

Learning to Provably Satisfy High Relative Degree Constraints for Black-Box Systems*

Jean-Baptiste Bouvier, Kartik Nagpal and Negar Mehr

Abstract—In this paper, we develop a method for learning a control policy guaranteed to satisfy an affine state constraint of high relative degree in closed loop with a black-box system. Previous reinforcement learning (RL) approaches to satisfy safety constraints either require access to the system model, or assume control affine dynamics, or only discourage violations with reward shaping. Only recently have these issues been addressed with POLICEd RL, which guarantees constraint satisfaction for black-box systems. However, this previous work can only enforce constraints of relative degree 1. To address this gap, we build a novel RL algorithm explicitly designed to enforce an affine state constraint of high relative degree in closed loop with a black-box control system. Our key insight is to make the learned policy be affine around the unsafe set and to use this affine region to dissipate the inertia of the high relative degree constraint. We prove that such policies guarantee constraint satisfaction for deterministic systems and are agnostic to the choice of the RL training algorithm. Our results demonstrate the capacity of our approach to enforce hard constraints in the Gym inverted pendulum and on a space shuttle landing simulation. Website: <https://iconlab.negarmehr.com/CDC-POLICEd-RL/>

I. INTRODUCTION

The lack of safety guarantees in reinforcement learning (RL) has been impeding its wide deployment in real-world settings [1]. Safety in RL is traditionally captured by state constraints preventing the system from entering unsafe regions [2]. This issue has been investigated by numerous approaches and most commonly under the framework of constrained Markov decision processes (CMDPs) [3]. CMDPs only *encourage* policies to respect safety constraints by penalizing the expected violations [4]; however, they do not provide any satisfaction guarantees [2]. For safety-critical tasks, such as autonomous driving or human-robot interactions, safety guarantees are primordial and require the learned policy to maintain constraint respect.

A few RL attempts at learning provably safe policies have involved control barrier functions (CBFs) [5], backward reachable sets [6], and projection of control inputs onto safe sets [7]. However, all these methods require precise knowledge of the system dynamics, which is usually not available in RL. To circumvent this issue and study black-box systems without an analytical model of the dynamics, the common approach has been to learn safety certificates [8], [9]. Yet, the formal safety guarantees of these methods hinge upon the quality of their CBF approximations. More

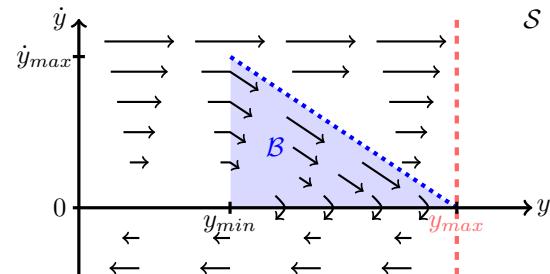


Fig. 1: Phase portrait of constrained output y illustrating our High Relative Degree POLICEd RL method on a system of relative degree 2. To prevent states from violating constraint $y \leq y_{max}$ (red dashed line), our policy guarantees that trajectories entering buffer region \mathcal{B} (blue) cannot leave it through its upper bound (blue dotted line). Our policy makes \dot{y} sufficiently negative in buffer \mathcal{B} to bring \dot{y} to 0 in all trajectories entering \mathcal{B} . Once $\dot{y} < 0$, trajectories cannot approach the constraint. Due to the states' inertia, it is physically impossible to prevent all constraint violations. For instance, $y = y_{max}$, $\dot{y} > 0$ will yield $y > y_{max}$ at the next timestep. Hence, we only aim at guaranteeing the safety of trajectories entering buffer \mathcal{B} .

reliable safety guarantees have been established by recent work [10], whose POLICEd RL approach designs a repulsive buffer to enforce constraint satisfaction in closed-loop with a black-box system. However, work [10] along with most other safe RL works such as [7], [8], [9] are limited to constraints of relative degree 1. In contrast, our approach enforces inviolable *constraints of high relative degree* in closed-loop with a learned control policy while exclusively using a *black-box model* of the system dynamics.

The higher the *relative degree of a constraint*, the more inertia it has, and the more challenging its satisfaction is [11]. We propose a backstepping inspired approach, which is compatible with systems non-affine in control and with black-box systems contrary to these CBF methods.

