Learning Hybrid Control Barrier Functions from Data

Lars Lindemann
KTH Royal Institute of Technology

Haimin Hu
Princeton University

Alexander Robey
University of Pennsylvania

Hanwen Zhang
University of Pennsylvania

Dimos V. Dimarogonas
KTH Royal Institute of Technology

Stephen Tu
Google Brain Robotics

Nikolai Matni
University of Pennsylvania

Abstract: Motivated by the lack of systematic tools to obtain safe control laws for hybrid systems, we propose an optimization-based framework for learning certifiably safe control laws from data. In particular, we assume a setting in which the system dynamics are known and in which data exhibiting safe system behavior is available. We propose hybrid control barrier functions for hybrid systems as a means to synthesize safe control inputs. Based on this notion, we present an optimization-based framework to learn such hybrid control barrier functions from data. Importantly, we identify sufficient conditions on the data such that feasibility of the optimization problem ensures correctness of the learned hybrid control barrier functions, and hence the safety of the system. We illustrate our findings in two simulations studies, including a compass gait walker.

Keywords: Control Barrier Functions, Hybrid Systems, Imitation Learning.

1 Introduction

Consider the following *safety-critical* scenarios: autonomous vehicles in urban areas [1], exoskeletons for improving mobility of lower-body impaired users [2], and robots navigating a warehouse using semantic logic [3]. These systems are all described by *hybrid dynamics*, i.e., states and transitions are both continuous and discrete, due to either their physics, or to higher level logical decision making and importantly, share that: 1) *data* exhibiting safe behavior is readily available or easily collected, and 2) in most cases, their hybrid system dynamics are well understood and can be identified. Based on these observations, we propose an optimization-based approach to learning *provably safe controllers* for hybrid systems using *hybrid control barrier functions* (HCBF).

Related work: *Barrier functions* (BFs) were introduced in [4] to certify the safety of continuous-time systems. Nonsmooth [5] and hybrid [6] BFs have also been defined. Different to our work, these focus on BFs with jumps for discontinuous systems. *Control barrier functions* (CBFs) for continuous-time control systems appeared in [7] for synthesizing feedback control laws that ensure safety by enforcing forward invariance of a desired safe set. Reciprocal [8] and zeroing [9] CBFs were proposed as less restrictive alternatives that do not enforce forward invariance on subsets of the safe set. Such CBFs can be used as a safety guard via convex quadratic programs [8, 9], a feature that has been used for safe learning [10, 11]. CBFs for discrete-time control systems can be found in [12, 13, 14]. Related to this paper, *hybrid BFs* were proposed in [15, 16] in the hybrid systems framework of Goebel et al. [17] as a means of certifying the safety of hybrid systems.

The underlying challenge and bottleneck in ensuring safety using (control) BFs is the construction of these functions. In [4], the authors propose a sum-of-squares programming approach for finding polynomial (hybrid) barrier functions for polynomial (hybrid) systems and semi-algebraic sets. For finding CBFs, bilinear sum-of-squares programming and analytic approaches are proposed in [18, 19, 20]; such methods, however, only apply to a restrictive class of systems and have limited

scalability. Recent approaches attempt to circumvent these limitations by treating the CBF synthesis task as a machine learning problem. In [21], a deep neural network is trained to imitate the control law obtained from a CBF, whereas in [22], a CBF is synthesized from safe and unsafe data samples using support vector machines. In Saveriano and Lee [23], the authors cluster observed data and learn a linear CBF for each cluster. While the aforementioned works [21, 22, 23] present detailed empirical validation of their methods, no formal correctness proofs are provided. In [24], a Lyapunov, barrier, and a policy function is learned from data: the validity of the learned certificates are then verified post-hoc using Lipschitz arguments. In [25], a method is proposed that learns a provably correct neural net safety guards for kinematic bicycle models. Finally, Robey et al. [26] propose a data-driven approach for learning CBFs for smooth nonlinear systems assuming known Lipschitz dynamics, as well as availability of expert demonstrations illustrating safe behavior. Sufficient conditions ensuring correctness of the learned CBF using a Lipschitz argument are also given.

Contributions: For a class of hybrid control systems, we define *hybrid control BFs* as a means to enforce forward invariance of a safe set. We provide sufficient conditions in terms of a HCBF that, if satisfied by a control law, ensure safety. We then show that learning HCBFs from data can be cast as a constrained optimization problem, and provide conditions under which a feasible solution is a valid HCBF. Finally, we present simulations showcasing the benefits of our framework.

2 Preliminaries and Problem Formulation

Notation: Let $\mathrm{dom}(z) := \{(t,j) \in \mathbb{R}_{\geq 0} \times \mathbb{N} | \exists \zeta \in \mathbb{R}^{n_z} \text{ s.t. } z(t,j) = \zeta\}$ be the domain of a function $z : \mathbb{R}_{\geq 0} \times \mathbb{N} \to \mathbb{R}^{n_z}$. A function $\alpha : \mathbb{R} \to \mathbb{R}$ is an extended class $\mathcal{K}$ function if α is strictly increasing and $\alpha(0) = 0$. For $\epsilon > 0$ and $p \geq 1$, let $\mathcal{B}_{\epsilon,p}(z^i) := \{z \in \mathbb{R}^{n_z} \, \big| \, \|z - z^i\|_p \leq \epsilon\}$ be the closed p-norm ball around $z^i \in \mathbb{R}^{n_z}$. Let $\mathrm{bd}(\mathcal{C})$ and $\mathrm{int}(\mathcal{C})$ be the boundary and interior of a set $\mathcal{C}$.

Control Barrier Functions At time $t \in \mathbb{R}_{\geq 0}$, let $x(t) \in \mathbb{R}^n$ be the state of the control system

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

where $f : \mathbb{R}^n \to \mathbb{R}^n$ and $g : \mathbb{R}^n \to \mathbb{R}^{n \times m}$ are continuous functions. Let the solutions to (1) under a 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. We do not assume forward completeness of (1) under u here, i.e., $\mathcal{I}$ may be bounded.

Consider next a continuously differentiable function $h : \mathbb{R}^n \to \mathbb{R}$ and define the set $\mathcal{C} := \{x \in \mathbb{R}^n \, \big| \, h(x) \geq 0\}$, 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. Note that $\mathcal{C}$ is closed and further assume that $\mathcal{C}$ is not the empty set. Now, let $\mathcal{D}$ be an open set that is such that $\mathcal{D} \supseteq \mathcal{C}$. The function $h(x)$ is said to be a *valid control barrier function* on $\mathcal{D}$ if there exists a locally Lipschitz continuous extended class $\mathcal{K}$ function $\alpha : \mathbb{R} \to \mathbb{R}$ such that

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

holds for all $x \in \mathcal{D}$, where $\mathcal{U} \in \mathbb{R}^n$ 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 $K_{\mathrm{CBF}}(x) := \{u \in \mathbb{R}^m \, \big| \, \langle \nabla h(x), f(x) + g(x)u \rangle \geq -\alpha(h(x))\}$. The next result follows mainly from Ames et al. [9].[1]

Lemma 1. *Assume that $h(x)$ is a valid control barrier function on $\mathcal{D}$ and that $u : \mathcal{D} \to \mathcal{U}$ is a continuous function with $u(x) \in K_{CBF}(x)$. Then $x(0) \in \mathcal{C}$ implies $x(t) \in \mathcal{C}$ for all $t \in \mathcal{I}$. If $\mathcal{C}$ is compact, it follows that $\mathcal{C}$ is forward invariant under $u(x)$, i.e., $\mathcal{I} = [0, \infty)$.*

Hybrid Systems We model and analyze hybrid systems using the formalism of [17, 28].

Definition 1. *A hybrid system [17] is a tuple $\mathcal{H} := (C, F, D, G)$ where $C \subseteq \mathbb{R}^{n_z}$, $D \subseteq \mathbb{R}^{n_z}$, $F : \mathbb{R}^{n_z} \to \mathbb{R}^{n_z}$, and $G : \mathbb{R}^{n_z} \to \mathbb{R}^{n_z}$ are the flow and jump sets and the continuous flow and jump maps, respectively. At the hybrid time $(t,j) \in \mathbb{R}_{\geq 0} \times \mathbb{N}$, let $z(t,j) \in \mathbb{R}^{n_z}$ be the hybrid state with initial condition $z(0,0) \in C \cup D$ and the hybrid system dynamics*

$$\dot{z}(t,j) = F(z(t,j)) \text{ for } z(t,j) \in C, \quad z(t,j+1) = G(z(t,j)) \text{ for } z(t,j) \in D. \tag{2}$$

---

[1]We provide a slightly modified version to account for Ames et al. [18, Remark 5] and to not require that $\nabla h(x) \neq 0$ when $x \in \mathrm{bd}(\mathcal{C})$. The proof (in Appendix A) uses the Comparison Lemma [27, Lemma 3.4].

Solutions to (2) are parameterized by (t, j), where t indicates continuous *flow* according to $F(z)$ and j indicates discontinuous *jumps* according to $G(z)$. Now let $\mathcal{E} \subseteq \mathbb{R}_{\geq 0} \times \mathbb{N}$ be a *hybrid time domain* [17, Ch. 2.2], i.e., $\mathcal{E}$ is an infinite union of intervals of the form $[t_j, t_{j+1}] \times \{j\}$ or a finite union of intervals of the form $[t_j, t_{j+1}] \times \{j\}$ where the last interval, if it exists, has the form $[t_j, t_{j+1}] \times \{j\}$, $[t_j, t_{j+1}) \times \{j\}$, or $[t_j, \infty) \times \{j\}$. We now formally define a *hybrid solution* $z : \mathcal{E} \to C \cup D$ to $\mathcal{H}$.

Definition 2. *A function $z : \mathcal{E} \to C \cup D$ is a hybrid solution to $\mathcal{H}$ if $z(0, 0) \in C \cup D$ and*

- *for each $j \in \mathbb{N}$ such that $I_j := \{t \in \mathbb{R}_{\geq 0} | (t, j) \in dom(z)\}$ is not a singleton, $z(t, j) \in C$ and $\dot{z}(t, j) = F(z(t, j))$ for all $t \in [\min I_j, \sup I_j)$*
- *for each $(t, j) \in dom(z)$ s.t. $(t, j + 1) \in dom(z)$, $z(t, j) \in D$ and $z(t, j + 1) = G(z(t, j))$.*

Example 1. *Consider the bouncing ball example from Goebel et al. [17, Example 1.1]. Let $x \in \mathbb{R}$ be the (vertical) position and $v \in \mathbb{R}$ be the (vertical) velocity of a point mass. Let $z := \begin{bmatrix} x & v \end{bmatrix}^T$ and*

$$\dot{z}(t, j) = \begin{bmatrix} v(t, j) & -\rho \end{bmatrix}^T \text{ for } z(t, j) \in C, \quad z(t, j + 1) = \begin{bmatrix} x(t, j) & -\kappa v(t, j) \end{bmatrix}^T \text{ for } z(t, j) \in D,$$

where $\rho > 0$ (gravity), $\kappa \in (0, 1)$ (damping), $C := \{z \in \mathbb{R}^2 | x > 0 \text{ or } x = 0, v \geq 0\}$, and $D := \{z \in \mathbb{R}^2 | x = 0, v < 0\}$. For an initial condition $z(0, 0) := \begin{bmatrix} x(0, 0) & v(0, 0) \end{bmatrix}^T$ with $x(0, 0) > 0$, the hybrid time domain $\mathcal{E}$ is an infinite union of intervals of the form $[t_j, t_{j+1}] \times \{j\}$ where, as $j \to \infty$, t_j converges to some constant value and where $\|z(t, j)\|$ converges to zero.

