

Domain Adaptive Safety Filters via Deep Operator Learning

Lakshmidheepakreddy Manda¹, Shaoru Chen², Mahyar Fazlyab¹

Abstract—Learning-based approaches for constructing Control Barrier Functions (CBFs) are increasingly being explored for safety-critical control systems. However, these methods typically require complete retraining when applied to unseen environments, limiting their adaptability. To address this, we propose a self-supervised deep operator learning framework that learns the mapping from environmental parameters to the corresponding CBF, rather than learning the CBF directly. Our approach leverages the residual of a parametric Partial Differential Equation (PDE), where the solution defines a parametric CBF approximating the maximal control invariant set. This framework accommodates complex safety constraints, higher relative degrees, and actuation limits. We demonstrate the effectiveness of the method through numerical experiments on navigation tasks involving dynamic obstacles.

I. INTRODUCTION

A crucial criterion for control design in safety-critical domains is to ensure that the system's trajectories will always remain in a safe region. One popular approach to enforce safety constraints is through control barrier functions (CBFs), see e.g., [1]. CBF-based safety filters minimally correct a legacy controller to ensure constraint satisfaction. Specifically, for continuous-time control-affine systems, this involves deploying a convex quadratic program, referred to as CBF-QP, that can be solved efficiently in real-time.

Despite their simplicity, there is no general methodology for constructing CBFs for arbitrary nonlinear systems. The main challenges in CBF synthesis include bounded actuation [2], high relative degree [3], [4], complex safety specifications [5], and the inherent conservatism of CBFs [6]. To address these issues, machine learning techniques have increasingly been employed to improve the synthesis of CBFs for more complex systems and environments [7]–[10]. However, a key limitation of data-driven approaches, which motivates the current work, is that CBFs learned in one environment may become invalid in another. This environment dependence contrasts with Model Predictive Control (MPC), which can adapt to changing environments.

Our Contribution: Motivated by the challenges in CBF synthesis and adaptation to changing environments, in this work, we propose a deep operator learning framework that learns the mapping between the environment and the CBF, as opposed to learning the CBF itself. Following [6], our methodology starts with a Hamilton-Jacobi-Bellman

(HJB) type parametric Partial Differential Equation (PDE), in which the finite-dimensional parameters characterize the environment (e.g., the location of obstacles), and the unique solution of the PDE yields a CBF that approximates the maximal control invariant set. Using the residual of this PDE as a self-supervised training loss function, we learn the solution operator of the PDE, dubbed the *CBF operator*, using a continuously differentiable Multi-Layer Perceptron (MLP) architecture. The resulting CBF and the corresponding safety filter are an explicit function of the environment parameters, enabling it to adapt to novel or changing environments during deployment. The overall framework takes bounded actuation, high relative degree, complex safety specifications, and conservatism of CBF into account. We demonstrate the efficacy of our method on various safe control tasks with moving obstacles.

A. Literature Review

1) *CBF synthesis:* Several key challenges in constructing a CBF have been investigated in the literature. For multiple safety constraints that are possibly associated with logical operations, Grottel et al. [11] compose multiple CBFs accordingly through the non-smooth min/max operators. In [12], a smooth composition of logical operations on constraints were introduced and later extended to handle actuation and state constraints simultaneously [13], [14] using integral CBFs [15]. Molnar et al. [5] propose an algorithmic way to create a single smooth CBF arbitrarily composing both min and max operators. Such smooth bounds have been used in changing environments [16]. Notably, [17] ensures the input-constrained feasibility of the CBF condition while composing multiple CBFs.

High-order CBF (HOCBF) [4] and exponential CBF [3] systematically approach CBF construction when the safety constraints have high relative degree. However, controlling the conservatism of these approaches is a challenge. To reduce conservatism or improve the performance of HOCBF, learning frameworks have emerged [18]–[20] allowing tuning of the class \mathcal{K} functions used in the CBF condition.

Learning frameworks [7]–[10] building on [21] have tackled the challenge of limited actuation. Liu et al. [2] explicitly consider input constraints in learning a CBF by finding counterexamples on the 0-level set of the CBF for training, while Dai et al. [22] propose a data-efficient, prioritized sampling method.

In [23], an adaptive sampling method for training HJB NN models favoring sharp gradient regions is proposed. Drawing tools from reachability analysis, the recent work [24] iteratively expands the volume of the control invariant

¹Department of Electrical and Computer Engineering, Johns Hopkins University, 3400 N Charles St, Baltimore, MD, 21218. lmanda1@jhu.edu

²Microsoft Research, 300 Lafayette Street, New York, NY, 10027. shaoruchen@microsoft.com

¹Department of Electrical and Computer Engineering, Johns Hopkins University, 3400 N Charles St, Baltimore, MD, 21218. mahyarfazlyab@jhu.edu

set by learning the policy value function and improving the performance through policy iteration. Similarly, Dai et al. [25] expand conservative hand-crafted CBFs by learning on unsafe and safe trajectories. In [26], an actor-critic reinforcement learning framework is applied to learn a CBF value function. Qin et al. [27] learn a control policy and CBF together for black-box constraints and dynamics. We exploit the connection between CBFs and Hamilton-Jacobi (HJ) reachability analysis investigated in [28], [29] for CBF learning in this work.