To learn our safe controller, we draw inspiration from the POLICEd RL method of [10] and transform the state space surrounding the affine state constraint into a buffer region that cannot be crossed. We overcome the major limitation of [10] by extending POLICEd RL to constraints of high relative degree. To dissipate the inertia of this high relative degree constraint, our key insight is to extend the buffer into the dimensions of the unactuated derivatives of the state constraint. While this task appears arduous, these derivatives are usually accessible since related to the states. In this buffer region, we train the policy to dissipate the state's inertia to progressively slow its progression towards the constraint, as illustrated in Fig. 1. This controller guarantees that trajectories entering the buffer do not violate the constraint. Since inertia cannot be dissipated instantly, some constraint violations are physically impossible to prevent. Inspired by

*This work is supported by the National Science Foundation, under grants ECCS-2145134 CAREER Award, CNS-2218759, and CCF-2211542.

Jean-Baptiste Bouvier, Kartik Nagpal and Negar Mehr are with the Department of Mechanical Engineering, University of California Berkeley, Berkeley, CA 94709, USA {bouvier3, kartiknagpal, negar}@berkeley.edu

[10], to easily verify the dissipative character of the controller in the buffer, we use the POLICE algorithm [12] to generate an affine policy over the buffer region.

In summary, our contributions in this work are as follows.

- 1) We introduce High Relative Degree POLICEd RL, a novel RL framework to guarantee satisfaction of an affine state constraint of high relative degree using a black-box model of the system in closed-loop with a trained policy.
- 2) We provide comprehensive proof, and we detail our method to evaluate our trained policy while directly guaranteeing constraint satisfaction.
- 3) We demonstrate the safety guarantees of our approach in a number of simulation studies involving an inverted pendulum and a space shuttle landing.

The remainder of this work is organized as follows. In Section II, we introduce our framework. In Section III, we establish the theoretical guarantees of our approach in enforcing the satisfaction of a high relative degree constraint. In Section IV, we illustrate our method on the Gym inverted pendulum and on a space shuttle landing scenario. More details can be found at <https://arxiv.org/abs/2407.20456>

Notation: We denote the positive integer interval from $a \in \mathbb{N}$ to $b \in \mathbb{N}$ inclusive by $\llbracket a, b \rrbracket$. We denote the component $i \in \llbracket 1, n \rrbracket$ of a vector $x \in \mathbb{R}^n$ by x_i and the vector composed of components i to $j > i$ by $x_{i:j}$. The set of nonnegative real numbers is \mathbb{R}^+ . We denote the k^{th} time derivative of a function y by $y^{(k)} = \frac{d^k y}{dt^k}$. If $x, y \in \mathbb{R}^n$, then $x \leq y$ denotes the element-wise inequalities $x_i \leq y_i$ for all $i \in \llbracket 1, n \rrbracket$.

II. FRAMEWORK

We consider a black-box deterministic system

$$\dot{x}(t) = f(x(t), u(t)), \quad u(t) \in \mathcal{U}, \quad x(0) \in \mathcal{X}, \quad (1)$$

with state space $\mathcal{X} \subseteq \mathbb{R}^n$ and admissible control set $\mathcal{U} \subseteq \mathbb{R}^m$. We consider dynamics (1) to be an implicit black-box, meaning that we can evaluate f but we do not have any explicit knowledge or analytical form of f . This is similar to the online RL setting where f is a simulator or a robot.

We assume that the system safety constraint is captured by a single affine inequality on output

$$y(t) := Cx(t) \leq y_{max}, \quad \text{for all } t \geq 0, \quad (2)$$

with $C \in \mathbb{R}^{1 \times n}$ and $y_{max} \in \mathbb{R}$. We model our deterministic feedback policy $u(t) = \pi_\theta(x(t)) \in \mathcal{U}$ by a deep neural network parameterized by θ . Our objective is to train policy π_θ to respect constraint (2) and maximize the following expected discounted reward

$$\max_{\theta} \mathcal{G}(\pi_\theta) := \mathbb{E}_{x_0 \sim \rho_0} \int_0^\infty \gamma^t R(x(t), \pi_\theta(x(t))) dt \quad \text{s.t. (2)}, \quad (3)$$

where $\gamma \in (0, 1]$ is a discount factor, R a reward function, and ρ_0 the distribution of initial states. The only stochasticity in our setting comes from the initial state sampling $x_0 \sim \rho_0$. We emphasize that constraint (2) is a *hard constraint* to be respected at all times. Contrary to previous work [10], we assume that constraint (2) has a *relative degree* at least 2.