Problem Formulation The class of *hybrid control systems* that we consider is

$$\begin{cases} \dot{z}(t, j) = f_c(z(t, j)) + g_c(z(t, j))u_c(z(t, j)) & \text{for } z(t, j) \in C \\ z(t, j + 1) = f_d(z(t, j)) + g_d(z(t, j))u_d(z(t, j)) & \text{for } z(t, j) \in D \end{cases} \quad (3)$$

where $u_c : C \to \mathcal{U}_c \subseteq \mathbb{R}^{m_c}$ and $u_d : D \to \mathcal{U}_d \subseteq \mathbb{R}^{m_d}$ are continuous control laws during flows and jumps, respectively. Let $f_c : \mathbb{R}^{n_z} \to \mathbb{R}^{n_z}$, $f_d : \mathbb{R}^{n_z} \to \mathbb{R}^{n_z}$, $g_c : \mathbb{R}^{n_z} \to \mathbb{R}^{n_z \times n_c}$, and $g_d : \mathbb{R}^{n_z} \to \mathbb{R}^{n_z \times m_d}$ be locally Lipschitz continuous functions.[2]

We are given a set of *expert trajectories* consisting of N_c and N_d discretely sampled data-points along flows and jumps as $Z_{\text{dyn}}^c := \{(z^i, u_c^i)\}_{i=1}^{N_c}$ for $z^i \in C$ and $Z_{\text{dyn}}^d := \{(z^i, u_d^i)\}_{i=1}^{N_d}$ for $z^i \in D$, as illustrated in Figure 1 (left).[3] It is assumed that each $z^i \in \text{int}(\mathcal{S})$ where $\mathcal{S} \subseteq \mathbb{R}^{n_z}$ is the *geometric safe set*, i.e., the set of safe states as naturally specified on a subset of the system configuration space (e.g., to avoid collision, vehicles must maintain a minimum separating distance). Our goal is now to learn, from Z_{dyn}^c and Z_{dyn}^d, a twice continuously differentiable function $h : \mathbb{R}^{n_z} \to \mathbb{R}$ such that

$$\mathcal{C} := \{z \in \mathbb{R}^{n_z} \mid h(z) \geq 0\} \quad (4)$$

is a subset of the geometric safe set $\mathcal{S}$, and that $\mathcal{C}$ can be made forward invariant by appropriate control actions $u_c(z) \in \mathcal{U}_c$ and $u_d(z) \in \mathcal{U}_d$.

3 Learning Hybrid Control Barrier Functions from Data

We begin by defining a suitable notion of *hybrid control barrier functions*, and show that they provide a mechanism for safe control of the hybrid control system (3). We then show how such HCBFs can be learned from data via a constrained optimization problem, and provide sufficient conditions under which a feasible solution is a valid HCBF.

Hybrid Control Barrier Functions Let $h : \mathbb{R}^{n_z} \to \mathbb{R}$ be a twice continuously differentiable function for which the set $\mathcal{C}$ in (4) is not empty and such that $\mathcal{C} \subseteq C \cup D$. Such an assumption is natural as we are only interested in regions where the system (3) is defined. Consider now the sets $\mathcal{D}_C \subseteq C$ and $\mathcal{D}_D \subseteq D$ that are such that $\mathcal{C} \cap C \subseteq \mathcal{D}_C$ and $\mathcal{C} \cap D \subseteq \mathcal{D}_D$ so that $\mathcal{C} \subseteq \mathcal{D}_C \cup \mathcal{D}_D$[4], ensuring that the set $\mathcal{D} := \mathcal{D}_C \cup \mathcal{D}_D$ fully covers $\mathcal{C}$ – see Fig. 1(middle) and (right). The sets $\mathcal{D}_C$ and $\mathcal{D}_D$ are the equivalent to the set $\mathcal{D}$ in Section 2, but considered separately for flows and jumps.

---

[2]We again do not assume completeness of the system (3) under u_c and u_d. Completeness here means that the hybrid time domain $\text{dom}(z)$ is unbounded (see Goebel et al. [17, Ch. 2.2] for a formal definition).

[3]We refer to the collection of data points Z_{dyn}^c and Z_{dyn}^d as expert trajectories to emphasize that this is a natural way of collecting the $\{(z^i, u_c^i)\}$ and $\{(z^i, u_d^i)\}$ pairs from the system (3). We note, however, that our method simply requires a collection of state-action pairs $\{(z^i, u_c^i)\}$ and $\{(z^i, u_d^i)\}$ demonstrating safe behavior.

[4]This follows as we assume that $\mathcal{C} \subseteq C \cup D$, which results in $(\mathcal{C} \cap C) \cup (\mathcal{C} \cap D) = \mathcal{C} \cap (C \cup D) = \mathcal{C}$, and since $(\mathcal{C} \cap C) \cup (\mathcal{C} \cap D) \subseteq \mathcal{D}_C \cup \mathcal{D}_D$ by the choices of the sets $\mathcal{D}_C$ and $\mathcal{D}_D$.
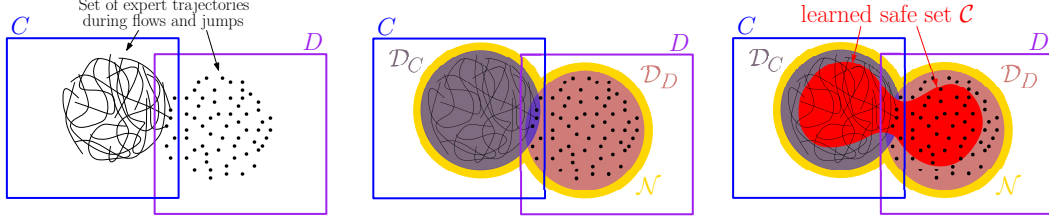
Figure 1: Problem setup (left): The flow and jump sets C and D (blue and purple boxes) and the set of safe expert trajectories during flows and jumps (black lines and dots). Set definitions (middle): The sets $\mathcal{D}_C$ and $\mathcal{D}_D$ (black and light red balls) are the union of ϵ balls around the expert trajectories during flows and jumps. The set $\mathcal{N}$ (golden rings), defined around $\mathcal{D}_C$ and $\mathcal{D}_D$, ensures that the learned safe set $\mathcal{C}$ is such that $\mathcal{C} \subset \mathcal{D} \subseteq \mathcal{S}$. Note that the geometrical safe set $\mathcal{S}$ is not depicted here. Desired result (right): The learned safe set $\mathcal{C}$ (red region) is defined via the learned valid HCBF $h(z)$.

Definition 3. *The function $h(z)$ is said to be a* valid hybrid control barrier function *on $\mathcal{D}$ if there exists a locally Lipschitz continuous extended class $\mathcal{K}$ function $\alpha : \mathbb{R} \to \mathbb{R}$ such that*

$$\sup_{u_c \in \mathcal{U}_c} \langle \nabla h(z), f_c(z) + g_c(z)u_c \rangle \geq -\alpha(h(z)) \text{ for all } z \in \mathcal{D}_C, \tag{5}$$

$$\sup_{u_d \in \mathcal{U}_d} h(f_d(z) + g_d(z)u_d) \geq 0 \text{ for all } z \in \mathcal{D}_D. \tag{6}$$

We define the sets of HCBF consistent inputs to be $K_{\text{HCBF},c}(z) := \{u_c \in \mathcal{U}_c \,\big|\, \langle \nabla h(z), f_c(z) + g_c(z)u_c \rangle \geq -\alpha(h(z))\}$ and $K_{\text{HCBF},d}(z) := \{u_d \in \mathcal{U}_d \,\big|\, h(f_d(z) + g_d(z)u_d) \geq 0\}$. Assume for the remainder of the paper that the set $\mathcal{D}_C$ is open.[5] This will allow us, in the next result, to establish forward invariance of the set $\mathcal{C}$ under control laws $u_c(z)$ and $u_d(z)$ when the set $\mathcal{C}$ is compact.

Theorem 1. *Assume that $h(z)$ is a valid hybrid control barrier function on $\mathcal{D}$ and that $u_c : \mathcal{D}_C \to \mathcal{U}_c$ and $u_d : \mathcal{D}_D \to \mathcal{U}_d$ are continuous functions with $u_c(z) \in K_{HCBF,c}(z)$ and $u_d(z) \in K_{HCBF,d}(z)$. Then $z(0, 0) \in \mathcal{C}$ implies $z(t, j) \in \mathcal{C}$ for all $(t, j) \in dom(z)$. If $\mathcal{C}$ is compact and satisfies $\mathcal{C} \subseteq C \cup D$, then the set $\mathcal{C}$ is forward invariant under $u_c(z)$ and $u_d(z)$, i.e., $dom(z)$ is unbounded.*

Proof. The proof is based on showing that $h(t, j) \geq 0$ for all $(t, j) \in \text{dom}(z)$ due to the constraints (5) and (6) and that $z(t, j)$ can not leave $\mathcal{D}_C \cup \mathcal{D}_D$. The complete proof is given in Appendix A. $\square$

Example 2. *Consider again the bouncing ball example from Example 1, here artificially equipped with continuous and discrete control inputs for illustrative purposes. The dynamics are*

$$\dot{z}(t, j) = [v(t, j) \quad u_c - \rho]^T \text{ for } z(t, j) \in C, \; z(t, j+1) = [x(t, j) \quad -u_d v(t, j)]^T \text{ for } z(t, j) \in D,$$

and the geometric safe set is $\mathcal{S} := \{z \in \mathbb{R}^2 | v^2 \leq \zeta\}$. The artificial control inputs u_c and u_d can be thought of as controllable wind resistance and floor damping, respectively. Define the candidate HCBF $h(z) := \zeta - v^2$, such that $\mathcal{C} = \mathcal{S}$. Evaluating constraint (5) with $\alpha(r) := r$ yields $-2vu_c + 2v\rho \geq -\zeta + v^2$ and for which we can always find a suitable $u_c \in \mathbb{R}$. Similarly, evaluating constraint (6) yields $\zeta - (u_d v)^2 \geq 0$ and for which we can, again, always find a suitable $u_d \in \mathbb{R}$.

An Optimization Based Approach First, define the data sets $Z_{\text{safe}}^c := \{z^i : (z^i, u_c^i) \in Z_{\text{dyn}}^c\}$ and $Z_{\text{safe}}^d := \{z^i : (z^i, u_d^i) \in Z_{\text{dyn}}^d\}$. We next define, for $\epsilon_c, \epsilon_d > 0$ and $p \geq 1$, the sets

$$\mathcal{D}_D := D \cap \bigcup_{z^i \in Z_{\text{safe}}^d} \mathcal{B}_{\epsilon_d, p}(z^i) \text{ and } \mathcal{D}_C := \mathcal{D}_C' \backslash \text{bd}(\mathcal{D}_C') \text{ where } \mathcal{D}_C' := C \cap \bigcup_{z^i \in Z_{\text{safe}}^c} \mathcal{B}_{\epsilon_c, p}(z^i) \tag{7}$$

that need to be such that $\mathcal{D} = \mathcal{D}_C \cup \mathcal{D}_D \subseteq \mathcal{S}$, which can be easily achieved even when data-points z^i are close to $\text{bd}(\mathcal{S})$ by adjusting ϵ_c and ϵ_d or by omitting z^i. Note that the set $\mathcal{D}_C$ is open by definition. For $\sigma > 0$, define $\mathcal{N} := \{\text{bd}(\mathcal{D}) \oplus \mathcal{B}_{\sigma, p}(0)\} \backslash \mathcal{D}$, where $\mathcal{N}$ is a ring of diameter σ that surrounds the set $\mathcal{D}$ (see the golden ring in Figure 1). We will use the set $\mathcal{N}$ to enforce that the value of the learned HCBF $h(z)$ is negative on $\mathcal{N}$ to ensure that the set $\mathcal{C}$ is contained within the set $\mathcal{D}$.

---

[5]If $\mathcal{D}_C$ is not open, one can instead assume that $\mathcal{C} \setminus \mathcal{D}_D$ is strictly contained within $\mathcal{D}_C \cup \mathcal{D}_D$.

