



BIROn - Birkbeck Institutional Research Online

Zhao, H. and Zeng, X. and Chen, Taolue and Liu, Z. and Woodcock, J. (2021) Learning safe neural network controllers with barrier certificates. *Formal Aspects of Computing* 33 (3), pp. 437-455. ISSN 0934-5043.

Downloaded from: <https://eprints.bbk.ac.uk/id/eprint/48301/>

Usage Guidelines:

Please refer to usage guidelines at <https://eprints.bbk.ac.uk/policies.html> or alternatively contact lib-eprints@bbk.ac.uk.

Learning Safe Neural Network Controllers with Barrier Certificates

Hengjun Zhao¹, Xia Zeng¹, Taolue Chen², Zhiming Liu¹ and Jim Woodcock^{1,3}

¹School of Computer and Information Science, Southwest University, Chongqing, China

²Department of Computer Science, University of Surrey, UK

³Department of Computer Science, University of York, UK

Abstract. We provide a new approach to synthesize controllers for nonlinear continuous dynamical systems with control against safety properties. The controllers are based on neural networks (NNs). To certify the safety property we utilize barrier functions, which are represented by NNs as well. We train the controller-NN and barrier-NN simultaneously, achieving a verification-in-the-loop synthesis. We provide a prototype tool `nncontroller` with a number of case studies. The experiment results confirm the feasibility and efficacy of our approach.

Keywords: Continuous dynamical systems; Controller synthesis; Neural networks; Safety verification; Barrier certificates

1. Introduction

Controller design and synthesis is one of the most fundamental problems in control theory. In recent years, especially with the boom of deep learning, there has been considerable research activities in the use of neural networks (NNs) for control of nonlinear systems [LHP⁺16, DCH⁺16]. NNs feature the versatile representational ability of nonlinear maps and fast computation, making them an ideal candidate for sophisticated control tasks [PEY01]. Typical examples include self-driving cars, drones, and smart cities. It is noteworthy that many of these applications are safety-critical systems, where safety refers to, in a basic form, that the system cannot reach a dangerous or unwanted state. For control systems in a multitude of Cyber-Physical-System domains, designing *safe* controllers which can guarantee safety behaviors of the controlled systems is of paramount importance [BT SK17, RBK18, DJST18a, RAA19, TSYA19, COMB19, CCTS20, YFS20, ICW⁺20, TYML⁺20].

Typically, when a controller is given, formal verification is required to certify its safety. Our previous work [ZZCL20] has dealt with the *verification* of continuous dynamical systems by the aid of neural networks. In a nutshell, we follow a deductive verification methodology therein by synthesizing a barrier function, the existence of which suffices to show the safety of the controlled dynamical system. The crux was to use neural

networks to represent the barrier functions, spurred by the well-known universal approximation theorem [LLPS93] which assures the expressibility of NNs.

It is imperative to realize that verification or certification of an existing controller does not lend itself to effective and efficient *construction* of controllers, which is the main focus of the current work. Following a correctness-by-design methodology, we aim to synthesize controllers which can guarantee that the controlled system is safe. This question is considerably more challenging and perhaps more interesting from a system engineering perspective. To this end we adopt a data-driven approach for the design of controllers which are to be represented as an NN. A key issue of controller synthesis is to provide a formal guarantee of the quality for the obtained controller, of which safety is arguably the most fundamental. A common practice is to first come up with a controller and then to verify it against desired properties. An interesting innovation of our work is, however, to integrate the synthesis and verification in a unified, data-driven framework, which is enabled by our earlier work by using NNs as a certification mechanism. At a high level, our approach for the controller synthesis will produce two neural networks simultaneously, i.e., one is used to represent the controller (henceforth referred to as controller-NN), and the other is used to represent the barrier function (henceforth referred to as barrier-NN). The synergy of the two NNs, supported by an additional verification procedure to make sure the learned barrier-NN is indeed a barrier certificate, provides the desired safety guarantee for the synthesized controller.

Our method follows a data-driven framework in the sense that both NNs are trained from datasets. For that purpose, we generate training sets and propose specifically designed loss functions which are the key towards the application of standard learning algorithms for NNs. In terms of the learned NN controllers, we find that they usually respect safety constraints, but may exhibit poor performance in terms of, e.g., stability. To further improve the synthesized controllers, we propose a number of approaches such as imposing a larger safety region, stability-aware loss functions, and bounded control inputs (via the Hardtanh activation function).

In general, the advantages of our approach are threefold: (1) the approach is data-driven, requiring considerably less control theory expertise; (2) the approach can support non-linear control systems and safety properties, owing to the representation power of neural networks; and (3) the approach can achieve verification-in-the-loop synthesis, owing to the co-synthesis of controller and barrier functions, which can be seamlessly integrated to provide a correctness-by-design controller as well as its certification.

The main contributions of the paper are summarized as follows:

- We put forward a learning-based framework to synthesize controllers as well as the associated safety certification. This is largely a data-driven approach, with little prior knowledge required, and enjoys great flexibility to effectively handle nonlinear (beyond polynomial) dynamics of ODEs.
- We instantiate the framework by using new class of activation functions. Moreover, we demonstrate how to generate training set, and to construct loss functions of neural networks. We also provide practical methods to formally verify the learnt barrier certificates represented as neural networks.
- We carry out proof-of-concept case studies to showcase the efficacy of the approach.

1.1. Related Work

Our work on learning and verifying NN controllers with barrier certificates is closely related to two categories of research, i.e. *safety critical control by machine learning* and *formal verification of neural networks*. Note that the discussions below are necessarily non-exhaustive as a reasonably detailed discussion requires an independent survey.

Safety Critical Control by Machine Learning. Research work in this category has been emerging in the past years. They differ in: (1) the overall learning framework, e.g., reinforcement learning (RL) or supervised learning; (2) the component to be learned (especially by NN), e.g., the system model, the feedback control policy, or the safety certificate; (3) the type of safety certificate, e.g., control Lyapunov function (CLF) or control barrier function (CBF) [ACE⁺19]. A verification-in-the-loop RL algorithm was proposed in [DKYP19] to learn safe NN controllers for known system dynamics using CBFs; an end-to-end safe RL architecture was developed by combining model-free RL control, model-based CBF control, and model learning in [COMB19]; CLFs and CBFs are integrated into the episodic learning framework and RL framework with an emphasis on model uncertainties in [TDL⁺19, TSYA19, CCTS20]; CBFs are integrated with imitation

learning to train safe NN controllers in [YFS20]. For all the above work, CLFs or CBFs are assumed to be given, at least in a parametric form. For CLFs or CBFs synthesis, a demonstrator-learner-verifier framework was proposed in [RS19] to learn polynomial CLFs for polynomial nonlinear dynamical systems; a special type of neural network was designed in [RBK18] as candidates for learning Lyapunov functions; a supervised learning approach was proposed in [CRG19] to learn neural network Lyapunov functions and linear control policies; data-driven model predictive control (MPC) exploiting neural Lyapunov function and neural network dynamics model was proposed in [DJST18a, MGQ⁺20]. For multi-agent systems, barrier functions have recently been applied for safe policy synthesis on POMDP models [ASBA19]. The computer science community has dealt with the issue of safe controller learning in different ways. For example, a proof-based approach was proposed in [FP18] towards safe RL; a synthesis framework capable of synthesizing deterministic programs from neural network policies was proposed in [ZXMJ19] which enables application of formal verification techniques for traditional software systems can be applied. Compared with the above work, our approach has the following distinguished features:

- the controller and safety certificate are both represented and learned by NNs of general structure; no prior knowledge or initial guess is required;
- training data generation is based on state space sampling, and therefore trajectory simulation is not needed;

Formal Verification of Neural Networks. This has attracted considerable research efforts in recent years, and the general problem is NP-hard [KBD⁺17]. A large body of research focuses on the robustness issue of neural networks. In particular, given an input subject to (adversarial) perturbations, one intends to determine whether the output of the neural network (e.g., the classification result) is invariant to these perturbations. Essentially, this is to estimate the output range of a given neural network on a compact set. There are now a wide range of methods including constraint-solving based approaches [KBD⁺17], optimization based approaches [DJST18b, WZC⁺18, XTJ18], abstract interpretation based approaches [PT10, LLY⁺19], etc. The underlying techniques have also been adopted in the study of continuous or hybrid systems [SL20, RS10, RS07]. By combing the verification of neural networks and continuous dynamical systems, work has been done recently for verification of control systems with neural network feedback components [DCS19, IWA⁺19, SKS19, DFG⁺19, TYML⁺20]. The main technique is reachability analysis of the closed-loop system, either by finite-state abstraction [SKS19], or by reachable set approximation based on interval or other abstract domain [DCS19, IWA⁺19, TYML⁺20]. Usually the reachable set computation can only verify safety up to a finite time horizon, and the approximation error of reachable set may explode. In contrast, we adopt a deductive approach based on barrier certificate, following and improving the line of work in [TKID18].

1.2. Outline

The rest of this paper is organized as follows: some preliminary knowledge is provided in Section 2 for self-containedness; the main steps of our approach is presented in Section 3 with a running example for demonstration; various improvements of the synthesized controllers are discussed in Section 4; implementation and experiment details are given in Section 5; the paper is concluded by Section 6. We note that a preliminary version is accepted by SETTA 2020 as a short paper under the same title [ZZC⁺20].

2. Preliminaries

Throughout this paper, \mathbb{R} denotes the set of real numbers. For any natural number n , let $[n] = \{1, \dots, n\}$.