Definition 1. The relative degree r of output y (2) for dynamics (1) is the order of its input-output relationship, i.e., $r := \min \{p \in \mathbb{N} : \frac{\partial}{\partial u} \frac{\partial^p y}{\partial t^p}(t) \neq 0 \text{ for all } x \in \mathcal{X}\}$ [13].

In simpler words, the relative degree is the minimal number of times output y has to be differentiated until control input u appears. As argued in [14], relative degree r can be obtained by first-order principles without the knowledge of dynamics f . Hence, knowing r is compatible with our black-box model of f .

Having $r \geq 2$ means that u does not appear in the expression of \dot{y} . Thus, a change in control input will not immediately modify y . We follow [11] and refer to the unactuated derivatives of y ($y^{(k)}$ for $k \in \llbracket 1, r-1 \rrbracket$) as generalized inertia, by analogy to inertia in kinematic systems. To enforce constraint (2), we need to dissipate this generalized inertia before reaching constraint line $y = y_{max}$. To easily assess this generalized inertia, we make the following assumption.

Assumption 1. There exists an invertible map T between state $x \in \mathbb{R}^n$ and $s \in \mathbb{R}^n$, whose first r components are $s_1 = y$, $s_2 = \dot{y}$, ..., $s_r = y^{(r-1)}$, where y is output (2). Transformation $T(x) = s$ gives rise to an equivalent state space $\mathcal{S} := T(\mathcal{X})$.

Note that this is a rather mild assumption. Indeed, a transformation $s = T(x)$ always exists since $y = Cx$ and thus $s_{k+2} = y^{(k+1)} = C \frac{\partial^k \dot{x}}{\partial t^k} = C \frac{\partial^k f}{\partial t^k}(x, u)$. Assumption 1 is required for the invertibility of T and the fact that T can be determined without knowledge of black-box dynamics f . Following Assumption 1, we now have two equivalent state representations: x denotes the original state of system (1), while s denotes the transformed state composed of the iterated derivatives of constrained output y .

We can now formally define our problem of interest.

Problem 1. Given:

- 1) black-box control system (1);
- 2) state space $\mathcal{X} \subseteq \mathbb{R}^n$;
- 3) admissible input set $\mathcal{U} \subseteq \mathbb{R}^m$;
- 4) affine constraint (2) of relative degree $r \geq 2$;
- 5) neural network policy $\pi_\theta(x)$ parameterized by θ ;
- 6) invertible transformation T of Assumption 1;

Our goal is to solve $\theta^* = \arg \max_{\theta} \mathcal{G}(\pi_\theta)$ s.t. (1) and (2).

III. CONSTRAINED REINFORCEMENT LEARNING

We will solve Problem 1 by designing a buffer preventing trajectories from breaching the constraint, as shown in Fig. 1.

A. Buffer Design

We formalize the concept of Fig. 1 and build buffer \mathcal{B} as the state space region where our controller will dissipate generalized inertia $\dot{y}, \dots, y^{(r-1)}$ to prevent violation of constraint (2). To adapt \mathcal{B} to this task, we design it in state space \mathcal{S} , because Assumption 1 states that the coordinates of \mathcal{S} are the generalized inertia components $y^{(k)}$.

Assume we can choose $s \in \mathcal{B}$ and let us investigate what bounds the components of s should satisfy to remain in \mathcal{B} .

Following Assumption 1, the first component of s is $s_1 = y$, and thus should satisfy $s_1 \leq y_{max}$ to respect constraint (2). We choose a lower bound for s_1 as $y_{min} < y_{max}$.

Following Assumption 1, the second component of s is $s_2 = \dot{y}$. To maintain $y \leq y_{max}$, we need $\dot{y} \leq 0$ when $y = y_{max}$. Requiring $\dot{y} \leq 0$ for all $s \in \mathcal{B}$ is the approach of [10] but restricts \mathcal{B} to only include states already moving away from upper bound y_{max} , whereas we want to slow down and stop trajectories going towards y_{max} . Thus, we must allow states with $\dot{y} > 0$ in \mathcal{B} . Let $\dot{y}_{max} > 0$ be the maximal velocity in \mathcal{B} . Our controller will later require \mathcal{B} to be a polytope. Thus, we naturally define the upper bound on $\dot{y} = s_2$ as