Hence, also assume that points $Z_N = \{z^i\}_{i=1}^{N_u}$ are sampled from $\mathcal{N}$, i.e., $z^i \in \mathcal{N}$. Note that only state data $z^i \in Z_N$ are sampled from $\mathcal{N}$, while no inputs u_c^i or u_d^i are needed, i.e., unsafe data can be obtained by gridding or sampling. While the set $\mathcal{C}$ defined in (4) considers all $z \in \mathbb{R}^{n_z}$ such that $h(z) \geq 0$, we modify this definition slightly by restricting the domain to the set $\mathcal{N} \cup \mathcal{D}$. This is a natural restriction as we are learning a HCBF from data sampled over the domain $\mathcal{N} \cup \mathcal{D}$, and we therefore instead consider learning a valid *local* HCBF $h(z)$ over $\mathcal{D}$ with respect to the set

$$\mathcal{C} := \{z \in \mathcal{N} \cup \mathcal{D} \,|\, h(z) \geq 0\}. \tag{8}$$

The optimization problem: We now propose an optimization problem for learning a valid local HCBF, and then prove its correctness. We solve

$$\min_{h \in \mathcal{H}} \ \|h\|$$

$$\text{s.t. } h(z^i) \geq \gamma_{\text{safe}}, \ \forall z^i \in Z_{\text{safe}}^c \cup Z_{\text{safe}}^d \qquad\qquad : \text{safe set with margin } \gamma_{\text{safe}} > 0 \quad (9a)$$

$$\qquad h(z^i) \leq -\gamma_{\text{unsafe}}, \ \forall z^i \in Z_N \qquad\qquad : \text{unsafe set with margin } \gamma_{\text{unsafe}} > 0 \quad (9b)$$

$$\qquad \text{Lip}(h(z^i), \bar{\epsilon}) \leq L_h, \ \forall z^i \in Z_N \qquad\qquad : \text{Lipschitz constraint on } h(z) \quad (9c)$$

$$\qquad q_c(z^i, u_c^i) := \langle \nabla h(z^i), f_c(z^i) + g_c(z^i)u_c^i \rangle$$
$$\qquad\qquad\qquad + \alpha(h(z^i)) \geq \gamma_{\text{dyn}}^c \qquad : \text{grad. constraint (5) with margin } \gamma_{\text{dyn}}^c > 0 \quad (9d)$$

$$\qquad \text{Lip}(q_c(z^i, u_c^i), \epsilon_c) \leq L_q^c, \ \forall (z^i, u_c^i) \in Z_{\text{dyn}}^c \qquad : \text{Lipschitz constraint on } q_c(z^i, u_c^i) \quad (9e)$$

$$\qquad q_d(z^i, u_d^i) := h(f_d(z^i) + g_d(z^i)u_d^i) \geq \gamma_{\text{dyn}}^d \quad : \text{jump constraint (6) with margin } \gamma_{\text{dyn}}^d > 0 \quad (9f)$$

$$\qquad \text{Lip}(q_d(z^i, u_d^i), \epsilon_d) \leq L_q^d, \ \forall (z^i, u_d^i) \in Z_{\text{dyn}}^d \qquad : \text{Lipschitz constraint on } q_d(z^i, u_d^i) \quad (9g)$$

where $\mathcal{H}$ is a normed function space and where the positive constants γ_{safe}, γ_{unsafe}, $\gamma_{\text{dyn}}^c, \gamma_{\text{dyn}}^d$, L_h, L_q^c, and L_q^d are *hyperparameters* determined by the data-sets Z_{safe}^c, Z_{safe}^d, and Z_N, which must be sufficiently dense, as quantified by $\bar{\epsilon}$, ϵ_c, and ϵ_d (conditions given below).

Lipschitz bounds: The constraints in (9c), (9e), and (9g) assume a function $\text{Lip}(\cdot, \epsilon)$ that returns an upper bound on the Lipschitz constant of its argument with respect to the state z in an ϵ neighborhood of z^i using the p-norm as detailed in Appendix A. It may be difficult to enforce the constraints (9c), (9e), and (9g) in the optimization problem (9). One can resort to bootstrapping the values of L_h, L_q^c, and L_q^d by iteratively solving the optimization problem (9), calculating the values of L_h, L_q^c, and L_q^d to verify if constraints (9c), (9e), and (9g) hold, and adjusting regularization hyperparameters accordingly. Appendix A explains how to compute Lipschitz constants for DNNs [29] and RKHS.

Should we trust the experts? Optimization problem (9) approximates the supremums over continuous and discrete inputs in the constraints (5) and (6) with the expert actions – while computationally expedient, this may prove conservative. We note that in many cases of interest, the supremum over the continuous input $u_c \in \mathcal{U}_c$ in constraint (5) admits a closed form expression. For example, when $\mathcal{U}_c$ is an $\|\cdot\|$-norm ball, the left hand side of constraint (5) reduces to $\langle \nabla h(z^i), f_c(z^i) \rangle + \|\nabla h(z^i)g_c^T(z^i)\|_\star + \alpha(h(z^i))$, for $\|\cdot\|_\star$ the dual norm. This in turn can be used to simplify constraint (9d) in optimization problem (9), and in particular eliminates the dependency on u_c^i. Nevertheless, the availability of expert demonstrations is still valuable as they indicate that a safe action exists, and we therefore expect a feasible HCBF h and control action u_c to exist.

Unconstrained relaxation: For general function classes $\mathcal{H} := \{h(z; \theta)|\theta \in \Theta\}$, e.g., when $h(z)$ is a DNN, optimization problem (9) is nonconvex: we propose an unconstrained relaxation that can be solved efficiently using stochastic first-order gradient methods such as Adam or stochastic gradient descent. Let $[r]_+ := \max\{r, 0\}$ for $r \in \mathbb{R}$ and let the unconstrained optimization problem be

$$\min_{\theta \in \Theta} \ \|\theta\|^2 + \lambda_s \sum_{z^i \in Z_{\text{safe}}^c \cup Z_{\text{safe}}^d} [\gamma_{\text{safe}} - h_\theta(z^i)]_+ + \lambda_u \sum_{z^i \in Z_N} [h_\theta(z^i) + \gamma_{\text{unsafe}}]_+ + \lambda_d \sum_{(z^i, u_c^i) \in Z_{\text{dyn}}^c} [\gamma_{\text{dyn}}^c$$

$$- \langle \nabla h_\theta(z^i), f_c(z^i) + g_c(z^i)u_c^i \rangle - \alpha(h_\theta(z^i))]_+ + \lambda_c \sum_{(z^i, u_d^i) \in Z_{\text{dyn}}^d} [\gamma_{\text{dyn}}^d - h_\theta(f_d(z^i) + g_d(z^i)u_d^i)]_+$$

where $\lambda_s, \lambda_u, \lambda_d, \lambda_c > 0$ are hyperparameters that tradeoff between the constraints of problem (9).

Guaranteeing Safety We next show correctness of the learned HCBF $h(z)$ obtained from (9) in two steps by: 1) showing that the certified safe set (8) is contained within the geometric safe set, i.e., that $\mathcal{C} \subset \mathcal{D} \subseteq \mathcal{S}$, and 2) proving that $h(z)$ is a valid local HCBF by ensuring that the set $\mathcal{C}$ is forward invariant under control laws $u_c(z) \in \mathcal{U}_c$ and $u_d(z) \in \mathcal{U}_d$. All proofs are in the Appendix.

1) Guaranteeing $\mathcal{C} \subset \mathcal{D} \subseteq \mathcal{S}$: Let $h(z)$ be Lipschitz continuous with local constant $L_h(z)$,[6] and Z_N be an $\bar{\epsilon}$-net of $\mathcal{N}$, i.e., for all $z \in \mathcal{N}$, there exists $z^i \in Z_N$ such that $\|z^i - z\|_p \leq \bar{\epsilon}$. Assume that $\bar{\epsilon} < \gamma_{\text{unsafe}}/L_h(z^i)$ for all $z^i \in Z_N$ and $\gamma_{\text{unsafe}} > 0$. Then a standard covering and Lipschitz argument (see Proposition 3 in Appendix A) shows that if $h(z)$ satisfies constraint (9b) for all $z^i \in Z_N$, we have that $h(z) < 0$ for all $z \in \mathcal{N}$. We can use a similar argument on constraint (9a), which ensures that the set $\mathcal{C}$ over which $h(z) \geq 0$, as defined in equation (8), has non-empty interior.

Proposition 1. *Let $h(z)$ be Lipschitz continuous with local constant $L_h(z)$, and Z_{safe}^c and Z_{safe}^d be ϵ_c- and ϵ_d-nets of $\mathcal{D}_C$ and $\mathcal{D}_D$, respectively, with $\max(\epsilon_c, \epsilon_d) \leq \gamma_{safe}/L_h(z^i)$ for all $z^i \in Z_{safe}^c \cup Z_{safe}^d$ and $\gamma_{safe} > 0$. Then, if $h(z)$ satisfies constraint (9a), we have that $h(z) \geq 0$ for all $z \in \mathcal{D}$.*

Propositions 3 and 1 then ensure that $\mathcal{C} \subset \mathcal{D} \subseteq \mathcal{S}$. Note that the constraints (9a) and (9b) may lead to infeasibility of (9). Similarly to Robey et al. [26, equation (3.4)], we instead enforce the constraint (9a) on smaller sets $\bar{Z}_{\text{safe}}^c$ and $\bar{Z}_{\text{safe}}^d$ with $\bar{Z}_{\text{safe}}^c \subset Z_{\text{safe}}^c$ and $\bar{Z}_{\text{safe}}^d \subset Z_{\text{safe}}^d$, i.e., replace (9a) by

$$h(z^i) \geq \gamma_{\text{safe}}, \ \forall z^i \in \bar{Z}_{\text{safe}}^c \cup \bar{Z}_{\text{safe}}^d. \tag{10}$$

2) Guaranteeing a valid local HCBF: For a state $z \in \mathcal{D}_C$, define the function $q_c(z, u_c^i) := \langle \nabla h(z), f_c(z) + g_c(z)u_c^i \rangle + \alpha(h(z))$ where the control sample u_c^i is associated with the sample z^i that is such that $\|z^i - z\|_p \leq \epsilon_c$. Note that such a pair $(z^i, u_c^i) \in Z_{\text{dyn}}^c$ is guaranteed to exist when the set Z_{safe}^c is an ϵ_c-net of $\mathcal{D}_D$. The function $q_c(z, u_c^i)$ is Lipschitz continuous in z with local constant denoted by $L_q^c(z, u_c^i)$, as we have assumed h to be twice continuously differentiable. Similarly, define $q_d(z, u_d^i) := h(f_d(z) + g_d(z)u_d^i)$ and note that $q_d(z, u_d^i)$ is Lipschitz continuous with local constant denoted by $L_q^d(z, u_d^i)$. We next provide conditions guaranteeing that the learned HCBF satisfies the constraint (5) for all $z \in \mathcal{D}_C$ and the constraint (6) for all $z \in \mathcal{D}_D$.

Proposition 2. *Suppose $q_c(z, u_c^i)$ and $q_d(z, u_d^i)$ are Lipschitz continuous with local constants $L_q^c(z, u_c^i)$ and $L_q^d(z, u_d^i)$, respectively. Let $\gamma_{dyn}^c, \gamma_{dyn}^d > 0$, and assume that (i) Z_{safe}^c is an ϵ_c-net of $\mathcal{D}_C$ with $\epsilon_c \leq \gamma_{dyn}^c/L_q^c(z^i, u_c^i)$ for all $(z^i, u_c^i) \in Z_{dyn}^c$, and (ii) Z_{safe}^d is an ϵ_d-net of $\mathcal{D}_D$ with $\epsilon_d \leq \gamma_{dyn}^d/L_q^d(z^i, u_d^i)$ for all $(z^i, u_d^i) \in Z_{dyn}^d$. Then, if $h(z)$ satisfies constraints (9d) and (9f), we have that $q_c(z, u_c^i) \geq 0$ for all $z \in \mathcal{D}_C$ and $q_d(z, u_d^i) \geq 0$ for all $z \in \mathcal{D}_D$.*

We can now ensure that the function $h(z)$ from (9) defines a set $\mathcal{C}$ that is such that $\mathcal{C} \subset \mathcal{D} \subseteq \mathcal{S}$ and that can be rendered forward invariant by control actions satisfying the constraints (5) and (6). The next theorem follows without proof and summarizes under which conditions $h(z)$ is a valid HCBF.