Unlike any of the mentioned learning approaches, our CBF learning approach considers both bounded actuation and complex safety constraints while aiming to learn the maximal control invariant set. Importantly, our method considers changing environments by learning a NN CBF operator.

2) *Adaptive CBF*: Adaptive CBFs (aCBF) [30] considers parametric model uncertainty and proposes an adaptive CBF that is robust against a bounded uncertainty set. Robust adaptive CBFs [31] improve the conservatism of aCBF and introduce set membership identification to reduce modeling errors. In the presence of system disturbances and measurement noises, Takano et al. [32] apply a robust CBF formulation for safety filter design. Xiao et al. [33] propose an adaptive high-order CBF that adapts to time-varying control inputs and noises in the dynamics. Close to our work are [34] which formulates an environmental CBF that incorporates the environmental parameters into CBF parameterization, and [20] which focuses on tuning the class \mathcal{K} function in CBF to achieve good performance across environments. Different from [34], which focuses on handling delay inputs and case studies assuming the availability of the environmental dynamics, and [20], which relies on the exponential CBF formulation, our work aims to provide a holistic, automated learning method for CBF construction across various environments with complex safety specifications.

B. Notation

An extended class \mathcal{K} function is a function $\alpha : (-b, a) \mapsto \mathbb{R}$ for some $a, b > 0$ that is strictly increasing and satisfies $\alpha(0) = 0$. We denote $L_r^+(h) = \{x \in \mathbb{R}^n \mid h(x) \geq r\}$ as the r -superlevel set of the function h . The positive and negative parts of a number $a \in \mathbb{R}$ are denoted by $(a)_+ = \max(a, 0)$ and $(a)_- = \min(a, 0)$, respectively. We denote the class of k -times continuously differentiable functions by C^k ($k \geq 0$). We denote the perspective of the log-sum-exp (LSE) function by $\text{LSE}(y; \beta) = \frac{1}{\beta} \log(\sum_{i=1}^M \exp(\beta y_i))$ for $\beta > 0$.

II. BACKGROUND

Consider a continuous time control-affine system:

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

where $f: \mathbb{R}^n \rightarrow \mathbb{R}^n$, $g: \mathbb{R}^n \rightarrow \mathbb{R}^{n \times m}$ are locally Lipschitz continuous, $x \in \mathcal{D} \subseteq \mathbb{R}^n$ is the state with domain \mathcal{D} , u is the control input, and $\mathcal{U} \subseteq \mathbb{R}^m$ is the control input constraint

set, which we assume to be a convex polyhedron. We denote the solution of (1) at time $t \geq 0$ by $x(t)$.

Given a set $\mathcal{X} \subseteq \mathcal{D}$ that represents a safe subset of the state space, the general objective of safe control design is to find a control law $\pi(x)$ that renders \mathcal{X} forward invariant under the closed-loop dynamics $\dot{x} = f(x) + g(x)\pi(x)$, i.e., if $x(t_0) \in \mathcal{X}$ for some $t_0 \geq 0$, then $x(t) \in \mathcal{X}$ for all $t \geq t_0$. In this paper, we assume that the safe set can be characterized by the 0-superlevel set of a continuous function $c(\cdot) \in C^0$,

$$\mathcal{X} = L_0^+(c) = \{x \mid c(x) \geq 0\}. \quad (2)$$

A. Control Barrier Functions (CBFs)

A continuously differentiable function $h: \mathbb{R}^n \rightarrow \mathbb{R}$ is a CBF if the following conditions hold:

(i): The 0-super level set of h is contained in \mathcal{X} , i.e.,

$$L_0^+(h) = \{x \mid h(x) \geq 0\} \subseteq \mathcal{X}. \quad (3)$$

(ii): A Lyapunov-like condition on the whole domain:

$$\sup_{u \in \mathcal{U}} \{\nabla h(x)^\top (f(x) + g(x)u)\} + \alpha(h(x)) \geq 0 \quad \forall x \in \mathcal{D}, \quad (4)$$

where α is an extended class \mathcal{K} function, typically chosen to be linear, $\alpha(h(x)) = \gamma h(x)$, $\gamma > 0$. If the system has unbounded actuation, $u \in \mathbb{R}^m$, (4) becomes equivalent to

$$\nabla h(x)^\top g(x) = 0 \implies \nabla h(x)^\top f(x) + \alpha(h(x)) \geq 0. \quad (5)$$

The above conditions together ensure that $L_0^+(h)$ is a control invariant subset of \mathcal{X} .

