

Safe control under input saturation with neural control barrier functions

Anonymous Author(s)

Affiliation

Address

email

1 **Abstract:** We propose new methods to synthesize control barrier function (CBF)
2 based safe controllers that avoid input saturation, which can cause safety viola-
3 tions. In particular, our method is created for high-dimensional, general nonlinear
4 systems, for which such tools are scarce. We leverage techniques from machine
5 learning, like neural networks and deep learning, to simplify this challenging prob-
6 lem in nonlinear control design. The method consists of a learner-critic architec-
7 ture, in which the critic gives counterexamples of input saturation and the learner
8 optimizes a neural CBF to eliminate those counterexamples. We provide empiri-
9 cal results on a 10D state, 4D input quadcopter-pendulum system. Our learned
10 CBF avoids input saturation and maintains safety over nearly 100% of trials.

11 **Keywords:** safety, control, learning

12 1 Introduction

13 In theory, control barrier functions are an appealing tool for safe control. However, it is difficult
14 to make the derived controllers respect input limits, which reduces their usage in practice. CBFs
15 target the *set invariance* class of safety problems, in which safety is defined as keeping a system’s
16 state to a prescribed region. A large part of their appeal is that they offer mathematical guarantees
17 of safety. Such assurances are essential for safety-critical robotics applications, like collision-free
18 drone flight [1, 2], manipulators that work safely around humans [3], and stable bipedal walking [4].
19 However, these safety guarantees break down when input saturation occurs, since that means the
20 system cannot exert the force required for an evasive maneuver. The system then becomes endan-
21 gered (or dangerous), with the possibility of expensive equipment failure or people getting harmed.
22 It is therefore imperative that we account for input limits when designing CBFs.

23 CBFs are *energy functions* which map states to an energy value, with safe states having lower en-
24 ergy. In *energy function methods*, an energy function is found and then a controller is crafted that
25 dissipates the energy. The core problem of these methods is constructing the energy function. A
26 valid energy function has to meet complex constraints that depend on the input limits, system dy-
27 namics, and safety specification. So far, this problem of designing CBFs around input limits has
28 been studied to a limited degree, mostly for small or simple systems. To our knowledge, nothing has
29 been proposed which handles the nonlinear and high-dimensional systems that are more realistic in
30 robotics. Examples of existing work include methods for designing CBFs using intuition for simple
31 systems [5]. Other works derive CBFs that are provably non-saturating for systems with special
32 structure, like kinematic bicycle, polynomial and Euler-Lagrange systems [6, 7, 8, 9, 10]. Another
33 approach is to frame this as an optimization problem, computing parameters for CBFs with simple
34 forms for use on low-dimensional systems [11, 12, 8]. A non-CBF safe control method with guar-
35 antees against saturation is HJ reachability [13]. It is generic (handles nonlinear systems), but has
36 not been able to scale past 6D¹, except in special cases [15, 16].

¹In general, synthesizing a safe controller for a nonlinear system with *formal guarantees* against satura-
tion is an NP-hard problem [13]. Hence, it is intrinsically difficult to scale Hamilton-Jacobi (HJ) reachability.

37 For the most part, existing works address narrow classes of systems, and some of them even require
 38 substantial human effort. In contrast, we envision a framework for *automating* CBF synthesis that
 39 has a *wide range of applicability*. Our goal is for it to be easy-to-use and generic, encompassing
 40 nonlinear systems of higher dimension. To achieve this, we borrow ideas from machine learning
 41 (ML), like neural networks, which are well suited to large-scale nonlinear optimization problems.

42 The related field of Lyapunov function (LF) synthesis has already taken inspiration from ML, with
 43 success. LFs are a different kind of energy function used for *certifying* stabilizing controllers, rather
 44 than constructing new controllers. Usually, it is assumed that the candidate stabilizing controllers
 45 already respect input limits, so input saturation is not considered. However, LFs are still hard to
 46 construct because they need to satisfy other complex constraints. Recent works have proposed
 47 representing LFs as neural networks (NNs), which are expressive and easy to optimize [17, 18,
 48 19, 20]. To enforce constraint satisfaction over input sets, these neural LFs are either trained on a
 49 grid of inputs [20], or better yet, violating inputs [19]. There has also been work on neural CBFs,
 50 but they have been used for the entirely distinct problem of inferring a safety criterion from expert
 51 trajectories [21, 22].

52 We propose using *neural CBFs* to simplify CBF synthesis and extend its range of applicability. This
 53 choice of functional representation enables us to handle high-dimensional, nonlinear systems. In
 54 particular, NNs can easily handle systems with large state spaces ($>6D$), since they can be efficiently
 55 trained on much larger inputs than that ($>100D$). Further, for a nonlinear system, a valid CBF can
 56 be arbitrarily nonlinear, which demands the rich expressiveness of NNs. Finally, learning a non-
 57 saturating CBF for a nonlinear system involves a highly nonlinear loss function. Thus, we also need
 58 a functional representation that can be optimized for good solutions very easily, like NNs.