$$s_2^{max} := \beta(y_{max} - y) \quad \text{with} \quad \beta := \frac{\dot{y}_{max}}{y_{max} - y_{min}}, \quad (4)$$

so that $s_2^{max}(y) = \dot{y}_{max}$ when $y = y_{min}$ and $s_2^{max} = 0$ when $y = y_{max}$, as illustrated in Fig. 1. We choose a lower bound $s_2^{min} \leq 0$ so that $s_2^{min} \leq s_2^{max}(y)$ for all y . Note that

$$s_2 = \dot{y} \leq s_2^{max}(y) = \beta(y_{max} - y) \quad (5)$$

is a differential inequality on y that we designed to maintain $y \leq y_{max}$. To enforce (5) we need a control input, but only $y^{(r)}$ is actuated. Our key idea is then to make our controller enforce a differential inequality on $y^{(r)}$, whose iterated integrations will lead to (5).

Working backwards, we differentiate (5) into $\ddot{y} \leq -\beta\dot{y}$, i.e., $s_3 \leq -\beta s_2$. Thus, we choose bounds $s_3^{min} < s_3^{max}(s) := -\beta s_2$. Iterating this process until $y^{(r)}$ leads to lower and upper bounds \underline{b} and \bar{b} on the first r components of $s \in \mathcal{B}$. Then, $s_{1:r} \in [\underline{b}, \bar{b}(s)]$ element-wise with

$$\underline{b} := [y_{min}, \quad s_2^{min}, \quad s_3^{min}, \dots, \quad s_r^{min}], \quad (6)$$

$$\bar{b}(s) := [y_{max}, \beta(y_{max} - s_1), -\beta s_2, \dots, -\beta s_{r-1}]. \quad (7)$$

The remaining $n-r$ components of $s \in \mathcal{B}$ are not derivatives of y and thus are not involved in the constraint enforcement process. As mentioned previously, we will need \mathcal{B} to be a polytope, hence we choose to bound the last $n-r$ components of $s \in \mathcal{B}$ by a polytope $\mathcal{P} \subset \mathbb{R}^{n-r}$ so that

$$\mathcal{B} := \{s \in \mathcal{S} : s_{1:r} \in [\underline{b}, \bar{b}(s)], s_{r+1:n} \in \mathcal{P}\}. \quad (8)$$

Note that the bounds we just derived only delimit region \mathcal{B} in \mathcal{S} , but without adequate control input, trajectories will not respect these bounds. Similarly, the differential inequalities obtained above only reflect the desired dynamics that we want to enforce with our controller. On the other hand, bounds $s_{1:r} \geq \underline{b}$ and $s_{r+1:n} \in \mathcal{P}$ will not be specifically enforced by the controller, but should be designed to encompass all trajectories to be safeguarded from y_{max} .

By design, buffer \mathcal{B} of (8) is then a compact convex polytope with a finite number N of vertices gathered in the set $\mathcal{V}(\mathcal{B}) := \{v^1, \dots, v^N\}$.

B. Controller Design

Let us now design a controller to maintain trajectories in buffer \mathcal{B} . Inspired by [10], we model our control policy with a POLICEd neural network $\mu_\theta := \pi_\theta \circ T : \mathcal{S} \rightarrow \mathcal{U}$, with continuous piecewise affine activation functions such as

ReLU [12]. This POLICEd architecture allows us to make the outputs of μ_θ affine over a polytopic region of state space \mathcal{S} , which we chose to be \mathcal{B} . Then, there exist matrices $D_\theta \in \mathbb{R}^{m \times n}$ and $e_\theta \in \mathbb{R}^m$ such that

$$\mu_\theta(s) = D_\theta s + e_\theta \quad \text{for all } s \in \mathcal{B}. \quad (9)$$

Since buffer \mathcal{B} and policy μ_θ are both in space \mathcal{S} and not \mathcal{X} , we need to calculate the state dynamics in \mathcal{S} . State $s = T(x)$ following controller μ_θ satisfies

$$\dot{s} = \frac{\partial T}{\partial t}(x) = \frac{\partial T}{\partial x} \frac{\partial x}{\partial t} = \frac{\partial T}{\partial x}(x) f\left(x, \mu_\theta(T(x))\right).$$