Theorem 2. *Let $h(z)$ be a twice continuously differentiable function and let the sets $\mathcal{S}$, $\mathcal{N}$, $\mathcal{D}_C$, $\mathcal{D}_D$, $\mathcal{D}$, $\mathcal{C}$, and the data-sets Z_{dyn}^c, Z_{dyn}^d, Z_{safe}^c, Z_{safe}^d, and Z_N be defined as above. Suppose that Z_N forms an $\bar{\epsilon}$-net of $\mathcal{N}$ satisfying $\bar{\epsilon} < \gamma_{unsafe}/L_h(z^i)$ for all $z^i \in Z_N$, and that Z_{safe}^c and Z_{safe}^d are ϵ_c- and ϵ_d-nets of $\mathcal{D}_C$ and $\mathcal{D}_D$, respectively, satisfying the conditions of Propositions 1 & 2. Let $h(z)$, $q_c(z, u_c^i)$, and $q_d(z, u_c^i)$ be Lipschitz continuous with local constants $L_h(z)$, $L_q^c(z, u_c^i)$, and $L_q^d(z, u_d^i)$, respectively. Then, if $h(z)$ satisfies constraints (10), (9b), (9d), and (9f), the set $\mathcal{C}$ is non-empty, $\mathcal{C} \subset \mathcal{D} \subseteq \mathcal{S}$, and the function $h(z)$ is a valid local HCBF on $\mathcal{D}$ with domain $\mathcal{N} \cup \mathcal{D}$.*

4   Case Studies

The code for both case studies is available at https://github.com/unstable-zeros/learning-hcbfs. In both our simulations, we numerically integrate the hybrid dynamics of interest using an integrator implementation inspired by Drake [30]. We build our implementation on top of `scipy.integrate.solve_ivp`'s event detection API. This allows us to get precise notifications for when a system makes a discrete jump.

---

[6]By local constant, we here mean a Lipschitz constant in an $\bar{\epsilon}$ neighborhood of z.

Bouncing ball   Our first experiment considers the controlled bouncing ball discussed in Example 2. Here we define a safe set $\mathcal{S} := \{z \in \mathbb{R}^2 \,|\, |v| \le 2\}$, and set $\rho = 9.81$. Details as to how expert data is generated can be found in Appendix B. At a high level, we use an exploratory controller that covers much of the state-space but does not guarantee safety, in combination with a safe controller composed of a continuous component based on the analytical HCBF in Example 2 and the CBF-QP problem [8], and a discrete component obtained by analytically solving the constraint described in Example 2. Using the above controllers, we obtain data-sets Z_{safe}^c and Z_{safe}^d containing 5000 safe states and expert demonstrations u_c^i and u_d^i. Those expert states, shown as the green and magenta dots in Fig. 2(left), are evenly gridded near the boundary of the safe set $\mathcal{S}$. We in addition sample 4560 unsafe samples by gridding along the boundary of $\mathcal{N}$ to form the data-set Z_N.

We parametrize the HCBF candidate h as a two-hidden-layer fully-connected DNN with tanh activation functions and 64 neurons in each hidden layer. The training is implemented using `jax` [31] and the Adam algorithm with a cosine decay learning rate. We craft a loss function by relaxing the constraints in (9) – hyperparameter choices and other training details can be found in Appendix B. The set $\mathcal{C}$ of the learned HCBF is shown as green dots in the middle subplot of Figure 2. One may clearly observe that $\mathcal{C}$ is strictly contained within the geometric safe set, i.e. $\mathcal{C} \subset \mathcal{S}$.
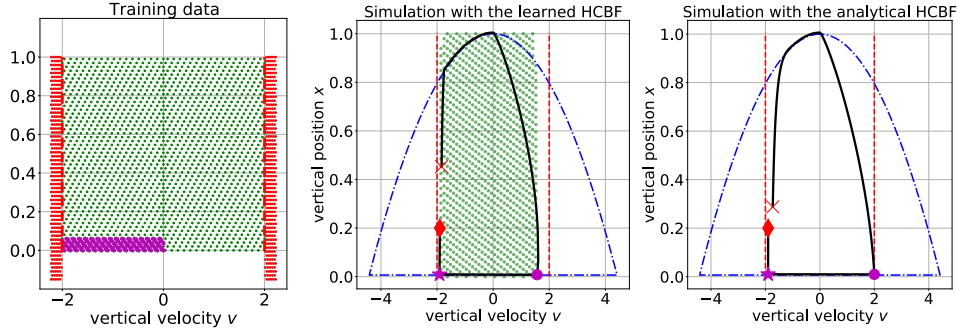

Figure 2: Left: Plot of the expert states Z_{safe}^c (green) and Z_{safe}^d (magenta), and unsafe samples Z_N (red). Dashed red lines indicate bounds on the velocity. Middle: Closed-loop simulation of the bouncing ball using the learned HCBF where the initial state, state before jump, state after jump and terminal state are indicated by a red diamond, magenta star, magenta circle and red cross, respectively. The solid black line and dash-dotted blue line represent the closed-loop trajectory and desired reference path. The green dots denote the safe invariant set (8) of the learned HCBF. Right: Closed-loop simulation of the bouncing ball using the analytical HCBF.

We now test if the learned HCBF is able to produce safe control inputs that keep the system within the set $\mathcal{C}$. We consider a nominal control law $u_{c,\text{nom}}$ that looks to track the reference path shown in dash-dotted blue in Fig. 2(middle) and (right): as the system has full control, it can exactly track the reference path, but without correction, this would lead to violation of the velocity constraint. Starting from an initial state $(0.2, -1.9)$, we simulate the controlled system for 2.25 seconds. During flow, we solve the CBF-QP problem for the continuous input u_c (see Appendix B.2). During jump, we perform line search with a decay factor 0.95 to find the scalar discrete input u_d that satisfies (6). The closed-loop trajectory produced by the learned HCBF is plotted in Figure 2 (middle). As a comparison, we also plot in Figure 2 (right) the trajectory obtained using the analytical HCBF. Albeit slightly more conservative, the learned HCBF keeps the ball within $\mathcal{C}$ at all times during both flows and jumps.

Compass gait   We consider the compass gait walker [32, 33]. Originally introduced by Goswami et al. [32], the compass gait dynamics describe a passive bipedal robot walking down an inclined plane at a constant velocity. This system is described by a four-dimensional continuous state $z = [\theta_{\text{stance}}, \theta_{\text{swing}}, \dot{\theta}_{\text{stance}}, \dot{\theta}_{\text{swing}}]$ consisting of the angle and angular velocity of each leg. In this notation, the "stance" foot corresponds to the foot that is in contact with the ground as the compass gait walker makes its descent down the ramp; hence, the "swing" foot refers to the foot that is not in contact with the ramp at a particular instant in time. In our simulations, the walker's initial stance leg is its left leg, and therefore its initial swing leg is its right leg. To improve the walking capabilities, we add actuation to hip joint and the ankle of the stance leg. To collect expert trajectories, we use the energy-based controller of Goswami et al. [34]. Implementation details can be found in Appendix C. Learning and analyzing a HCBF for the compass gait walker hybrid system poses
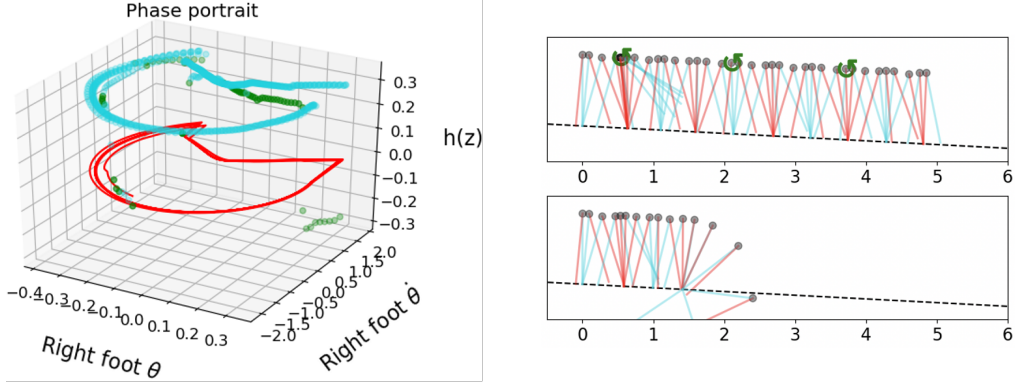
Figure 3: (Left) In red, we show the phase portrait of the right foot using the CBF-QP controller. On the vertical axis, we plot the values of $h(x)$ associated with this phase portrait. The regions where $u_{\text{HCBF}} \neq u_{\text{nom}}$, i.e., where the HCBF corrects the nominal controller, are shown in green; values of $h(x)$ where $u_{\text{HCBF}} = u_{\text{nom}}$ are shown in blue. Notice that when the value of $h(z)$ dips below the "safe" threshold $\gamma_{\text{safe}} = 0.3$, the HCBF-QP almost always intervenes. (Right) From a fixed initial state, we show a trajectory with the CBF-QP controller (top) and the nominal zero-valued controller (bottom). The states are down-sampled for clarity, the left and right feet are shown in red and blue respectively, and torques applied by the CBF-QP controller are shown as green arrows.

several challenges, including the well-known sensitivity of the compass gait walker to its initial conditions and the inherent difficulties in visualizing a four-dimensional state space. To facilitate a meaningful visualization of the four-dimensional state space, when collecting expert data, we fix the stance leg initial condition to the point $[\theta_{\text{stance}}, \dot{\theta}_{\text{stance}}] = [0, 0.4]$ on the passive limit cycle, and vary the initial condition of the swing leg by adding uniform noise to corresponding passive limit cycle state $[\theta_{\text{swing}}, \dot{\theta}_{\text{swing}}] = [0, 2.0]$ of the swing leg.

We again parameterize the candidate HCBF h with a two-hidden-layer fully-connected neural network with tanh activations, and as before, determine the hyperpameters by grid search. However, unlike the previous example, the task of identifying boundary points is complicated by the higher dimensionality of the state space. Therefore, we propose a novel algorithm that can be used to identify boundary points by sampling from the space of expert trajectories. The algorithm, described in Appendix C, identifies boundary points by computing the pairwise distances between all of the expert states and thresholding based on the number of neighbors a point has within an ϵ-norm ball.

In the left panel of Fig. 3, we visualize the learned HCBF. Starting from an initial condition with the same left leg state as the expert trajectories, we identify the values for $h(z)$ at which $u_{\text{HCBF}} \neq u_{\text{nom}} > 0$ in green, i.e. where the HCBF controller corrects the nominal controller. In red, we plot the phase portrait for the right leg corresponding to a trajectory using the CBF-QP controller. To demonstrate the physical interpretation of this learned HCBF, in the right top panel of Fig. 3, we show the motion of the compass gait walker down the ramp, marking in green where the CBF-based controller takes over; in the right lower panel of Fig. 3 we see the failure of a zero-valued controller from the same initial condition, showing that the learned HCBF preserves safety. We refer to Appendix C for further illustrative visualizations. Videos of both the safe HCBF and unsafe nominal controller trajectories can be found in the supplementary material.