59 In the rest of the paper, we expound on our two main contributions, the first of which is our aforemen-
 60 tioned neural CBF design (Sec. 3.1). The second contribution is our learning framework (Sec. 3.2).
 61 We begin by proposing a loss function that captures the severity of saturation. In our learner-critic
 62 paradigm, the critic finds states that maximize the loss (that is, worst saturators) and the learner
 63 updates the neural CBF to minimize saturation at those states. We also introduce a number of techni-
 64 ques that are required to make our learner stable and our critic efficient and powerful. Our method
 65 produces a CBF that ensures safety nearly 100% of the time for a input-limited, nonlinear, 10D
 66 quadcopter-pendulum system (Sec. 4).

67 2 Preliminaries

68 In this section, we provide a review of CBFs, mathematically define a non-saturating CBF, and
 69 explain the premise of CBF synthesis. First, some notation: for a function $c : \mathbb{R}^n \rightarrow \mathbb{R}$, let $\mathcal{C} \triangleq$
 70 $\{c\}_{\leq 0}$ be its zero-sublevel set, $\partial\mathcal{C} = \{c\}_{=0}$ the boundary of this set, and $\text{Int}(\mathcal{C}) = \{c\}_{<0}$ the interior.
 71 Now, we assume the following is given: (1) a control-affine system $\dot{x} = f(x) + g(x)u$, where
 72 $x \in \mathcal{D} \subset \mathbb{R}^n$, $u \in \mathcal{U} \subseteq \mathbb{R}^m$ and $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$, $g : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times m}$ are locally Lipschitz continuous
 73 on \mathbb{R}^n , (2) input set \mathcal{U} , a bounded convex polytope, and (3) a safety specification $\rho : \mathbb{R}^n \rightarrow \mathbb{R}$,
 74 which implicitly defines the *allowable set* $\mathcal{A} \triangleq \{\rho\}_{\leq 0}$. ρ is continuous and piecewise smooth.

75 Typically, in CBF-based safe control, a CBF and *safe set* are constructed based on the user-defined
 76 safety specification and then a controller is formed to keep the safe set *forward invariant* (FI). A set
 77 is forward invariant if all trajectories starting within it stay within it. As such, let r be the *relative*
 78 *degree* from ρ to u (i.e. the first derivative of ρ where u appears) and then a standard *limit-blind*
 79 *CBF* is:

$$\phi = \left[\prod_{i=1}^{r-1} \left(1 + c_i \frac{\partial}{\partial t} \right) \right] \rho \quad (\text{Limit-blind CBF})$$

However, recent advances in HJ reachability have abandoned formal guarantees to create potential for scalability [14], which we have also done in our work. We note, however, that despite the lack of guarantees, our method is very effective in practice.

80 , where $c_i < 0$ [5]. This CBF has a corresponding safe set $\mathcal{S} \subseteq \mathcal{A}$; see the Appendix for its exact
 81 definition. Then, as usual, we can straightforwardly construct a safe controller. Any controller
 82 which can repel the system back into the safe set whenever it reaches the safe set boundary is a safe
 83 controller. The following theorem formalizes this idea:

84 **Theorem 1** (Taken from [5]). *Given a CBF ϕ and safe set \mathcal{S} , any feedback controller $k(x) : \mathbb{R}^n \rightarrow$
 85 \mathbb{R}^m satisfying*

$$\dot{\phi}(x, k(x)) \triangleq \underbrace{\nabla\phi(x)^\top f(x)}_{L_f\phi(x)} + \underbrace{\nabla\phi(x)^\top g(x)k(x)}_{L_g\phi(x)} \leq 0 \quad \forall x \in \partial\mathcal{S} \quad (1)$$

86 *renders the system forward invariant over \mathcal{S} .*

87 This theorem requires the controller to repel the system (decrease ϕ) from the boundary ($x \in \partial\mathcal{S}$).
 88 It allows us to follow an arbitrary nominal policy, k_{nom} , as long as we project the input to the half-
 89 space satisfying Eqn. 1. Thus, the standard CBF safe controller simply applies this projection at
 90 every timestep of control execution [5, 7]. This is implemented as a quadratic program (QP):

$$k(x) = \arg \min_u \frac{1}{2} \|u - k_{nom}(x)\|_2^2 \quad (\text{CBF-QP})$$

$$\text{s.t.} \quad L_f\phi(x) + L_g\phi(x)u \leq \begin{cases} 0 & \text{if } x \in \partial\mathcal{S} \\ \infty & \text{o.w.} \end{cases} \quad (2)$$

$$u \in \mathcal{U} \quad (3)$$

91 The significant problem with this paradigm is that this controller can saturate. Specifically, saturation
 92 occurs when no $u \in \mathcal{U}$ exists that satisfies Eqn. 2, causing the loss of safety guarantees.² This can
 93 happen because we have not yet accounted for input limits in our design of ϕ . What we propose is
 94 to craft a different CBF that does not lead to saturation down the line, which is the process of *CBF*
 95 *synthesis*. We will call this a *non-saturating CBF*:

