

Simultaneous Synthesis and Verification of Neural Control Barrier Functions through Branch-and-Bound Verification-in-the-Loop Training

Xinyu Wang¹, Luzia Knoedler¹, Frederik Baymler Mathiesen², and Javier Alonso-Mora¹

Abstract—Control Barrier Functions (CBFs) that provide formal safety guarantees have been widely used for safety-critical systems. However, it is non-trivial to design a CBF. Utilizing neural networks as CBFs has shown great success, but it necessitates their certification as CBFs. In this work, we leverage bound propagation techniques and the Branch-and-Bound scheme to efficiently verify that a neural network satisfies the conditions to be a CBF over the continuous state space. To accelerate training, we further present a framework that embeds the verification scheme into the training loop to synthesize and verify a neural CBF simultaneously. In particular, we employ the verification scheme to identify partitions of the state space that are not guaranteed to satisfy the CBF conditions and expand the training dataset by incorporating additional data from these partitions. The neural network is then optimized using the augmented dataset to meet the CBF conditions. We show that for a non-linear control-affine system, our framework can efficiently certify a neural network as a CBF and render a larger safe set than state-of-the-art neural CBF works. We further employ our learned neural CBF to derive a safe controller to illustrate the practical use of our framework.

I. INTRODUCTION

Safety is a critical element of autonomous systems, such as self-driving cars and manipulators that interact with humans. As autonomous systems grow more complex, determining whether they operate safely becomes challenging.

Safety can be formulated via invariance, in the sense that any trajectory originating within an invariant set will never traverse beyond the boundaries of that set. Lately, the use of Control Barrier Functions (CBFs) to derive a forward invariant set has received significant attention in the control and learning community [1]. However, there exists no general and scalable technique for designing CBFs. Therefore, recent works [2], [3] synthesize continuous CBFs using Neural Networks (NNs) as a function template, which are referred to as Neural Control Barrier Functions (nCBFs). Yet, these works rely on an initial guess of the forward invariant set or the function structure of the CBF to synthesize the nCBF. An improper initial guess usually results in a suboptimal

This paper has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No. 101017008. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

¹Xinyu Wang, Luzia Knoedler, and Javier Alonso-Mora are with the Cognitive Robotics Department, Delft University of Technology, 2628 CD Delft, The Netherlands x.wang-55@student.tudelft.nl l.knoedler@tudelft.nl j.alonsomora@tudelft.nl

²Frederik Baymler Mathiesen is with the Delft Center for Systems and Control, Delft University of Technology, 2628 CD Delft, The Netherlands f.b.mathiesen@tudelft.nl




Fig. 1: A schematic overview of the presented Branch-and-Bound Verification-in-the-Loop Training. The framework comprises of two key components: the learner and the verifier, which operate sequentially. The learner optimizes the nCBF using a fixed dataset and a counterexample dataset. The verifier leverages bound propagation techniques and the Branch-and-Bound scheme to refine a partition of the state space until the CBF conditions are satisfied or counterexamples are generated.

nCBF. Constructing an optimal CBF that renders a maximum forward invariant set is challenging. A recent work [4] introduced the Control Barrier-Value Function (CBVF) which is a safe value function and renders the maximum forward invariant set for a chosen time span. In this work, we synthesize a continuous nCBF that approximates the infinite-horizon CBVF and renders a safe set that is close to the maximum forward invariant set.

Although utilizing NNs as CBFs offers universal approximation capabilities, it necessitates their certification as CBFs to provide safety guarantees. Verifying the NN as an nCBF in the continuous state space presents a significant challenge. Specifically, since the NN is trained using a finite set of data points, it will only be verified on those points. Outside the certified points, safety is no longer guaranteed. There are works [5], [6] that use Satisfiability Modulo Theory (SMT) to verify their NNs. However, they are restricted to very simple NNs due to expensive computation. In this work, we leverage bound propagation techniques [7] and the Branch-and-Bound scheme (BBS) to efficiently verify nCBFs. In particular, we partition the state space and utilize linear bound propagation techniques to provide lower and upper bounds of the NN and its Jacobian. These bounds are used to verify if the NN satisfies the conditions to be a CBF. The BBS is applied

to refine the partition to improve scalability and achieve less conservative bounds. We refer to the above verification scheme as Branch-and-Bound Verification (BBV). This approach is similar to [8], however, we verify CBFs instead of barrier functions. To accelerate training, we embed the BBV into the training loop to synthesize an nCBF simultaneously, which we refer to as Branch-and-Bound Verification-in-the-Loop Training (BBVT), see Fig. 1. We show the efficiency of our method and the practical use of nCBFs on an inverted pendulum and a 2D navigation task in a simulation environment.