5 Conclusion

We presented a framework for learning safe control laws for hybrid systems. We introduced HCBFs and sufficient conditions under which an HCBF based control law guarantees safety, i.e., a desired safe set is made forward invariant. We then showed how to learn such HCBFs from data using an optimization-based framework. We gave sufficient conditions under which feasible solutions to the optimization problem are valid HCBFs, and illustrated our methods on two case studies. Future work will look to extend this approach to hybrid systems described by differential *inclusions*, so as to be applicable to systems with friction and stiction, as well as to develop statistical guarantees of correctness to alleviate the sampling burden of our method.

8

References

[1] W. Schwarting, J. Alonso-Mora, and D. Rus. Planning and decision-making for autonomous vehicles. *Annual Review of Control, Robotics, and Autonomous Systems*, 2018.

[2] M. Tucker, E. Novoseller, C. Kann, Y. Sui, Y. Yue, J. Burdick, and A. D. Ames. Preference-based learning for exoskeleton gait optimization. *arXiv preprint arXiv:1909.12316*, 2019.

[3] H. Kress-Gazit, G. E. Fainekos, and G. J. Pappas. Temporal-logic-based reactive mission and motion planning. *IEEE Trans. Robot.*, 25(6):1370–1381, 2009.

[4] S. Prajna, A. Jadbabaie, and G. J. Pappas. A framework for worst-case and stochastic safety verification using barrier certificates. *IEEE Trans. Autom. Control*, 52(8):1415–1428, 2007.

[5] P. Glotfelter, J. Cortés, and M. Egerstedt. Nonsmooth barrier functions with applications to multi-robot systems. *IEEE Control Syst. Lett.*, 1(2):310–315, 2017.

[6] P. Glotfelter, I. Buckley, and M. Egerstedt. Hybrid nonsmooth barrier functions with applications to provably safe and composable collision avoidance for robotic systems. *IEEE Robotics and Automation Letters*, 4(2):1303–1310, 2019.

[7] P. Wieland and F. Allgöwer. Constructive safety using control barrier functions. In *Proc. IFAC Symp. Nonlin. Control Syst.*, pages 462–467, Pretoria, South Africa, August 2007.

[8] A. D. Ames, J. W. Grizzle, and P. Tabuada. Control barrier function based quadratic programs with application to adaptive cruise control. In *Proceedings of the Conference on Decision and Control (CDC)*, pages 6271–6278, Los Angeles, CA,, December 2014.

[9] 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*, 62(8):3861–3876, 2017.

[10] M. J. Khojasteh, V. Dhiman, M. Franceschetti, and N. Atanasov. Probabilistic safety constraints for learned high relative degree system dynamics. *arXiv preprint arXiv:1912.10116*, 2019.

[11] 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.

[12] A. Agrawal and K. Sreenath. Discrete control barrier functions for safety-critical control of discrete systems with application to bipedal robot navigation. In *Robotics: Science and Systems*, 2017.

[13] M. Ohnishi, L. Wang, G. Notomista, and M. Egerstedt. Barrier-certified adaptive reinforcement learning with applications to brushbot navigation. *IEEE Transactions on robotics*, 35(5):1186–1205, 2019.

[14] M. Cavorsi, M. Khajenejad, R. Niu, Q. Shen, and S. Z. Yong. Tractable compositions of discrete-time control barrier functions with application to lane keeping and obstacle avoidance. *arXiv preprint arXiv:2004.01858*, 2020.

[15] M. Maghenem and R. G. Sanfelice. Characterizations of safety in hybrid inclusions via barrier functions. In *Proceedings of the 22nd ACM International Conference on Hybrid Systems: Computation and Control*, pages 109–118, 2019.

[16] A. Bisoffi and D. V. Dimarogonas. A hybrid barrier certificate approach to satisfy linear temporal logic specifications. In *2018 Annual American Control Conference (ACC)*, pages 634–639. IEEE, 2018.

[17] R. Goebel, R. G. Sanfelice, and A. R. Teel. *Hybrid Dynamical Systems: modeling, stability, and robustness*. Princeton University Press, Princeton, NJ, 1 edition, 2012. ISBN 9781400842636.

[18] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada. Control barrier functions: Theory and applications. In *2019 18th European Control Conference (ECC)*, pages 3420–3431, Naples, Italy, June 2019.

[19] 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*, 15(3):1216–1229, 2017.

[20] L. Wang, D. Han, and M. Egerstedt. Permissive barrier certificates for safe stabilization using sum-of-squares. In *2018 Annual American Control Conference (ACC)*, pages 585–590. IEEE, 2018.

[21] 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.

[22] 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.

[23] M. Saveriano and D. Lee. Learning barrier functions for constrained motion planning with dynamical systems. In *IEEE International Conference on Intelligent Robots and Systems*, 2019.

[24] W. Jin, Z. Wang, Z. Yang, and S. Mou. Neural certificates for safe control policies. *arXiv preprint arXiv:2006.08465*, 2020.

[25] J. Ferlez, M. Elnaggar, Y. Shoukry, and C. Fleming. Shieldnn: A provably safe nn filter for unsafe nn controllers. *arXiv preprint arXiv:2006.09564*, 2020.

[26] A. Robey, H. Hu, L. Lindemann, H. Zhang, D. V. Dimarogonas, S. Tu, and N. Matni. Learning control barrier functions from expert demonstrations. *arXiv preprint arXiv:2004.03315*, 2020.

[27] H. K. Khalil. *Nonlinear Systems*. Prentice-Hall, Englewood Cliffs, NJ, 2nd edition, 1996.

[28] R. Goebel, R. G. Sanfelice, and A. R. Teel. Hybrid dynamical systems. *IEEE Control Systems*, 29(2):28–93, 2009.

[29] M. Fazlyab, A. Robey, H. Hassani, M. Morari, and G. Pappas. Efficient and accurate estimation of lipschitz constants for deep neural networks. In *Advances in Neural Information Processing Systems*, pages 11427–11438, 2019.

[30] R. Tedrake and the Drake Development Team. Drake: Model-based design and verification for robotics, 2019. URL https://drake.mit.edu.

[31] 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. *URL http://github. com/google/jax*, page 18, 2020.

[32] A. Goswami, B. Espiau, and A. Keramane. Limit cycles and their stability in a passive bipedal gait. In *Proceedings of IEEE international conference on robotics and automation*, volume 1, pages 246–251. IEEE, 1996.

[33] K. Byl and R. Tedrake. Approximate optimal control of the compass gait on rough terrain. In *2008 IEEE International Conference on Robotics and Automation*, pages 1258–1263. IEEE, 2008.

[34] A. Goswami, B. Espiau, and A. Keramane. Limit cycles in a passive compass gait biped and passivity-mimicking control laws. *Autonomous Robots*, 4(3):273–286, 1997.

A Appendix A

A.1 Computing the Lipschitz constants L_h, L_q^c, and L_q^c

To verify (9c), (9e), and (9g), the Lipschitz constant L_h of the learned function $h(z)$ as well the Lipschitz constant $L_{\nabla h}$ of its gradient $\nabla h(z)$ need to be calculated first. The following discussion is mainly taken from Robey et al. [26, Section 3.4.2]. When $\mathcal{H}$ consists of twice continuously differentiable functions, as is the case when $h \in \mathcal{H}$ is parametrized by a Deep Neural Net with twice continuously differentiable activation functions or a Reproducing Kernel Hilbert Space, the following is shown to hold.

Reproducing Kernel Hilbert Space: In the case of random Fourier features with l random features where $h(z) := \langle \phi(z), \theta \rangle$ with $\theta \in \Theta$ where Θ is a convex set and with $\phi : \mathbb{R}^{n_z} \to \mathbb{R}^l$ where

$$\phi(z) = \sqrt{2/l}(\cos(\langle z, w_1 \rangle + b_1), ..., \cos(\langle z, w_l \rangle + b_l)),$$

an upper bound on the Lipschitz constant of $h(z)$ is provided as $\sqrt{2\sigma^2}(1 + \sqrt{n_z/l} + \sqrt{(2/l)\log(1/\delta)})\|\theta\|_2$ with probability at least $1 - \delta$ where σ is as explained in Robey et al. [26, Section 3.4.2]. An upper bound on the Lipschitz constant of $\nabla h(z)$ can be derived by the bound $\|\nabla^2 h(z)\| \leq 3\sqrt{2}\|\theta\|_\infty \sigma^2(l + n_z + 2\log(1/\delta))/\sqrt{l}$ that holds with probability at least $1 - \delta$.

Deep Neural Net: When $h(z)$ is a DNN, the problem of exactly computing the Lipschitz constant of $h(z)$ is known to be NP-hard. 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 $h(z)$ is 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 [29]. Recently, the authors of Fazlyab et al. [29] proposed a semidefinite-programming based approach to efficiently compute an accurate upper bound on L_h. On the other hand, there are relatively few results that provide accurate upper bounds for the Lipschitz constant of the gradient of $h(z)$. The only general method for computing an upper bound on $L_{\nabla h}$ is through post-hoc sampling.

Now, using these Lipschitz constants L_h and $L_{\nabla h}$ of $h(z)$ and $\nabla h(z)$, respectively, it can be seen that the functions $q_c(z, u_c^i)$ and $q_d(z, u_d^i)$ in (9d) and (9f) are locally Lipschitz continuous in z since $\nabla h(z)$, $f_c(z)$, $f_d(z)$, $g_c(z)$, $g_d(z)$, and $\alpha(h(z))$ are locally Lipschitz continuous and since function composition preserves Lipschitz continuity. Upper bounds of the Lipschitz constants L_q^c and L_q^d follow immediately.

A.2 Proof of Lemma 1

Proof of Lemma 1. First note that $\dot{v}(t) = -\alpha(v(t))$ with $v(0) \geq 0$ admits a unique solution $v(t)$ that is such that $v(t) \geq 0$ for all $t \geq 0$ [27, Lemma 4.4]. Each solution $x(t)$ to (1) under $u(x)$ is now, due to the chain rule and since $u(x) \in K_{\text{CBF}}(x)$, such that $\dot{h}(x(t)) \geq -\alpha(h(x(t)))$ for all $t \in \mathcal{I}$. Using the Comparison Lemma [27, Lemma 3.4] and assuming that $h(x(0)) \geq 0$, it follows that $h(x(t)) \geq v(t) \geq 0$ for all $t \in \mathcal{I}$, i.e., $x(0) \in \mathcal{C}$ implies $x(t) \in \mathcal{C}$ for all $t \in \mathcal{I}$. Note next that (1) is defined on $\mathcal{D}$ since $u(x)$ is only defined for $x \in \mathcal{D}$. Since $x \in \mathcal{C}$ for all $t \in \mathcal{I}$ and when $\mathcal{C}$ is compact, it follows by Khalil [27, Theorem 3.3] that $\mathcal{I} = [0, \infty)$, i.e., $\mathcal{C}$ is forward invariant. $\square$

A.3 Proof of Theorem 1

Proof of Theorem 1. During flows with I_j not being a singleton, and if $h(z(\min(I_j), j)) \geq 0$, we infer that $h(z(t, j)) \geq 0$ for all $t \in [\min I_j, \sup I_j)$ due to (5) and as in the proof of Lemma 1. By continuity of $h(z)$ and $z(t, j)$ and since $\mathcal{C}$ is closed, it also holds that $h(z(\sup I_j, j)) \geq 0$ if $I_j = [t_j, t_{j+1}] \times \{j\}$, i.e., the right end point is included in I_j. After each jump, it holds that $h(z(t_{j+1}, j + 1)) \geq 0$ as a consequence of (6). Consequently, $z(0, 0) \in \mathcal{C}$ implies $z(t, j) \in \mathcal{C}$ for all $(t, j) \in \text{dom}(z)$.