96 **Definition 1** (*Non-saturating CBF*). *A function $\phi : \mathbb{R}^n \rightarrow \mathbb{R}$ is a non-saturating CBF over a set \mathcal{S}*
 97 *if for all $x \in \partial\mathcal{S}$:*

$$\inf_{u \in \mathcal{U}} \dot{\phi}(x, u) \leq 0 \quad (4)$$

98 Intuitively, this is just requiring that there exist a feasible control input to decrease ϕ (return the
 99 system to the interior of \mathcal{S}) at every state on the boundary, $\partial\mathcal{S}$. It is possible to find a non-saturating
 100 CBF ϕ^* that defines a smaller safe set ($\mathcal{S}^* \subseteq \mathcal{S}$) than the original saturating CBF. The idea is
 101 that a smaller set can exclude all the irrecoverable states from the original set. Formally, we want
 102 $\rho^* : \mathbb{R}^n \rightarrow \mathbb{R}$ such that $\{\rho^*\}_{\leq 0} \subseteq \{\rho\}_{\leq 0}$. Then, our new CBF is just the following:

$$\phi^* = \left[\prod_{i=1}^{r-1} \left(1 + c_i \frac{\partial}{\partial t} \right) \right] \rho - \rho + \rho^* \quad (\text{Limit-aware CBF})$$

103 3 Learning Non-Saturating Control Barrier Functions

104 In this section, we present our neural CBF design (Sec. 3.1) and then discuss the learning framework
 105 which trains it to be non-saturating (Sec. 3.2). Learning is a min-max optimization problem of the
 106 following form, where θ denotes the parameters of ϕ^* :

$$\min_{\theta} \max_{x \in \partial\mathcal{S}^*} \mathcal{L}(\theta, x) \quad (5)$$

107 We call this loss $\mathcal{L}(\theta, x)$ the *saturation risk*. This min-max problem is solved using a learner-critic
 108 algorithm (Alg. 1), where the critic and learner alternately find where the worst saturation occurs
 109 and then update the CBF to reduce saturation there. We also propose some strategies for training
 110 stably, boosting critic efficiency, and for increasing the volume of the safe set. We visualize the
 111 training process for a simple running example (Fig. 1), where we learn a CBF to prevent an inverted
 112 pendulum from tipping.

²In practice, to avoid an unsolvable QP when saturation occurs, we add a slack variable to Eqn. 2. However, we will still violate safety guarantees.

Algorithm 1 Learning non-saturating CBF

```
1: function LEARN( $\mathbb{X}_{ce}, \theta$ )
2:   Set learning rate  $\beta = 0.001$ 
3:    $\theta \leftarrow \theta - \beta \cdot \nabla_{\theta} [\sum_{x \in \mathbb{X}_{ce}} \text{softmax}(\mathcal{L}(\theta, x)) + \mathcal{R}(\theta)]$   $\triangleright$  From Eqn. 6, 7
4:   return  $\theta$ 
5: end function
6: function COMPUTECE( $\mathbb{X}_{ce}, \theta$ )
7:   Set learning rate  $\alpha = 0.001$ , number of gradient steps  $N = 20$ 
8:    $\mathbb{X}_{ce} \leftarrow$  project  $\mathbb{X}_{ce}$  to new boundary  $\triangleright$  See Alg. 3
9:    $\mathbb{X}_{rand} \leftarrow$  uniformly sample a set on the boundary  $\triangleright$  See Alg. 5
10:   $\mathbb{X} \leftarrow \mathbb{X}_{rand} \cup \mathbb{X}_{ce}$ 
11:  for  $i$  in  $[0, \dots, N]$  do
12:     $\mathbb{G} \leftarrow \nabla_{\mathbb{X}} \mathcal{L}(\theta, \mathbb{X})$   $\triangleright$  Batch gradient
13:     $\mathbb{P} \leftarrow$  project  $\mathbb{G}$  along boundary
14:     $\mathbb{X} \leftarrow \mathbb{X} + \alpha \cdot \mathbb{P}$   $\triangleright$  Batch update
15:     $\mathbb{X} \leftarrow$  project  $\mathbb{X}$  to boundary  $\triangleright$  See Alg. 3
16:  end for
17:   $\mathbb{X}_{ce} \leftarrow$  worst saturating states in  $\mathbb{X}$ 
18:  return  $\mathbb{X}_{ce}$ 
19: end function
20: function MAIN()
21:   Input: dynamical system  $\dot{x}$ , safety specification  $\rho$ 
22:   Randomly initialize neural CBF parameters  $\theta$ 
23:   Repeat:
24:      $\mathbb{X}_{ce} \leftarrow$  COMPUTECE( $\mathbb{X}_{ce}, \theta$ )  $\triangleright \mathbb{X}_{ce}$  is a set of counterexamples
25:      $\theta \leftarrow$  LEARN( $\mathbb{X}_{ce}, \theta$ )
26:   Until convergence
27:   return  $\theta$ 
28: end function
```