II. RELATED WORK

Many works use CBFs to ensure the safety of a system [9]–[11]. However, it is non-trivial to construct CBFs. In recent years, new techniques emerged to automatically synthesize CBFs. For a system with polynomial dynamics, a CBF can be obtained by solving a sum-of-squares (SOS) optimization problem [12]. Unfortunately, SOS scales poorly to higher dimensional systems [13]. To address this shortcoming, NNs have been employed to approximate CBFs. They are trained by supervised learning [2], [3], [14] or Reinforcement Learning (RL) with the Actor-Critic framework [15], [16]. However, the quality of the nCBF in those works depends on an initial guess of the forward invariant set, CBF candidate, or exploration strategy. An improper initial guess results in a conservative nCBF with a small forward invariant set. To address the conservativity, in this work, we learn a continuous nCBF that renders a safe set close to the maximum forward invariant set. Furthermore, the training does not require an initial guess.

Commonly, NNs are trained through backpropagation of the empirical loss on a finite set of data points. Therefore, it is important to note that even an empirical loss of zero does not guarantee that the certificate is valid everywhere in the state space. Only a few works have verified their NNs, such as [5], [6], [17], which leverage SMT to provide counterexamples (CEs) and guarantee the correctness of the synthesis procedure. However, SMT is limited to simple NNs with around 20 neurons in one or two hidden layers due to the need for expensive computation. In contrast to using SMT for exact verification, several efficient NN verification methods using linear bound propagation techniques have been developed [7], [18]. These bounding methods provide a new direction to verify neural certificates. The work in [8] partitions the state space with a BBS and verifies the property of the discrete-time stochastic barrier function for each partition leveraging the method in [7]. Our work extends the BBS of [8] to CBFs for continuous-time deterministic control-affine systems where the control input constraints must be considered and uses the BBV scheme to verify the learned continuous nCBF.

III. PROBLEM FORMULATION

Given the following continuous-time control-affine system

$$\dot{x} = f(x) + g(x)u; \quad x(0) = x_0; \quad (1)$$

where $x \in X \subset \mathbb{R}^n$, $u \in U \subset \mathbb{R}^m$, $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ denotes the autonomous dynamics, and $g : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times m}$ denotes the input dynamics. We assume that f and g are Lipschitz continuous and X, U are compact sets.

The safety requirement for the system in (1) is encoded via a state admissible set $X_a \subset X$ and a convex input admissible set $U_a \subset U$. A safe system stays in the state admissible set for all time. To formally define safety, we use $x(t; x_0)$ to refer to a trajectory of the system in (1) at time t with initial condition x_0 and control policy $u = u(x)$. Safety is then defined as:

Definition 1 (Safety) The system in (1) is safe if $x(t; x_0) \in X_a$ and $u = u(x(t; x_0)) \in U_a; \forall t \in [0, 1]$.

However, it should be noted that X_a is not safe everywhere. There may not exist a control input that transitions a state close to the boundary towards the interior of X_a . A safe set should have the property that if the system starts in the safe set, it stays inside for all time. Towards formally defining this property, let a set C be defined as the 0-superlevel set of a continuously differentiable function $h : \mathbb{R}^n \rightarrow \mathbb{R}$, i.e.,

$$C = \{x \in X : h(x) \geq 0\};$$

$$\partial C = \{x \in X : h(x) = 0\}.$$

Then forward invariance and a safe set are defined as follows.

Definition 2 (Forward invariance) The set C is forward invariant if for every $x_0 \in C$, there exists a control policy $u = u(x) \in U_a$ such that the trajectory of system in (1) $x(t; x_0) \in C; \forall t \in [0, 1]$.

Definition 3 (Safe Set) The set C is a Safe Set if C is forward invariant and $C \subset X_a$.

A CBF renders a safe set and can be used to derive safe control inputs. Before defining CBFs, we must introduce extended class \mathcal{K}_1 functions. An extended class \mathcal{K}_1 function is a mapping $\alpha : \mathbb{R} \rightarrow \mathbb{R}$ that is strictly increasing and for which $\alpha(0) = 0$ holds. We define a continuous CBF as:

Definition 4 (Control Barrier Function) Let $C \subset X_a$ be the 0-superlevel set of a continuously differentiable function $h : \mathbb{R}^n \rightarrow \mathbb{R}$, then h is a CBF in X_a for system in (1) if there exists an extended class \mathcal{K}_1 function α such that

$$\sup_{u \in U_a} [L_f h(x) + L_g h(x)u] \geq \alpha(h(x)) \quad (2)$$

for all $x \in X_a$, where L_f, L_g represent Lie derivatives.