We next show that $\mathcal{C}$ is forward invariant under $u_c(z)$ and $u_d(z)$, i.e., $\text{dom}(z)$ is unbounded, if $\mathcal{C}$ is compact and if $\mathcal{C} \subseteq C \cup D$. First recall that $\mathcal{C} \subseteq \mathcal{D}_C \cup \mathcal{D}_D$ due to $\mathcal{C} \subseteq C \cup D$ and since $\mathcal{C} \cap C \subseteq \mathcal{D}_C$ and $\mathcal{C} \cap D \subseteq \mathcal{D}_D$. Assume now that the hybrid solution $z(t, j)$ to the system (3) under control laws

$u_c(z)$ and $u_d(z)$ is maximal[7] and that the hybrid time domain $\mathrm{dom}(z)$ is bounded. By defining $(T, J) := \sup_{(t,j)} \mathrm{dom}(z)$, we distinguish between the two cases: 1) $z(T, J) \notin \mathrm{dom}(z)$ and 2) $z(T, J) \in \mathrm{dom}(z)$. 1) Note that $z(T, J) \notin \mathrm{dom}(z)$ can only happen when I_J is not a singleton and when $z(t, j)$ has left the set $\mathcal{D}_C$ by flowing without entering $\mathcal{D}_D$, which would enable a successive jump that is not possible since it is assumed that the solution $z(t, j)$ is maximal. Since the set $\mathcal{C}$ is compact and since $\mathcal{D}_C$ is open, there has to exist a time $t < T$ such that $z(t, J) \in \mathcal{D}_C \setminus \mathcal{C}$ according to Khalil [27, Theorem 3.3], which does not hold since $(t, J) \in \mathrm{dom}(z)$ and since $z(t, J) \in \mathcal{C}$ as shown previously. Since $z(t, j)$ is maximal, it follows by contradiction that $\mathrm{dom}(z)$ is unbounded. 2) If $z(T, J) \in \mathrm{dom}(z)$, it holds that $z(T, J) \in \mathcal{C}$ so that either $z(T, J) \in \mathcal{D}_D \cap \mathcal{C}$ or $z(T, J) \in (\mathcal{D}_C \cap \mathcal{C}) \setminus (\mathcal{D}_D \cap \mathcal{C})$. In the former case, the solution $z(t, j)$ can be extended by a jump. In the latter case, note that $z(T, J)$ is strictly contained within the set $\mathcal{D}_C \cup \mathcal{D}_D$ so that the solution $z(t, j)$ can be extended into $\mathcal{D}_C$ or into $\mathcal{D}_D$ by flowing. By contradiction, it again follows that $\mathrm{dom}(z)$ is unbounded. $\qquad\square$

A.4 Proposition 3

Proposition 3. *Let $h(z)$ be Lipschitz continuous with local constant $L_h(z)$[8], $\gamma_{unsafe} > 0$ and Z_N be an $\bar{\epsilon}$-net of $\mathcal{N}$ with $\bar{\epsilon} < \gamma_{unsafe}/L_h(z^i)$ for all $z^i \in Z_N$. Then, if $h(z)$ satisfies constraint (9b), we have that $h(z) < 0$ for all $z \in \mathcal{N}$.*

Proof of Proposition 3. Note first that, for all $z \in \mathcal{N}$, it follows that there exists a point $z^i \in Z_N$ satisfying $\|z - z^i\|_p \le \bar{\epsilon}$ due to the assumption that Z_N is an $\bar{\epsilon}$-net of $\mathcal{N}$. For any $z \in \mathcal{N}$, we can now select a point $z^i \in Z_N$ satisfying $\|z - z^i\|_p \le \bar{\epsilon}$ for which it follows that

$$h(z) = h(z) - h(z^i) + h(z^i) \overset{(a)}{\le} |h(z) - h(z^i)| - \gamma_{\text{unsafe}}$$
$$\overset{(b)}{\le} L_h(z^i)\|z - z^i\|_p - \gamma_{\text{unsafe}} \overset{(c)}{\le} L_h(z^i)\bar{\epsilon} - \gamma_{\text{unsafe}} \overset{(d)}{<} 0$$

In particular, inequality (a) follows from the constraint (9b) which says that $h(z^i) \le -\gamma_{\text{unsafe}}$ for all $z^i \in Z_N$. Inequality (b) follows by the local Lipschitz constant $L_h(z)$ on $h(z)$ in an $\bar{\epsilon}$ neighborhood of z^i, while inequality (c) follows again by the assumption that Z_N forms an $\bar{\epsilon}$-net of $\mathcal{N}$. The strict inequality (d) follows simply by the assumption that $\bar{\epsilon} < \gamma_{\text{unsafe}}/L_h(z^i)$ for all $z^i \in Z_N$. $\qquad\square$

A.5 Proof of Proposition 1

Proof of Proposition 1. Note first that, for all $z \in \mathcal{D}_C$, it again follows that there exists a point $z^i \in Z_{\text{safe}}^c$ satisfying $\|z - z^i\|_p \le \epsilon_c$ due to the assumption that Z_{safe}^c is an ϵ_c-net of $\mathcal{D}_C$. For any $z \in \mathcal{D}_C$, we can now select a point $z^i \in Z_{\text{safe}}^c$ satisfying $\|z - z^i\|_p \le \epsilon_c$ for which it follows that

$$0 \overset{(a)}{\le} h(z^i) - \gamma_{\text{safe}} = h(z^i) - h(z) + h(z) - \gamma_{\text{safe}} \le |h(z^i) - h(z)| + h(z) - \gamma_{\text{safe}}$$
$$\overset{(b)}{\le} L_h(z^i)\|z^i - z\|_p + h(z) - \gamma_{\text{safe}} \overset{(c)}{\le} L_h(z^i)\epsilon_c + h(z) - \gamma_{\text{safe}} \overset{(d)}{\le} h(z)$$

In particular, inequality (a) follows from the constraint (9a) which says that $h(z^i) \ge \gamma_{\text{safe}}$ for all $z^i \in Z_{\text{safe}}^c$. Inequality (b) follows by the local Lipschitz constant $L_h(z)$ on $h(z)$ in an ϵ_c neighborhood of z^i, while inequality (c) follows again by the assumption that Z_{safe}^c is an ϵ_c-net of $\mathcal{D}_C$. The inequality (d) follows simply by the assumption that $\max(\epsilon_c, \epsilon_d) \le \gamma_{\text{safe}}/L_h(z^i)$ for all $z^i \in Z_{\text{safe}}^c \cup Z_{\text{safe}}^d$. The same analysis holds for all $z \in \mathcal{D}_D$, so that $h(z) \ge 0$ for all $z \in \mathcal{D}$. $\qquad\square$

A.6 Proof of Proposition 2

Proof of Proposition 2. Note first that, for all $z \in \mathcal{D}_C$, it follows that there exists a pair $(z^i, u_c^i) \in Z_{\text{dyn}}^c$ satisfying $\|z - z^i\|_p \le \epsilon_c$ due to the assumption that Z_{safe}^c is an ϵ_c-net of $\mathcal{D}_C$. For any $z \in \mathcal{D}_C$,

---

[7]A hybrid solution $z(t, j)$ is maximal if there exists no other hybrid solution $z'(t, j)$ with $\mathrm{dom}(z) \subset \mathrm{dom}(z')$ and with $z(t, j) = z'(t, j)$ for all $(t, j) \in \mathrm{dom}(z)$.

[8]By local constant, we here mean a Lipschitz constant in an $\bar{\epsilon}$ neighborhood of z.

we can now select a pair $(z^i, u_c^i) \in Z_{\text{dyn}}^c$ satisfying $\|z - z^i\|_p \le \epsilon_c$ for which it follows that

$$0 \stackrel{(a)}{\le} q_c(z^i, u_c^i) - \gamma_{\text{dyn}}^c = q_c(z^i, u_c^i) - q_c(z, u_c^i) + q_c(z, u_c^i) - \gamma_{\text{dyn}}^c$$

$$\le |q_c(z^i, u_c^i) - q_c(z, u_c^i)| + q_c(z, u_c^i) - \gamma_{\text{dyn}}^c \stackrel{(b)}{\le} L_q^c(z^i, u_c^i)\|z^i - z\|_p + q_c(z, u_c^i) - \gamma_{\text{dyn}}^c$$

$$\stackrel{(c)}{\le} L_q^c(z^i, u_c^i)\epsilon_c + q_c(z, u_c^i) - \gamma_{\text{dyn}}^c \stackrel{(d)}{\le} q_c(z, u_c^i)$$

In particular, inequality (a) follows from the constraint (9d) which says that $q_c(z^i, u_c^i) \ge \gamma_{\text{dyn}}^c$ for all $(z^i, u_c^i) \in Z_{\text{dyn}}^c$. Inequality (b) follows by the local Lipschitz constant $L_q^c(z, u_c^i)$ on $q_c(z, u_c^i)$ in an ϵ_c neighborhood of z^i, while inequality (c) follows again by the assumption that Z_{safe}^c is an ϵ_c-net of $\mathcal{D}_C$. The inequality (d) follows simply by the assumption that $\epsilon_c \le \gamma_{\text{dyn}}^c / L_q^c(z^i, u_c^i)$ for all $z^i \in Z_{\text{safe}}^c$. This implies that $q_c(z, u_c^i) \ge 0$ for all $z \in \mathcal{D}_C$. The same analysis holds for all $z \in \mathcal{D}_D$, so that $q_d(z, u_d^i) \ge 0$ for all $z \in \mathcal{D}$. $\qquad\square$

B   Appendix B: Bouncing Ball

The code for the bouncing ball case study can be found in `https://github.com/unstable-zeros/learning-hcbfs/tree/main/bouncing_ball_for_share`.

B.1   Generating and tracking the reference path

The control goal of the bouncing ball case study is to track a desired reference path shown as dash-dotted blue lines in Fig. 2 (middle and right). This reference path is obtained by dropping the ball from $(1, 0)$ without applying any control input, i.e., $u_c := 0$, until it hits the ground. The discrete control input u_d is set to 1.0 when this happens. Reversing the velocity in sign while keeping the position unchanged gives the reference path in the right half plane.

In order to track the reference path, we design a linear error-feedback control law $u_{c,\mathrm{nom}} = K(z(t, j) - z_{\mathrm{ref}}(t))$ with $K = [10 \quad 5.48]^T$ via solving the Riccati equation. The reference state $z_{\mathrm{ref}}(t)$ is selected as the state on the reference path that is the closet to $z(t, j)$ in Euclidean distance. We set the nominal discrete input to be $u_{d,\mathrm{nom}} = 1.2$ such that it amplifies the velocities after collisions and thus is more likely to cause constraint violation. Note that neither the reference path nor the closed-loop trajectory obtained by applying the nominal controllers $u_{c,\mathrm{nom}}$ and $u_{d,\mathrm{nom}}$ complies with the geometric safe set $\mathcal{S}$.

B.2   Training data

To obtain training data during flows, we use the above reference controller $u_{c,\mathrm{nom}}$ together with the analytical HCBF in Example 2, which, per solving the CBF-QP problem [8], gives a safe controller u_c during flow. The CBF-QP basically consists of solving a convex quadratic program with decision variable u_c, cost function $\|u_c - u_{c,\mathrm{nom}}\|$, and the constraint $\langle \nabla h(z), f_c(z) + g_c(z)u_c \rangle \geq -\alpha(h(z))$. The safe control input u_d used to obtain training data during jumps is essentially a thresholding function and can be found analytically as mentioned in Example 2. Based on these safe controllers u_c and u_d, we obtain data-sets Z_{safe}^c and Z_{safe}^d containing 5000 safe states z^i and associated expert demonstrations u_c^i and u_d^i, which are shown as the green and magenta dots in the left subplot of Figure 2.

B.3   Hyperparameters for training the neural network

We train the neural network for 1.5×10^3 epochs using the loss given by unconstrained relaxation of Problem (9) in Section 3. The hyperparameters, shown in Table 1, are found by grid search.

