333v(\theta)$ and $x(t_f) = 0.2333v(\theta)$. We grid across both $\theta \in [0, 2\pi]$ and $t \in$ $[0, t_f]$ to ensure a densely sampled set of points. All sample points Z_{dyn} are shown in Figure 2(left). We consider the x_i corresponding to the circular trajectories (green in Fig. 2) as defining X_{safe} . We then set X_{unsafe} to be points sampled (red in Fig. 2) along the circle at r = -0.5 and r = -0.1. Our samples are specifically chosen to form a net over \mathcal{D} and \mathcal{N} , with $\epsilon = 0.01666$ and $\bar{\epsilon} = 0.0333$, respectively.

We parameterize \mathcal{H} using $\ell = 200$ random Fourier features corresponding to the Gaussian kernel with $\sigma = 1.2$. We set $\alpha(x) = x$ and then solve the optimization problem with $\gamma_{\text{safe}} = 0.1$, $\gamma_{\text{unsafe}} = 0.3$, $\gamma_{\text{dyn}} = 0.01$ using cvxpy with the MOSEK backend. Next, we verify that our specific choices of γ_{safe} , γ_{unsafe} , γ_{dyn} satisfied the necessary conditions, computing $\|\nabla h(x_i)\|$ and $\|\nabla_x q(x_i, u_i)\|$ to obtain $L_h(x_i)$ and $L_q(x_i)$, respectively. This verification is shown in Fig. 3. The level sets of the resulting CBF h are plotted in Fig. 2(right).



Fig. 2: Left: Plot of the expert trajectories (green), dynamic samples (black) and unsafe samples (red) used for training. Right: Level set plot of the learned CBF for the two dimensional planar example. The dotted black line represents the boundary of the safe set S.



Fig. 3: Safe, derivative, and unsafe slacks which verify the sufficient conditions given in Theorem III.4. (a) the safe slack plot, for each constraint, shows the value $h(x_i) - L_h(x_i)\epsilon$, which needs to be positive. (b) the derivative slack plot, for each constraint, shows the value of $q(x_i, u_i) - L_q(x_i)\epsilon$, which also needs to be positive (c) the unsafe slack plot, for each constraint, shows the value of $h(x_i) + L_h(x_i)\overline{\epsilon}$, which needs to be negative.

2) Aircraft Collision Avoidance: In this subsection, we apply the control barrier filter technique (Sec. III-B) to the aircraft collision avoidance problem in [20]. The joint state vector of the two aircraft, indexed with a and b, is $x = [p_{x,a} \ p_{y,b} \ \theta_a \ p_{x,b} \ p_{y,b} \ \theta_b]^{\mathsf{T}} \in \mathbb{R}^6$, denoting positions in the (x, y)-plane and orientations. The controls $u = [v_a \ \omega_a \ v_b \ \omega_b]^{\mathsf{T}} \in \mathbb{R}^4$ are the translational and angular velocities with constraints $0.1 \le v_a \le 1.0$ and $-1.0 \le \omega_a \le 1.0$. The control goal is to reach the target states $x_g = [-5 \ 0 \ \pi \ 5 \ 0 \ 0]^{\mathsf{T}}$ if $p_{x,a}(0) \ge 0$, $p_{x,b}(0) \le 0$ or $x_g = [5 \ 0 \ 0 \ -5 \ 0 \ \pi]^{\mathsf{T}}$ if $p_{x,a}(0) \le 0$, $p_{x,b}(0) \ge 0$. The safety specification is that the two airplanes should maintain a minimal distance $D_s = 0.5$ to avoid collisions. To this end, we define the geometric safe set as,

$$\mathcal{S} := \left\{ x \in \mathbb{R}^6 \, \big| \, p_{x,r}^2 + p_{y,r}^2 \ge D_s^2 \right\}.$$
(14)

where $p_{i,r} = p_{i,a} - p_{i,b}, i \in \{x, y\}$ is the relative position.

Generating training data: We consider two ways of generating expert demonstrations. First, we used a standard tracking MPC as the nominal controller equipped with the closed form constructive CBF in [20] for collision avoidance (which we refer to as CBF-MPC). To generate the expert trajectories, we started the system from 400 randomly generated initial conditions inside the set S. Each run terminated when the airplanes were sufficiently far away from each other. Furthermore, we uniformly sampled safe and unsafe states; these states as well as the expert trajectories are shown in Figure 4 in relative coordinates.