113 3.1 Neural CBF Design

114 We consider the design of ρ^* . Let $\text{nn} : \mathbb{R}^n \rightarrow \mathbb{R}$ be a multilayer perceptron with tanh activations.
115 Then, we define

$$\rho^*(x) = (\text{nn}(x) - \text{nn}(x_e))^2 + \rho(x) \quad (\text{Design 1})$$

116 if there is a state, x_e , that should belong to any valid \mathcal{S}^* . Otherwise, let

$$\rho^*(x) = \text{softplus}(\text{nn}(x)) + \rho(x) \quad (\text{Design 2})$$

117 The reason for these designs is that a valid ρ^* needs to obey three constraints: *Constraint 1*: ρ^* is
118 piecewise smooth. This is required to preserve the piecewise smoothness of ϕ^* , which allows us to
119 make existence and uniqueness arguments for the closed-loop system. Our ρ^* satisfies this because
120 nn has smooth tanh activations and ρ is assumed piecewise smooth. *Constraint 2*: The 0-sublevel
121 set of ρ^* is contained within the 0-sublevel set of ρ : $\{\rho^*\}_{\leq 0} \subseteq \{\rho\}_{\leq 0}$. The allowable set can be
122 shrunk but not enlarged; otherwise, dangerous states may be incorporated. Our design meets this
123 criterion because $\rho^* \geq \rho$. *Constraint 3*: \mathcal{S}^* is nonempty, where \mathcal{S}^* is defined by ρ^* . In some cases, it
124 makes sense to enforce some state x_e to lie within \mathcal{S}^* (Design 1), which means \mathcal{S}^* will be nonempty
125 by design. One such case is stabilization-type safety problems, where safety is defined as bounded
126 deviation from an equilibrium. Then, the equilibrium would be a natural choice for x_e . However,
127 if this is not the case, then we can choose Design 2 and encourage non-emptiness by regularizing
128 the safe set volume (Sec. 3.2). Our CBF forms have the advantage of satisfying constraints 1-3 by
129 design, while still offering the nonlinear expressiveness that we need to find a non-saturating CBF.

130 3.2 Learning framework

131 In the following sections, we present a loss function that encourages satisfaction of Eqn. 1 and the
132 algorithm for solving the min-max problem on this loss. First, we define θ as the parameters of ϕ^* ,

133 which include the weights of m and the c_i coefficients. Our loss function, called *saturation risk*, is
 134 defined as:

$$\mathcal{L}(\theta, x) \triangleq \inf_{u \in \mathcal{U}} \dot{\phi}_\theta^*(x, u) \quad (\text{Saturation risk})$$

135 It is a measure of the best-case saturation at a given state x . When $\mathcal{L}(\theta, x) \leq 0$, then no saturation
 136 occurs at x ; when $\mathcal{L}(\theta, x) > 0$, it measures how severe the saturation is. Thus, our min-max problem
 137 (Eqn. 5) is to minimize the *worst best-case saturation* over the boundary. When the worst best-case
 138 is negative, i.e.

$$\max_{x \in \partial \mathcal{S}^*} \mathcal{L}(\theta, x) \leq 0 \quad (\text{Training goal})$$

139 , then we have successfully found a non-saturating CBF.

140 To solve the min-max problem on $\mathcal{L}(\theta, x)$ (Eqn. 5), we propose a learner-critic algorithm (Alg. 1).
 141 Essentially, the algorithm alternates between the critic computing counterexamples (maximization
 142 with respect to x) and the learner updating the CBF (minimization with respect to θ). The critic
 143 uses projected gradient descent to produce an approximate maximizer, \hat{x}^* , and then the learner uses
 144 gradient descent to minimize the saturation loss at \hat{x}^* . See below for more details on the critic. Since
 145 both learner and critic perform gradient descent on $\mathcal{L}(\theta, x)$, it is useful to re-express it as an analytic
 146 function, rather than a continuous minimization. To find the analytic expression, observe that this
 147 $\mathcal{L}(\theta, x)$ a minimization where the objective is affine in the variable u (from Eqn. 1) and the variable
 148 is constrained to a convex polyhedron set \mathcal{U} . This means the minimizing u^* is one of the vertices
 149 $v \in \mathcal{V}(\mathcal{U})$ of the constraint set. Thus, $\mathcal{L}(\theta, x)$ can be computed as a discrete minimization, which is
 150 analytic:

$$\mathcal{L}(\theta, x) \triangleq \min_{v \in \mathcal{V}(\mathcal{U})} \dot{\phi}_\theta^*(x, v) \quad (\text{Analytic saturation risk})$$

151 In practice, we get more stable convergence if we use a *weighted average* of the loss over all coun-
 152 terexamples \mathbb{X}_{ce} instead of the loss on the single worst counterexample. That is, the learner update
 153 becomes:

$$\theta = \theta - \alpha \cdot \nabla_\theta \left[\sum_{x \in \mathbb{X}_{ce}} \text{softmax}(\mathcal{L}(\theta, x)) \right] \quad (6)$$

154 This avoids two problems which hinder convergence: (1) deadlock (when improvement at one coun-
 155 terexample causes saturation at another) and (2) noisy gradient ($\nabla_\theta \mathcal{L}(\theta, \hat{x}^*)$ only *approximates* the
 156 true gradient $\nabla_\theta \mathcal{L}(\theta, x^*)$, since \hat{x}^* is a suboptimal maximizer). A weighted average loss both en-
 157 courages progress on many counterexamples at once and averages out the effect of noisy gradients.

