

Differentiable Safe Controller Design Through Control Barrier Functions

Shuo Yang^{ID}, *Graduate Student Member, IEEE*, Shaoru Chen^{ID}, *Graduate Student Member, IEEE*, Victor M. Preciado^{ID}, *Senior Member, IEEE*, and Rahul Mangharam^{ID}, *Senior Member, IEEE*

Abstract—Learning-based controllers, such as neural network (NN) controllers, can show high empirical performance but lack formal safety guarantees. To address this issue, control barrier functions (CBFs) have been applied as a safety filter to monitor and modify the outputs of learning-based controllers in order to guarantee the safety of the closed-loop system. However, such modification can be myopic with unpredictable long-term effects. In this letter, we propose a safe-by-construction NN controller which employs differentiable CBF-based safety layers and relies on a set-theoretic parameterization. We compare the performance and computational complexity of the proposed controller and an alternative projection-based safe NN controller in learning-based control. Both methods demonstrate improved closed-loop performance over using CBF as a separate safety filter in numerical experiments.

Index Terms—Safety-critical control, control barrier functions, neural network controller, safe learning control.

I. INTRODUCTION

LEARNING-BASED control has become increasingly popular for controlling complex dynamical systems [1] since it requires little expert knowledge and can be carried out in an automatic, data-driven manner. However, due to the black-box nature of learning models, learning-based controllers such as neural network (NN) controllers lack formal guarantees which significantly limits their deployment in safety-critical applications.

The integration of control-theoretical approaches and machine learning has provided a promising solution to safe learning control, where trainable machine learning modules are embedded into a control framework that guarantees the safety or stability of the dynamical system [2], [3]. Of wide applicability is the control barrier function (CBF) framework [4], [5]

Manuscript received 16 September 2022; revised 23 November 2022; accepted 17 December 2022. Date of publication 30 December 2022; date of current version 11 January 2023. This work was supported in part by the U.S. Department of Transportation’s Mobility21 National University Transportation Center and in part by NSF CCRI under Grant 1925587. Recommended by Senior Editor S. Dey. (*Shuo Yang and Shaoru Chen contributed equally to this work.*) (*Corresponding author: Shuo Yang.*)

The authors are with the Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104 USA (e-mail: yangs1@seas.upenn.edu; srchen@seas.upenn.edu; preciado@seas.upenn.edu; rahulm@seas.upenn.edu).

Digital Object Identifier 10.1109/LCSYS.2022.3233322

which explicitly specifies a safe control set and guards the system inside a safe invariant set. This is achieved by constructing a CBF-based safety filter that projects any reference control input (possibly generated by a NN controller) onto the safe control set online. When a continuous-time control-affine system is considered, such projection reduces to a convex quadratic program (QP) which is referred to as CBF-QP. Due to its simplicity, flexibility, and formal safety guarantees, CBFs have been applied in safe learning control with many successful applications [6], [7].

Compared with model predictive control (MPC) [8], which needs to solve a nonconvex optimization problem in the face of nonlinear dynamical systems, CBF-QP is computationally efficient to solve online. However, unlike MPC, the QP-based safety filter only operates in a minimally invasive manner, i.e., it generates the safe control input closest (in the Euclidean norm) to the reference control input, unaware of the long-term effects of its action. This indicates that the effects of the safety filter on the performance of the closed-loop system are hard to predict. Therefore, the application of the safety filter may give rise to myopic controllers [9] that induce sub-par performance in the long term.

To address the issue of myopic CBF-based safety filters, in this letter we propose to utilize CBF to construct safe-by-construction NN controllers that allow end-to-end learning. Incorporating safety layers in the NN controller allows the learning agent to take the effects of safety filters into account during training in order to maximize long-term performance. Inspired by [10], we design a differentiable safety layer using the gauge map which establishes a bijection between the polytopic output set of a NN (e.g., an ℓ_∞ norm ball) and the safe control set characterized by CBFs. We denote the proposed architecture as NN-gauge (Fig. 1(a)). We compare NN-gauge with an alternative safe NN controller, NN-diff-QP, which consists of a NN followed by a differentiable CBF-QP layer (Fig. 1(b)). In the online execution, NN-gauge requires closed-form evaluation or solving a linear program (LP) while NN-diff-QP solves a quadratic program. Both methods are significantly cheaper to run online than MPC.