Then, by defining the map

$$\tilde{f}(s; \mu_\theta) := \frac{\partial T}{\partial x}(T^{-1}(s)) f(T^{-1}(s), \mu_\theta(s)),$$

we can write $\dot{s}(t) = \tilde{f}(s(t); \mu_\theta)$. We are mostly interested in the dynamics of the actuated derivative of output y :

$$y^{(r)}(t) = \frac{\partial y^{(r-1)}}{\partial t}(t) = \frac{\partial s_r}{\partial t}(t) = \tilde{f}_r(s(t); \mu_\theta), \quad (10)$$

where s_r and \tilde{f}_r denote the r^{th} component of s and \tilde{f} respectively. If dynamics (10) were known and affine, their coupling with affine policy μ_θ on \mathcal{B} would lead to a simple constraint enforcement process. However, dynamics (10) are a black-box and possibly nonlinear. We will thus use an affine approximation of (10) inside using the following definition.

Definition 2. An approximation measure ε of dynamics (10) in buffer (8) is any $\varepsilon \geq 0$ for which there exists any matrices $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times m}$, and $c \in \mathbb{R}^n$ such that for all $s \in \mathcal{B}$

$$|\tilde{f}_r(s; \mu_\theta) - C(As + B\mu_\theta(s) + c)| \leq \varepsilon. \quad (11)$$

Intuitively, the value of ε quantifies how far from affine is function \tilde{f}_r over buffer \mathcal{B} . Having access to map T , controller μ_θ and to a black-box model of f , we can evaluate \tilde{f}_r and compute ε using linear least square approximation [15]. Since ε is estimated from data, it might not verify (11) for some $s \in \mathcal{B}$ absent from the dataset. To satisfy Definition 2, we need to over-approximate ε since any upper bound will verify (11). With such an upper bound, we will guarantee the satisfaction of constraint (2) with actual dynamics (1).

We now establish our central result demonstrating how to guarantee satisfaction of constraint (2) by black-box environment (1) armed only with an approximation measure ε and without knowing A , B , c , f , or \tilde{f}_r .

Theorem 1. Assume that for some approximation measure ε , dissipation condition

$$\tilde{f}_r(v; \mu_\theta) \leq -2\varepsilon - \beta v_r, \quad (12)$$

holds for all $v \in \mathcal{V}(\mathcal{B})$, where v_r is the r^{th} component of v and β comes from (4). If a trajectory s steered by μ_θ verifies

$$s_{1:r}(t_0) < \bar{b}(s(t_0)) \quad (13)$$

for some $t_0 \geq 0$, and satisfies

$$s_{1:r}(t) \geq \underline{b} \quad \text{and} \quad s_{r+1:n}(t) \in \mathcal{P} \quad (14)$$

for all $t \in [t_0, t_1]$, then $s_{1:r}(t) < \bar{b}(s(t))$ for all $t \in [t_0, t_1]$.

In simpler words, Theorem 1 guarantees that trajectories entering buffer \mathcal{B} below upper bound \bar{b} cannot exit \mathcal{B} through \bar{b} as long as dissipation condition (12) is satisfied. Theorem 1 generates the bent arrows of the flow illustrated in Fig. 1 which prevent trajectories from violating constraint (2). The major strength of our approach is that dissipation condition (12) only needs to be enforced at the vertices of \mathcal{B} and thus does not require knowledge of f or \tilde{f}_r .

Proof of Theorem 1. The intuition behind this proof is to use the convexity of buffer \mathcal{B} and affine approximation (11) to extend condition (12) to the entire \mathcal{B} . By combining this condition with the specific design of upper bound (7), we can derive bounds on the output derivatives $y, \dot{y}, \dots, y^{(r-1)}$ and show that they cannot cross upper bound \bar{b} .

We divide the proof into three lemmas. First, we show in Lemma 1 that condition (12) yields $\dot{s}_r \leq -\beta s_r$ for all s in \mathcal{B} . This condition combined in Lemma 2 with (14) yields the differential inequalities to be respected by $y, \dot{y}, \dots, y^{(r-1)}$ as long as trajectory s remains in \mathcal{B} . In Lemma 3 these differential equations are then paired with initial conditions (13) to obtain $s_{1:r}(t) < \bar{b}(s(t))$ for all $t \in [t_0, t_1]$. \square

C. Supporting Lemmata

Lemma 1. *If for some approximation measure ε , condition (12) holds for all $v \in \mathcal{V}(\mathcal{B})$, then controller μ_θ yields*

$$\dot{s}_r(t) \leq -\beta s_r(t), \quad \text{i.e.,} \quad y^{(r)}(t) \leq -\beta y^{(r-1)}(t), \quad (15)$$

for all $s(t) \in \mathcal{B}$.