2.1. Constrained Continuous Dynamical System

A continuous dynamical system is modeled by a system of first-order ordinary differential equations (ODEs) $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x})$, where

- $\mathbf{x} = (x_1, x_2, \dots, x_n)^T \in \mathbb{R}^n$ is a column vector, $\dot{\mathbf{x}}$ denotes the derivative of \mathbf{x} with respect to the time variable t , and
- $\mathbf{f}(\mathbf{x}) : \Omega \rightarrow \mathbb{R}^n$ is a vector field $\mathbf{f}(\mathbf{x}) = (f_1(\mathbf{x}), \dots, f_n(\mathbf{x}))^T$ defined on an open subset $\Omega \subseteq \mathbb{R}^n$.

We assume that \mathbf{f} satisfies the *local Lipschitz condition*, which ensures that, given $\mathbf{x} = \mathbf{x}_0 \in \Omega$, there exists a time $\mathcal{T} > 0$ and a unique time trajectory $\mathbf{x}(t) : [0, \mathcal{T}] \rightarrow \mathbb{R}^n$ such that $\frac{d(\mathbf{x}(t))}{dt} = \mathbf{f}(\mathbf{x}(t))$ for any $t \in [0, \mathcal{T}]$ and $\mathbf{x}(0) = \mathbf{x}_0$. In the sequel, the trajectory is denoted by $\mathbf{x}(t, \mathbf{x}_0)$.

A *constrained continuous dynamical systems* (CCDS) is represented by $\Gamma = (\mathbf{f}, X_D, X_I, X_U)$, where

- $\mathbf{f} : \Omega \rightarrow \mathbb{R}^n$ is the vector field,
- $X_D \subseteq \Omega$ is an evolution constraint (or system domain),
- $X_I \subseteq X_D$ is the set of initial states, and
- $X_U \subseteq X_D$ is the set of unsafe states.

For CCDSs, the following problem is widely investigated in safety critical applications.

Definition 2.1 (Safety Verification). A CCDS $\Gamma = (\mathbf{f}, X_D, X_I, X_U)$ is safe if for all $\mathbf{x}_0 \in X_I$, there does not exist $t_1 > 0$ such that

$$\mathbf{x}(t_1, \mathbf{x}_0) \in X_U \quad \text{and} \quad \forall t \in [0, t_1]. \mathbf{x}(t, \mathbf{x}_0) \in X_D,$$

that is, the system's trajectory never reaches X_U from X_I as long as it remains in X_D .

Remark 2.1. According to Definition 2.1, the case that the system's trajectory from X_I can first leave X_D and then enter X_U does not affect the safety property of the system.

2.2. Controlled CCDS

In this paper, we consider a *controlled CCDS* $\Gamma = (\mathbf{f}, X_D, X_I, X_U)$ with continuous dynamics defined by

$$\begin{cases} \dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{u}) \\ \mathbf{u} = \mathbf{g}(\mathbf{x}) \end{cases}, \quad (1)$$

where $\mathbf{x} \in X_D \subseteq \mathbb{R}^n$ is the system state, $\mathbf{u} \in U \subseteq \mathbb{R}^m$ is the control input, and $\mathbf{f} : X_D \times U \rightarrow \mathbb{R}^n$ and $\mathbf{g} : X_D \rightarrow U$ are the locally Lipschitz continuous vector field and feedback controller function, respectively. The problem we considered in this paper is defined as follows.

Definition 2.2 (Safe Controller Synthesis). Given a controlled CCDS $\Gamma = (\mathbf{f}, X_D, X_I, X_U)$ with \mathbf{f} defined by (1), design a locally Lipschitz continuous feedback control law \mathbf{g} such that the closed-loop system Γ with $\mathbf{f} = \mathbf{f}(\mathbf{x}, \mathbf{g}(\mathbf{x}))$ is safe as per Definition 2.1.

2.3. Barrier Certificate

Given a system Γ , a barrier certificate is a real-valued function $B(\mathbf{x})$ over the states of the system satisfying the condition that $B(\mathbf{x}) \leq 0$ for any reachable state \mathbf{x} and $B(\mathbf{x}) > 0$ for any state in the unsafe set X_U . If such a function $B(\mathbf{x})$ exists, one can easily deduce that the system can *not* reach a state in the unsafe set from the initial set [PJP07]. In this paper, we will certify the safety of a synthesized controller by generating barrier certificates.

There are several different formulations of barrier certificates without explicit reference to the solutions of the ODEs [PJP07, KHS⁺13, DGXZ17, SGT18]. We will adopt what are called *strict barrier certificate* [SPW12] conditions.

Theorem 2.1 (Strict barrier certificate). Given a system $\Gamma = (\mathbf{f}, X_D, X_I, X_U)$, if there exists a continuously differentiable function $B : X_D \rightarrow \mathbb{R}$ s.t.

1. $B(\mathbf{x}) \leq 0$ for $\forall \mathbf{x} \in X_I$
2. $B(\mathbf{x}) > 0$ for $\forall \mathbf{x} \in X_U$
3. $\mathcal{L}_{\mathbf{f}}B(\mathbf{x}) < 0$ for $\forall \mathbf{x} \in \{\mathbf{x} \in X_D \mid B(\mathbf{x}) = 0\}$,




Fig. 1. The structure of a multilayer feed-forward artificial neural network

then the system Γ is safe, and such B is a barrier certificate.

Note that in the above third condition, $\mathcal{L}_{\mathbf{f}}B$ is the *Lie derivative* of B w.r.t. \mathbf{f} , that is, the inner product of \mathbf{f} and the gradient of B :

$$\mathcal{L}_{\mathbf{f}}B(\mathbf{x}) = (\nabla B) \cdot \mathbf{f}(\mathbf{x}) = \sum_{i=1}^n \left(\frac{\partial B(\mathbf{x})}{\partial x_i} \cdot f_i(\mathbf{x}) \right). \quad (2)$$

Proof. We prove the theorem by contradiction. Suppose that the theorem does not hold, that is, there exists a function B satisfying the three conditions in the premise but Γ is not safe. Then by Definition 2.1 there exists $\mathbf{x}_0 \in X_I$ and $t_1 > 0$ such that

$$\mathbf{x}(t_1, \mathbf{x}_0) \in X_U \quad \text{and} \quad \forall t \in [0, t_1]. \mathbf{x}(t, \mathbf{x}_0) \in X_D. \quad (3)$$

Denote $\mathbf{x}(t, \mathbf{x}_0)$ by $\mathbf{x}(t)$ for short for any $t \geq 0$. Then by Condition 1 and 2 of Theorem 2.1, we have $B(\mathbf{x}(0)) \leq 0$ and $B(\mathbf{x}(t_1)) > 0$. Noting that both B and $\mathbf{x}(t)$ are continuous, it follows from the Intermediate Value Theorem that there exists $t_2 \in [0, t_1]$ s.t. $B(\mathbf{x}(t_2)) = 0$. Let $\mathcal{Z} = \{t \in [0, t_1] \mid B(\mathbf{x}(t)) = 0\}$. Then it is obvious that \mathcal{Z} is a nonempty bounded set since $t_2 \in \mathcal{Z}$. By the Completeness of Reals, \mathcal{Z} has a supremum, denoted by $t_{\text{sup}} = \sup \mathcal{Z}$. Again by the continuity of $B(\mathbf{x}(t))$, it is not difficult to show that t_{sup} has the following properties: (i) $t_{\text{sup}} \in [0, t_1]$; (ii) $B(\mathbf{x}(t_{\text{sup}})) = 0$; (iii) $B(\mathbf{x}(t)) > 0, \forall t \in (t_{\text{sup}}, t_1]$. Then by (ii) and (iii) the right-hand derivative of $B(\mathbf{x}(t))$ at t_{sup} , i.e. $\lim_{h \rightarrow 0^+} \frac{B(\mathbf{x}(t_{\text{sup}}+h)) - B(\mathbf{x}(t_{\text{sup}}))}{h}$, is non-negative. However, according to (i), (ii), (3) and Condition 3 of Theorem 2.1, we have that the derivative of $B(\mathbf{x}(t))$ at t_{sup} equals $\mathcal{L}_{\mathbf{f}}B(\mathbf{x}(t_{\text{sup}}))$ and thus is negative, which is a contradiction. Therefore the theorem holds. \square

Remark 2.2. The converse of Theorem 2.1 does not hold in general. However, under moderate assumptions, various converse theorems have been established [WS16, Rat18], which suggests that practically a barrier certificate does exist for a broad class of safe controlled CCDSs. Thus the crux in applying Theorem 2.1 is to find the barrier certificate effectively.

2.4. Neural Networks

In this paper, both the synthesized control law \mathbf{g} and the barrier certificate B are represented by (feed-forward) neural networks (NNs). We introduce some basic notions here. A typical NN consists of a number of interconnected neurons which are organized in a layered structure. Each neuron is a single processing element that responds to the weighted inputs received from other neurons (cf. Fig. 1.)