We denote the class of all admissible CBFs by \mathcal{H} :

$$\mathcal{H} = \{h \in C^1 \mid h \text{ satisfies (3) and (4)}\}. \quad (6)$$

The CBF can be used as a safety filter, framed as a convex QP, to ensure safety for a given baseline control policy $u_r(x)$,

$$\begin{aligned} u_s(x) &:= \operatorname{argmin}_{u \in \mathcal{U}} \|u - u_r(x)\|^2 \\ \text{s.t. } &\nabla h(x)^\top (f(x) + g(x)u) + \alpha(h(x)) \geq 0, \end{aligned} \quad (7)$$

B. Neural Network-based CBF

The function $c(\cdot)$, which defines the boundary of the safe set \mathcal{X} , often fails to qualify as a CBF, either due to a lack of differentiability or because it does not satisfy the invariance condition (4). Thus, we must construct a CBF, either on a case-by-case basis or using function approximation techniques.

Specifically, let $h_\theta(x): \mathbb{R}^n \rightarrow \mathbb{R}$ be the parameterized function approximator. The parameters θ can be found by solving the following self-supervise risk minimization problem, in which the violation of both conditions (3) and (4) is penalized:

$$\min_{\theta} \mathbb{E}_{x \sim \rho} [(-\mathcal{I}_\gamma(h_\theta(x)))_+^2 + \lambda (h_\theta(x) - c(x))_+^2], \quad (8)$$

Here ρ is typically a uniform distribution over \mathcal{D} and

$$\mathcal{I}_\gamma(h) := \sup_{u \in \mathcal{U}} \{\nabla h(x)^\top (f(x) + g(x)u)\} + \gamma h(x). \quad (9)$$

This approach, however, focuses solely on satisfying the invariance conditions and does not incorporate notions of performance. Consequently, it may result in an overly conservative CBF in practice.

C. Performance-Oriented CBF via HJ PDE

Ideally, we want to find a CBF h whose zero super level set $L_0^+(h)$ has the maximal volume within the safe set \mathcal{X} . To achieve this, we can characterize the maximal control invariant set inside \mathcal{X} through an HJ PDE. Consider the dynamics (1),

$$\dot{x}(\tau) = f(x(\tau)) + g(x(\tau))u(\tau), \quad \tau \in [t, 0], \quad x(t) = x, \quad (10)$$

on a time interval $[t, 0]$ where $t \leq 0$ is the initial time, and x is the initial state. Let $\mathcal{U}_{[t, 0]}$ be the set of Lebesgue measurable functions $u: [t, 0] \rightarrow \mathcal{U}$, and $\psi(\tau) := \psi(\tau; x, t, u(\cdot)) : [t, 0] \mapsto \mathbb{R}^n$ denote the unique solution of (10) given x and $u(\cdot) \in \mathcal{U}_{[t, 0]}$. Given a bounded Lipschitz continuous function $c: \mathcal{D} \mapsto \mathbb{R}$, the viability kernel of $L_0^+(c) = \{x \mid c(x) \geq 0\}$ is defined as

$$\mathcal{V}(t) := \{x \in L_0^+(c) \mid \exists u(\cdot) \in \mathcal{U}_{[0, t]}, \text{ s.t. } \forall \tau \in [t, 0], \psi(\tau) \in L_0^+(c)\}. \quad (11)$$

Taking $t \rightarrow -\infty$, the viability kernel gives us the *maximal* control invariant set contained in $L_0^+(c)$ [28]. For $t \leq 0$, the viability kernel $\mathcal{V}(t)$ can be expressed as the super level set of the control barrier-value function $B_\gamma: \mathcal{D} \times [0, \infty) \mapsto \mathbb{R}$ defined as

$$B_\gamma(x, t) := \max_{u(\cdot) \in \mathcal{U}_{[t, 0]}} \min_{\tau \in [t, 0]} e^{\gamma(\tau-t)} c(\psi(\tau)), \quad (12)$$

where $\gamma \geq 0$. In other words, for each $\gamma \geq 0$ and all $t \leq 0$, we have $\{x \mid B_\gamma(x, t) \geq 0\} = \mathcal{V}(t)$ [28]. Furthermore, B_γ is a unique Lipschitz continuous viscosity solution of the HJ PDE:

$$\begin{aligned} & \min\{c(x) - B_\gamma(x, t), \frac{\partial}{\partial t} B_\gamma(x, t) + \\ & \max_{u \in \mathcal{U}} \nabla_x B_\gamma(x, t)^T (f(x) + g(x)u) + \gamma B_\gamma(x, t)\} = 0, \end{aligned} \quad (13)$$

with the boundary condition $B_\gamma(x, t) = c(x)$. We obtain the steady-state HJ PDE by taking $t \rightarrow -\infty$:

$$\begin{aligned} & \mathcal{N}_\gamma(c, B_\gamma) := \min\{c(x) - B_\gamma(x), \\ & \max_{u \in \mathcal{U}} \nabla_x B_\gamma(x)^T (f(x) + g(x)u) + \gamma B_\gamma(x)\} = 0. \end{aligned} \quad (14)$$

For a given input function $c(\cdot)$ to this PDE, the unique solution $B_\gamma(\cdot)$, when differentiable everywhere, is a valid CBF, satisfying both conditions (3) and (4). In the case where not fully differentiable, for certain representative cases, e.g., when $c(\cdot)$ is a signed distance function and the system dynamics is Lipschitz continuous, $B_\gamma(\cdot)$ will be differentiable almost everywhere [35].