A. Related Works

Safe controller design: The use of gauge map in safe learning control was proposed in [10], [11] which only consider

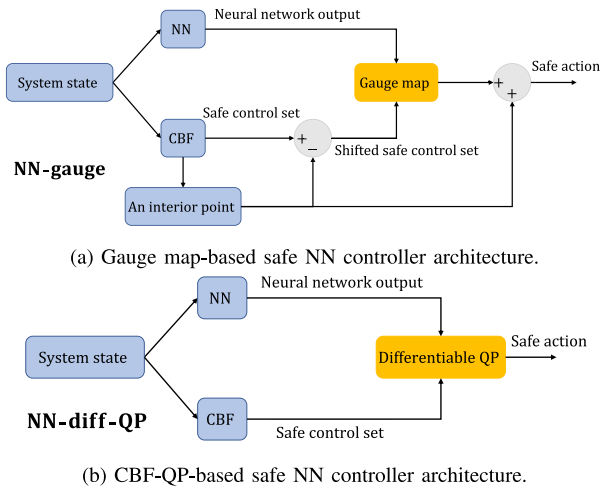


Fig. 1. Safe-by-construction NN controllers that utilize CBFs to construct differentiable safety layers (yellow blocks).

linear dynamics for which control invariant sets and an interior safe control policy are achievable. In this letter, by proposing NN-gauge, we significantly extend the scope of this framework to handle nonlinear dynamics with CBFs and overcome the arising computational difficulties by applying an implicit interior policy parameterization. The construction of NN-diff-QP naturally follows from the use of differentiable optimization layers [12], [13], [14]. NN-diff-QP is applied in [15] and [16] for reinforcement learning tasks. In [17], CBFs are applied as penalty functions to promote the safety of NN controllers. However, unlike NN-gauge or NN-diff-QP, the resulting NN controllers do not have formal safety guarantees.

Differential CBF: Introducing learning modules in the parameterization of CBFs can improve the feasibility and performance of the safety filter [18], [19] even in the face of changing environments [20], [21]. These works focus on learning or improving CBFs such that the safe control set is enlarged and the safety filter can work better with a reference controller. Instead, in this letter, we consider NN controller synthesis with a given CBF.

The contributions of this letter are summarized below:

- 1) We propose a novel, differentiable, safe-by-construction NN controller NN-gauge as shown in Fig. 1(a) using CBFs. To the best of our knowledge, NN-gauge is the only alternative to the projection-based safe NN controller NN-diff-QP (Fig. 1(b)). Compared with NN-diff-QP, NN-gauge enjoys more efficient training and online evaluation.
- 2) We provide detailed case studies to evaluate the performance-complexity trade-off of NN-gauge and NN-diff-QP.
- 3) We demonstrate that learning safe-by-construction NN controllers leads to better long-term closed-loop performance than filtering a trained, possibly unsafe NN controller.

In Section II, background on CBF is provided and the control problem is introduced. The construction of NN-gauge and

NN-diff-QP is shown in Section III, followed by numerical examples in Section IV. Section V concludes this letter.

II. PRELIMINARY AND PROBLEM FORMULATION

A. System Model

In this letter we are interested in a continuous-time nonlinear control-affine system:

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

where f and g are locally Lipschitz, $x \in D \subseteq \mathbb{R}^n$ is the state, D denotes a compact set in \mathbb{R}^n , and $u \in \mathbb{R}^m$ is the control input subject to bounded polytopic control constraints:

$$u \in U = \{u \in \mathbb{R}^m \mid A_u u \leq b_u\}. \quad (2)$$

B. Control Barrier Functions

Safety can be framed in the context of enforcing set invariance in the state space, i.e., the state cannot exit a safe set. The safe set \mathcal{C} is represented by the superlevel set of a continuously differentiable function $h(x)$. The algebraic expressions of the safe set \mathcal{C} , boundary of the safe set $\partial\mathcal{C}$, and interior of the safe set $\text{Int}(\mathcal{C})$ are given by:

$$\begin{aligned} \mathcal{C} &= \{x \in \mathbb{R}^n : h(x) \geq 0\}, \\ \partial\mathcal{C} &= \{x \in D \subset \mathbb{R}^n : h(x) = 0\}, \\ \text{Int}(\mathcal{C}) &= \{x \in D \subset \mathbb{R}^n : h(x) > 0\}. \end{aligned} \quad (3)$$

For a locally Lipschitz continuous control law $u = k(x)$, we have that $\dot{x} = f(x) + g(x)k(x)$ is locally Lipschitz continuous. Thus, for any initial condition $x_0 \in D$, there exists a maximum time interval of existence $I(x_0) = [0, \tau_{max})$, such that $x(t)$ is the unique solution to the ordinary differential equation (1) on $I(x_0)$. We frame the safety of system (1) in terms of set invariance as shown below.

Definition 1 (Forward Invariance and Safety): The set \mathcal{C} is forward invariant if for every $x_0 \in \mathcal{C}$, $x(t) \in \mathcal{C}$ holds for $x(0) = x_0$ and all $t \in I(x_0)$. If \mathcal{C} is forward invariant, we say system (1) is safe.