In general, an NN represents a function $\mathcal{N}(\mathbf{x})$ on the input \mathbf{x} and can be represented as a composition of its layers. We normally reserve 0 and L for the indices of the input and the output layer respectively, and all of the other layers in between are hidden layers. In this paper, we use superscripts to index layer-specific variables. In particular, the layer l comprises neurons $n_i^{(l)}$ for $i \in [d^{(l)}]$, where $d^{(l)}$ is the dimension of the layer l . Neuron $n_j^{(l-1)}$ of the layer $l-1$ is connected with neuron $n_i^{(l)}$ of layer l by a directed edge with weight $w_{ij}^{(l)} \in \mathbb{R}$. Each neuron $n_i^{(l)}$ of layer $l \in [L]$ is associated with a *bias* $b_i^{(l)} \in \mathbb{R}$ and an activation function

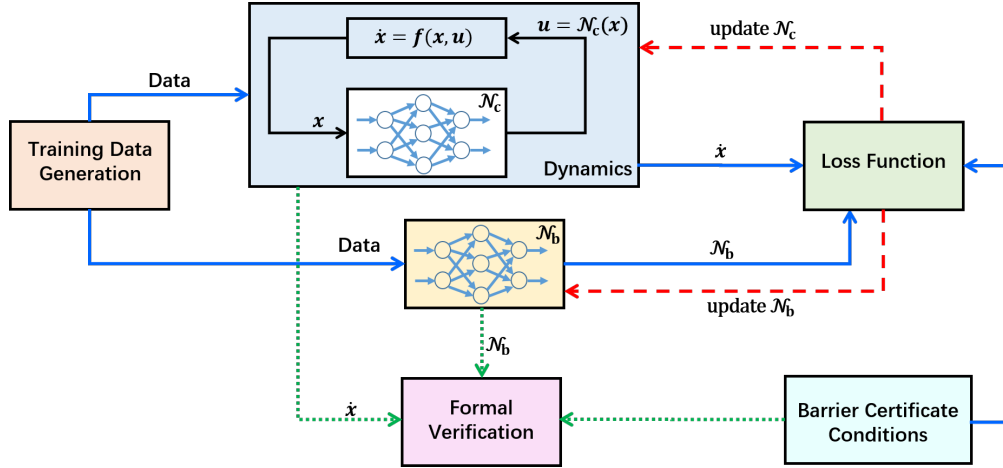


Fig. 2. The framework of safe neural network controller synthesis

$a_i^{(l)} : \mathbb{R} \rightarrow \mathbb{R}$. Usually the neurons in the same layer has identical activation functions, denoted by $a^{(l)}$. Commonly used activation functions include ReLU (rectified linear unit, i.e., $\max(0, x)$ for $x \in \mathbb{R}$), sigmoid, hyperbolic tangent, etc.

Denote the input vector to the NN by $\mathbf{x} \in \mathbb{R}^{d^{(0)}}$. Let the output vector of the l -th layer be $\mathbf{x}^{(l)}$. Then $\mathbf{x}^{(0)} = \mathbf{x}$. We introduce the vector variable $\mathbf{z}^{(l)}$ to denote the input vector to the l -th layer for $l \in [L]$. Thus the *forward propagation equations* of an NN can be defined as

$$\begin{cases} \mathbf{x}^{(0)} &= \mathbf{x} \\ \mathbf{z}^{(l)} &= \mathbf{W}^{(l)} \cdot \mathbf{x}^{(l-1)} + \mathbf{b}^{(l)} \quad \text{for } l \in [L] \\ \mathbf{x}^{(l)} &= a^{(l)}(\mathbf{z}^{(l)}) \quad \text{for } l \in [L] \\ \mathbf{y} &= \mathcal{N}(\mathbf{x}) = \mathbf{x}^{(L)} \end{cases}, \quad (4)$$

where $\mathbf{W}^{(l)}$ is a matrix of dimension $d^{(l)} \times d^{(l-1)}$, $\mathbf{b}^{(l)}$ is a $d^{(l)}$ -dimensional column vector, and $a^{(l)}$ is taken as an element-wise function for a vector input.

Training of NNs is usually performed through *backward propagation*, during which the parameters \mathbf{W} 's and \mathbf{b} 's are updated through an optimization algorithm (e.g., stochastic gradient descent, SGD for short) applied on the training set [GBC16].

3. Methodology

The framework of our safe controller learning approach is demonstrated in Fig. 2. Given a controlled CCDS $\Gamma = (\mathbf{f}, X_D, X_I, X_U)$, the basic idea of the proposed approach is to represent the controller function \mathbf{g} as well as the safety certificate function B by two NNs, i.e. \mathcal{N}_c and \mathcal{N}_b respectively. Then we formulate the barrier certificate conditions as per Theorem 2.1 w.r.t. \mathcal{N}_b and the closed-loop dynamics $\mathbf{f}(\mathbf{x}, \mathcal{N}_c(\mathbf{x}))$ into a loss function, and then train the two NNs simultaneously on a generated training data set until the loss is reduced to 0. The resulting two NNs are the controller and barrier certificate candidates, which satisfy the conditions of Theorem 2.1 on the sampled data set. To overcome the limitation of data-driven approaches, i.e., the generalization issue of the learned NNs on non-sampled data, formal verification (by SMT solvers in this paper) is performed on the synthesized candidates to show that the barrier certificate conditions are indeed satisfied. The blue (solid), red (dashed), and green (dotted) arrows in Fig. 2 show the information flow of forward propagation, backward propagation, and formal verification, respectively.

Next, before giving more detailed steps of our approach, we first introduce a running example.

Example 3.1 (Dubins' Car [TKID18, DKYP19]). The control objective is to steer a car with constant velocity 1 to track a path, here the X -axis in the positive direction. The states of the car are the x, y position and the driving direction θ , which can be transformed to the distance error d_e and angle error θ_e between




Fig. 3. States of Dubins' car: $d_e = y$, $\theta_e = \frac{\pi}{2} - \theta$




Fig. 4. Simulated car trajectories with learned NN controller

the current position and the target path (cf. Fig. 3). The controlled CCDS $\Gamma = (\mathbf{f}, X_D, X_I, X_U)$ is:

$$\mathbf{f} : \begin{bmatrix} \dot{d}_e \\ \dot{\theta}_e \end{bmatrix} = \begin{bmatrix} \sin(\theta_e) \\ -u \end{bmatrix}, \quad \text{where } u \text{ is the scalar control input}$$

- X_D : $\{(d_e, \theta_e) \in \mathbb{R}^2 \mid -6 \leq d_e \leq 6, -7\pi/10 \leq \theta_e \leq 7\pi/10\}$;
- X_I : $\{(d_e, \theta_e) \in \mathbb{R}^2 \mid -1 \leq d_e \leq 1, -\pi/16 \leq \theta_e \leq \pi/16\}$;
- X_U : the complement of $\{(d_e, \theta_e) \in \mathbb{R}^2 \mid -5 \leq d_e \leq 5, -\pi/2 \leq \theta_e \leq \pi/2\}$ in X_D .

Figure 4 shows 50 simulated trajectories on the x - y plane from random initial states in X_I using our learned NN controller u . The two red horizontal lines are the safety upper and lower bounds (± 5) for y (the same bounds as d_e). In the rest of this paper, we will use Example 3.1 to demonstrate our safe controller synthesis approach.

3.1. The Structure of \mathcal{N}_c and \mathcal{N}_b

We first fix the structure of \mathcal{N}_c and \mathcal{N}_b as follows, assuming that in the controlled CCDS Γ , \mathbf{x} and \mathbf{u} are of n and m dimension respectively, e.g. $n = 2, m = 1$ for Example 3.1.

- **Input layer** has n neurons for both \mathcal{N}_c and \mathcal{N}_b ;
- **Output layer** has m neurons for \mathcal{N}_c and one single neuron for \mathcal{N}_b ;
- **Hidden layer**: there is no restriction on the number of hidden layers or the number of neurons in each hidden layer; for Example 3.1, the structures are fixed such that \mathcal{N}_c has one hidden layer with 5 neurons, and \mathcal{N}_b has one hidden layer with 10 neurons;
- **Activation function**: considering the inherent requirement of local Lipschitz continuity for \mathcal{N}_c and the inherent requirement of differentiability for \mathcal{N}_b , and for ease of formal verification, we adopt ReLU, i.e. $a(x) = \max(0, x)$, and *Bent-ReLU* [ZZCL20], i.e.,

$$a(x) = 0.5 \cdot x + \sqrt{0.25 \cdot x^2 + 0.0001} \quad (5)$$

as activation functions for hidden layers of \mathcal{N}_c and \mathcal{N}_b respectively. (The Lipschitz continuity of ReLU is by [JD20].) The activation function of the output layer is the identity map for both \mathcal{N}_c and \mathcal{N}_b .

Remark 3.1. According to Remark 2.2, provided that the barrier certificate for the considered system exists, the success of our approach relies on choosing NN architectures that are sufficiently expressive for representing the sought controller and barrier functions. The relation between the NN architecture and its approximation ability is a hard theoretical problem and there has been much recent progress. For example, it was shown [Tel17] that for any *rational* function there is a ReLU network of size (number of neurons) $\mathcal{O}(\text{poly} \log(1/\epsilon))$ which is ϵ -close.

3.2. Training Data Generation

In our training algorithm, training data are generated by sampling points from the domain X_D , initial set X_I , and unsafe region X_U of the considered system Γ . No simulation of the continuous dynamics is needed. The simplest sampling method is to grid the super-rectangles bounding X_D , X_I , X_U with a fixed mesh size, and then filter out those points not satisfying the constraints of X_D , X_I , X_U . For example, we generate a mesh with $2^8 \times 2^8$ points from X_D for Example 3.1. The obtained three finite data sets are denoted by S_D , S_I , and S_U .