Secondly, we built a web-based simulator that allows a user to control two simulated aerial vehicles. As before, the goal of the simulation is to control the two aerial vehicles such that they do not collide. We emphasize that these trajectories were solely by human guidance; no nominal controller was used. The data is plotted in Figure 4.



Fig. 4: Left: Plot of the expert trajectories Z_{dyn} generated by CBF-MPC (green), safe samples X_S (blue) and unsafe samples X_N (red). Right: Plot of the expert trajectories obtained from human demonstrations with the safe and unsafe samples.

Training procedure: We parametrized the CBF candidate h(x) with a two-hidden-layer fully-connected neural network with 64 neurons in each layer and tanh activation functions. The training procedure was implemented using jax [29] and the Adam algorithm with a cosine decay learning rate. We trained the neural network for 10^5 epochs using the loss in (12) with $\lambda_s = 2.0$, $\lambda_u = 2.0$, $\lambda_d = 15.0$, $\gamma_{safe} = 3.0$, $\gamma_{unsafe} = 0.5$ and $\gamma_{dyn} = 0.05$. Each of these hyperparameters was chosen via grid-search. The learned CBFs and the closed form CBF [20] evaluated at the training points are plotted in Figure 5 in relative coordinates.



Fig. 5: Left: The CBF learned from CBF-MPC (blue) and the closed form CBF in [20] (red) evaluated at the states in training data-set. Right: The CBF learned from human demonstrations.

Closed-loop control with learned CBF: To demonstrate the efficacy of the CBF learned from expert demonstrations, we used it in the aircraft collision avoidance problem with the same control goal and safety specification as in (14). The two airplanes were initialized at various symmetric initial positions on the circle $p_x^2 + p_y^2 = 1$ such that they were facing each other. In this way, if both airplanes used the nominal MPC controller, they would collide.

The closed-loop state trajectories using our learned CBF are shown in Figure 6. The CBFs learned on both data-sets

successfully steer the airplanes away from each other for all initial states, which experimentally validates the forward invariance of S. As a comparison, we also plotted the state trajectories produced by the CBF from [20] under the same settings in Figure 6. Since this CBF is derived analytically, it appears to render more aggressive control actions which manage to separate the airplanes at a closer distance.



Fig. 6: Closed-loop control of the two airplanes starting from multiple initial conditions using closed form CBF from [20] (top left), the CBF learned from CBF-MPC (top right), and the CBF learned from human demonstrations (bottom). Trajectories of the same run are marked with an identical color. The initial states of agent a are marked with diamonds, and those of agent b with circles.

V. CONCLUSION

We proposed and analyzed an optimization based approach to learning CBFs from expert demonstrations for known nonlinear control affine dynamical systems. We showed that under suitable assumptions of smoothness on the underlying dynamics and the learned CBF and under sufficiently fine sampling, the learned CBF is provably valid, guaranteeing safety. This work provides a firm theoretical foundation for future exploration that will look to leverage tools from statistical learning theory to reduce the sample complexity burden of the proposed method by focusing on guaranteeing safety for "typical" behaviors, as opposed to uniform coverage of the state-space.

REFERENCES

- S. Prajna, A. Jadbabaie, and G. J. Pappas, "A framework for worstcase and stochastic safety verification using barrier certificates," *IEEE Trans. Autom. Control*, vol. 52, no. 8, pp. 1415–1428, 2007.
- [2] P. A. Parrilo, "Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization," Ph.D. dissertation, California Institute of Technology, 2000.
- [3] P. Wieland and F. Allgöwer, "Constructive safety using control barrier functions," in *Proc. IFAC Symp. Nonlin. Control Syst.*, Pretoria, South Africa, August 2007, pp. 462–467.
- [4] A. D. Ames, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs with application to adaptive cruise control," in *Proc. Conf. Decis. Control*, Los Angeles, CA, December 2014, pp. 6271–6278.
- [5] 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, 2017.