To verify invariance of \mathcal{C} under the control input constraints (2), a control barrier function is constructed as a certificate which characterizes the admissible set of control inputs that render \mathcal{C} forward invariant.

Definition 2 (Extended Class \mathcal{K}_∞ Function): A continuous function $\alpha: \mathbb{R} \rightarrow \mathbb{R}$ is said to be an extended \mathcal{K}_∞ if it is strictly increasing and $\alpha(0) = 0$.

Definition 3 (Control Barrier Function): Let $\mathcal{C} \subset D \subset \mathbb{R}^n$ be the superlevel set of a continuously differentiable function $h: D \rightarrow \mathbb{R}$, then h is a control barrier function if there exists an extended class \mathcal{K}_∞ function $\alpha(\cdot)$ such that for the control system (1):

$$\sup_{u \in U} [L_f h(x) + L_g h(x)u] \geq -\alpha(h(x)), \quad (4)$$

for all $x \in D$ where $L_f h(x)$ and $L_g h(x)$ denote the Lie derivatives.

Given the CBF $h(x)$, the set of all control values that render \mathcal{C} safe is given by:

$$K_{cbf}(x) = \{u \in U : L_f h(x) + L_g h(x)u + \alpha(h(x)) \geq 0\} \quad (5)$$

which we denote as the safe control set. The following theorem shows that the existence of a control barrier function implies that the control system (1) is safe:

Theorem 1 [4, Th. 2]: Assume $h(x)$ is a CBF on $D \supset \mathcal{C}$ and $\frac{\partial h}{\partial x}(x) \neq 0$ for all $x \in \partial \mathcal{C}$. Then any Lipschitz continuous controller $u(x)$ such that $u(x) \in K_{cbf}(x)$ for all $x \in \mathcal{C}$ will render the set \mathcal{C} forward invariant.

One important feature of $K_{cbf}(x)$ is that it is a polytope¹ for all states. This enables the construction of a QP-based safety filter that modifies any given reference controller $k(x)$ in a minimally invasive fashion [5] as follows:

$$\begin{aligned} u(x) = \underset{u}{\operatorname{argmin}} \quad & \|u - k(x)\|_2^2 \\ \text{subject to} \quad & L_f h(x) + L_g h(x)u + \alpha(h(x)) \geq 0, \\ & u \in U. \end{aligned} \quad (6)$$

Although the safety filter (6) guarantees the forward invariance of the safe set, it does not take into account the consequences of the projection on future states and the performance of the closed-loop system. This issue is inevitable when the reference controller $k(x)$ and the safety filter are designed separately, and we propose to fix it by designing safe-by-construction controllers that are amenable to any learning or optimization framework. Particularly, in this letter, we consider optimizing NN controllers using modern machine learning solvers (such as stochastic gradient descent (SGD) and Adam [22]) with known system dynamics.

C. Problem Formulation

Following the definition of safe control set (5), we define the set of safe control policies as $\Pi := \{\pi : \mathbb{R}^n \mapsto \mathbb{R}^m | \pi(x) \in K_{cbf}(x), \forall x \in \mathcal{C}\}$. Our task is to design a controller for system (1) such that a performance objective is optimized and the closed-loop system always stays inside the safe set \mathcal{C} . In other words, we consider finding a policy $\pi(x)$ that solves the following optimal control problem within a horizon $T < \tau_{max}$ ²:

$$\begin{aligned} \underset{\pi \in \Pi}{\operatorname{minimize}} \quad & \mathbb{E}_{x(0) \in \mathcal{C}} \left[\frac{1}{T} \int_0^T c(x(t), u(t)) dt \right] \\ \text{subject to} \quad & \dot{x} = f(x) + g(x)u, \\ & u(t) = \pi(x(t)), \end{aligned} \quad (7)$$

where $\mathbb{E}_{x(0) \in \mathcal{C}}$ is the expectation with respect to the initial state, $c(x(t), u(t))$ is the cost associated with occupying state $x(t)$ and action $u(t)$. The cost function $c(x(t), u(t))$ could be of any form and is problem-specific, e.g., it can formulate the penalty or barrier functions of constraints on the state $x(t)$.

Despite the complex dynamics, cost functions, and safe policy constraints, problem (7) can still be effectively approached by parameterizing an NN policy and applying SGD/Adam

¹The polytopic safe control set $K_{cbf}(x)$ is non-empty for all $x \in \mathcal{C}$ by definition of CBF. In this letter, we assume a valid CBF for system (1) is given and use it for NN controller design, although synthesizing CBFs can be challenging itself and is an active area of research.