With the definition of a CBF, we may derive sufficient conditions for a safe system. According to the main result in [19], the following theorem holds:

Theorem 1 ([19, Theorem 2]) If function h is a CBF for the system in (1) and $\frac{\partial h}{\partial x}(x) \in \mathcal{K}_1$ for all $x \in \partial C$, then any Lipschitz continuous controller $u \in K_{cbf}(x)$ with

$$K_{cbf}(x) = \{u \in U_a : L_f h(x) + L_g h(x)u + \alpha(h(x)) \geq 0\} \quad (3)$$

renders the set \mathcal{C} safe. Additionally, the set \mathcal{C} is asymptotically stable in X_a .

With Theorem 1, we are able to ensure the safety of the system in (1) as long as a CBF is found and its gradient does not vanish on \mathcal{C} . We show in Section V-D how to obtain a safe policy satisfying (3) using CBFs.

The objective of this work is to automatically synthesize an nCBF and verify it for the continuous state space. The problem is defined as follows.

Problem 1. Given the system (1), state admissible set \mathcal{X}_a , convex input admissible set \mathcal{U}_a and $\dot{x} = f(x) + g(x)u$ where δ is a positive constant, synthesize an nCBF that is denoted by $\hat{h}_w(x)$, where w are the parameters of the NN, and renders set \mathcal{C} safe for the system (1). This is equivalent to

$$\hat{C} \supseteq X_a; \quad (4a)$$

$$\text{inequality (2) holds in } X_a; \quad (4b)$$

where $\hat{C} = \{x \in X : \hat{h}_w(x) \geq 0\}$ is the 0-superlevel set of the nCBF.

Remark 1. The condition $\frac{\partial \hat{h}_w}{\partial x}(x) \neq 0$ for all $x \in \mathcal{C}$ is omitted since it generally holds in our setting as we consider Tanh-based Fully-Connected Neural Network (FCNN). More specifically, since $\frac{\partial \tanh}{\partial x}(x) \in (0; 1]$ for all x , the condition is only violated if either $w = 0$ or catastrophic cancellation occurs in the linear layers, which will almost surely never happen.

IV. NEURAL CONTROL BARRIER FUNCTION TRAINING AND VERIFICATION

In this work, we design a new empirical loss to synthesize an nCBF, which is introduced in Section IV-A. As the training set only contains a finite set of data points, the CBF conditions may not hold in the continuous state space. Therefore, in Section IV-B, we present the BBV to verify nCBFs. Nevertheless, it is often necessary to iterate through multiple training and verification cycles before successfully learning an nCBF. Thus, we introduce BBVT in Section IV-C, which embeds BBV in the training loop to accelerate training for certifiability.

A. Learning a Neural Control Barrier Function

The primary goal of this work is to train an NN $\hat{h}_w(x)$ until it satisfies conditions (4a) and (4b) and render a large forward invariant set. Towards this end, we leverage the main result in [4, Theorem 3], where a CBF is shown to recover the maximum safe set subject to safety constraints. Conversely, we are interested in finite-horizon properties. Thus we extend the time-dependent Control Barrier-Value Function Variational Inequality (CBVF-VI) to the finite-horizon. Let $h(x)$ denote the finite-horizon CBVF and $\rho(x) : X \rightarrow \mathbb{R}$ denote the signed-distance function for the set X_a , which is defined as $\rho(x) = \inf_{y \in X_a} \|y - x\|$ if $x \in X_a$ and $\rho(x) = -\inf_{y \in X_a} \|y - x\|$ if $x \in X \setminus X_a$. The finite-horizon CBVF-VI is defined as

$$0 = \min_{x \in X} f(x) - h(x); \quad (5)$$

$$\max_{u \in \mathcal{U}_a} L_f h(x) + L_g h(x)u + \delta h(x)g; \quad (5)$$

We use an NN $\hat{h}_w(x)$ to approximate the infinite-horizon CBVF $h(x)$. Then, the empirical loss is defined as follows:

$$L = \frac{1}{N_1} \sum_{x \in X_a} \min_{u \in \mathcal{U}_a} [L_f \hat{h}_w(x) + L_g \hat{h}_w(x)u + \delta \hat{h}_w(x) - g]; \quad (6a)$$

$$\sup_{u \in \mathcal{U}_a} L_f \hat{h}_w(x) + L_g \hat{h}_w(x)u + \delta \hat{h}_w(x) - g; \quad (6a)$$

$$+ \frac{1}{N_2} \sum_{x \in X \setminus X_a} \max_{u \in \mathcal{U}_a} \hat{h}_w(x) + \delta; \quad (6b)$$