Proof. Since ε is an approximation measure, there exist A , B and c verifying (11) which we evaluate at $s = v \in \mathcal{V}(\mathcal{B})$

$$\begin{aligned} C(Av + B\mu_\theta(v) + c) \\ \leq |C(Av + B\mu_\theta(v) + c) - \tilde{f}_r(v; \mu_\theta)| + \tilde{f}_r(v; \mu_\theta) \\ \leq \varepsilon - 2\varepsilon - \beta v_r, \end{aligned} \quad (16)$$

where the first inequality is a triangular inequality, the second follows from (11) and (12). Using the convexity of polytope \mathcal{B} of vertices $\mathcal{V}(\mathcal{B}) = \{v^1, \dots, v^N\}$, for any $s \in \mathcal{B}$, there exist $\alpha^1, \dots, \alpha^N \in \mathbb{R}^+$ such that $\sum_{k=1}^N \alpha^k = 1$ and $s = \sum_{k=1}^N \alpha^k v^k$. Controller (9) applied at $s \in \mathcal{B}$ yields

$$\begin{aligned} C(As + B\mu_\theta(s) + c) &= C(As + B(D_\theta s + e_\theta) + c) \\ &= C(A + BD_\theta)s + C(Be_\theta + c) \\ &= C(A + BD_\theta) \sum_{k=1}^N \alpha^k v^k + C(Be_\theta + c) \sum_{k=1}^N \alpha^k \\ &= \sum_{k=1}^N \alpha^k C((A + BD_\theta)v^k + Be_\theta + c) \\ &= \sum_{k=1}^N \alpha^k C(Av^k + B\mu_\theta(v^k) + c) \\ &\leq \sum_{k=1}^N \alpha^k (-\varepsilon - \beta v_r) = -\varepsilon \sum_{k=1}^N \alpha^k - \beta \sum_{k=1}^N \alpha^k v_r \\ &= -\varepsilon - \beta s_r, \end{aligned} \quad (17)$$

where the only inequality comes from (16) on each vertex v^k and the last equality stems from the linear decomposition of component r of state s between component r of vertices v^k . For any state $s \in \mathcal{B}$, (10) yields

$$\begin{aligned} y^{(r)} &= \tilde{f}_r(s, \mu_\theta) \\ &\leq |\tilde{f}_r(s, \mu_\theta) - C(As + B\mu_\theta(s) + c)| + C(As + B\mu_\theta(s) + c) \\ &\leq \varepsilon - \varepsilon - \beta s_r = -\beta y^{(r-1)}, \end{aligned}$$

where we first use the triangular inequality, then (11) and (17), and the last equality comes from the definition of state s in Assumption 1. \square

Lemma 1 uses the convexity of \mathcal{B} and affine approximation (11) to extend (12), valid *only* at the vertices of \mathcal{B} , into (15), valid all over \mathcal{B} . Without the POLICE algorithm [12], μ_θ would not be affine over \mathcal{B} , and dissipation condition (12) would need to be enforced everywhere on the buffer at a prohibitive computational cost.

Lemma 2. *If (15) holds for all $s \in \mathcal{B}$, $s(t_0) \in \mathcal{B}$, and (14) holds for all $t \in [t_0, t_1]$ for some $t_1 > t_0 \geq 0$, then*

$$y(t) \leq (y(t_0) - y_{max})e^{-\beta(t-t_0)} + y_{max} \quad (18)$$

and

$$y^{(k)}(t) \leq y^{(k)}(t_0)e^{-\beta(t-t_0)} \quad (19)$$

for all $k \in [1, r-1]$ and all $t \in [t_0, t_1]$.

Proof. We apply the comparison lemma of [16] to differential inequality (15), which yields $y^{(r-1)}(t) \leq y^{(r-1)}(t_0)e^{-\beta(t-t_0)}$.