²The horizon T is a tuning parameter for NN controller design. While a larger T is always preferred to improve the closed-loop performance of the trained NN controller, it necessarily increases the computational complexity of training.

which is empowered by automatic differentiation and modern machine learning solvers [23]. This procedure is simple and has been shown effective to synthesize high-performance NN controllers [24], [25]. Next, we study different parameterizations of safe NN policies $\pi \in \Pi$ for solving (7) and evaluate their performances. Notably, with safe NN policies, the effects of the CBF-based safety layer are automatically considered.

III. SAFE CONTROLLER DESIGN

A natural way to construct a safe NN controller $\pi(x)$ is to restrict the output of the controller into the safe control set (5) for all states. NN-diff-QP (Fig. 1(b)) achieves this by concatenating a differentiable projection layer which can be implemented in a NN using toolboxes such as `cvxpylayers` [14] and `qpth` [12]. In this section, we propose a different parameterization of a safe NN controller which achieves improved online computational efficiency using gauge maps.

A. Gauge Map

Tabas and Zhang [10] observe that, while it is challenging to directly restrict the output of a NN inside a general polytope, adding a hyperbolic tangent activation layer easily constrains the NN output into the unit ℓ_∞ -norm ball \mathbb{B}_∞ . This motivates the application of the gauge map that establishes a bijection between \mathbb{B}_∞ and a general polytope which in this letter we consider as $K_{cbf}(x)$. The notion of gauge map is facilitated by the concept of C-set.

Definition 4 (C-Set [26]): A C-set is a convex and compact set including the origin as an interior point.

The gauge function (or Minkowski function) of a vector $v \in \mathbb{R}^m$ with respect to a C-set $Q \subset \mathbb{R}^m$ is given by

$$\gamma_Q(v) = \inf\{\lambda \geq 0 | v \in \lambda Q\}. \quad (8)$$

When Q is a polytopic C-set defined by $\{w \in \mathbb{R}^m | F_i^T w \leq g_i, i = 1, 2, \dots, r\}$, the gauge function can be written in closed-form as $\gamma_Q(v) = \max_i\{F_i^T v / g_i\}$. For any $v \in \mathbb{B}_\infty$, the gauge map is defined as

$$G(v|Q) = \frac{\|v\|_\infty}{\gamma_Q(v)} \cdot v, \quad (9)$$

which constructs a bijection (see [10, Lemma 1]) between the unit ball \mathbb{B}_∞ and the C-set Q . As shown in Fig. 2, all points $v \in \mathbb{B}_\infty$ are mapped “proportionally” to the same level set of the polytope Q by the gauge map, while their projections onto Q tend to concentrate on the boundary of Q .

B. Interior Safe Policy

To apply the gauge map as a safety layer that maps the output of an NN to the safe control set $K_{cbf}(x)$, we have to first find an interior safe policy $\pi_{int}(x)$ such that $\pi_{int}(x) \in \operatorname{Int}(\mathcal{C})$ for all $x \in \mathcal{C}$. This is necessary since $K_{cbf}(x)$ may not be a C-set, and we need to shift $K_{cbf}(x)$ by $\pi_{int}(x)$ such that it is recentered around the origin. When $\pi_{int}(x)$ is available, we have that the shifted safe control set $\hat{K}_{cbf}(x) := K_{cbf}(x) - \pi_{int}(x)$ is a C-set for which we can apply the gauge map.




Fig. 2. Illustration of the gauge map from the ℓ_∞ ball B_∞ to a polytopic set Q .

An explicit construction of $\pi_{int}(x)$ is achievable for linear dynamical systems through multi-parametric programming [11]. The recent work [27] proposes an algorithm to extract an explicit or closed-form safe policy from CBFs for system (1). Such an explicit interior policy is desirable since in this case the gauge map can be evaluated in closed-form, making both the training and online evaluation of NN-gauge computationally efficient. However, for general nonlinear systems, the proposed algorithm in [27] can be complex to apply in practice. To address this issue, we propose an alternative method that implicitly constructs an interior safe policy by choosing $\pi_{int}(x)$ as the Chebyshev center [28] of $K_{cbf}(x)$. Specifically, we have $\pi_{int}(x) = u^*$ where u^* is the solution to the following LP:

$$\begin{aligned} & \underset{u, R \geq 0}{\text{maximize}} && R \\ & \text{subject to} && -L_g h(x)u + R \| -L_g h(x) \|_2 \\ & && \leq L_f h(x) + \alpha(h(x)), \\ & && a_{u,i}^\top u + R \| a_{u,i} \|_2 \leq b_{u,i}, \quad i = 1, \dots, d_u, \end{aligned} \quad (10)$$

where $a_{u,i}$ denotes the i -th row of A_u , $b_{u,i}$ denotes the i -th entry of b_u , and d_u represents the number of linear constraints defining U . The Chebyshev center formulation of $\pi_{int}(x)$ pushes $\pi_{int}(x)$ away from the boundary of $K_{cbf}(x)$. It also facilitates training of the upstream NN, since it makes the target function that the NN needs to learn smoother. By the validity of CBF, the safe control set $K_{cbf}(x)$ is a non-empty polytope for all $x \in \mathcal{C}$. By choosing $\pi_{int}(x)$ as the Chebyshev center from (10), we readily have that $\pi_{int}(x) \in \text{Int}(K_{cbf}(x))$ for $R > 0$ and the shifted safe control set $\hat{K}_{cbf}(x)$ is a C-set.