158 Design of the critic

159 In this section, we discuss the design of a fast and powerful critic, which is key to the success
 160 of the training algorithm (see function COMPUTECE in Alg. 1). The critic needs to return good
 161 counterexamples, otherwise the CBF will not train properly. However, it also needs to accomplish
 162 this efficiently, since it needs to recompute counterexamples repeatedly throughout training. Another
 163 challenge to speed is systems with large state spaces, since state space size is also the size of the
 164 critic’s optimization space. To this end, we propose strategies for boosting critic efficacy cheaply.

165 The critic is in charge of finding counterexamples, which are states along the safe set boundary
 166 where saturation would occur at control-time. This is a constrained optimization problem, which
 167 we handle using projected gradient descent (PGD). We project the saturation risk’s gradient along
 168 the boundary, use this to update x , and finally project x back to the boundary. However, in practice,
 169 this basic critic produces unsatisfactory counterexamples. There are many poor local optima in this
 170 optimization problem that can trap PGD. We explore two strategies for circumventing this issue:
 171 taking the best of a batch and warmstarting.

172 For the first strategy, we apply batch PGD to a *set* of initial states and take the best of the result-
 173 ing candidate counterexamples. This batch optimization is a cheap way to improve our chances of




Figure 1: Learned safe sets for toy inverted pendulum problem, at four points during training. The black dots represent candidate counterexamples; the critic correctly identifies that the states where angle and angular velocity have the same sign are the most likely to saturate (angular velocity acting to destabilize pendulum). In (d), green denotes the largest safe set, approximated by model predictive control (MPC) (see Appendix for details). Our regularization is so effective that there is virtually no difference between our learned set and the MPC set.

174 evading local optima. A portion of the initial states is sampled *uniformly* on the boundary using
 175 the fast routine from [23]. The other initial states are “warmstarted”: we take them to be the best
 176 counterexamples from the previous critic call. This strategy is based on the idea that old counterex-
 177 amples can still be good candidates for the current call. The current call should only have a slightly
 178 different optimization objective and constraint than the previous call, due to the incremental update
 179 to θ . Overall, this reuse of work is a big boon to efficiency. Together, these two techniques help the
 180 critic perform efficiently and effectively, even for high-dimensional state spaces, as in Sec. 4.

181 **Enlarging the safe set** In this section, we introduce a regularization term that we add to the training
 182 objective to help enlarge the safe set. In general, one benefit of optimization-driven approaches
 183 to synthesis is that we can augment the objective with terms that target control *performance*. One
 184 measure of a good safe controller is the size of its safe set: a larger safe set means the controller can
 185 prevent more states from evolving toward danger. Thus, we devise a regularization term:

$$\mathcal{R}(\theta) \triangleq \sum_{x \in \mathbb{X}_{reg}} \text{sigmoid}(s_{\theta}^*(x)) \quad (7)$$

186 , where $s_{\theta}^*(x)$ defines inclusion in the safe set ($\mathcal{S}^* = \{s_{\theta}^* \leq 0\}$) and \mathbb{X}_{reg} is a set sampled uniformly
 187 in \mathcal{D} . The idea behind this term is that we want states near the safe set boundary to move inside the
 188 safe set (or further inside, if they are already). Mathematically, we want to decrease $s_{\theta}^*(x)$ toward
 189 negative values (move into \mathcal{S}^*) for x with $s_{\theta}^*(x) \approx 0$ (near boundary). In order to only affect x such
 190 that $s_{\theta}^*(x) \approx 0$, we apply the sigmoid transform. To test the efficacy of this term, we compare the
 191 volume of safe sets learned with and without it, at iterations where they attain similar saturation risk.
 192 It turns out that without this term, the safe set tends to shrink more than necessary during training, as
 193 smaller safe sets have a lower possibility of saturating. We find that inclusion of the term provides a
 194 nearly 200% increase in volume.

195 4 Experiments

196 In this section, we train a neural CBF on a challenging input-limited robotic system. We pose two
 197 experimental questions:

198 **Q1.** How well do we achieve our training objective?

199 **Q2.** Does the CBF-based safe controller ensure FI?