Building on this, we can adopt a similar approach to the previous risk minimization formulation (8). However, this time, we focus on minimizing the residual of the PDE (14)

$$\min_{\theta} \mathbb{E}_{x \sim \rho} [(\mathcal{N}_\gamma(c(x), h_\theta(x)))^2]. \quad (15)$$

An appealing feature of the above training loss function is that it admits an “optimal” CBF in the sense that $L_0^+(h)$ is approximately maximal.

III. MAIN IDEA: DOMAIN ADAPTATION VIA OPERATOR LEARNING

The learned solution of problem (15) will be valid only for a specific environment defined by the function $c(\cdot)$. As a result, it cannot be deployed in novel environments without retraining. This limits its generalization and usefulness in real-time applications.

Instead of retraining a new CBF for each different environment, we propose to learn the *solution operator* of the PDE. Formally, given any $c(\cdot) \in C^0$ and the corresponding solution $h(\cdot) \in C^1$ that satisfies $\mathcal{N}_\gamma(c(x), h(x)) = 0 \forall x$, we define the solution operator $\mathcal{G}: C^0 \rightarrow C^1$ as

$$h(x) = \mathcal{G}(c(x)). \quad (16)$$

$\mathcal{G}(\cdot)$ is a nonlinear operator that maps between infinite dimensional Banach spaces. To approximately learn this map, we assume that the input function $c(\cdot)$ is parameterized by a *finite-dimensional* vector $e \in \mathbb{R}^{n_e}$ that represents the parameters of the environment, e.g., the location and sizes of obstacles. Under this assumption, we can abstract the safe set as $\mathcal{X} = \{x \mid c(x, e) \geq 0\}$.

By viewing the environment parameters as “virtual” states without dynamics, i.e., $\dot{e} = 0$, we define the augmented state as $\xi = (x, e)$, obeying the following dynamics,

$$\dot{\xi} = f^{\text{aug}}(\xi) + g^{\text{aug}}(\xi)u, \quad (17)$$

where

$$f^{\text{aug}}(\xi) = \begin{bmatrix} f(x) \\ 0 \end{bmatrix}, \quad g^{\text{aug}}(\xi) = \begin{bmatrix} g(x) \\ 0 \end{bmatrix}. \quad (18)$$

Using this augmentation, our approach is to learn a neural CBF in the *joint space* of the states and the environment parameters using a loss akin to (15),

$$\min_{\theta} \mathbb{E}_{\xi \sim \rho} [(\mathcal{N}_\gamma(c(\xi), h_\theta(\xi)))^2]. \quad (19)$$

In the next subsection, we provide the details of the proposed training loss.

A. Training Loss Function

To train the CBF operator, $h_\theta(x, e)$, we adopt a similar procedure as [6] which for each environment approximates the maximal control invariant set characterized by the HJ PDE (14) and enhances the feasibility of the safety filter with a CBF condition penalty term. We uniformly sample the environmental parameters $E = \{e_i\}_{i=1}^M$ and the states $X_e = \{x_e^j\}_{j=1}^N$ for each environment. Then, we construct the total dataset of joint states as $\Xi = \bigcup_{i=1}^M \{e_i\} \times X_{e_i}$, where \times denotes the Cartesian product of sets. We train the CBF operator $h_\theta(x, e)$ using two loss functions \mathcal{L}_{HJ} and \mathcal{L}_{CBF} .

\mathcal{L}_{HJ} for Maximal Control Invariant Set Approximation: This loss guides the CBF $h_\theta(\cdot, e)$ to approximate the maximal control invariant set for each environment

$$\mathcal{L}_{\text{HJ}}(\theta) = \frac{1}{|\Xi|} \sum_{\xi \in \Xi} (\mathcal{N}_\gamma(c(\xi), h_\theta(\xi)))^2. \quad (20)$$

When the control input set \mathcal{U} is polyhedral, the maximum over $u \in \mathcal{U}$ in the PDE is achieved at one of the vertices and can be obtained in closed form.

\mathcal{L}_{CBF} for Feasibility Enhancement: The HJ PDE is derived for the policy that achieves the maximal control invariant set; hence, the safe control set $\{u \in \mathcal{U} \mid \nabla_x h_\theta(x, e)^\top (f(x) + g(x)u) + \gamma h_\theta(x, e) \geq 0\}$ at state x can be small or even be a singleton. This tends to give rise to infeasibility of the CBF-QP (7) in practice. Therefore, we also explicitly enforce the second CBF condition (4) with \mathcal{L}_{CBF} ,

$$\mathcal{L}_{\text{CBF}}(\theta) = \frac{1}{|\Xi|} \sum_{\xi \in \Xi} [-\mathcal{I}_\gamma(h_\theta(\xi))_+]^2, \quad (21)$$

which improves the feasibility of the CBF-QP. We combine \mathcal{L}_{HJ} and \mathcal{L}_{CBF} into a single objective with scalarization hyperparameter $\lambda > 0$,

$$\min_{\theta} \mathcal{L}(\theta) = \mathcal{L}_{\text{HJ}}(\theta) + \lambda \mathcal{L}_{\text{CBF}}(\theta). \quad (22)$$

We use standard first-order optimization methods for this objective. We note that the first CBF condition (3), $c(\xi) \geq h_\theta(\xi)$, does not appear explicitly in the above loss function. Our specific parameterization will guarantee this property. The details of this parameterization, which plays a crucial role in improving the learning process, will be discussed in the following subsection.