- [6] E. D. Sontag, "A 'universal' construction of artstein's theorem on nonlinear stabilization," *Systems & control letters*, vol. 13, no. 2, pp. 117–123, 1989.
- [7] X. Xu, P. Tabuada, J. W. Grizzle, and A. D. Ames, "Robustness of control barrier functions for safety critical control," in *Proc. Conf. Analys. Design Hybrid Syst.*, Atlanta, GA, October 2015, pp. 54–61.
- [8] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: Theory and applications," in *Proc. European Control Conf.*, Naples, Italy, June 2019, pp. 3420– 3431.
- [9] X. Xu, J. W. Grizzle, P. Tabuada, and A. D. Ames, "Correctness guarantees for the composition of lane keeping and adaptive cruise control," *IEEE Transactions on Automation Science and Engineering*, vol. 15, no. 3, pp. 1216–1229, 2017.
- [10] L. Wang, D. Han, and M. Egerstedt, "Permissive barrier certificates for safe stabilization using sum-of-squares," in *Proc. American Control Conf.*, Milwaukee, WI, June 2018, pp. 585–590.
- [11] 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. Artificial Intel.*, Honolulu, HI, February 2019, pp. 3387–3395.
- [12] L. Wang, E. A. Theodorou, and M. Egerstedt, "Safe learning of quadrotor dynamics using barrier certificates," in *Proc. Conf. Robot. Automat.*, Brisbane, Australia, May 2018, pp. 2460–2465.
- [13] A. Taylor, A. Singletary, Y. Yue, and A. Ames, "Learning for safetycritical control with control barrier functions," in *Proc. Conf. Learning for Dynamics and Control*, June 2020, pp. 708–717.
- [14] A. J. Taylor, A. Singletary, Y. Yue, and A. D. Ames, "A control barrier perspective on episodic learning via projection-to-state safety," *arXiv* preprint arXiv:2003.08028, 2020.
- [15] S. Yaghoubi, G. Fainekos, and S. Sankaranarayanan, "Training neural network controllers using control barrier functions in the presence of disturbances," *arXiv preprint arXiv:2001.08088*, 2020.
- [16] W. Jin, Z. Wang, Z. Yang, and S. Mou, "Neural certificates for safe control policies," arXiv preprint arXiv:2006.08465, 2020.
- [17] N. M. Boffi, S. Tu, N. Matni, J.-J. E. Slotine, and V. Sindhwani, "Learning stability certificates from data," *arXiv preprint* arXiv:2008.05952, 2020.
- [18] M. Srinivasan, A. Dabholkar, S. Coogan, and P. Vela, "Synthesis of control barrier functions using a supervised machine learning approach," arXiv preprint arXiv:2003.04950, 2020.
- [19] M. Saveriano and D. Lee, "Learning barrier functions for constrained motion planning with dynamical systems," in *Proc. Conf. Intelligent Robots Systems*, Macau, China, November 2019.
- [20] E. Squires, P. Pierpaoli, and M. Egerstedt, "Constructive barrier certificates with applications to fixed-wing aircraft collision avoidance," in *Proc. Conf. Control Techn. Appl.*, Copenhagen, Denmark, August 2018, pp. 1656–1661.
- [21] U. Rosolia and F. Borrelli, "Learning model predictive control for iterative tasks. a data-driven control framework," *IEEE Trans. Autom. Control*, vol. 63, no. 7, pp. 1883–1896, 2017.
- [22] A. Rahimi and B. Recht, "Random features for large-scale kernel machines," in *Proc. Advances Neur. Inform. Proc. Syst.*, Vancouver, Canada, December 2008, pp. 1177–1184.
- [23] M. Fazlyab, A. Robey, H. Hassani, M. Morari, and G. Pappas, "Efficient and accurate estimation of lipschitz constants for deep neural networks," in *Proc. Advances Neur. Inform. Proc. Syst.*, Vancouver, Canada, December 2019, pp. 11423–11434.
- [24] F. Blanchini, "Set invariance in control," *Automatica*, vol. 35, no. 11, pp. 1747–1767, 1999.
- [25] R. Vershynin, High-dimensional prob.: An introduction with applications in data science. Cambridge university press, 2018, vol. 47.
- [26] A. Virmaux and K. Scaman, "Lipschitz regularity of deep neural networks: analysis and efficient estimation," in *Proc. Advances Neur. Inform. Proc. Syst.*, Montréal, Canada, December 2018, pp. 3835– 3844.
- [27] G. Wood and B. Zhang, "Estimation of the lipschitz constant of a function," *Journ. of Global Opt.*, vol. 8, no. 1, pp. 91–103, 1996.
- [28] S. Kolathaya and A. D. Ames, "Input-to-state safety with control barrier functions," *IEEE Control Systems Letters*, vol. 3, no. 1, pp. 108–113, 2019.
- [29] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, and S. Wanderman-Milne, "JAX: composable transformations of Python+NumPy programs," 2018. [Online]. Available: http://github.com/google/jax