Initial condition $s(t_0) \in \mathcal{B}$ yields

$$s_r(t_0) = y^{(r-1)}(t_0) \leq -\beta s_{r-1}(t_0) = -\beta y^{(r-2)}(t_0).$$

Define function $g(t) := y^{(r-1)}(t) + \beta y^{(r-2)}(t)$. Then, (15) is equivalent to $\dot{g}(t) \leq 0$ and our initial condition is $g(t_0) \leq 0$. Therefore, $g(t) \leq 0$ for all $t \in [t_0, t_1]$, i.e., $y^{(r-1)}(t) \leq -\beta y^{(r-2)}(t)$. Using the comparison lemma of [16], we can solve this differential inequality and obtain $y^{(r-2)}(t) \leq y^{(r-2)}(t_0)e^{-\beta(t-t_0)}$.

We can iterate this process for $k \in \{r-3, \dots, 1\}$ and obtain $y^{(k+1)}(t) \leq -\beta y^{(k)}(t)$, which yields $y^{(k)}(t) \leq y^{(k)}(t_0)e^{-\beta(t-t_0)}$ for all $t \in [t_0, t_1]$.

For $k = 1$, we thus have $\ddot{y}(t) + \beta \dot{y}(t) \leq 0$. Initial condition $s(t_0) \in \mathcal{B}$ yields $\dot{y}(t_0) \leq \beta(y_{max} - y(t_0))$, or equivalently $\dot{y}(t_0) + \beta y(t_0) \leq \beta y_{max}$. As previously, with $g(t) := \dot{y}(t) + \beta y(t)$, we have $\dot{g}(t) \leq 0$ and $g(t_0) \leq \beta y_{max}$. Thus, $g(t) \leq \beta y_{max}$, i.e., $\dot{y}(t) + \beta y(t) \leq \beta y_{max}$ for all $t \in [t_0, t_1]$. We solve this differential inequality using the comparison lemma of [16] and obtain (18) for all $t \in [t_0, t_1]$. \square

Lemma 2 used (15) and upper bound \bar{b} of (7) to obtain the differential equations verified by the output derivatives in \mathcal{B} . We will now use initial condition (13) to show that trajectories cannot leave \mathcal{B} through its upper bound \bar{b} .

Lemma 3. *If (15) holds for all $s \in \mathcal{B}$, (18) and (19) hold for all $t \in [t_0, t_1]$, then (13) implies $s_{1:r}(t) \leq \bar{b}(s(t))$ for all $t \in [t_0, t_1]$.*

Proof. Condition $s(t_0) < \bar{b}(s(t_0))$ yields $y(t_0) < y_{max}$. This initial condition combined with (18) leads to $y(t) < y_{max}$ for all $t \in [t_0, t_1]$, i.e., $s_1(t) < \bar{b}(s(t))_1$.

Initial condition $s_2(t_0) < \bar{b}(s(t_0))_2$ yields $\dot{y}(t_0) < \beta(y_{max} - y(t_0))$. Starting from (19) for $k = 1$, we have

$$\dot{y}(t) \leq \dot{y}(t_0)e^{-\beta(t-t_0)} < \beta(y_{max} - y(t_0))e^{-\beta(t-t_0)} \quad (20)$$

for all $t \in [t_0, t_1]$. Reorganizing (18) leads to $-y(t_0)e^{-\beta(t-t_0)} \leq -y(t) + y_{max}(1 - e^{-\beta(t-t_0)})$, which can be combined with (20) into

$$\begin{aligned} \dot{y}(t) &< \beta(y_{max}e^{-\beta(t-t_0)} - y(t) + y_{max}(1 - e^{-\beta(t-t_0)})) \\ &< \beta(y_{max} - y(t))e^{-\beta(t-t_0)}, \end{aligned}$$

i.e., $s_2(t) < \bar{b}(s(t))_2$ for all $t \in [t_0, t_1]$.

Similarly for $k \in [2, r-1]$, (19) combined with initial condition $y^{(k)}(t_0) < -\beta y^{(k-1)}(t_0)$ leads to

$$y^{(k)}(t) \leq y^{(k)}(t_0)e^{-\beta(t-t_0)} < -\beta y^{(k-1)}(t_0)e^{-\beta(t-t_0)} \quad (21)$$

for all $t \in [t_0, t_1]$. Reversing (19) at $k-1$ leads to $-y^{(k-1)}(t_0)e^{-\beta(t-t_0)} \leq -y^{(k-1)}(t)$, which can be combined with (21) into $y^{(k)}(t) < -\beta y^{(k-1)}(t)$, i.e., $s_{k+1}(t) < \bar{b}(s(t))_{k+1}$ for all $t \in [t_0, t_1]$. \square

Remark 1. Control set \mathcal{U} of (1) might prevent the existence of an admissible stabilizing policy μ_θ . That is why we use RL to find policy μ_θ and we verify its safety with Theorem 1.