B. CBF Parameterization

Instead of learning the CBF operator h directly, we can learn the difference function $c - h$. To formalize this, we introduce a parameterized CBF operator as

$$h_\theta(x, e) = c(x, e) - \delta_\theta(x, e), \quad (23)$$

where $\delta_\theta(\cdot, e) : \mathbb{R}^{n+n_e} \mapsto \mathbb{R}$ is a non-negative valued function approximator. Importantly, by construction, we have $h_\theta(x, e) \leq c(x, e)$ which implies that the 0-superlevel set of $h_\theta(\cdot, e)$ is always contained in the safe region \mathcal{X} for all environments. We parameterize the difference function $\delta_\theta(x, e)$ as the following MLP,

$$\begin{aligned} z_0 &= [x^\top \ e^\top]^\top, \\ z_{k+1} &= \sigma(W_k z_k + b_k), k = 0, \dots, L-1, \\ \delta_\theta(x, e) &= \sigma_+(W_L z_L + b_L), \end{aligned} \quad (24)$$

where W_k, b_k denote the weights and bias of the $(k+1)$ -th linear layer, $\sigma(\cdot)$ is a smooth activation, and $\sigma_+(\cdot)$ is a smooth non-negative valued activation such as the Softplus function. Note that $\delta_\theta(x, e)$ is a function of all the states and so is the CBF operator h_θ . Therefore, we alleviate the issue of high relative degree [4] through the parameterization of the MLP CBF operator.

However, the above parameterization is valid only when $c(x, e)$ is continuously differentiable. In the next subsection, we address cases where $c(x, e)$ is not differentiable or where the environment includes multiple safety constraints.

C. Complex Environments

The PDE (14) takes a single constraint function $c(\cdot, \cdot)$ as input. For complex environments with multiple constraints, we propose to build this function by an arbitrary composition of conjunctions, disjunctions, and negations [36] [5]. To illustrate

this, consider the sets $\mathcal{S}_i = L_0^+(s_i), i \in I$. The following relations hold for any e :

$$\begin{aligned} \bigcup_{i \in I} \mathcal{S}_i &= \{x \in \mathcal{D} \mid \max_{i \in I} s_i(x, e) \geq 0\} \\ \bigcap_{i \in I} \mathcal{S}_i &= \{x \in \mathcal{D} \mid \min_{i \in I} s_i(x, e) \geq 0\} \\ \bar{\mathcal{S}}_i &= \{x \in \mathcal{D} \mid -s_i(x, e) \geq 0\}. \end{aligned} \quad (25)$$

Similarly, arbitrary combinations of unions and intersections can be captured by composing the max and min operators.

Example 1: Consider the Dubins car dynamics $\dot{x}_1 = V \cos(\psi), \dot{x}_2 = V \sin(\psi), \dot{\theta} = u$ where (x_1, x_2) denote the position coordinates on the 2D plane, ψ denotes the angle of the unicycle, V is a constant speed, and u is the angular velocity considered as control input. Suppose there are two circular obstacles in the environment which respectively give the following safety constraints:

$$\mathcal{S}_i = \{x \in \mathbb{R}^2 \mid \|x - o_i\|_i^2 \geq r_i^2\} \quad i = 1, 2,$$

The parameter e in this example is given by $e = [o_1^\top \ r_1 \ o_2^\top \ r_2]^\top \in \mathbb{R}^6$, and the safe set $\mathcal{X} = \mathcal{S}_1 \cap \mathcal{S}_2$ is represented as

$$\mathcal{X} = \{x \mid c(x, e) := \min(\|x - o_1\|_1^2 - r_1^2, \|x - o_2\|_2^2 - r_2^2) \geq 0\}.$$