Parameter name	Value
λ_{s}	4.0
λ_{u}	5.0
λ_{d}	1.0
λ_{c}	1.0
γ_{safe}	0.0025
γ_{unsafe}	0.075
γ_{dyn}^c	0.055
γ_{dyn}^d	0.055

Table 1: Hyperparameters used for training the neural network based HCBF for the bouncing ball.

B.4   Visualization of the learned HCBF

We visualize the learned and analytical HCBFs in $(v, x, h(z))$-space in Figure 4 by plotting their level sets. The green dots indicate where $h(z) \geq 0$ on $C \cup D$, while the purple dots additionally indicate where $h(z) \geq 0$ on $\{z \in C \mid x = 0, \ 0 \leq v \leq 2\}$, i.e. the possible set of states after an impact with the ground.
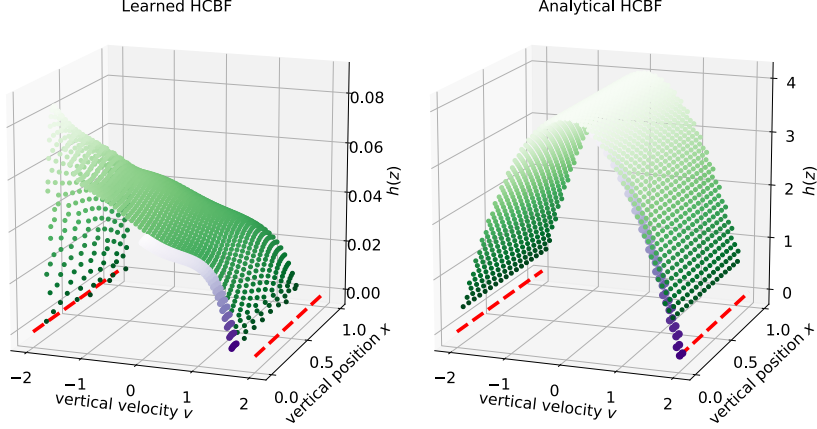
Figure 4: Left: Level sets of the learned HCBF. Right: Level sets of the analytical HCBF. For both plots, green dots indicate where $h(z) \geq 0$ on $C \cup D$, while the purple dots additionally indicate where $h(z) \geq 0$ on $\{z \in C \mid x = 0 \, , \, 0 \leq v \leq 2\}$. Deeper color indicates a lower value of $h(z)$.

B.5 Validness of the learned HCBF

We performed post-verification of the learned HCBF satisfying Propositions 1, 2 and 3, which serve as sufficient conditions for learning a valid HCBF as stated in Theorem 2. Since all of our data sets were gridded evenly on the state space, densities of the ϵ-nets are equal to the gridding resolutions, which are $\bar{\epsilon} = 0.01$ and $\epsilon_c = \epsilon_d = 0.02$. The local Lipschitz constants $L_h(\cdot), L_q^c(\cdot), L_q^d(\cdot)$ were approximated using the norm of their gradients, e.g. $L_h(z) \approx \|\nabla h(z)\|_2$, which gives a tight estimation of the Lipschitz constants due to our dense gridding scheme. The percentages of data points that satisfy those conditions are summarized in Table 2.

Condition	Involved datasets	Satisfaction rate
(9a)	$Z_{\text{safe}}^c, Z_{\text{safe}}^d$	95.38%
(9b)	Z_N	96.6%
(9d)	$Z_{\text{safe}}^c, Z_{\text{dyn}}^c$	100%
(9f)	$Z_{\text{safe}}^d, Z_{\text{dyn}}^d$	100%
Proposition 1	$Z_{\text{safe}}^c, Z_{\text{safe}}^d$	98.81%
Proposition 2	$Z_{\text{safe}}^c, Z_{\text{safe}}^d, Z_{\text{dyn}}^c, Z_{\text{dyn}}^d$	98.74%
Proposition 3	Z_N	100%

Table 2: Satisfaction rates of sufficient conditions for learning a valid HCBF for the bouncing ball.

C Appendix C: Compass Gait Walker

The code for the compass gait walker case study can be found in `https://www.dropbox.com/sh/4dzborvnav4su59/AABHfU7mgc5ACqQyM9NxKZ7Ea?dl=0`.

C.1 Sampling boundary points

Algorithm 1 Neighbor-based Unsafe Trajectory Sampling (NUTS)

Input: Tuple of n_z-dimensional vectors $Z := Z_{\text{safe}}^c \cup Z_{\text{safe}}^d$, minimum number of neighbors $N \in \mathbb{Z}_+$, neighbor threshold $\eta > 0$
Output: Binary vector $y := \{y_1, \ldots, y_n\} \subset \{0, 1\}^n$

1: $tree \leftarrow \text{KDTree}(Z)$
2: $num\_nbrs \leftarrow tree.\text{query\_radius}(Z, \text{radius} = \eta)$
3: $y = \mathbb{1}(num\_nbrs < N)$

The approach we have described crucially relies on being able to sample from "unsafe" regions of state space, i.e. to sample states $z^i \in Z_N$. For low-dimensional state spaces, it is often possible to take advantage of the underlying geometry of the hybrid system to sample from this region as for instance in the bouncing ball case study. However, for higher-dimensional state spaces such as the compass gait walker, it may not be possible to easily leverage the underlying geometry to obtain these states.

To collect unsafe states $z^i \in Z_N$ for dynamical or hybrid systems with high dimensional state spaces, we propose an algorithm for obtaining these unsafe states by sampling the expert trajectories. The pseudocode for this algorithm is given in Algorithm 1. This algorithm takes two input parameters: a positive integer $N \in \mathbb{Z}_+$ and a small nonnegative constant $\eta > 0$. In line 1, we first compute the pairwise distances between each of the $n := N_c + N_d$ states in our expert trajectories $Z_{\text{safe}}^c \cup Z_{\text{safe}}^d$; the result is a symmetric $n \times n$ matrix M, where the element at position (i, j) in M represents the pairwise distance between states z^i and z^j. Next, we threshold M so that all values in M that are greater than η are set to zero in $\tilde{M}$ and all values in M that are less than η are set to one. In line 2, we use the specific notation $\mathbb{1}(M_{i,j} \leq \eta) := 1$ if $M_{i,j} \leq \eta$ and $\mathbb{1}(M_{i,j} \leq \eta) := 0$ otherwise. Intuitively, the idea in this step is to identify the neighbors of each data point. More specifically, the ith row of $\tilde{M}$ corresponds to a state z^i; then the indices j in this row with $M_{ij} < \eta$ correspond to states z^j which we deem neighbors of z^i. In line 3, we sum along the rows of $\tilde{M}$ to create the vector $\tilde{y} \in \mathbb{R}^n$, which counts the number of neighbors for each state z^i for $i \in \{1, 2, \ldots, n\}$. Finally in line 4, we threshold this vector $\tilde{y}$ based on the minimum number of neighbors N. Note here that $\mathbb{1}(\tilde{y} \leq N)$ operates element-wise. The result is a binary vector $y \in \mathbb{R}^n$, where the indices i that are set to one correspond to states z^i that the algorithm identifies as boundary points. In this way, the total number of boundary points identified by the algorithm is $\sum_{i=1}^{n} y_i$.

Before demonstrating how we used this algorithm in the compass gait case study, we provide some simple examples to illustrate the efficacy of this algorithm toward identifying the boundaries of different sets of points. In Figure 5, we show the result of running Algorithm 1 on mixtures of two Gaussians. In Figure 6, we show a three-dimensional example in the unit cube.

Finally, we show the result of running Algorithm 1 on ten rollouts from the compass gait walker in Figure 7. In this case, the majority of the collected states are located in the first bin, which is corroborated by the histogram of swing foot angular velocity in Figure 8. In the following sections, we will use these sampled unsafe states for training the HCBF function.

C.2 Collecting expert trajectories for the compass gait walker

Our implementation of the compass gait walker relies heavily on the C++ implementation in the Drake package [30]. In particular, we re-implement the compass gait walker in Python with Jax [31]: our implementation is publicly available on the project Github repository. We use the same default settings for the compass gait walker as are used in the Drake implementation. Specifically, for the walker itself we set the hip and leg masses to be $m_H = 10.0$ kg and leg mass $m_L = 5.0$ kg
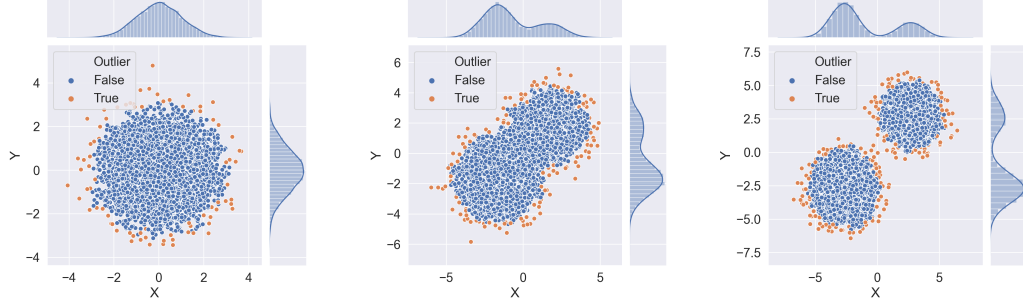
(a) We generate 20000 data points according to $z \sim \mathcal{N}(0, I)$. The marginals of the distribution are shown above and to the right of the plot.

(b) We mix two Gaussians. In this case, we generate 20000 data points z according to $z|y \sim \mathcal{N}(1.7y\mathbb{1}, I)$ with $x \sim$ Bernoulli$(1/2)$ and $y = 2x - 1$

(c) Again, we mix two Gaussians. In this case, we generate 20000 data points z according to $z|y \sim \mathcal{N}(2.7y\mathbb{1}, I)$ with $x \sim$ Bernoulli$(1/2)$ and $y = 2x - 1$

Figure 5: We run Algorithm 1 on data generated from two-dimensional Gaussian distributions with $\eta = 0.3$ and $N = 5$. All boundary points are marked in orange, and the non-boundary points are marked in blue.
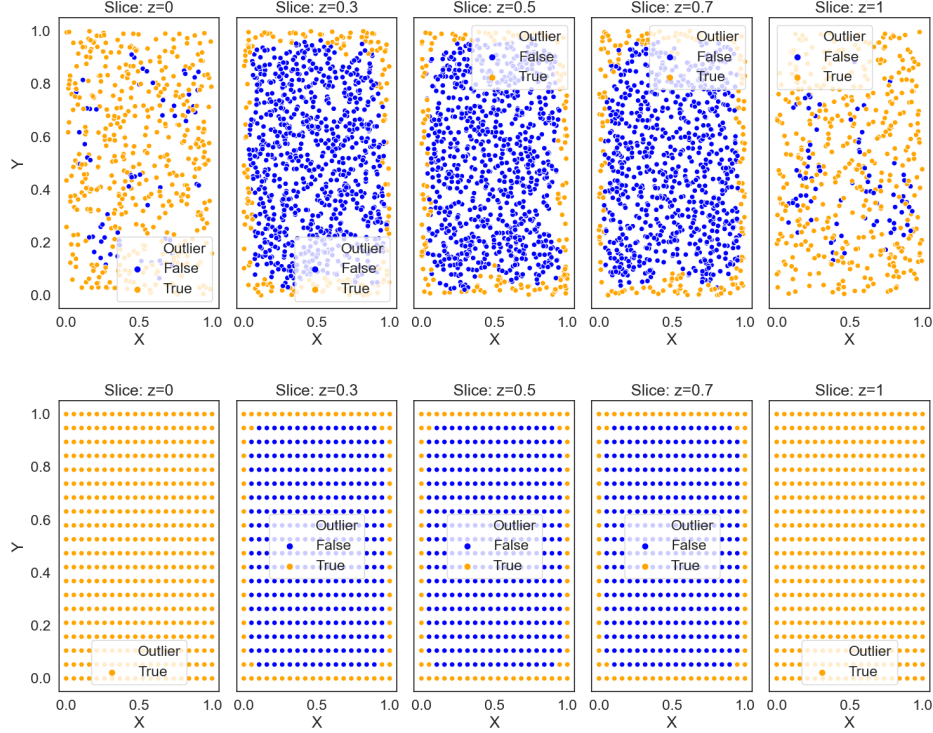