C. Control Policy Architecture

With an interior safe policy $\pi_{int}(x)$, we can now construct a safe NN controller using the gauge map, as shown in the following theorem.

Theorem 2: Let $\phi_\theta : \mathcal{C} \rightarrow \mathbb{B}_\infty$ be a neural network parameterized by θ and $\pi_{int}(x)$ be an interior safe policy. Then, for any system state x in the set \mathcal{C} , the policy:

$$\pi_\theta(x) := G(\phi_\theta(x))\hat{K}_{cbf}(x) + \pi_{int}(x) \quad (11)$$

has the following properties:

- 1) π_θ is safe.
- 2) The policy $\pi_\theta(\cdot)$ is trainable with respect to the NN parameters θ .

Proof: 1) By the construction of the gauge map, we have $G(\phi_\theta(x))\hat{K}_{cbf}(x) \in \hat{K}_{cbk}(x)$ which is the safe control set $K_{cbf}(x)$ shifted by the interior safe control input $\pi_{int}(x)$. Therefore, we have $\pi_\theta(x) \in K_{cbf}(x)$ for all $x \in \mathcal{C}$, and conclude that $\pi_\theta(x)$ is safe.

2) As shown in [10, Th. 1], automatic differentiation can be applied to compute the subgradients of π_θ with respect to θ , making it possible to train π_θ . ■

The NN ϕ_θ can embed any architecture and learns the residual control policy added to the interior safe policy $\pi_{int}(x)$. By the construction of the gauge map, the applied controller $\pi_\theta(x)$ shown in (11) belongs to the safe control set $K_{cbf}(x)$, and the performance of $\pi_\theta(x)$ is no worse than $\pi_{int}(x)$ after training. The online evaluation of π_θ can be done in closed-form if an explicit interior safe policy $\pi_{int}(x)$ is given, or by solving an LP if the implicit construction (10) is used.

Remark 1: Our analysis of safe NN controllers can be readily applied to incorporate high order CBFs [29] in which case the safe control set $K_{cbf}(x)$ is still a polytope. In addition to the CBF-based safe control sets, NN-gauge and NN-diff-QP can easily encode other forms of polytopic safe control sets such as $u \in \{u \mid Gu \leq b\}$ and $u \in \{u \mid Fx + Gu \leq b\}$.

IV. NUMERICAL EXPERIMENTS

In this section, we demonstrate the application of safe NN controllers in adaptive cruise control (ACC) and aircraft collision avoidance. The following control methods are considered:

- 1) *MPC:* Model predictive control could guarantee the safety of the closed-loop system with good performance, but it is computationally costly to run online when the horizon is large or the dynamics is nonlinear.
- 2) *NN:* A feedforward NN controller is trained to optimizes (7) with regularizers penalizing violations of safety constraints. No safety filter is applied in the online evaluation of this NN controller. This method enjoys fast online evaluation but suffers the risk of safety violations.
- 3) *NN-QP:* The above NN controller is equipped with the CBF-QP safety filter during online evaluation.
- 4) *NN-diff-QP and NN-gauge:* The safe-by-construction NN controllers introduced in Section III that are trained directly to optimize the control performance (7). We also include the interior policy π_{int} given by (10) as a special case of NN-gauge for comparison.

For NN-diff-QP, we use `cvxpylayers` [14] to construct the differentiable QP layer. For NN-gauge, the implicit interior policy (10) is applied. All training is performed using PyTorch on Google Colab with Adam [22] as the optimizer.

A. Adaptive Cruise Control

Adaptive cruise control is a common example to validate safe control strategies [4], [30]. The control goal of ACC is to let the ego car achieve the desired cruising speed while maintaining a safe distance from the leading car. We consider the scenario where the ego car tries to follow the leading car on a straight road. The dynamics of the problem is given by

(model adapted from [30]):

$$\dot{x}(t) = \begin{bmatrix} 0 & 1 & 0 \\ 0 & -0.1 & 0 \\ 0 & -1 & 0 \end{bmatrix} x(t) + \begin{bmatrix} 0 \\ 2.5 \\ 0 \end{bmatrix} u(t), \quad (12)$$

where $x = [x_1 \ x_2 \ x_3]^\top = [p_f \ v_f \ d]^\top$, p_f is the position, v_f is the velocity of the ego car, d is the distance between the ego and leading cars, and $u \in [-1, 1]$ is the control input denoting the acceleration. To prevent collision between the two cars, the CBF is chosen as $h(x) = x_3 - 1.8x_2$ and the trajectory cost for the ego car is given by

$$\text{Cost} = \int_0^T (0.01(x_2(t) - v_f^*)^2 + 0.05u(t)^2) dt \quad (13)$$

with a desired speed $v_f^* = 30\text{m/s}$. The leading car travels at a constant speed of 16m/s , so we expect the ego car's speed to converge to 16m/s at the steady state since a speed greater than 16m/s will lead to a violation of safe distance.