3.3. Loss Function Encoding

Given S_I , S_U , and S_D , the loss function for training \mathcal{N}_c and \mathcal{N}_b can be expressed as

$$L(S_D, S_I, S_U) = c_1 \cdot \sum_{\mathbf{x} \in S_I} L_1(\mathbf{x}) + c_2 \cdot \sum_{\mathbf{x} \in S_U} L_2(\mathbf{x}) + c_3 \cdot \sum_{\mathbf{x} \in S_D} L_3(\mathbf{x}) \quad (6)$$

with

$$\begin{aligned} L_1(\mathbf{x}) &= \text{ReLU}(\mathcal{N}_b(\mathbf{x}) + \varepsilon_1) && \text{for } \mathbf{x} \in S_I, \\ L_2(\mathbf{x}) &= \text{ReLU}(-\mathcal{N}_b(\mathbf{x}) + \varepsilon_2) && \text{for } \mathbf{x} \in S_U, \\ L_3(\mathbf{x}) &= \text{ReLU}(\mathcal{L}_f \mathcal{N}_b(\mathbf{x}) + \varepsilon_3) && \text{for } \mathbf{x} \in \{\mathbf{x} \in S_D : |\mathcal{N}_b(\mathbf{x})| \leq \varepsilon_4\} \end{aligned} \quad (7)$$

denoting the sub-loss functions encoding the three conditions of Theorem 2.1, and c_1, c_2, c_3 three positive constant weight coefficients for the sub-losses L_1, L_2, L_3 respectively. The basic idea is to impose a positive (resp., zero) penalty to those sampled points that violate (resp., satisfy) barrier certificate conditions. $\varepsilon_1, \varepsilon_2, \varepsilon_3$ in (7) are three small non-negative tolerances, the role of which is to increase the generalizability of the learned NNs, i.e., to enforce zero loss on the non-sampled data points. ε_4 in (7) is a small positive constant characterizing a narrow tube around the zero-level set of \mathcal{N}_b , since it is hard to sample data on the level set exactly. Note that in the above expression L_3 , \mathbf{f} is $\mathbf{f}(\mathbf{x}, \mathcal{N}_c(\mathbf{x}))$.

3.4. The Training Process

We adopt a modified SGD optimization technique for training the two NNs \mathcal{N}_c and \mathcal{N}_b . That is, we partition the training data sets S_D, S_I, S_U into mini-batches and shuffle the list of batches to gain some randomness effect, rather than shuffling the whole training data set. For each mini-batch of data, the loss is calculated according to (6) and the weights and biases of the two NNs are updated by a gradient descent step through backward propagation. To start the training, we must first specify the ε_1 to ε_4 in the loss function, as well as hyper-parameters such as number of restarts n_{restart} , number of epoches n_{epoch} , number of mini-batches n_{batch} , and learning rate l_r , etc. For Example 3.1, we set $n_{\text{restart}} = 5$, $n_{\text{epoch}} = 100$, $n_{\text{batch}} = 4096$ and $l_r = 0.1$. The choices of ε_1 to ε_4 will be presented in the following subsection. The training process terminates when the loss is reduced to 0 on all mini-batches or the number of restarts exceeds n_{restart} .

3.5. Formal Verification

The rigorousness of the NNs resulted from 0 training loss is not guaranteed since our approach is data-driven and the learned NNs may lack generalization property, that is, the three conditions in Theorem 2.1 are not necessarily satisfied by \mathcal{N}_c and \mathcal{N}_b on non-sampled data. Therefore we resort to formal verification to guarantee the correctness our synthesized controllers. To conduct the verification, we replace the occurrences of \mathbf{f} and B in Theorem 2.1 by $\mathbf{f}(\mathbf{x}, \mathcal{N}_c(\mathbf{x}))$ and \mathcal{N}_b , and try to show that the negation of the conjunction of the three conditions, i.e.

$$\begin{aligned} &\exists \mathbf{x}. \mathbf{x} \in X_I \wedge \mathcal{N}_b(\mathbf{x}) > 0 \\ \vee &\exists \mathbf{x}. \mathbf{x} \in X_U \wedge \mathcal{N}_b(\mathbf{x}) \leq 0 \\ \vee &\exists \mathbf{x}. \mathbf{x} \in X_D \wedge \mathcal{N}_b(\mathbf{x}) = 0 \wedge \mathcal{L}_{\mathbf{f}(\mathbf{x}, \mathcal{N}_c(\mathbf{x}))} \mathcal{N}_b(\mathbf{x}) \geq 0 \end{aligned} \quad (8)$$




Fig. 5. Learned and verified NN controller and barrier certificate for Example 3.1: the inner (green) and outer (red) shaded areas are the initial and unsafe regions, black arrows in the white area are the closed-loop vector fields $\mathbf{f}(\mathbf{x}, \mathcal{N}_c(\mathbf{x}))$, and the blue curve surrounding the inner shaded box is the zero-level set of \mathcal{N}_b

is UNSATISFIABLE. Due to the high degree of nonlinearity in \mathbf{f} and \mathcal{N}_b of (8), its satisfiability is resolved by the interval-propagation based, nonlinear SMT solver iSAT3.¹ To speed up the verification process, we compute piecewise linear approximations (with interval error bounds) of Bent-ReLU function and its derivative, and replace their occurrences in \mathcal{N}_b and $\mathcal{L}_f \mathcal{N}_b$ by the linear approximations. In this way, the efficiency and effectiveness of formal verification are relevant to the following three issues:

- The tolerances chosen for loss function encoding in (6) and (7);
- The piece-wise linear approximation error of Bent-ReLU function and its derivative;
- The interval splitting width for iSAT3.

For the third issue, we usually set the minimal splitting width option `--msw` to 0.001 for iSAT3. The first and second issues are addressed in the following two sub-sections.

3.5.1. Pre-training and Fine-tuning

The success of synthesis and formal verification heavily relies on the choices of the four constants ε_1 to ε_4 in (6) and (7). Generally, small tolerances are preferred for faster training, while larger tolerances are preferred for formal verification to compensate for the errors caused by activation function linearization and interval arithmetic computation. In practice, we adopt a pre-training and fine-tuning combination strategy. That is, we start with small positive ε_4 and zero ε_1 to ε_3 to perform the initial training. If the pre-trained NNs failed formal verification, they are iteratively refined by gradually increasing the tolerances. For Example 3.1, the first controller and barrier certificate are synthesized with $\varepsilon_4 = 0.01$ and $\varepsilon_1 = \varepsilon_2 = \varepsilon_3 = 0$, for which the formal verification fails, while the fine-tuned controller and barrier certificate are successfully verified when ε_3 was increased to 0.01 (cf. Fig. 5).

3.5.2. Adding Normalized Lie Derivative in Loss Encoding

Larger tolerances in the loss function (6) and (7) are not always useful for formal verification. To see this, consider checking unsatisfiability of the third condition of (8). Note that $\mathcal{L}_f B = \nabla B \cdot \mathbf{f} = \|\nabla B\| \|\mathbf{f}\| \cos \theta_{\nabla B, \mathbf{f}}$, where $\|\cdot\|$ denotes the Euclidean norm and $\theta_{\nabla B, \mathbf{f}}$ denotes the angle between ∇B and \mathbf{f} . Fig. 6(a) illustrates a situation that a point \mathbf{x} on the zero-level set of a barrier candidate B has negative Lie derivative, as $\theta_{\nabla B, \mathbf{f}}$ is slightly larger than $\frac{\pi}{2}$ at \mathbf{x} . Moreover, it can be concluded that $\mathcal{L}_f B(\mathbf{x}) < -\varepsilon_3$ for very large ε_3 since $\|\mathbf{f}\|$ is large. However, formal verification of the negative Lie derivative condition would be very hard at \mathbf{x} , where the direction of ∇B has a large approximation error due to piecewise linearization. For instance, if the approximated $\nabla B(\mathbf{x})$ ranges from $\underline{\nabla B}$ to $\overline{\nabla B}$, then formal verification becomes impossible since $\theta_{\overline{\nabla B}, \mathbf{f}} < \frac{\pi}{2}$

¹ <https://projects.informatik.uni-freiburg.de/projects/isat3/>




Fig. 6. The sign of normalized Lie derivative is robust to Bent-ReLU linearization errors




Fig. 7. Simulations of Dubins' car from $(-1, -0.19)$ with different NN controllers for comparison of stability performance

which makes the Lie derivative positive. The reason for such a phenomenon is that negative $\mathcal{L}_f B$ does not necessarily force the span angle of ∇B and \mathbf{f} to be large, so the sign of $\mathcal{L}_f B$ is not robust to approximation noises of ∇B . The problem can be resolved by introducing additional sub-loss function specifying normalized Lie derivative into the loss function (6) as follows:

$$L_4(\mathbf{x}) = \text{ReLU}\left(\frac{\mathcal{L}_f \mathcal{N}_b(\mathbf{x})}{\|\nabla \mathcal{N}_b\| \cdot \|\mathbf{f}\|} + \varepsilon_5\right), \quad \text{for } \mathbf{x} \in \{\mathbf{x} \in S_D : |\mathcal{N}_b(\mathbf{x})| \leq \varepsilon_4\} \quad (9)$$

where ε_4 are defined in (7) and ε_5 is a non-negative constant. By (9), if a barrier certificate is synthesized with zero L_4 value and sufficiently large ε_5 , then the angle between $\nabla \mathcal{N}_b$ and \mathbf{f} would be large enough to tolerant approximation errors, which leads to successful verification (cf. Fig. 6(b)).

4. Improvement of the Learned Controllers

The controller synthesized and verified in the last section is guaranteed to be safe. However, it may perform poorly regarding properties such as stability. As an illustration, we simulate the Dubins' car system from initial state $d_e = -1, \theta_e = -0.19$ using the NN controller corresponding to Fig. 5. The changes of d_e and θ_e within 60 time units are shown in Fig. 7 by *-marked dashed (d_e) or solid (θ_e) lines. It is obvious that the car has a large distance error although it is still within safety bounds (± 5). We therefore propose a series of ways to improve the performance of synthesized controllers in this section.