where δ is a small positive constant to encourage the strict satisfaction of the conditions. The loss term (6a) shapes the NN to be the solution of the finite-horizon CBVF-VI introduced in (5), which encourages the satisfaction of condition (4b). The loss term (6b) encourages that the nCBF is negative in the inadmissible area $X \setminus X_a$, which is equivalent to condition (4a). Since the system is control-affine, the optimal solution u^* for $\sup_{u \in \mathcal{U}_a} [L_f \hat{h}_w(x) + L_g \hat{h}_w(x)u]$ must be one of the vertices \mathcal{U}_a^V . Let \mathcal{U}_a^V denote the vertices of the input admissible set, we choose control input $u^* = \arg \max_{u \in \mathcal{U}_a^V} L_g \hat{h}_w(x)u$. However, the Lie derivative of $\hat{h}_w(x)$ in the early training stage may not align with the Lie derivative of the true CBVF $h(x)$. This results in an undesirable optimization path and the NN can get stuck at a deadlock. The occurrence of a deadlock situation signifies that improvements at certain data points cause constraint violations at other data points, as noted in [20]. To facilitate the training process and avoid deadlocks, we borrow ideas from [2], [6], which use a nominal controller to guide the training. Here, we train another neural network \hat{h}_w with the same structure as \hat{h}_w based on the loss of [21] and choose $u^* = \arg \max_{u \in \mathcal{U}_a^V} \hat{h}_w(x + (f(x) + g(x)u) \cdot t)$ by simulating one step ahead to guide the training of the nCBF. To further guide the training, we may integrate the verification procedure with a so-called Counterexample Guided Inductive Synthesis (CEGIS) approach, as described in Section IV-C.

B. Verifying the learned Neural Control Barrier Function

Since the NN is trained on finite data points, one must note that the NN may not satisfy the CBF conditions everywhere in the state space, even if the empirical loss decreases to zero. In fact, condition (4b) may be violated almost everywhere, which means the NN may fail to render a forward invariant set and the safety guarantee no longer exists. In this section, we propose to use the BBV to verify the learned nCBF in the continuous state space. Specifically, our primary goal is to verify the satisfaction of conditions (4a) and (4b). Before we explain our verification scheme in detail, we introduce some notations first. Let the partition of the state space be denoted as hyperrectangle $\mathcal{R}_i = \{x : x_j \in [x_{j-1}, x_{j+1}]\}$ centered at point $x_i \in X$ with radius $r_i \in \mathbb{R}^n$,

Similarly to condition (4a), the BBS starts from a coarse partition and refines each hyperrectangle when $q_i < 0$ until $q_i \geq 0$ or $t_i \geq t_{\text{gap}}$. If condition (9) holds for all hyperrectangles $B(x_i; i)$, then condition (4b) holds in the continuous state space.

However, the challenge arises in the computation of q_i . The computation of q_i can be reframed as an optimization problem within the hyperrectangle $B(x_i; i)$

$$q_i = \min_x q(x) \quad (10a)$$

$$\text{st: } x \in B(x_i; i) \quad (10b)$$

Fig. 2: An example of the BBV in a 2D state space. The scheme starts with a coarse partition $B(x_0; 0)$ and refines it using the Branch-and-Bound scheme. For each hyperrectangle $B(x_i; i)$, $i = 0; 1; 2; \dots$, upper bounds for the neural network are computed. In this case, the hyperrectangles $B(x_0; 0)$ and $B(x_2; 2)$ are refined as $\hat{h}_{u0} > 0$, $\hat{h}_{u2} > 0$.

The term $q(x)$ is a complex function containing nonlinear dynamic functions f, g , the NN \hat{h}_w as well as its Jacobian, which renders a constrained Nonlinear Program (NLP) in (10a). The state-of-the-art NLP solver [22] requires gradients of the objective function, which involves computation of the Hessian of the NN. The expensive computation makes it impractical to solve (10a) directly.

Although computing the lower bound $q(x)$ is quite complex, computing the bound of the components $q(x)$ separately is much simpler. We can compute the bound of the NN using CROWN [7] and its Jacobian leveraging a recent result in [23] or [24]:

$$\hat{h}_{li} \leq \hat{h}_w(x) \leq \hat{h}_{ui}; \quad x \in B(x_i; i); \quad (11)$$

$$J_{li} \leq J_{\hat{h}_w(x)} \leq J_{ui}; \quad x \in B(x_i; i); \quad (12)$$

see Fig. 2. Initially, all hyperrectangles have the same radius $r_i = r_{\text{init}}$. Let $B = \{B(x_0; 0); \dots; B(x_N; N)\}$ denote the set of all hyperrectangles $B(x; i)$, $B_{X=X_a}$ denote the set of hyperrectangles that covers the inadmissible area X_a , and $B_{X_a} \setminus B$ denote the set of hyperrectangles that covers the admissible area.