All NN controllers, namely NN, NN-gauge and NN-diff-QP, are trained to optimize (7) with horizon $T = 1\text{s}$ by Adam with randomly sampled initial states. The system dynamics (12) is discretized with sampling rate $\Delta t = 0.1\text{s}$ during training. The trainable NN modules ϕ_θ in NN, NN-gauge, and NN-diff-QP have the same architecture.

We test all NN controllers on 5 randomly sampled initial states over a horizon $T = 20\text{s}$ in order to evaluate their long-term performance. The results are reported in Table I. For one of the testing initial condition $x(0) = [0 \ 30 \ 100]^\top$, we plot the values of the CBF function and the velocity along the trajectory of the ego car in Fig. 3(b) and Fig. 3(a), respectively.

We observe that both NN-gauge and NN-diff-QP achieve similar closed-loop performance, comparable to MPC.³ While NN achieves reasonable performance, it violates safety constraints (Fig. 3(a)). Directly applying the CBF-QP safety filter on it enforces safe control, but deteriorates the long-term closed-loop performance of the NN controller as shown in Fig. 3(b) where the optimal behavior of the ego car is supposed to have a steady-state velocity of 16m/s . NN-gauge has an edge over NN-diff-QP in training and online evaluation time due to its use of LP-based safety layers. We also observe a large performance improvement of NN-gauge compared with π_{int} which means training the NN module in NN-Gauge can greatly improve the performance of the policy.

B. Aircraft Collision Avoidance

We apply our framework to the aircraft collision avoidance problem which is adapted from [31]. Specifically, we consider a dynamical system with states $x = [x_a^\top \ x_b^\top]^\top$, where $x_a = [p_{a,x} \ p_{a,y} \ \theta_a]^\top \in \mathbb{R}^3$ is the state of aircraft a with $p_{a,x}$ and $p_{a,y}$ denoting the position and θ_a denoting the orientation. The state x_b of aircraft b is defined similarly. The control inputs are $u = [v_a \ w_a \ v_b \ w_b]^\top \in \mathbb{R}^4$ where v_a and w_a are speed and turning rate of aircraft a , respectively, and v_b , w_b are defined

³While MPC solves a finite horizon optimal control problem to optimality, its closed-loop performance is not guaranteed to be optimal. In Table I, NN-diff-QP achieved better performance than MPC in the long term.

TABLE I
CLOSED-LOOP PERFORMANCE COMPARISON OF CONTROLLERS IN THE ADAPTIVE CRUISE CONTROL EXAMPLE. AVERAGE VALUES FROM 5 RANDOMLY SAMPLED INITIAL STATES ARE REPORTED

	Safety	Trajectory cost	Training time per epoch	Solve time
MPC	Safe	269.3	N/A	3.11s
NN	Unsafe	640.3	0.03s	0.04s
NN-QP	Safe	818.2	0.03s	0.36s
NN-diff-QP	Safe	258.3	11.0s	0.35s
NN-gauge	Safe	270.9	8.3s	0.28s
π_{int}	Safe	734.8	N/A	0.26s

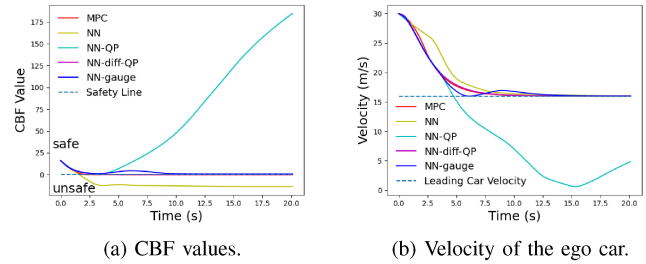


Fig. 3. CBF values (left) and velocity of the ego car (right) under different controllers are evaluated in closed-loop for 20s. A CBF value below zero indicates unsafety, and the optimal behavior of the ego car is expected to have a steady state velocity of 16m/s same as the leading car.