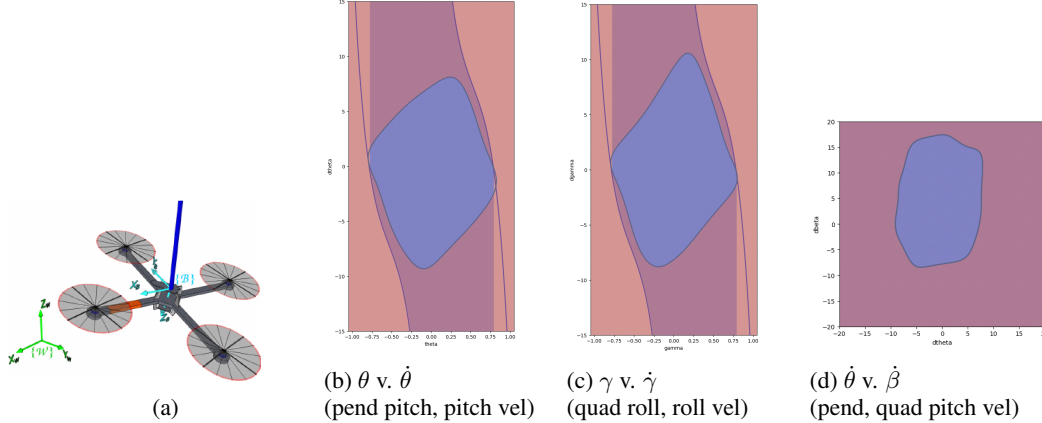


Figure 2: (Left) quadcopter-pendulum system (image from [24]). (Others) Various axis-aligned 2D slices of the 10D safe set (learned in blue, baseline in purple). Note that for each slice, the unvisualized states have been set to 0.

200 Our system is a pendulum on top of a quadcopter (Fig. 2). This is a coupled system, with the dynam-
 201 ics of both components found by the Euler-Lagrangian method [25]. The states are quadcopter posi-
 202 tion and roll-pitch-yaw orientation (x, y, z and γ, α, β) and roll-pitch pendulum orientation (ϕ_p, θ_p),
 203 as well as the first derivatives of these states. The inputs are thrust and torque ($F, \tau_\gamma, \tau_\beta, \tau_\alpha$), which
 204 are limited to a bounded convex polytope set. See the Appendix for the system dynamics. The safety
 205 specification is to prevent the pendulum from tipping and the quadcopter from rolling:

$$\rho = \gamma^2 + \beta^2 + \delta_p^2 - (\pi/4)^2 \quad (8)$$

206 , where $\delta_p = \arccos(\cos(\phi_p) \cos(\theta_p))$ is the pendulum’s angle from the vertical. Since the quad-
 207 copter position does not impact safety and the position dynamics can be decoupled, we exclude
 208 position and consider the resulting 10D state, 4D input system. Finally, we note that we have a
 209 stabilization-type safety problem (the goal is to limit deviation from the equilibrium, $x_e = \vec{0}$), so we
 210 use Design 1 for our CBF.

211 Next, we propose metrics to answer each of our experimental questions:

212 **M1.** % of non-saturating states on ∂S^* : we address Q1 by measuring how well we satisfy Eqn. 4.
 213 We compute this by uniformly sampling 10k states on ∂S^* and calculating what percentage of them
 214 are non-saturating. We also use these samples to compute the mean and variance of the severity of
 215 saturation, $\mathbb{E}[\mathcal{L}(\theta, x)]$ and $\sigma[\mathcal{L}(\theta, x)]$. To compute the severity of the worst saturation, $\mathcal{L}(\theta, x^*)$, we
 216 apply our critic and allow it to use more samples and computation time than during CBF training.

217 **M2.** % of simulated rollouts that are FI: Q2 considers the in-the-loop control performance of our
 218 learned CBF. For this metric, we initialize 5k rollouts uniformly randomly inside S^* , simulate their
 219 trajectories under the safe controller (CBF-QP) until just after they reach the boundary. Then,
 220 we record whether the system exited or remained inside S^* after arriving at the boundary. The
 221 value of this metric depends on the choice of nominal controller, k_{nom} . We try two different ones:
 222 $k_{nom,unact}(x) = 0$, which yields an unactuated system, and $k_{nom,LQR}$, which is a linear-quadratic
 223 regulator (LQR) stabilizing controller.

224 We choose a non-neural CBF baseline. For nonlinear functions, the usual approach to synthesis is
 225 to hand-select a function form and then optimize the parameters. We select a simple but often used
 226 form [6, 11]:

$$\rho_{baseline}^* = (\gamma^2 + \beta^2 + \delta_p^2)^{a_1} - (\pi/4)^{2 \cdot a_1} + a_2 \quad (9)$$

227 , where $a_1 > 0, a_2 \geq 0$ are parameters that are optimized using our formulated loss (Saturation
 228 risk). See Appendix for details of this optimization.

229 **Training details.** The learning framework and metrics were implemented using Python and Py-
 230 Torch [26]. Training took around 11 hours on a single NVIDIA GeForce RTX 2080 Ti GPU.

method	Q1			Q2: % FI rollouts	
	% non-sat. states	mean, std dev of sat.	worst sat.	$k_{nom,unact}$	$k_{nom,LQR}$
Baseline	78.68	4.99 ± 3.78	29.50	79.53	49.50
Ours	99.00	1.75 ± 2.40	15.19	98.90	97.00

Table 1: Metrics computed on our CBF and baseline CBF. The “mean, std dev of sat.” is $\mathbb{E}[\mathcal{L}(\theta, x)] \pm \sigma[\mathcal{L}(\theta, x)]$ and “worst sat.” is $\mathcal{L}(\theta, x^*)$.