To verify condition (4a), which is equivalent to $\hat{h}_w(x) < 0; \forall x \in X \setminus X_a$, we rely on the linear bounds of the NN computed using CROWN [7]. The linear bounds are defined as follows:

$$\hat{h}_{li} \leq \hat{h}_w(x) \leq \hat{h}_{ui}; \quad x \in B(x_i; i); \quad (7)$$

We use these linear bounds to certify the satisfaction of condition (4a). In particular, the upper bound \hat{h}_{ui} can be used to check for non-positivity

$$\hat{h}_w(x) \leq \hat{h}_{ui} < 0; \quad x \in B(x_i; i); \quad B(x_i; i) \in B_{X \setminus X_a}; \quad (8)$$

However, this upper bound tends to be conservative when $B(x_i; i)$ covers a large area. Therefore, we leverage the BBS that starts from the coarse partition and refines each hyperrectangle when $\hat{h}_{ui} > 0$ until $\hat{h}_{ui} \leq 0$ or $t_i \geq t_{\text{gap}}$ where $t_{\text{gap}} > 0$ is the minimum partition size, see Fig. 2. If condition (8) holds for all hyperrectangles $B(x; i)$, then the condition (4a) holds in the continuous state space.

Although verifying condition (4a) is simple, verifying condition (4b) $\sup_{u \in U_a} [L_f \hat{h}_w(x) + L_g \hat{h}_w(x)u]$

$\hat{h}_w(x); \forall x \in X_a$ is challenging. For improved readability, we denote $q(x) = \sup_{u \in U_a} [L_f \hat{h}_w(x) + L_g \hat{h}_w(x)u + \hat{h}_w(x)]$. Hence, verifying condition (4b) is equivalent to verifying $q(x) \leq 0; \forall x \in X_a$. Let q_i denote a lower bound of $q(x)$ for $x \in B(x_i; i)$. Then the following condition has to hold:

$$q(x) \leq q_i \leq 0; \quad x \in B(x_i; i); \quad B(x_i; i) \in B_{X_a}; \quad (9)$$

Furthermore, we can approximate the nonlinear dynamic functions f and g using Taylor Models as done in [25] or sampling:

$$x_{ji} \leq f(x) + g(x)u \leq x_{ui}; \quad x \in B(x_i; i); \quad (13)$$

In (10a), the objective function depends on the variable x and is constrained within the feasible region $B(x_i; i)$. We simplify (10a) by considering three independent variables subject to independent constraints. This results in

$$q_i^0 = \min_{h, J; x} q^0(h; J; x) = hJ; xi + h \quad (14a)$$

$$\text{st: } \hat{h}_{li} \leq h \leq \hat{h}_{ui}; \quad (14b)$$

$$J_{li} \leq J \leq J_{ui}; \quad (14c)$$

$$x_{ji} \leq x \leq x_{ui}; \quad (14d)$$

where x denotes the value of $f(x) + g(x)u$, h denotes the value of $\hat{h}_w(x)$ and J denotes the value of $J_{\hat{h}_w(x)}$. When (11), (12), and (13) are over-approximations of the true intervals, it is clear that the optimal solution q_i^0 from (14a) is an over-approximation of the optimal solution q_i from (10a), which means $q_i^0 \geq q_i$. To efficiently solve (14a), we may compute the optimal solution independently for each term, taking the minimum over the set of vertices.

Although the theoretical complexity of the BBV is still exponential in the dimension of the state space, it improves the scalability in practice. One must note that our method

is a sound verification method instead of a complete one, which means the failure to obtain $B_{x=x_a}$ and B_{x_a} that satisfy condition (8), (9) does not imply the invalidation of the nCBF, as we over-approximate the conditions. We want to emphasize that the chosen over-approximation method, CROWN, has been the winning strategy at the Verification of Neural Networks Competition for multiple years [26].

C. Branch and Bound Verification-in-the-Loop Training

Although the BBV provides a practical way to certify the NN as nCBF, it requires several training and verification processes until an nCBF is obtained. Therefore, leveraging the information from the verification and ensuring the satisfaction of conditions (4a) and (4b) becomes the task of BBVT. This type of method is also known as CEGIS [27]. See Fig. 1 for an overview of the framework.

We start with the initial fixed training dataset D that contains a number of uniformly sampled points. During the training procedure, we optimize the NN to decrease the loss in (6) using D . After k epochs, the verifier starts with a coarse partition of the state space. The upper bound $B(x_i; i) \supseteq B_{x=x_a}$ and lower bound $B(x_i; i) \subseteq B_{x_a}$ are computed. The hyperrectangles, whose $q_i^0 = 0$ or $q_i^1 = 0$, are split until t_{gap} . After reaching the minimum partition size t_{gap} , the hyperrectangles whose $q_i^0 = 0$ or $q_i^1 = 0$ are treated as the violation areas. The center points are added to the CE dataset and the training procedure is repeated until the verifier returns satisfaction or the maximum number of iterations n_{max} is reached.