Fig. 8. NN controller learned and verified for Example 3.1 with larger safety margin: $\varepsilon_1 = 0.02$, $\varepsilon_2 = 0.8$, $\varepsilon_3 = 0.01$, $\varepsilon_4 = 0.05$; the inner (green) and outer (red) shaded areas are the initial and unsafe regions, black arrows in the white area are the closed-loop vector fields $\mathbf{f}(\mathbf{x}, \mathcal{N}_c(\mathbf{x}))$, and the blue curve surrounding the inner shaded box is the zero-level set of \mathcal{N}_b

4.1. Larger Safety Margin

The first improvement is to gradually increase the safety margin specified by the ε_2 constant in the loss function (6) and (7) by iterative fine-tuning. For example, when ε_2 is increased to 0.8, an NN controller \mathcal{N}_c and the corresponding barrier certificate \mathcal{N}_b are synthesized and shown in Fig. 8. The simulation performance of \mathcal{N}_c is shown in Fig. 7 by \circ -marked dashed (d_e) or solid (θ_e) lines. It is obvious that the distance error is reduced compared to the controller of Fig. 5.

4.2. Asymptotic Stability

Figure 7 shows that using the NN controller with larger safety margin, the distance error of the Dubins' car stabilizes at a value larger than 0.5, which is not desirable. To further reduce the distance error in the long run, we introduce additional loss terms into the loss function to express asymptotic-stability-like properties. Suppose that \mathbf{x}_o is an expected equilibrium point of the system, that is, $\mathbf{f}(\mathbf{x}_o, \mathcal{N}_c(\mathbf{x}_o)) = \mathbf{0}$. For example, the system in Example 3.1 is expected to stabilize with 0 distance and angle errors and so \mathbf{x}_o is $(0, 0)$. Then we define the sub-loss functions for asymptotic stability as:

$$\begin{aligned} L_5(\mathbf{x}) &= \text{ReLU}(-\|\mathbf{f}(\mathbf{x}, \mathcal{N}_c(\mathbf{x}))\| + \varepsilon_6) & \text{for } \mathbf{x} \in \{\mathbf{x} \in S_D : \|\mathbf{x} - \mathbf{x}_o\| > \varepsilon_7\}, \\ L_6(\mathbf{x}) &= \text{ReLU}(\|\mathbf{f}(\mathbf{x}, \mathcal{N}_c(\mathbf{x}))\| - \varepsilon_8) & \text{for } \mathbf{x} = \mathbf{x}_o \end{aligned} \quad (10)$$

where $\varepsilon_6, \varepsilon_7, \varepsilon_8$ are three small positive constants. The basic idea of L_5, L_6 is to impose such constraints that the closed-loop vector field $\mathbf{f}(\mathbf{x}, \mathcal{N}_c(\mathbf{x}))$ has negligible norm at the asymptotically stable point \mathbf{x}_o , and strictly positive norm outside a neighborhood of \mathbf{x}_o with radius ε_7 . By choosing $\varepsilon_7 = 0.1$, $\varepsilon_6 = 0.05$, $\varepsilon_8 = 0.001$ we obtain a fine-tuned \mathcal{N}_c whose simulation performance is shown in Fig. 7 by \square -marked dashed (d_e) or solid (θ_e) lines, which demonstrate good asymptotic stability property. We also fix $\varepsilon_8 = 0.001, \varepsilon_6 = 0.05$ and compare the performances of \mathcal{N}_c obtained from different ε_7 values. The simulation results are shown in Fig. 9. It can be roughly concluded that decreasing ε_7 will have an effect of increasing the *overshoot* and decreasing the *settling time* of the simulated traces. An intuitive explanation of such effects is that by L_5 , shrinking ε_7 increases $\|\mathbf{f}\|$ near \mathbf{x}_o , and thus trajectories approaches \mathbf{x}_o quickly but may overshoot.

Comparison with LQR Controllers. To further evaluate the performance of synthesized NN controllers, we linearize the Dubins' car system near $\mathbf{x}_o = (0, 0)$ and then compute the classic LQR (linear quadratic regulator [Hes18]) controllers for the linearized system. Preliminary experiment shows that for fixed Q and R matrices in the LQR controller computation, by tuning the values of ε_6 and ε_7 , we can obtain NN controllers with comparable performances to LQR controllers (cf. Fig. 10).

Remark 4.1. NN controllers are in principle much more expressive than linear controllers such as LQR,




Fig. 9. Comparison of NN controllers learned using L_5 and L_6 losses with $\varepsilon_6 = 0.05$, $\varepsilon_8 = 0.001$ for Example 3.1: all simulations are from initial state $(-1, -0.19)$; dashed and solid lines represent d_e and θ_e traces respectively; simulations corresponding to controllers learned with $\varepsilon_7 = 0.3, 0.1, 0.05$ are marked by $*$, \circ , and \square respectively




Fig. 10. Simulation of NN and LQR controllers with initial state $(-1, -0.19)$ for Example 3.1: the NN controller is synthesized with $\varepsilon_6 = \varepsilon_7 = 0.05$, $\varepsilon_8 = 0.001$, and the LQR controller is synthesized with Q the 2-dimensional identity matrix and $R = 1$; dashed and solid lines represent d_e and θ_e traces respectively, and traces simulated with LQR and NN controllers are marked by $*$ and \circ respectively

and so it is interesting to investigate better ways of loss function encoding and controller tuning to synthesize NN controllers superior to linear controllers (e.g., LQR) in future.

4.3. Bounded Control Inputs

In practice, the control input \mathbf{u} to system (1) cannot take arbitrary values (cf. Fig. 11(a)) but are bounded within a compact set U . Therefore it is necessary to consider how to synthesize bounded NN controllers for practical applications. Actually this can be achieved simply by replacing the identity activation function in the output layer of \mathcal{N}_c (cf. Section 3.1) by any activation with bounded range, say hyperbolic tangent function. For ease of formal verification, we adopt a piece-wise linear activation *Hardtanh* for the output




Fig. 11. Plotting of surfaces of unbound or bounded NN controllers for Example 3.1 over X_D

layer of \mathcal{N}_c , that is,

$$a^{(L)}(x) = c \cdot \max(-1, \min(1, x))$$

with c a positive constant, which restricts the output of \mathcal{N}_c to be within $[-c, c]$ for each dimension. For Example 3.1, by choosing $c = 3$ we obtained a bounded NN controller as shown in Fig. 11(b). In our experiment, the Hardtanh activation can either be applied in the pre-training or fine-tuning phase.

5. Implementation and Experiments

Given a controlled CCDS $\Gamma = (\mathbf{f}, X_D, X_I, X_U)$ and generated training data set S_D, S_I, S_U , in the most general form, the loss function we adopted for training safe NN controllers is:

$$L(S_D, S_I, S_U) = c_1 \sum_{\mathbf{x} \in S_I} L_1(\mathbf{x}) + c_2 \sum_{\mathbf{x} \in S_U} L_2(\mathbf{x}) + \sum_{\mathbf{x} \in S_D} (c_3 L_3(\mathbf{x}) + c_4 L_4(\mathbf{x}) + c_5 L_5(\mathbf{x})) + c_6 L_6(\mathbf{x}_o) \quad (11)$$

where \mathbf{x}_o is the equilibrium point, $L_1, L_2, L_3, L_4, L_5, L_6$ are defined in (7), (9) and (10), c_1, c_2, c_3 are defined in (6), and c_4, c_5, c_6 are non-negative constant sub-loss weights. Thus there are totally 6 sub-loss weights denoted by $\mathbf{c} = (c_1, c_2, \dots, c_6)$ for short; besides, there are 8 tolerances in (11) denoted by $\boldsymbol{\varepsilon} = (\varepsilon_1, \varepsilon_2, \dots, \varepsilon_8)$ for short. Our implementation and experiments are conducted based on (11) and related notations.

5.1. The Training Algorithm

The main algorithm for training a safe NN controller is presented in Algorithm 1, which can be explained as follows:

- $n_{\text{restart}}, n_{\text{epoch}}, n_{\text{batch}}$ and l_r are hyper-parameters for training (cf. Section 3.4); in all our case studies, n_{restart} and n_{batch} are fixed at 5 and 4096 respectively;
- `nn_construct()` in Line 1 is to construct the structure of \mathcal{N}_c and \mathcal{N}_b (cf. Section 3.1); in all our case studies, \mathcal{N}_c has one hidden layer with 5 neurons, and \mathcal{N}_b has one hidden layer with 10 neurons;
- `data_gen()` in Line 2 is to generate batches of training data (cf. Section 3.2);
- `initialize()` in Line 4 is to initialize weights and biases of \mathcal{N}_c and \mathcal{N}_b by Gaussian distribution;
- `compute_batch_loss()` in Line 8 is to compute the loss value on each batch of data using the input $\mathbf{c}, \boldsymbol{\varepsilon}$ (cf. Section 3.3 and (11));
- `update()` in Line 9 is to update \mathcal{N}_c and \mathcal{N}_b using *gradient descent* with step size l_r ;
- `decide_success()` in Line 11 is to decide the termination condition, which involves checking whether the epoch loss L_{epoch} reaches 0.