Figure 6: (Top) We uniformly sample 1000 points $(x, y, z) \in \mathbb{R}^3$ in the unit cube. We then run Algorithm 1 with $\eta = 0.13$ and $N = 70$ and show slices of this state space $\{(x, y, z)|z_0 - 0.05 \leq z \leq z_0 + 0.05\}$ for $z_0 \in \{0, 0.3, 0.5, 0.7, 1\}$. Notice that on the boundary slices $z_0 = 0$ and $z_0 = 1$, the algorithm identifies the majority of the slice as a boundary points. On the other hand, for the slices in the middle of the unit cube, generally only the outline of the slice contains boundary points. In this way, the algorithm successfully identifies the boundary of the cube. (Bottom) We repeat the same experiment as the top panel, except in this case we grid the unit cube with 8000 points and rerun the algorithm with the same parameters. Notice that the boundary of the cube is identified exactly in this case.
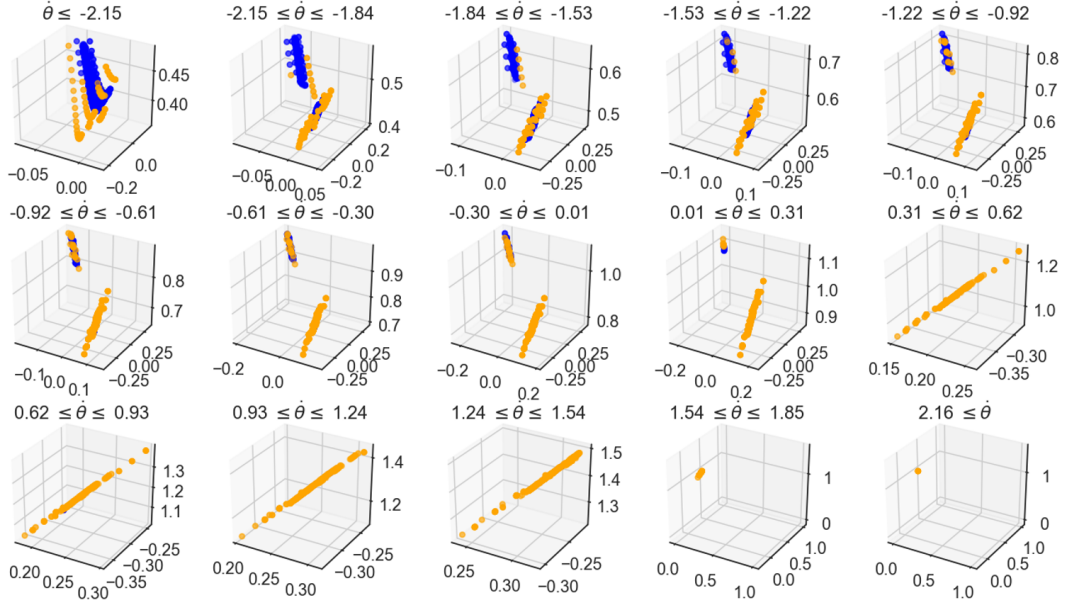
Figure 7: We show the result of running Algorithm 1 with $\eta = 0.035$ and $N = 50$ on 10 rollouts for the compass gait walker from different initial conditions sampled in the neighborhood of a point $[0.0, 0.0, 0.4, -2.0]$ on the passive limit cycle. As the state space is four-dimensional, we bin the states in these rollouts according to the angular velocity of the swing leg and plot the other components of the state in each subplot. In particular, the x, y, and z axes of each plot are angle of the stance leg, the angle of the swing leg, and the angular velocity of the stance leg respectively. Boundary points are shown in orange, whereas the interior points are shown in blue.
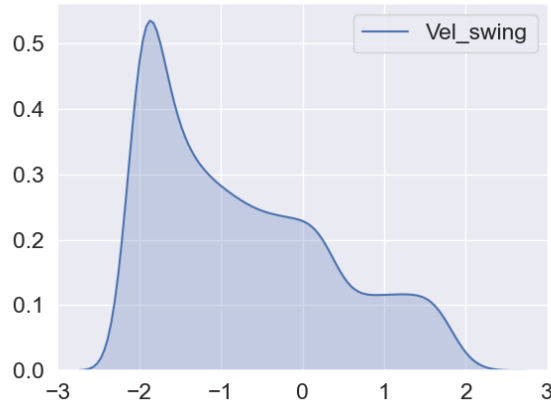

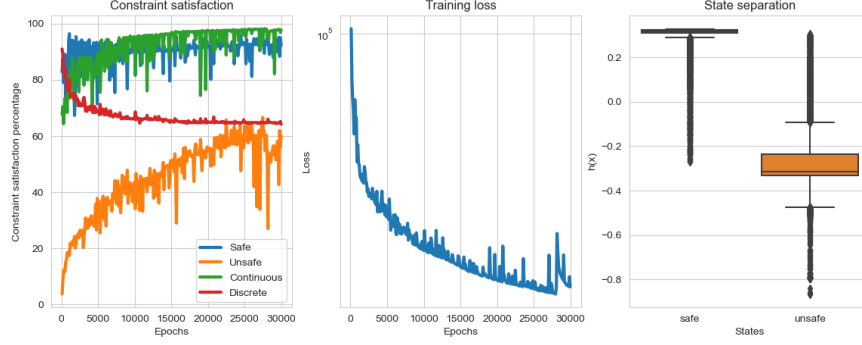**Figure 8:** Histogram of swing foot angular velocity in Figure 7.

Figure 9: Left: We show the constraint satisfaction rates for the soft constraints imposed in the unconstrained relaxation of (9) as described in the paper while training the HCBF. Middle: We plot the training loss as a function of the training epoch. Right: We show boxplots corresponding to the values of the trained HCBF on the safe and unsafe states. Notice that for states marked safe, we generally have $h(z) > 0$, whereas for states marked unsafe, we generally have $h(z) < 0$.

respectively. Furthermore, we use a leg length $\ell = 1.0$ meters, and we set the center of mass of each leg to be 0.5 meters from the ground. Finally, we use a ramp slope angle of $\gamma = 0.0525$ radians, and we set gravity to be $g = 9.81$ m/s^2.

To collect expert trajectories, we use the energy-based controller described in eq. (24) in [34]. This controller applies actuation to the hip joint, leaving the ankle joints unactuated. In particular, this controller takes the following form

$$u_H = -\frac{\lambda(E - E^*)}{\dot{\theta}_{\text{stance}} - \dot{\theta}_{\text{swing}}}$$

where we set $\lambda = 0.3$ and where E^* is the reference energy from the passive limit cycle, and E is the current total mechanical energy of the compass gait walker. Throughout, we use a reference energy of $E^* = 153.244 J$ as suggested by [34].

As described in the main text, we collect expert trajectories by fixing the initial condition of the left leg to $[\theta_{\text{stance}}, \dot{\theta}_{\text{stance}}] = [0, 0.4]$ and varying the initial condition of the right foot $[\theta_{\text{swing}}, \dot{\theta}_{\text{swing}}] = [0 + u_1, 2.0 + u_2]$ where $u_1 \sim \text{Uniform}([-0.2, 0.2])$ and $u_2 \sim \text{Uniform}([-0.5, 0.5])$. In Figure 10 (right), we show the successful right foot initial conditions for the energy-based controller corresponding to this initial condition. In particular, to train the HCBF described in the next section, we collected 500 rollouts using this scheme. For each rollout, we used a horizon of $T = 750$ steps and a time interval of $\Delta t = 0.01$.

C.3 Training the HCBF

To train the neural-network based HCBF, we relax (9) into the unconstrained problem of Section 3. The hyperparameters used for this problem are given in Table 3. In Figure 10, we show the constraint satisfaction rates, the training loss, and the state separation attained by the learned HCBF.

In the main text, we showed that our HCBF controller can be used in conjunction with a nominal control law that does not actuate the system. In particular, in the experiments shown in Figure 3, we used trajectories obtained by simulating the compass gait walker with an energy-based controller to train the HCBF. Following this, we used *no control* as the nominal controller for the HCBF-based controller. In Figure 10, we show that despite the fact that our nominal controller is not providing any control inputs, the HCBF-based controller has a similar safe set to that of the energy-based controller. In this way, the HCBF causes the uncontrolled system to match the safety characteristics of the expert energy-based controller. In particular, it can be observed that the safe set $\mathcal{C}$, which is described by the zero superlevel set of $h(z)$, under-approximates the safe expert behavior. As can further be observed, there are regions in the state space where our HCBF-based controller provides safe system trajectories while the energy-based controller results in unsafe system trajectories (e.g., the top left corner or also the bottom right corner in Figure 10). We suspect, at this point, that the

19

Parameter name	Value
λ_{s}	5.0
λ_{u}	5.0
λ_{d}	0.5
λ_{c}	0.5
γ_{safe}	0.3
γ_{unsafe}	0.3
γ_{dyn}^{c}	0.05
γ_{dyn}^{d}	0.05

Table 3: Hyperparameters used for training the neural network based HCBF for the compass gait.
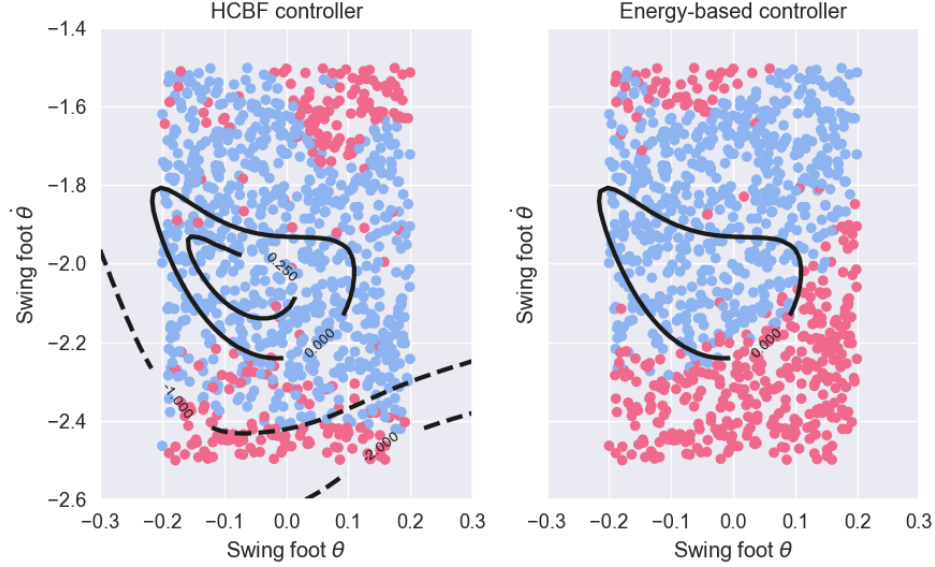

Figure 10: For a fixed standing foot initial condition of $[\theta_{\text{stance}}, \dot{\theta}_{\text{stance}}] = [0, 0.4]$, we vary the swing foot initial condition and simulate the compass gait walker with the HCBF controller on top of a nominal zero-valued controller (left) and the energy-based controller (right). Each blue dot denotes a swing foot initial condition that corresponds to a trajectory in which the walker travels eight or more steps; red dots denote trajectories in which the walker travels fewer than eight steps. We also plot the level sets of the learned HCBF for reference.

safe set $\mathcal{C}$ is attractive similarly to the continuous case [9] and provides hence more safety in some regions. A more detailed investigation of this observation is, however, subject to future work.