The above composition procedure may not preserve differentiability, resulting in a non-smooth c in the parameterization (23). To address this, we can approximate these operators via the perspective of the LSE. These approximations for the max operator are as follows,

$$\begin{aligned} \text{LSE}(s; \beta) - \frac{\log(|I|)}{\beta} &\leq \max_{i \in I} s_i(x, e) \leq \text{LSE}(s; \beta), \\ -\text{LSE}(-s; \beta) &\leq \min_{i \in I} s_i(x, e) \leq -\text{LSE}(-s; \beta) + \frac{\log(|I|)}{\beta}, \end{aligned} \quad (26)$$

becoming arbitrarily tight as $\beta \rightarrow +\infty$. Now for a given $c(\cdot, e)$ composed of max and min operators, we can construct a single smooth function $\underline{c}(\cdot, e)$ such that $\underline{c}(x, e) \leq c(x, e)$ for all $x \in \mathcal{D}$. By substituting $\underline{c}(x, e)$ for $c(x, e)$ in (23), we obtain the following CBF parameterization,

$$h_\theta(x, e) = \underline{c}(x, e) - \delta_\theta(x, e). \quad (27)$$

It directly follows that $L_0^+(h_\theta(\cdot, e)) \subseteq L_0^+(\underline{c}(\cdot, e)) \subseteq L_0^+(c(\cdot, e)) = \mathcal{X}$ for any $e \in \mathbb{R}^{n_e}$.

D. Pipeline

Here, we summarize the training pipeline. For each environmental parameter e , the safe set is defined as $L_0^+(c(x, e))$. We first fix $\beta > 0$ and obtain a smooth lower bound $\underline{c}(x, e)$ of $c(\cdot)$ through (26). Next, we parameterize the neural CBF operator as $h_\theta(x, e) = \underline{c}(x, e) - \delta_\theta(x, e)$. Finally, we optimize the training loss in (22) over the parameters θ .

IV. EXPERIMENTS

We evaluate the performance of the CBF operator learning framework and demonstrate its effectiveness as an adaptive safety filter through numerical experiments. In all experiments, the MLP δ_θ consists of four hidden layers, each with 50




Fig. 1. ● start state, ▲ target state, — zero contour of the CBF instance $h_\theta(\cdot, e)$, horizontal-axis (position) and vertical-axis (velocity). The circular obstacles are marked black. For each environment, the CBF value $h_\theta(\cdot, e)$ is denoted by the color, with red denoting high values and blue denoting low values.

neurons and tanh activation functions. The output layer uses a softplus activation to ensure differentiability and non-negativity. The CBF operator is trained in JAX [37] using the ADAM optimizer [38] on a single T4 GPU.

A. Double Integrator

In the first example, we consider the double integrator dynamics:

$$\begin{bmatrix} \dot{x} \\ \dot{v} \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} x \\ v \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \end{bmatrix} u, \quad (28)$$

where $u \in [-1, 1]$, x and v denote the position and velocity, respectively. The safe set is restricted to the set $[0, 10] \times [-5, 5]$ with two additional varying circular obstacles parameterized as $e = [r_1, x_{c1}, v_{c1}, r_2, x_{c2}, v_{c2}]$. The training data is composed of 1000 uniformly randomly sampled obstacle configurations. The radii were sampled from the interval $[1, 2]$. For each configuration, 10,000 state samples are drawn uniformly from $[-1, 11] \times [-6, 6]$.

In Figure 1, we apply the learned CBF operator to filter a proportional-derivative reference controller. The underlying shape of $L_0^+(h_\theta(\cdot, e))$ is robust across environments and is close to the maximal control invariant set of the double integrator in obstacle-free environments. Furthermore, the closed-loop system trajectories under the adaptive safety filter demonstrate the efficacy of the learned CBF operator.

B. Unicycle

Next, we consider the unicycle system given by

$$\begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{\psi} \end{bmatrix} = A(\psi)u, \quad A(\psi) = \begin{bmatrix} \cos(\psi) & 0 \\ \sin(\psi) & 0 \\ 0 & 1 \end{bmatrix}, \quad u = \begin{bmatrix} v \\ \omega \end{bmatrix}, \quad (29)$$

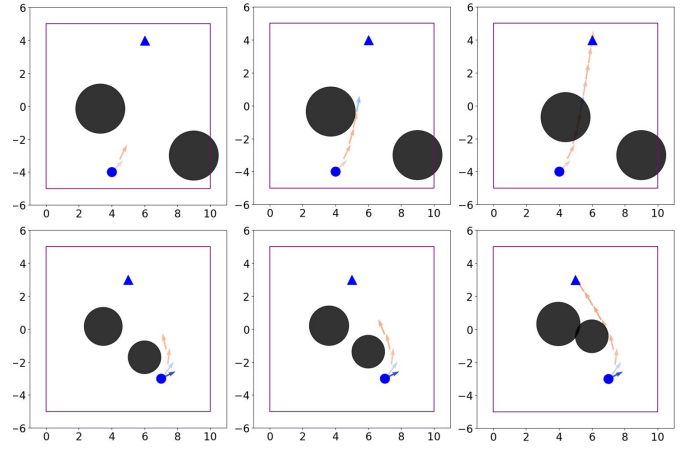


Fig. 2. ● start state, ▲ target state. The horizontal and vertical axes denote the position of the unicycle, and the arrow denotes its orientation. Top row: An obstacle moves from left to right trying to block the agent. Bottom row: the smaller obstacle chases the agent while the bigger obstacle moves slowly and grows over time. During motion, the CBF operator value changes in response to the moving obstacles. The heatmap on the arrows indicates the CBF values along the trajectory, with warmer colors denoting higher values.