Algorithm 1 Safe NN-Controller Training Algorithm

Input: $\Gamma = (\mathbf{f}, X_D, X_I, X_U)$, n_{restart} , n_{epoch} , n_{batch} , l_r , \mathbf{c} , ε ;

Output: $\mathcal{N}_c, \mathcal{N}_b$;

```

1:  $\mathcal{N}_c, \mathcal{N}_b = \text{nn\_construct}(\Gamma)$ ;
2:  $\text{data\_gen}(\Gamma)$ ;
3: for  $i = 1$  to  $n_{\text{restart}}$  do
4:    $\text{initialize}(\mathcal{N}_c, \mathcal{N}_b)$ ;
5:   for  $j = 1$  to  $n_{\text{epoch}}$  do
6:      $L_{\text{epoch}} = 0$ ;
7:     for  $k = 1$  to  $n_{\text{batch}}$  do
8:        $L_{\text{epoch}} += \text{compute\_batch\_loss}(\mathbf{c}, \varepsilon)$ ;
9:        $\text{update}(\mathcal{N}_c, \mathcal{N}_b, l_r)$ ;
10:    end for
11:    if  $\text{decide\_success}(L_{\text{epoch}})$  then
12:      return  $\mathcal{N}_c, \mathcal{N}_b$ ;
13:    end if
14:  end for
15: end for

```

We have implemented a prototype tool `nncontroller`² based on the Pytorch³ platform. Given a problem description and a set of user-specified parameters (cf. Algorithm 1), `nncontroller` automatically learns a safe NN controller candidate with an NN barrier certificate, and generates script files as the input to iSAT3 for formal verification. We have applied `nncontroller` to a number of cases in the literature [TKID18, DKYP19, ZXMJ19]. All experiments are performed on a laptop workstation running Ubuntu 18.04 with Intel i7-8550u CPU and 32GB memory. The details of cases studies are presented in the following sub-section.

5.2. Experiment Results

In addition to the running example, we have synthesized and verified NN controllers using `nncontroller` for the following cases.

Example 5.1 (Inverted Pendulum [ZXMJ19]). The controlled CCDS $\Gamma = (\mathbf{f}, X_D, X_I, X_U)$ is:

$$\mathbf{f} : \begin{bmatrix} \dot{\theta} \\ \dot{\omega} \end{bmatrix} = \begin{bmatrix} \omega \\ \frac{g}{l}(\theta - \frac{\theta^3}{6}) + \frac{1}{ml^2}u \end{bmatrix},$$

where $m = 1$ and $l = 1$ denote the pendulum mass and length respectively, $g = 9.8$ is the gravitational acceleration, u is the scalar control input maintaining the pendulum upright, and

- $X_D: \{(\theta, \omega) \in \mathbb{R}^2 \mid -\pi/2 \leq \theta \leq \pi/2, -\pi/2 \leq \omega \leq \pi/2\}$;
- $X_I: \{(\theta, \omega) \in \mathbb{R}^2 \mid -\pi/9 \leq \theta \leq \pi/9, -\pi/9 \leq \omega \leq \pi/9\}$;
- X_U : the complement of $\{(\theta, \omega) \in \mathbb{R}^2 \mid -\pi/6 \leq \theta \leq \pi/6, -\pi/6 \leq \omega \leq \pi/6\}$ in X_D .

Example 5.2 (Duffing Oscillator [ZXMJ19]). The controlled CCDS $\Gamma = (\mathbf{f}, X_D, X_I, X_U)$ is:

$$\mathbf{f} : \begin{bmatrix} \dot{x} \\ \dot{y} \end{bmatrix} = \begin{bmatrix} y \\ -0.6y - x - x^3 + u \end{bmatrix},$$

where u is the scalar control input that regulates the system's trajectories to $(0, 0)$, and

- $X_D: \{(x, y) \in \mathbb{R}^2 \mid -6 \leq x \leq 6, -6 \leq y \leq 6\}$;
- $X_I: \{(x, y) \in \mathbb{R}^2 \mid -2.5 \leq x \leq 2.5, -2 \leq y \leq 2\}$;
- X_U : the complement of $\{(x, y) \in \mathbb{R}^2 \mid -5 \leq x \leq 5, -5 \leq y \leq 5\}$ in X_D .

² Publicly available at: <https://github.com/zhaohj2017/FAoC-tool>

³ <https://pytorch.org/>

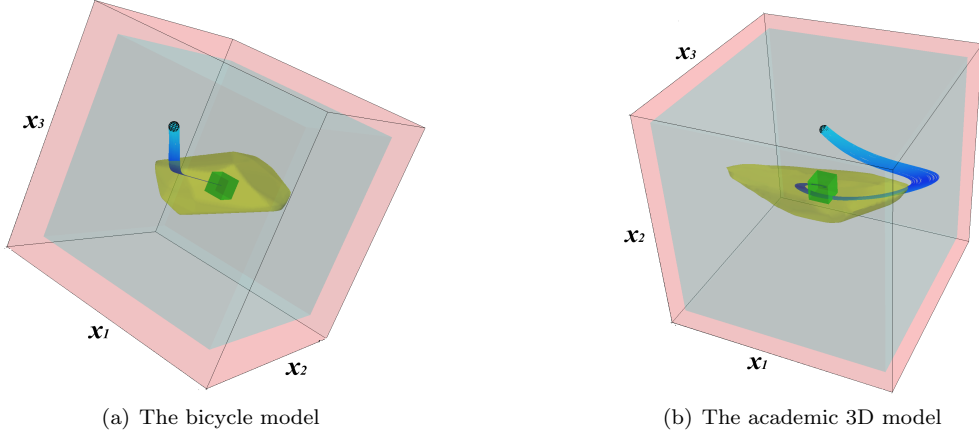


Fig. 12. Learned and verified NN controllers and barrier certificates for Example 5.3 and 5.4: for both cases, the innermost cube (green) represents the initial set, the outermost cube (pink) represents the system domain, and the space between the outermost and the middle cube (grey) is the unsafe region; the irregular surface (yellow) surrounding the innermost cube is the zero-level set of synthesized NN barrier certificates; the curves (blue) approaching the origin are simulated system trajectories

Example 5.3 (Bicycle Steering [DKYP19]). The control objective is to balance a bicycle. The states of the bicycle are (x_1, x_2, x_3) which denote the tilt angle, the angular velocity of tilt, and the handle bar angle with body respectively. The controlled CCDS $\Gamma = (\mathbf{f}, X_D, X_I, X_U)$ is:

$$\mathbf{f} : \begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \\ \dot{x}_3 \end{bmatrix} = \begin{bmatrix} \frac{ml}{J}(g \sin x_1 + \frac{x_2}{b} \cos x_1 \tan x_3) \\ 0 \\ 1 \end{bmatrix} + \begin{bmatrix} 0 \\ \frac{amlv}{Jb} \cdot \frac{\cos x_1}{\cos^2 x_3} \\ 1 \end{bmatrix} u,$$

where u is the scalar control input, $m = 20$ is the mass, $l = 1$ is the height, $b = 1$ is the wheel base, $J = \frac{mb^2}{3}$ is the moment of inertia, $v = 10$ is the velocity, $g = 10$ is the acceleration of gravity, $a = 0.5$ is the distance between the rear wheel and the line passing through the center of mass, and

- $X_D: \{(x_1, x_2, x_3) \in \mathbb{R}^3 \mid -\pi/2.5 \leq x_1 \leq \pi/2.5, -\pi/2.5 \leq x_2 \leq \pi/2.5, -\pi/2.5 \leq x_3 \leq \pi/2.5\}$;
- $X_I: \{(x_1, x_2, x_3) \in \mathbb{R}^3 \mid -\pi/30 \leq x_1 \leq \pi/30, -\pi/30 \leq x_2 \leq \pi/30, -\pi/30 \leq x_3 \leq \pi/30\}$;
- X_U : the complement of $\{(x_1, x_2, x_3) \in \mathbb{R}^3 \mid -\pi/3 \leq x_1 \leq \pi/3, -\pi/3 \leq x_2 \leq \pi/3, -\pi/3 \leq x_3 \leq \pi/3\}$ in X_D .

By introducing \tilde{u} such that $u = \tilde{u} \cos^2 x_3 - 20 \cos x_3 \sin x_3$, the original \mathbf{f} is transformed equivalently into

$$\tilde{\mathbf{f}} : \begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \\ \dot{x}_3 \end{bmatrix} = \begin{bmatrix} x_2 \\ 30 \sin x_1 + 15\tilde{u} \cos x_1 \\ \tilde{u} \cos^2 x_3 - 20 \cos x_3 \sin x_3 \end{bmatrix}.$$

An NN controller representing \tilde{u} was learned and verified for the transformed system $(\tilde{\mathbf{f}}, X_D, X_I, X_U)$ (cf. Fig. 12(a)).

Example 5.4 (Academic 3D [DKYP19]). The controlled CCDS $\Gamma = (\mathbf{f}, X_D, X_I, X_U)$ is:

$$\mathbf{f} : \begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \\ \dot{x}_3 \end{bmatrix} = \begin{bmatrix} x_3 + 8x_2 \\ -x_2 + x_3 \\ -x_3 - x_1^2 \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} u, \quad \text{where } u \text{ is the scalar control input}$$

- $X_D: \{(x_1, x_2, x_3) \in \mathbb{R}^3 \mid -2.2 \leq x_1 \leq 2.2, -2.2 \leq x_2 \leq 2.2, -2.2 \leq x_3 \leq 2.2\}$;
- $X_I: \{(x_1, x_2, x_3) \in \mathbb{R}^3 \mid -0.2 \leq x_1 \leq 0.2, -0.2 \leq x_2 \leq 0.2, -0.2 \leq x_3 \leq 0.2\}$;
- X_U : the complement of $\{(x_1, x_2, x_3) \in \mathbb{R}^3 \mid -2 \leq x_1 \leq 2, -2 \leq x_2 \leq 2, -2 \leq x_3 \leq 2\}$ in X_D .