Note that although the universal approximation theorem in [28] guarantees the existence of $f(x)$ to be an nCBF that renders maximum forward invariant set, this is under the assumption that the NN has a sufficient number of neurons. The training procedure is not guaranteed to converge to an nCBF, but if the verifier returns satisfaction, the NN is an nCBF for the given system in the continuous state space. It is possible to introduce adversarial training, i.e. training on the worst-case state in a region around each sample to improve the convergence to a verifiable nCBF [8].

V. RESULTS

In this section, we evaluate our proposed framework on two systems: an inverted pendulum and a 2D navigation task. The experimental setup is introduced in Section V-A. In Section V-B and Section V-C, we provide a comprehensive assessment on the inverted pendulum, addressing the verification efficiency and the size of the safe set, respectively. In Section V-D, we consider a 2D navigation task with nonconvex constraints to display the practical use of our framework and combine the nCBF with RL to achieve safe policy learning. Our code is available on GitHub

We consider the following baseline methods:

LST: The Level Set Toolbox (LST) [29] generates a safe value function by Hamilton-Jacobian-Issac Reachability Analysis (HJI-RA) over a discrete grid.

NeuralCLBF: Neural Control Lyapunov Barrier Function (NeuralCLBF) [2] parametrizes the CBF as an NN and optimizes it according to their empirical loss based on (2) and a nominal safe set.

SMT: [6] trains a neural Lyapunov function with SMT generating CEs and ensures the validation of the result. The constraints considered by SMT are conditions (4a) and (4b). To have a fair comparison, the training loss is chosen to be the same as in (6).

A. Experimental Setup

1) Inverted Pendulum: Let $s = [\ ; \] \in \mathbb{R}^2$ be the state variable and $u \in \mathbb{R}$ be the control input. We consider the state space $\mathcal{X} = \{s : s \in [-2, 2] \times [-5, 5]g\}$ and the input space $\mathcal{U} = \{u : u \in [-12, 12]g\}$. The dynamics of the inverted pendulum are given by:

$$\dot{s} = \begin{bmatrix} s_2 \\ -\frac{3g}{2l} \sin(s_1) - \frac{3}{ml^2} s_2 + \frac{3}{ml^2} u \end{bmatrix} \quad (15)$$

where $m = 1$, $b = 0.1$, $g = 9.81$, and $l = 1$. The state admissible set is $\mathcal{X}_a = \{s : s \in [-\frac{5}{6}, \frac{5}{6}] \times [-2, 4]g\}$ and the input admissible set is $\mathcal{U}_a = \mathcal{U}$, see Fig. 3.

2) 2D navigation task: We consider a 2D navigation task in which a point robot should reach a goal position while avoiding obstacles, see Fig. 6a. Let $[x; y; \dot{x}; \dot{y}] \in \mathbb{R}^4$ be the state variable and $a = [a_x; a_y] \in \mathbb{R}^2$ be the control input representing the acceleration along the x-axis and y-axis. The dynamics of the point robot are:

$$\begin{bmatrix} \dot{x} \\ \dot{y} \\ \ddot{x} \\ \ddot{y} \end{bmatrix} = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} s + \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} a_x \\ a_y \end{bmatrix} \quad (16)$$

We consider the admissible position set $\mathcal{X}^1 = \{s : x \in [0, 4]; y \in [0, 4]g\}$ except the obstacle set $\mathcal{X}^2 = \{s : x \in [1.5, 2.5]; y \in [0, 2]g\}$, together with velocity constraints $\mathcal{X}^3 = \{s : \dot{x} \in [-1, 1]; \dot{y} \in [-1, 1]g\}$. Thus, the state admissible set is $\mathcal{X}_a = \mathcal{X}^1 \setminus (\mathcal{X}^2 \cup \mathcal{X}^3)$, where $(\)^c$ represents the complement of a set. See Fig. 6 for a pictorial representation of the set. The input admissible set is $\mathcal{U}_a = \{a : a_x \in [-1, 1]; a_y \in [-1, 1]g\}$.

¹<https://github.com/tud-amr/ncbf-simultaneous-synthesis-and-verification>

Fig. 3: The workspace of the considered inverted pendulum.