where (x, y) denotes the position, ψ denotes the orientation angle, v is the velocity, and ω is the angular velocity. The controls are restricted to $\omega \in [-1, 1]$ and $v \in [0.2, 2]$. The safe set is given by $(x, y, \psi) \in [0, 10] \times [-5, 5] \times [-\infty, \infty]$ with two additional varying circular obstacles in x, y (see Figure 2). Importantly, the obstacles are moving which requires the system to apply different CBFs over time. As in the double integrator experiment, we sample 1000 environmental parameters but only 30,000 states to demonstrate sample efficiency, as opposed to the 1 million states predicted by dimensional scaling $(10^4)^{\frac{3}{2}}$. The observed sample efficiency exceeds expectations based on simple scaling, highlighting the potential of CBF operator learning for higher-dimensional systems.

The reference controller for the unicycle is given by [39]:

$$\text{rerr}(x, y) = R(\psi)^{-1} \begin{bmatrix} x_f - x \\ y_f - y \end{bmatrix}, \quad (30)$$

$$d = \text{atan2}(\text{rerr}[1], \text{rerr}[0]), \quad \omega = k_\omega d, \quad v = k_v \cos(d),$$

with $R(\psi)$ being a rotation matrix and $k_\omega, k_v > 0$ arbitrary gains. We apply the learned CBF operator as the safety filter in Figure 2 with two different obstacle setups. In both of them, the unicycle system is able to make timely response and reach the target without collision.

V. CONCLUSION

We proposed a deep operator learning framework to learn domain-adaptive control barrier functions that can adapt to novel environments without retraining. Leveraging domain randomization, the proposed framework learns the solution operator of a Hamilton-Jacobi PDE that maps the environment parameters to a control barrier function that approximates the maximal control invariant set. Our numerical experiments

show the effectiveness of the learned CBF to ensure safe navigation in new environments. For future work, we will incorporate a falsification/verification module to refine the learned model based on counter-examples.