An NN controller was successfully learned and verified for Γ (cf. Fig. 12(b)).

Table 1. Key parameters for pre-training and fine-tuning using `nncontroller` (cf. Algorithm 1 and Remark 5.1)

E.g.	n_e	l_r	\mathbf{c}	$\boldsymbol{\epsilon}$	\mathbf{c}^v	$\boldsymbol{\epsilon}^v$
3.1	100	0.1	(1, 1, 1, 0, 0, 0)	(0, 0, 0, 0.01, ·, ·, ·, ·)	(1, 1, 1, 0, 0, 0)	(0, 0, 0.01, 0.01, ·, ·, ·, ·)
5.1	100	0.1	(1, 1, 1, 0, 0, 0)	(0, 0, 0, 0.01, ·, ·, ·, ·)	(1, 1, 1, 0, 0, 0)	(0.01, 0, 0.02, 0.01, ·, ·, ·, ·)
5.2	100	0.01 \rightsquigarrow 0.1	(1, 1, 1, 0, 0, 0)	(0, 0, 0, 0.05, ·, ·, ·, ·)	(1, 1, 1, 0, 0, 0)	(0, 0, 0, 0.05, ·, ·, ·, ·)
5.3	200	0.01 \rightsquigarrow 0.2	(1, 1, 0.1, 0.1, 0, 0)	(0, 0, 0, 0.02, 0, ·, ·, ·)	(1, 1, 0.1, 0.1, 0.01, 0.01)	(0, 0, 0.35, 0.02, 0.35, 0.1, 0.1, 0.01)
5.4	200	0.01 \rightsquigarrow 0.2	(1, 1, 0.1, 0.1, 0, 0)	(0, 0, 0, 0.02, 0, ·, ·, ·)	(1, 1, 0.1, 0.1, 0.01, 0.01)	(0.01, 0.01, 0.15, 0.02, 0.1, 0.1, 0.2, 0.01)

Table 2. Time costs of synthesis and verification by `nncontroller` and `iSAT3` (cf. Remark 5.2)

E.g.	run 1		run 2		run 3		run 4		run 5		learning avg. cost	verification cost
	time	n_r	time	n_r	time	n_r	time	n_r	time	n_r		
3.1	21.11	0	15.04	0	14.98	0	65.25	0	15.37	0	26.35	8.27
5.1	478.29	1	168.75	0	292.96	0	111.55	0	43.89	0	219.09	15.24
5.2	60.59	0	72.47	0	64.64	0	48.08	0	851.49	1	219.45	4.71
5.3	752.63	1	1528.07	2	499.83	0	122.64	0	924.41	1	765.52	1344.50
5.4	240.94	0	301.22	0	2522.14	3	1001.66	1	390.25	0	891.24	6070.83

The key parameters used by `nncontroller` for our experiments are summarized in Table 1, and the time costs of synthesis and verification by `nncontroller` and `iSAT3` are summarized in Table 2.

Remark 5.1. In Table 1, n_e is a shorthand for n_{epoch} , \cdot means the corresponding parameter is not applicable, \rightsquigarrow means we adopt a self-adaptive learning rate scheduling strategy, and the superscript v means that the weight coefficients \mathbf{c}^v and parameters $\boldsymbol{\epsilon}^v$ are for the fine-tuned controllers, which are formally verified.

Remark 5.2. In Table 2, all time costs are measured in seconds; the time cost of NN controller training is not deterministic since the NN models are initialized randomly and the batches of training data are shuffled during the training process, and therefore we record the time costs of 5 separate runs of the training algorithm and compute the averaged cost; n_r denotes how many times we restart the algorithm when no NN controller is learned within the specified number of training epochs, i.e. n_{epoch} ; the last column corresponds to time costs of formal verification for the NN controllers and barrier certificates obtained with the \mathbf{c}^v and $\boldsymbol{\epsilon}^v$ parameters in Table 1 for each case.

Remark 5.3. Comparison of time costs of our experiment with related work such as [DKYP19, ZXMJ19] is not straightforward since we train two NNs simultaneously, while [DKYP19] requires user-provided barrier functions and [ZXMJ19] requires pre-trained NN controllers as their inputs. However, considering the number of layers and neurons (we use one hidden layer with 5 neurons and ReLU activations for \mathcal{N}_c uniformly), it can be asserted that our synthesized NN controllers have much simpler structure than [DKYP19, ZXMJ19].

6. Conclusion

We have proposed a new approach to synthesize neural network controllers for nonlinear continuous dynamical systems with control against safety properties. Our approach features in verification-in-the-loop synthesis: we simultaneously train the controller and its certificate, which we use barrier functions, represented by an NN as well. We have provided a prototype tool `nncontroller` with a number of case studies. The experiment results have confirmed the feasibility and efficacy of our approach.

Future work includes experimenting on different sampling and training strategies to reduce the data set size and to improve the training efficiency, as well different verification methods/tools other than interval SMT solvers. In particular, we plan to combine the counter-example-driven framework for program analysis [NARH17] with our proposed approach. Recently, the counter-example-guided inductive synthesis procedure

(CEGIS) has been employed in NN barrier certificates generation for continuous and hybrid systems with no control input [PAA20]. We anticipate that these would potentially further improve the scalability of our approach. We also plan to extend our approach to other properties such as reachability coupled with cost/reward based optimality as what has been done in optimal control and reinforcement learning.

Acknowledgements. We thank the anonymous reviewers for their valuable comments on the earlier versions of this paper, and thank Prof. Jyotirmoy V. Deshmukh for the explanation on the bicycle model of Example 5.3. H. Zhao was supported partially by the National Natural Science Foundation of China (No. 61702425, 61972385); X. Zeng was supported partially by the National Natural Science Foundation of China (No. 61902325), and “Fundamental Research Funds for the Central Universities” (SWU117058); T. Chen is partially supported by NSFC grant (No. 61872340), and Guangdong Science and Technology Department grant (No. 2018B010107004), the Overseas Grant of the State Key Laboratory of Novel Software Technology (No. KFKT2018A16), the Natural Science Foundation of Guangdong Province of China (No. 2019A1515011689); Z. Liu was supported partially by the National Natural Science Foundation of China (No. 62032019, 61672435, 61732019, 61811530327), and Capacity Development Grant of Southwest University (SWU116007); J. Woodcock was partially supported by the research grant from Southwest University.

References

- [ACE⁺19] Aaron D. Ames, Samuel Coogan, Magnus Egerstedt, Gennaro Notomista, Koushil Sreenath, and Paulo Tabuada. Control barrier functions: Theory and applications. In *2019 18th European Control Conference (ECC)*, pages 3420–3431, 2019.
- [ASBA19] Mohamadreza Ahmadi, Andrew Singletary, Joel W Burdick, and Aaron D. Ames. Safe policy synthesis in multi-agent POMDPs via discrete-time barrier functions. In *2019 IEEE 58th Conference on Decision and Control (CDC)*, pages 4797–4803. IEEE, 2019.
- [BTSK17] Felix Berkenkamp, Matteo Turchetta, Angela P. Schoellig, and Andreas Krause. Safe model-based reinforcement learning with stability guarantees. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, pages 908–919, Red Hook, NY, USA, 2017. Curran Associates Inc.
- [CCTS20] Jason Choi, Fernando Castañeda, Claire J. Tomlin, and Koushil Sreenath. Reinforcement learning for safety-critical control under model uncertainty, using control Lyapunov functions and control barrier functions. <https://arxiv.org/abs/2004.07584>, 2020.
- [COMB19] Richard Cheng, Gábor Orosz, Richard M. Murray, and Joel W. Burdick. End-to-end safe reinforcement learning through barrier functions for safety-critical continuous control tasks. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, pages 3387–3395. AAAI Press, 2019.
- [CRG19] Ya-Chien Chang, Nima Roohi, and Sicun Gao. Neural Lyapunov control. In *Advances in Neural Information Processing Systems 32*, pages 3245–3254. Curran Associates, Inc., 2019.
- [DCH⁺16] Yan Duan, Xi Chen, Rein Houthoofd, John Schulman, and Pieter Abbeel. Benchmarking deep reinforcement learning for continuous control. In *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, volume 48 of *JMLR Workshop and Conference Proceedings*, pages 1329–1338. JMLR.org, 2016.
- [DCS19] Souradeep Dutta, Xin Chen, and Sriram Sankaranarayanan. Reachability analysis for neural feedback systems using regressive polynomial rule inference. In *Proceedings of the 22nd ACM International Conference on Hybrid Systems: Computation and Control, HSCC*, pages 157–168, 2019.
- [DFG⁺19] Tommaso Dreossi, Daniel J. Fremont, Shromona Ghosh, Edward Kim, Hadi Ravanbakhsh, Marcell Vazquez-Chanlatte, and Sanjit A. Seshia. VerifAI: A toolkit for the formal design and analysis of artificial intelligence-based systems. In *Computer Aided Verification*, pages 432–442. Springer International Publishing, 2019.
- [DGXZ17] Liyun Dai, Ting Gan, Bican Xia, and Naijun Zhan. Barrier certificates revisited. *Journal of Symbolic Computation*, 80:62–86, 2017.
- [DJST18a] Souradeep Dutta, Susmit Jha, Sriram Sankaranarayanan, and Ashish Tiwari. Learning and verification of feedback control systems using feedforward neural networks. *IFAC-PapersOnLine*, 51(16):151 – 156, 2018. 6th IFAC Conference on Analysis and Design of Hybrid Systems ADHS 2018.
- [DJST18b] Souradeep Dutta, Susmit Jha, Sriram Sankaranarayanan, and Ashish Tiwari. Output range analysis for deep feedforward neural networks. In *NASA Formal Methods*, pages 121–138. Springer International Publishing, 2018.
- [DKYP19] Jyotirmoy V. Deshmukh, James Kapinski, Tomoya Yamaguchi, and Danil Prokhorov. Learning deep neural network controllers for dynamical systems with safety guarantees: Invited paper. In *2019 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, pages 1–7, 2019.
- [FP18] Nathan Fulton and André Platzer. Safe reinforcement learning via formal methods: Toward safe control through proof and learning. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 6485–6492. AAAI Press, 2018.
- [GBC16] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. The MIT Press, 2016.
- [Hes18] João P. Hespanha. *Linear Systems Theory*. Princeton University Press, second edition, 2018.
- [ICW⁺20] Radoslav Ivanov, Taylor J. Carpenter, James Weimer, Rajeev Alur, George J. Pappas, and Insup Lee. Case study: verifying the safety of an autonomous racing car with a neural network controller. In *HSCC ’20: 23rd ACM International Conference on Hybrid Systems: Computation and Control, Sydney, New South Wales, Australia, April 21-24, 2020*, pages 28:1–28:7. ACM, 2020.