(a) Training without verification-in-the-loop (b) Distribution of counterexamples (c) Training with verification-in-the-loop

Fig. 4: Shapes of 0-superlevel sets of NNs trained with and without BBVT for the inverted pendulum. In Fig. 4a the NN is trained with a xed dataset and evaluated on a denser testing dataset to showcase that condition (4b) is not satisfied for the continuous state space. Figure 4b shows the CEs added to the dataset according to BBVT. Figure 4c showcases that, after training the NN with BBVT, no validations are detected since the NN is an nCBF.

TABLE I: Hyper-parameter for nCBF Training.

	0.5		0.05
learning rate	10^{-3}	learning rate decay	0.995
verify after every k epochs	20	minimum partition gap t_{gap}	0.005
initial radius r_{init} (inverted pendulum)	[0.2, 0.2]	initial radius r_{init} (2D navigation)	[0.2, 0.2, 0.2, 0.2]
Num. xed points (inverted pendulum)	10^5	Num. xed points (2D navigation)	10^6
n_{max}	100		

different grid gaps, which are 0.2 and 0.05, respectively. It is evident that an increased grid density leads to improved accuracy at the cost of longer computation time. However, a dense grid map is not always possible, since the memory space of LST grows exponentially, which is referred to as the Curse of Dimensionality. With a grid gap of 0.05, LST requires 24.41kB memory space, while we only need to store the parameters of the nCBF, which is 1.2kB. This is important for embedded devices such as the control unit on drones.

3) Training Configuration: For both systems, we train the nCBF using Pytorch on NVIDIA A40, and Stochastic Gradient Descent is used as the optimizer to avoid local minima. The used hyper-parameters can be found in Table I. For the inverted pendulum, we choose a Tanh-based FCNN with one hidden layer which consists of 36 neurons. For the 2D navigation task, a larger Tanh-based FCNN is required since the shape of the environment is more complex. Here we choose a Tanh-based FCNN with two hidden layers, each of which consists of 256 neurons.

B. Verification and Efficiency

In this section, we use the inverted pendulum to discuss the certification of the trained NN as a CBF. To showcase the disadvantage of training without verification, we train the nCBF with a xed data set and stop training after 100 epochs. We then examine the satisfaction of condition (4b) with a denser testing dataset. The 0-superlevel set of the trained NN is shown in blue in Fig. 4a. The orange area indicates the testing data points that violate condition (4b).

We resume the training with the same dataset and use BBVT to augment the training dataset with CEs every epochs until the verifier returns satisfaction. Figure 4b shows the distribution of the CEs after the first verification loop. As we augment the dataset, the verifier returns satisfaction after 240 epochs, see Fig. 4c.

To highlight the efficacy of BBVT, we evaluate the training time, verification time, and the ratio of violation areas for our framework and the baseline methods. The results are shown in Table II. We first compare our method with LST [29]. The table shows the results of LST for two

Then, we compare our method with NeuralCLBF. Due to the lack of a verification process and CE data set, the xed data set for NeuralCLBF contains 10^6 data points in order to have a fair comparison with our method. Since NeuralCLBF learns an nCBF based on a nominal safe set, the training process is assisted by prior knowledge and results in less training time, see Table II. However, there are sparse areas that violate the conditions as discussed in [2] and how these sparse areas grow with the complexity of the system has not been studied yet.

We also compare our method with SMT. However, SMT did not return satisfaction until the maximum number of iterations n_{max} was reached, see Table II. Although there exist some works [5], [6] that use SMT to verify a neural controller, they only use a very simple FCNN with around 9 neurons. In our case, the computation time of SMT grows dramatically since the NN is more complex. Also, SMT can only generate several CEs at each iteration, while BBVT generates all the CEs in state space, which is more efficient than SMT.

C. Size of Safe Set

We will compare the size of the forward invariant set derived using our framework and the baseline methods in this section. Since SMT failed to verify the nCBF and LST with a grid gap of 0.2 has a large violation area, we compare our framework only against LST with a grid gap of 0.05 and NeuralCLBF with the nominal safe set being $\mathcal{H}_n = \{s : s_k < \frac{3}{4}g\}$.

The forward invariant sets derived by the different methods are illustrated in Fig. 5. We see that the size of the forward

TABLE II: Verification and Efficiency Comparison for the inverted pendulum. BBVT is compared against LST, NeuralCLBF, and SMT to synthesize an nCBF. LST and NeuralCLBF do not verify their safe value function, which is represented by '-' in columns 3 and 4. To validate the verification process, we calculate the ratio of points that violate condition (4b) on a uniform grid with a size of $10^3 \times 10^3$ within the state space.