231 **Discussion.** As we can see in Table 1, for our learned CBF, almost 100% of the boundary states
232 are non-saturating and almost 100% of the rollouts are FI across the nominal policies. This shows
233 that our neural CBF and learning framework are an effective combination and capable of handling
234 systems of high complexity. We are not able to attain 100% non-saturating states and FI rollouts,
235 which is either due to suboptimality of training or limitations of the function class, as NNs are only
236 universal approximators when their size is taken to the infinite. The optimized non-neural CBF does
237 not do as well: only 80% of the boundary states are non-saturating and around 50 – 80% of rollouts
238 are FI. For our learned CBF, the severity of saturation at the saturating states is not negligible, but is
239 it still low, and the worst saturation is large, but rarely encountered.

240 Visualizing slices of the safe set can provide deeper insight into the results of training (Fig. 2). Recall
241 that a safe set contains only states that are *recoverable* from danger, given our input limits. We
242 observe that our CBF has diamond-shaped safe sets in slices B and C, which makes sense because
243 small angles can still be recoverable at larger speeds. Additionally, there should be a maximum safe
244 angular speed for quadcopter and pendulum; our CBF learns this, but the baseline does not. Next,
245 our CBF has a roughly rectangular set in slice D, indicating that $\dot{\theta}, \dot{\beta}$ (pendulum, quadcopter pitch
246 velocity) should both be limited to an interval. This makes sense, since higher angular velocities are
247 certainly harder to pull back from. We also observe a larger range for the quadcopter’s pitch than
248 the pendulum’s, which makes sense because the quadcopter is directly actuated and the pendulum is
249 only indirectly actuated. On the other hand, the baseline does not restrict $\dot{\theta}, \dot{\beta}$ in slice D. Due to the
250 functional form of the baseline CBF, $\dot{\theta}, \dot{\beta}$ are only restricted when $\dot{\theta}\theta > 0$ or $\dot{\beta}\beta > 0$ (that is, when
251 angular velocity is acting to destabilize). However, this safety criterion is too lenient, since large
252 angular velocities that are currently swinging the system to the origin can cause overshooting and
253 toppling shortly after. Overall, we can see that the non-neural CBF does not have the right function
254 form. In general, it can be hard to guess what the right form might be for a system of this complexity.
255 However, our algorithmic approach, combined with the NN functional representation, removes the
256 need for guessing.

257 **Limitations and future work** One limitation of our results is that we had to do some minor feature
258 engineering on the inputs to the neural CBF. Namely, we changed the pendulum’s angular velocity
259 to a linear velocity and the quadcopter’s angular velocity to the linear velocity of its vertical body
260 axis. (Note that the dynamics of these states is still nonlinear). Before this change, learning was
261 difficult. This might be because a valid CBF is a *highly* nonlinear function (possibly involving some
262 trigonometric terms) of the original states, and such functions are challenging for NNs to learn. This
263 problem might be able to be solved otherwise with a different network architecture. However, we
264 note that after we made this adjustment, the rest of the synthesis required no human intervention.
265 Additionally, in this work we have assumed known, deterministic systems, but real-world systems
266 are often unknown and stochastic. An interesting future direction is extending our work with more
267 realistic modeling assumptions.

268 5 Conclusion

269 In summary, we proposed a framework for facilitating CBF synthesis under input limits. Thanks
270 to our neural CBF representation and our effective and efficient learning framework, our method
271 scales to higher-dimensional nonlinear systems. We learned a virtually non-saturating CBF on one
272 such system, the quadcopter-pendulum. We hope that this safe control tool will makes CBFs more
273 accessible and of practical value to roboticists.