Fig. 4. Trajectories of the aircraft a and b under different controllers. The induced costs of each trajectory are labeled accordingly.

similarly. The dynamics of the aircraft a (and similarly for aircraft b) is given by:

$$\dot{x}_a(t) = \begin{bmatrix} v_a(t)\cos(\theta_a(t)) \\ v_a(t)\sin(\theta_a(t)) \\ w_a(t) \end{bmatrix}. \quad (14)$$

As shown in Fig. 4, our goal is to drive aircraft a to the left and aircraft b to the right while maintaining a minimum safe distance of 0.5 between them. Aircrafts a and b try to stay close to $p_{a,x} = -5$ and $p_{b,x} = 5$, respectively. A quadratic cost function is defined accordingly over a horizon T , and we adopt the constructive CBF developed in [31] to encode the safe set in the state space which also considers the input constraint $(v_a, w_a, v_b, w_b) \in [0.1, 1] \times [-1, 1] \times [0.1, 1] \times [-1, 1]$. Note that the minimum admissible velocity of aircraft a and b is 0.1 , so they cannot stop exactly at $p_{a,x} = -5$ or $p_{b,x} = 5$.

All NN controllers are trained similarly as in the ACC example with sampling rate $\Delta t = 0.1\text{s}$ and horizon $T = 2\text{s}$,

TABLE II

CLOSED-LOOP PERFORMANCE COMPARISON OF CONTROLLERS IN THE AIRCRAFT COLLISION AVOIDANCE EXAMPLE. AVERAGE VALUES FROM 5 RANDOMLY SAMPLED INITIAL STATES ARE REPORTED

	Safety	Trajectory cost	Training time per epoch	Solve time
MPC	Safe	1346.8	N/A	44.1s
NN	Unsafe	712.0	0.24s	0.06s
NN-QP	Safe	2557.6	0.24s	1.28s
NN-diff-QP	Safe	1900.7	133.0s	1.30s
NN-gauge	Safe	2157.8	101.5s	0.74s
π_{int}	Safe	3495.9	N/A	0.69s

and are tested on 5 randomly sampled initial states with horizon $T = 20$ s. The results are shown in Table II. One set of closed-loop system trajectories starting from the initial state $x(0) = [0.5 \ 0 \ \pi \ -0.5 \ 0 \ 0]^T$ under different controllers are plotted in Fig. 4 together with the induced costs. With this initial condition, aircrafts a and b start close to each other with orientations leading to a head-on collision.

From Table II and Fig. 4, we observe that NN achieves the best performance but is unsafe. Adding the CBF-QP as a safety filter (i.e., NN-QP) drastically deteriorates the performance of the NN controller. Among the NN controllers with safety guarantee, NN-diff-QP performs the best and NN-gauge achieves a similar level of performance. The MPC controller has the best performance with safety guarantee, but it has a significantly higher online solve time.

V. CONCLUSION

In this letter, we showed that CBF-based safety filters can degrade closed-loop performance if their long-term effects are not considered during learning. To address this issue, we proposed a novel safe-by-construction NN controller which utilizes CBF and gauge map to construct a differentiable safety layer. The proposed gauge map-based NN controller achieves comparable performances as the projection-based NN controller while being computationally more efficient to train and evaluate online. Both the gauge map-based and projection-based safe NN controllers demonstrate improved performance compared with filtered NN controllers in numerical examples.

ACKNOWLEDGMENT

The authors thank Nikolai Matni for useful discussions.