	Stop criteria	Total computation time (s)	Average verification time (s/epoch)	Average generation time (s/per counterexample)	Violation points/test points (%)
LST(0.2)	value converges	5.34	-	-	1.9
LST(0.05)	value converges	104.48	-	-	0.0064
NeuralCLBF	loss converges	584.6	-	-	0.0013
SMT	max # iter reached	5311.68	14.73	1.34	0.7742
BBVT(ours)	verified	1214.15	16.20	0.004	0.0

Fig. 5: The forward invariant set of safe value functions obtained by different methods for the inverted pendulum.

invariant set from NeuralCLBF is conservative, while our method approximates CBVF and renders a safe set that is close to the maximum forward invariant set. Since the forward invariant set of the CBF is always a subset of that from HJI-RA, which is discussed in [4], it is not surprising that LST renders a larger safe set than ours. We note that $\dot{V} > 0$ in (6a) encourages the satisfaction of the CBF conditions at the expense of rendering a smaller safe set.

D. Application of Neural Control Barrier Functions to Safe Policy Learning

In this section, we use RL to address the 2D navigation task introduced in Section V-A.2. Let $\mathbf{x}_g = [x_g; y_g; 0; 0]$ be the goal state. The step reward is defined as $r = 0.01$ if $\mathbf{x}_g \leq \mathbf{x}_k$, the terminal reward is $r_{\text{collision}} = 5$ when the robot collides with the obstacles and $r_{\text{goal}} = 10$ when the robot reaches the goal area $\mathbf{x}_g = \mathbf{x}_k$ if $\mathbf{x}_g \leq \mathbf{x}_k \leq \mathbf{x}_g + \mathbf{g}$ where $\mathbf{g} = 0.1$ is the goal tolerance. We use Proximal Policy Optimization [30] to train the agent and solve

$$u_{\text{safe}} = \arg \min_{u \in \mathcal{U}_a} \|u - u_{\text{RL}}\|_2^2 \quad (17)$$

$$\text{s.t. } L_f \hat{H}_w(x) + L_g \hat{H}_w(x)u + (\hat{H}_w(x)) \leq 0$$

to project the action u_{RL} of the RL policy to the safe action u_{safe} with the least modification.

We note that, theoretically, this controller guarantees safety with infinite control frequency. However, a continuous controller is not possible to implement on discrete control units. This limits the safety guarantees we may provide. How to address the gap between continuous controllers and their discrete implementations remains an open question.

Figure 6a shows all trajectories performed during the training, and we can see that several trajectories collide with the obstacles, while there are no unsafe trajectories in Fig. 6b with nCBF as a safety filter. However, we observe that the average reward with nCBF is larger than for nominal RL without nCBF in the very early stage but has a slower growth rate and converges to a lower reward level compared with nominal RL, see Fig. 6c. The reason is that nCBF provides prior knowledge about the environment and the agent could avoid exploring unsafe regions in the early stage and gain a higher reward than nominal RL. However, the forward invariant set is still suboptimal as discussed in Section V-C, which means only a suboptimal policy is learned and exploration is restricted. Nevertheless, we believe that provided safety guarantees are beneficial in safety-critical applications.

VI. CONCLUSION

In this work, we presented a framework that simultaneously synthesizes and verifies continuous Neural Control Barrier Functions (nCBFs). To this end, we leveraged bound propagation techniques and the Branch-and-Bound scheme to efficiently verify neural networks as Control Barrier Functions in the continuous state space. In experiments, we showed that our framework efficiently synthesizes an nCBF which renders a larger safe set than state-of-the-art methods without requiring an initial guess.

Since the memory requirements and computation time of the Branch-and-Bound Verification scheme increase exponentially with the system dimension, in future work, we may address the scalability of our framework.

REFERENCES

- [1] 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.
- [2] C. Dawson, Z. Qin, S. Gao, and C. Fan, "Safe nonlinear control using robust neural lyapunov-barrier functions," *Conference on Robot Learning* pp. 1724–1735, PMLR, 2022.
- [3] B. Dai, P. Krishnamurthy, and F. Khorrami, "Learning a better control barrier function," in *2022 IEEE 61st Conference on Decision and Control (CDC)* IEEE, 2022.
- [4] J. J. Choi, D. Lee, K. Sreenath, C. J. Tomlin, and S. L. Herbert, "Robust control barrier-value functions for safety-critical control," in *60th IEEE Conference on Decision and Control (CDC)* IEEE, 2021.
- [5] A. Peruffo, D. Ahmed, and A. Abate, "Automated and formal synthesis of neural barrier certificates for dynamical models," *Tools and Algorithms for the Construction and Analysis of Systems: 27th International Conference, TACAS 2023* Springer, 2021.
- [6] Y.-C. Chang, N. Roohi, and S. Gao, "Neural lyapunov control," in *Advances in Neural Information Processing Systems* vol. 32, Curran Associates, Inc., 2019.