References

- 301
- 302 [1] B. Xu and K. Sreenath. Safe teleoperation of dynamic uavs through control barrier functions. In
303 *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 7848–7855.
304 IEEE, 2018.
- 305 [2] A. Singletary, K. Klingebiel, J. Bourne, A. Browning, P. Tokumaru, and A. Ames. Comparative
306 analysis of control barrier functions and artificial potential fields for obstacle avoidance. In
307 *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages
308 8129–8136. IEEE, 2020.
- 309 [3] R. Liu, R. Chen, Y. Sun, Y. Zhao, and C. Liu. Jerk-bounded position controller with real-time
310 task modification for interactive industrial robots. 2022.
- 311 [4] A. D. Ames, E. A. Cousineau, and M. J. Powell. Dynamically stable bipedal robotic walking
312 with nao via human-inspired hybrid zero dynamics. In *Proceedings of the 15th ACM interna-*
313 *tional conference on Hybrid Systems: Computation and Control*, pages 135–144, 2012.
- 314 [5] C. Liu and M. Tomizuka. Control in a safe set: Addressing safety in human-robot interactions.
315 In *Dynamic Systems and Control Conference*, volume 46209, page V003T42A003. American
316 Society of Mechanical Engineers, 2014.
- 317 [6] W. Zhao, T. He, and C. Liu. Model-free safe control for zero-violation reinforcement learning.
318 In *5th Annual Conference on Robot Learning*, 2021.
- 319 [7] A. D. Ames, J. W. Grizzle, and P. Tabuada. Control barrier function based quadratic programs
320 with application to adaptive cruise control. In *53rd IEEE Conference on Decision and Control*,
321 pages 6271–6278. IEEE, 2014.
- 322 [8] Y. Lyu, W. Luo, and J. M. Dolan. Probabilistic safety-assured adaptive merging control for
323 autonomous vehicles. In *2021 IEEE International Conference on Robotics and Automation*
324 *(ICRA)*, pages 10764–10770. IEEE, 2021.
- 325 [9] W. S. Cortez and D. V. Dimarogonas. Safe-by-design control for euler-lagrange systems. *arXiv*
326 *preprint arXiv:2009.03767*, 2020.
- 327 [10] W. S. Cortez, X. Tan, and D. V. Dimarogonas. A robust, multiple control barrier function
328 framework for input constrained systems. *IEEE Control Systems Letters*, 6:1742–1747, 2021.
- 329 [11] T. Wei and C. Liu. Safe control with neural network dynamic models. *arXiv preprint*
330 *arXiv:2110.01110*, 2021.
- 331 [12] A. Clark. Verification and synthesis of control barrier functions. *arXiv preprint*
332 *arXiv:2104.14001*, 2021.
- 333 [13] S. Bansal, M. Chen, S. Herbert, and C. J. Tomlin. Hamilton-jacobi reachability: A brief
334 overview and recent advances. In *2017 IEEE 56th Annual Conference on Decision and Control*
335 *(CDC)*, pages 2242–2253. IEEE, 2017.
- 336 [14] S. Bansal and C. J. Tomlin. Deepreach: A deep learning approach to high-dimensional reach-
337 ability. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages
338 1817–1824. IEEE, 2021.
- 339 [15] M. Chen, S. Herbert, and C. J. Tomlin. Fast reachable set approximations via state decoupling
340 disturbances. *arXiv preprint arXiv:1603.05205*, 2016.
- 341 [16] M. Chen, S. L. Herbert, M. S. Vashishtha, S. Bansal, and C. J. Tomlin. Decomposition of
342 reachable sets and tubes for a class of nonlinear systems. *IEEE Transactions on Automatic*
343 *Control*, 63(11):3675–3688, 2018.

- 344 [17] S. M. Richards, F. Berkenkamp, and A. Krause. The lyapunov neural network: Adaptive
345 stability certification for safe learning of dynamical systems. In *Conference on Robot Learning*,
346 pages 466–476. PMLR, 2018.
- 347 [18] H. Ravanbakhsh and S. Sankaranarayanan. Learning control lyapunov functions from coun-
348 terexamples and demonstrations. *Autonomous Robots*, 43(2):275–307, 2019.
- 349 [19] Y.-C. Chang, N. Roohi, and S. Gao. Neural lyapunov control. *Advances in neural information*
350 *processing systems*, 32, 2019.
- 351 [20] C. Dawson, Z. Qin, S. Gao, and C. Fan. Safe nonlinear control using robust neural lyapunov-
352 barrier functions. In *Conference on Robot Learning*, pages 1724–1735. PMLR, 2022.
- 353 [21] A. Robey, H. Hu, L. Lindemann, H. Zhang, D. V. Dimarogonas, S. Tu, and N. Matni. Learn-
354 ing control barrier functions from expert demonstrations. In *2020 59th IEEE Conference on*
355 *Decision and Control (CDC)*, pages 3717–3724. IEEE, 2020.
- 356 [22] N. M. Boffi, S. Tu, N. Matni, J.-J. E. Slotine, and V. Sindhvani. Learning stability certificates
357 from data. *arXiv preprint arXiv:2008.05952*, 2020.
- 358 [23] H. Narayanan and P. Niyogi. Sampling hypersurfaces through diffusion. In *Approximation,*
359 *Randomization and Combinatorial Optimization. Algorithms and Techniques*, pages 535–548.
360 Springer, 2008.
- 361 [24] R. Figueroa, A. Faust, P. Cruz, L. Tapia, and R. Fierro. Reinforcement learning for balancing
362 a flying inverted pendulum. In *Proceeding of the 11th World Congress on Intelligent Control*
363 *and Automation*, pages 1787–1793. IEEE, 2014.
- 364 [25] M. Hehn and R. D’Andrea. A flying inverted pendulum. In *2011 IEEE International Confer-*
365 *ence on Robotics and Automation*, pages 763–770. IEEE, 2011.
- 366 [26] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin,
367 N. Gimselshein, L. Antiga, et al. Pytorch: An imperative style, high-performance deep learning
368 library. *Advances in neural information processing systems*, 32, 2019.