REFERENCES

- [1] H. A. Pierson and M. S. Gashler, "Deep learning in robotics: A review of recent research," *Adv. Robot.*, vol. 31, no. 16, pp. 821–835, 2017.
- [2] P. L. Donti, M. Roderick, M. Fazlyab, and J. Z. Kolter, "Enforcing robust control guarantees within neural network policies," in *Proc. Int. Conf. Learn. Represent.*, 2020, pp. 1–26.
- [3] L. Furieri, C. L. Galimberti, M. Zakwan, and G. Ferrari-Trecate, "Distributed neural network control with dependability guarantees: A compositional port-hamiltonian approach," in *Proc. Learn. Dyn. Control Conf.*, 2022, pp. 571–583.
- [4] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs for safety critical systems," *IEEE Trans. Autom. Control*, vol. 62, no. 8, pp. 3861–3876, Aug. 2017.
- [5] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: Theory and applications," in *Proc. 18th Eur. Control Conf. (ECC)*, 2019, pp. 3420–3431.
- [6] A. Anand, K. Seel, V. Gjørnum, A. Håkansson, H. Robinson, and A. Saad, "Safe learning for control using control lyapunov functions and control barrier functions: A review," *Procedia Comput. Sci.*, vol. 192, pp. 3987–3997, 2021.
- [7] C. Dawson, Z. Qin, S. Gao, and C. Fan, "Safe nonlinear control using robust neural Lyapunov-barrier functions," in *Proc. Conf. Robot Learn.*, 2022, pp. 1724–1735.
- [8] E. F. Camacho and C. B. Alba, *Model Predictive Control*. London, U.K.: Springer, 2013.
- [9] M. H. Cohen and C. Belta, "Approximate optimal control for safety-critical systems with control barrier functions," in *Proc. 59th IEEE Conf. Decis. Control (CDC)*, 2020, pp. 2062–2067.
- [10] D. Tabas and B. Zhang, "Computationally efficient safe reinforcement learning for power systems," in *Proc. Amer. Control Conf. (ACC)*, 2022, pp. 3303–3310.
- [11] D. Tabas and B. Zhang, "Safe and efficient model predictive control using neural networks: An interior point approach," 2022, *arXiv:2203.12196*.
- [12] B. Amos and J. Z. Kolter, "OptNet: Differentiable optimization as a layer in neural networks," in *Proc. Int. Conf. Mach. Learn.*, 2017, pp. 136–145.
- [13] B. Amos, I. Jimenez, J. Sacks, B. Boots, and J. Z. Kolter, "Differentiable MPC for end-to-end planning and control," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 31, 2018, pp. 1–12.
- [14] A. Agrawal, B. Amos, S. Barratt, S. Boyd, S. Diamond, and J. Z. Kolter, "Differentiable convex optimization layers," in *Proc. Adv. Neural Inf. Processing Syst.*, vol. 32, 2019, pp. 1–19.
- [15] M. Pereira, Z. Wang, I. Exarchos, and E. Theodorou, "Safe optimal control using stochastic barrier functions and deep forward-backward SDEs," in *Proc. Conf. Robot Learn.*, 2021, pp. 1783–1801.
- [16] Y. Emam, P. Glotfelter, Z. Kira, and M. Egerstedt, "Safe model-based reinforcement learning using robust control barrier functions," 2021, *arXiv:2110.05415*.
- [17] W. S. Cortez, J. Drgona, A. Tuor, M. Halappanavar, and D. Vrabie, "Differentiable predictive control with safety guarantees: A control barrier function approach," 2022, *arXiv:2208.02319*.
- [18] H. Parwana and D. Panagou, "Recursive feasibility guided optimal parameter adaptation of differential convex optimization policies for safety-critical systems," in *Proc. Int. Conf. Robot. Autom. (ICRA)*, 2022, pp. 6807–6813.
- [19] K. Xu, W. Xiao, and C. G. Cassandras, "Feasibility guaranteed traffic merging control using control barrier functions," 2022, *arXiv:2203.04348*.
- [20] H. Ma, B. Zhang, M. Tomizuka, and K. Sreenath, "Learning differentiable safety-critical control using control barrier functions for generalization to novel environments," 2022, *arXiv:2201.01347*.
- [21] W. Xiao, R. Hasani, X. Li, and D. Rus, "BarrierNet: A safety-guaranteed layer for neural networks," 2021, *arXiv:2111.11277*.
- [22] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. 3rd Int. Conf. Learn. Represent.*, 2015, pp. 1–13.
- [23] A. G. Baydin, B. A. Pearlmutter, A. A. Radul, and J. M. Siskind, "Automatic differentiation in machine learning: A survey," *J. Mach. Learn. Res.*, vol. 18, pp. 1–43, Apr. 2018.
- [24] J. Drgona, A. Tuor, and D. Vrabie, "Learning constrained adaptive differentiable predictive control policies with guarantees," 2020, *arXiv:2004.11184*.
- [25] S. Mukherjee, J. Drgoňa, A. Tuor, M. Halappanavar, and D. Vrabie, "Neural lyapunov differentiable predictive control," 2022, *arXiv:2205.10728*.
- [26] F. Blanchini and S. Miani, *Set-Theoretic Methods in Control*, vol. 78. Cham, Switzerland: Springer, 2008.
- [27] H. Wang, K. Margellos, and A. Papachristodoulou, "Explicit solutions for safety problems using control barrier functions," 2022, *arXiv:2204.09380*.
- [28] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press, 2004.
- [29] W. Xiao and C. Belta, "Control barrier functions for systems with high relative degree," in *Proc. IEEE 58th Conf. Decis. Control (CDC)*, 2019, pp. 474–479.
- [30] J. Zeng, B. Zhang, Z. Li, and K. Sreenath, "Safety-critical control using optimal-decay control barrier function with guaranteed point-wise feasibility," in *Proc. Amer. Control Conf. (ACC)*, 2021, pp. 3856–3863.
- [31] E. Squires, P. Pierpaoli, and M. Egerstedt, "Constructive barrier certificates with applications to fixed-wing aircraft collision avoidance," in *Proc. IEEE Conf. Control Technol. Appl. (CCTA)*, 2018, pp. 1656–1661.