REFERENCES

- [1] 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)*, pp. 3420–3431, IEEE, 2019.
- [2] S. Liu, C. Liu, and J. Dolan, "Safe control under input limits with neural control barrier functions," in *Conference on Robot Learning*, pp. 1970–1980, PMLR, 2023.
- [3] Q. Nguyen and K. Sreenath, "Exponential control barrier functions for enforcing high relative-degree safety-critical constraints," in *2016 American Control Conference (ACC)*, pp. 322–328, IEEE, 2016.
- [4] W. Xiao and C. Belta, "High-order control barrier functions," *IEEE Transactions on Automatic Control*, vol. 67, no. 7, pp. 3655–3662, 2021.
- [5] T. G. Molnar and A. D. Ames, "Composing control barrier functions for complex safety specifications," *IEEE Control Systems Letters*, 2023.
- [6] S. Chen and M. Fazlyab, "Learning performance-oriented control barrier functions under complex safety constraints and limited actuation," *arXiv preprint arXiv:2401.05629*, 2024.
- [7] C. Dawson, S. Gao, and C. Fan, "Safe control with learned certificates: A survey of neural lyapunov, barrier, and contraction methods for robotics and control," *IEEE Transactions on Robotics*, 2023.
- [8] A. Robey, H. Hu, L. Lindemann, H. Zhang, D. V. Dimarogonas, S. Tu, and N. Matni, "Learning control barrier functions from expert demonstrations," in *2020 59th IEEE Conference on Decision and Control (CDC)*, pp. 3717–3724, IEEE, 2020.
- [9] C. Dawson, Z. Qin, S. Gao, and C. Fan, "Safe nonlinear control using robust neural lyapunov-barrier functions," in *Conference on Robot Learning*, pp. 1724–1735, PMLR, 2022.
- [10] Z. Qin, K. Zhang, Y. Chen, J. Chen, and C. Fan, "Learning safe multi-agent control with decentralized neural barrier certificates," *arXiv preprint arXiv:2101.05436*, 2021.
- [11] P. Glotfelter, J. Cortés, and M. Egerstedt, "Nonsmooth barrier functions with applications to multi-robot systems," *IEEE control systems letters*, vol. 1, no. 2, pp. 310–315, 2017.
- [12] L. Lindemann and D. V. Dimarogonas, "Control barrier functions for signal temporal logic tasks," *IEEE control systems letters*, vol. 3, no. 1, pp. 96–101, 2018.
- [13] P. Rabiee and J. B. Hoagg, "Soft-minimum barrier functions for safety-critical control subject to actuation constraints," in *2023 American Control Conference (ACC)*, pp. 2646–2651, IEEE, 2023.
- [14] P. Rabiee and J. B. Hoagg, "Soft-minimum and soft-maximum barrier functions for safety with actuation constraints," *arXiv preprint arXiv:2305.10620*, 2023.
- [15] A. D. Ames, G. Notomista, Y. Wardi, and M. Egerstedt, "Integral control barrier functions for dynamically defined control laws," *IEEE control systems letters*, vol. 5, no. 3, pp. 887–892, 2020.
- [16] A. Safari and J. B. Hoagg, "Time-varying soft-maximum control barrier functions for safety in an a priori unknown environment," *arXiv preprint arXiv:2310.05261*, 2023.
- [17] J. Breeden and D. Panagou, "Compositions of multiple control barrier functions under input constraints," in *2023 American Control Conference (ACC)*, pp. 3688–3695, IEEE, 2023.
- [18] W. Xiao, T.-H. Wang, R. Hasani, M. Chahine, A. Amini, X. Li, and D. Rus, "BarrierNet: Differentiable control barrier functions for learning of safe robot control," *IEEE Transactions on Robotics*, 2023.
- [19] W. Xiao, C. G. Cassandras, and C. A. Belta, "Learning feasibility constraints for control barrier functions," *arXiv preprint arXiv:2303.09403*, 2023.
- [20] H. Ma, B. Zhang, M. Tomizuka, and K. Sreenath, "Learning differentiable safety-critical control using control barrier functions for generalization to novel environments," in *2022 European Control Conference (ECC)*, pp. 1301–1308, IEEE, 2022.
- [21] B. Djeridane and J. Lygeros, "Neural approximation of pde solutions: An application to reachability computations," in *Proceedings of the 45th IEEE Conference on Decision and Control*, pp. 3034–3039, IEEE, 2006.
- [22] B. Dai, H. Huang, P. Krishnamurthy, and F. Khorrami, "Data-efficient control barrier function refinement," in *2023 American Control Conference (ACC)*, pp. 3675–3680, IEEE, 2023.
- [23] T. Nakamura-Zimmerer, Q. Gong, and W. Kang, "Adaptive deep learning for high-dimensional hamilton-jacobi-bellman equations," *SIAM Journal on Scientific Computing*, vol. 43, no. 2, pp. A1221–A1247, 2021.
- [24] O. So, Z. Serlin, M. Mann, J. Gonzales, K. Rutledge, N. Roy, and C. Fan, "How to train your neural control barrier function: Learning safety filters for complex input-constrained systems," *arXiv preprint arXiv:2310.15478*, 2023.
- [25] B. Dai, P. Krishnamurthy, and F. Khorrami, "Learning a better control barrier function under uncertain dynamics," *arXiv preprint arXiv:2310.04795*, 2023.
- [26] D. C. Tan, F. Acero, R. McCarthy, D. Kanoulas, and Z. A. Li, "Your value function is a control barrier function: Verification of learned policies using control theory," *arXiv preprint arXiv:2306.04026*, 2023.
- [27] Z. Qin, D. Sun, and C. Fan, "Sablas: Learning safe control for black-box dynamical systems," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 1928–1935, 2022.
- [28] J. J. Choi, D. Lee, K. Sreenath, C. J. Tomlin, and S. L. Herbert, "Robust control barrier-value functions for safety-critical control," in *2021 60th IEEE Conference on Decision and Control (CDC)*, pp. 6814–6821, IEEE, 2021.
- [29] S. Tonkens and S. Herbert, "Refining control barrier functions through hamilton-jacobi reachability," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 13355–13362, IEEE, 2022.
- [30] A. J. Taylor and A. D. Ames, "Adaptive safety with control barrier functions," in *2020 American Control Conference (ACC)*, pp. 1399–1405, IEEE, 2020.
- [31] B. T. Lopez, J.-J. E. Slotine, and J. P. How, "Robust adaptive control barrier functions: An adaptive and data-driven approach to safety," *IEEE Control Systems Letters*, vol. 5, no. 3, pp. 1031–1036, 2020.
- [32] R. Takano and M. Yamakita, "Robust constrained stabilization control using control lyapunov and control barrier function in the presence of measurement noises," in *2018 IEEE Conference on Control Technology and Applications (CCTA)*, pp. 300–305, IEEE, 2018.
- [33] W. Xiao, C. Belta, and C. G. Cassandras, "Adaptive control barrier functions," *IEEE Transactions on Automatic Control*, vol. 67, no. 5, pp. 2267–2281, 2021.
- [34] T. G. Molnar, A. K. Kiss, A. D. Ames, and G. Orosz, "Safety-critical control with input delay in dynamic environment," *IEEE transactions on control systems technology*, 2022.
- [35] K. P. Wabersich, A. J. Taylor, J. J. Choi, K. Sreenath, C. J. Tomlin, A. D. Ames, and M. N. Zeilinger, "Data-driven safety filters: Hamilton-jacobi reachability, control barrier functions, and predictive methods for uncertain systems," no. 5, pp. 137–177, 2023.
- [36] P. Glotfelter, J. Cortes, and M. Egerstedt, "Nonsmooth barrier functions with applications to multi-robot systems," vol. 1, no. 2, pp. 310–315.
- [37] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang, "JAX: composable transformations of Python+NumPy programs," 2018.
- [38] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [39] P. Mestres, C. Nieto-Granda, and J. Cortés, "Distributed safe navigation of multi-agent systems using control barrier function-based optimal controllers."