- [IWA⁺19] Radoslav Ivanov, James Weimer, Rajeev Alur, George J. Pappas, and Insup Lee. Verisig: verifying safety properties of hybrid systems with neural network controllers. In *Proceedings of the 22nd ACM International Conference on Hybrid Systems: Computation and Control, HSCC 2019.*, pages 169–178, 2019.
- [JD20] Matt Jordan and Alexandros G. Dimakis. Exactly computing the local Lipschitz constant of ReLU networks. <https://arxiv.org/abs/2003.01219>, 2020.
- [KBD⁺17] Guy Katz, Clark Barrett, David L Dill, Kyle Julian, and Mykel J Kochenderfer. Reluplex: An efficient smt solver for verifying deep neural networks. In *International Conference on Computer Aided Verification*, pages 97–117. Springer, 2017.
- [KHS⁺13] Hui Kong, Fei He, Xiaoyu Song, William NN Hung, and Ming Gu. Exponential-condition-based barrier certificate generation for safety verification of hybrid systems. In *Proceedings of the 25th International Conference on Computer Aided Verification (CAV)*, pages 242–257. Springer, 2013.
- [LHP⁺16] Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016.
- [LLPS93] Moshe Leshno, Vladimir Ya. Lin, Allan Pinkus, and Shimon Schocken. Multilayer feedforward networks with a nonpolynomial activation function can approximate any function. *Neural Networks*, 6(6):861 – 867, 1993.
- [LLY⁺19] Jianlin Li, Jiangchao Liu, Pengfei Yang, Liqian Chen, Xiaowei Huang, and Lijun Zhang. Analyzing deep neural networks with symbolic propagation: Towards higher precision and faster verification. In *Static Analysis*, pages 296–319. Springer International Publishing, 2019.
- [MGQ⁺20] Mayank Mittal, Marco Gallieri, Alessio Quaglino, Seyed Sina Mirrazavi Salehian, and Jan Koutník. Neural Lyapunov model predictive control. <https://arxiv.org/abs/2002.10451>, 2020.
- [NARH17] ThanhVu Nguyen, Timos Antonopoulos, Andrew Ruef, and Michael Hicks. Counterexample-guided approach to finding numerical invariants. In *Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering, ESEC/FSE 2017*, page 605–615, New York, NY, USA, 2017. Association for Computing Machinery.
- [PAA20] Andrea Peruffo, Daniele Ahmed, and Alessandro Abate. Automated and formal synthesis of neural barrier certificates for dynamical models. <https://arxiv.org/abs/2007.03251>, 2020.
- [PEY01] Alex Poznyak, Sanchez EN, and Wen Yu. *Differential Neural Networks for Robust Nonlinear Control*. World Scientific, 2001.
- [PJP07] Stephen Prajna, Ali Jadbabaie, and George J. Pappas. A framework for worst-case and stochastic safety verification using barrier certificates. *IEEE Transactions on Automatic Control*, 52(8):1415–1429, 2007.
- [PT10] Luca Pulina and Armando Tacchella. An abstraction-refinement approach to verification of artificial neural networks. In *Computer Aided Verification*, pages 243–257, 2010.
- [RAA19] Alex Ray, Joshua Achiam, and Dario Amodei. Benchmarking safe exploration in deep reinforcement learning. <https://cdn.openai.com/safexp-short.pdf>, 2019.
- [Rat18] Stefan Ratschan. Converse theorems for safety and barrier certificates. *IEEE Transactions on Automatic Control*, 63(8):2628–2632, 2018.
- [RBK18] Spencer M. Richards, Felix Berkenkamp, and Andreas Krause. The Lyapunov neural network: Adaptive stability certification for safe learning of dynamic systems. <http://arxiv.org/abs/1808.00924>, 2018.
- [RS07] Stefan Ratschan and Zhikun She. Safety verification of hybrid systems by constraint propagation-based abstraction refinement. *ACM Trans. Embed. Comput. Syst.*, 6(1):1–23, 2007.
- [RS10] Stefan Ratschan and Zhikun She. Providing a basin of attraction to a target region of polynomial systems by computation of Lyapunov-like functions. *SIAM Journal on Control and Optimization*, 48(7):4377–4394, 2010.
- [RS19] Hadi Ravanbakhsh and Sriram Sankaranarayanan. Learning control Lyapunov functions from counterexamples and demonstrations. *Autonomous Robots*, 43(2):275–307, 2019.
- [SGTP18] Andrew Sogokon, Khalil Ghorbal, Yong Kiam Tan, and André Platzer. Vector barrier certificates and comparison systems. In *Formal Methods*, pages 418–437, 2018.
- [SKS19] Xiaowu Sun, Haitham Khedr, and Yasser Shoukry. Formal verification of neural network controlled autonomous systems. In *Proceedings of the 22nd ACM International Conference on Hybrid Systems: Computation and Control, HSCC 2019.*, pages 147–156, 2019.
- [SL20] Zhikun She and Meilun Li. Over- and under-approximations of reachable sets with series representations of evolution functions. *IEEE Transactions on Automatic Control*, 2020.
- [SPW12] Christoffer Sloth, George J. Pappas, and Rafael Wisniewski. Compositional safety analysis using barrier certificates. In *Proc. of the Hybrid Systems: Computation and Control (HSCC)*, pages 15–24. ACM, 2012.
- [TDL⁺19] Andrew J. Taylor, Victor D. Dorobantu, Hoang M. Le, Yisong Yue, and Aaron D. Ames. Episodic learning with control Lyapunov functions for uncertain robotic systems. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6878–6884, 2019.
- [Tel17] Matus Telgarsky. Neural networks and rational functions. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70, ICML’17*, page 3387–3393. JMLR.org, 2017.
- [TKID18] Cumhur Erkan Tuncali, James Kapinski, Hisahiro Ito, and Jyotirmoy V. Deshmukh. Invited: Reasoning about safety of learning-enabled components in autonomous cyber-physical systems. In *2018 55th ACM/ESDA/IEEE Design Automation Conference (DAC)*, pages 1–6, 2018.
- [TSYA19] Andrew Taylor, Andrew Singletary, Yisong Yue, and Aaron Ames. Learning for safety-critical control with control barrier functions. <https://arxiv.org/abs/1912.10099>, 2019.
- [TYML⁺20] Hoang-Dung Tran, Xiaodong Yang, Diego Manzananas Lopez, Patrick Musau, Luan Viet Nguyen, Weiming Xiang, Stanley Bak, and Taylor T. Johnson. NNV: The neural network verification tool for deep neural networks and

- learning-enabled cyber-physical systems. In *Computer Aided Verification*, pages 3–17. Springer International Publishing, 2020.
- [WS16] Rafael Wisniewski and Christoffer Sloth. Converse barrier certificate theorems. *IEEE Transactions on Automatic Control*, 61(5):1356–1361, 2016.
- [WZC⁺18] Tsui-Wei Weng, Huan Zhang, Hongge Chen, Zhao Song, Cho-Jui Hsieh, Luca Daniel, Duane S. Boning, and Inderjit S. Dhillon. Towards fast computation of certified robustness for relu networks. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018*, pages 5273–5282, 2018.
- [XTJ18] Weiming Xiang, Hoang-Dung Tran, and Taylor T. Johnson. Output reachable set estimation and verification for multilayer neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 29(11):5777–5783, 2018.
- [YFS20] Shakiba Yaghoubi, Georgios Fainekos, and Sriram Sankaranarayanan. Training neural network controllers using control barrier functions in the presence of disturbances. <https://arxiv.org/abs/2001.08088>, 2020.
- [ZXMJ19] He Zhu, Zikang Xiong, Stephen Magill, and Suresh Jagannathan. An inductive synthesis framework for verifiable reinforcement learning. In *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2019*, pages 686–701, New York, NY, USA, 2019. Association for Computing Machinery.
- [ZZC⁺20] Hengjun Zhao, Xia Zeng, Taolue Chen, Zhiming Liu, and Jim Woodcock. Learning safe neural network controllers with barrier certificates. In *Dependable Software Engineering. Theories, Tools, and Applications*, pages 177–185, Cham, 2020. Springer International Publishing.
- [ZZCL20] Hengjun Zhao, Xia Zeng, Taolue Chen, and Zhiming Liu. Synthesizing barrier certificates using neural networks. In *HSCC '20*, pages 25:1–25:11. ACM, 2020.