IV. NUMERICAL SIMULATIONS

A. Gym Inverted Pendulum

We consider the Inverted Pendulum Gym environment with the MuJoCo dynamics engine. The environment state x is composed, in that order, of the cart position p , the pole angle θ , and their derivatives \dot{p} and $\dot{\theta}$. The objective within this environment is to maintain the pole close to the vertical, i.e., $|\theta| \leq 0.2$ rad. Let us focus on enforcing the upper constraint, $y := \theta \leq 0.2$ rad. This constraint has a relative degree $r = 2$ since the control input is the force exerted on the cart, which directly impacts $\ddot{\theta}$.

Given $y = \theta$, we define state $s := [y, \dot{y}, s_3, s_4] = T(x) := [\theta, \dot{\theta}, p, \dot{p}]$ by simply reordering state x .

Following Section III-A, we will now design buffer \mathcal{B} , whose architecture should help dissipate the inertia of trajectories arriving at $\theta = y_{min}$ with velocities $\dot{\theta} \leq \dot{y}_{max}$. We choose $y_{max} := 0.2$ rad, $y_{min} = 0.1$ rad, $s_2^{min} = 0$ rad/s and $\dot{y}_{max} = 1$ rad/s, and define buffer $\mathcal{B} := \{s \in \mathcal{S} : s_1 = y = \theta \in [0.1, 0.2], s_2 = \dot{y} = \dot{\theta} \in [0, 2 - 10\theta], s_3 = p \in [-0.9, 0.9], s_4 = \dot{p} \in [-1, 1]\}$. This choice of \mathcal{B} allows only $\dot{\theta} = 0$ rad/s when $\theta = 0.2$ rad, hence preventing θ in \mathcal{B} from growing past 0.2 rad.

Following Definition 2, we sample states in \mathcal{B} and perform a linear regression on $\ddot{\theta}$ to obtain an approximation measure $\varepsilon = 0.53$. We model controller μ_θ with a deep neural network trained to stabilize the pole at $\theta = 0$ in a reinforcement learning fashion using proximal policy optimization (PPO) [17]. We train two such policies, one being a standard multi-layer perceptron (MLP) to form a baseline, and the other

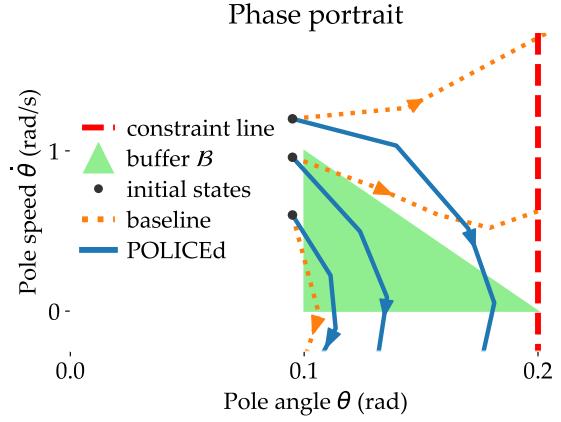


Fig. 2: Phase portrait of $(\theta, \dot{\theta})$ for the inverted pendulum. None of the POLICEd trajectories (blue) entering buffer \mathcal{B} (green) cross constraint line $\theta = 0.2$ rad (dashed red), whereas some of the baseline trajectories do (dotted orange). Our approach guarantees that a pole arriving at $\theta = 0.1$ rad with a velocity $\dot{\theta} < 1$ rad/s will satisfy $\theta \leq 0.2$ rad. We do not guarantee the safety of POLICEd trajectories not entering the buffer.

having POLICEd layers [12] enforcing affine condition (9). Both policies follow the same training and are encouraged to enforce dissipation condition (12) at the vertices of \mathcal{B} , which translates to $\dot{\theta} \leq -2\varepsilon - 10\dot{\theta}$.

To illustrate Theorem 1, assume that our POLICEd controller μ_θ enforces dissipation condition (12) and let us consider a trajectory s entering buffer \mathcal{B} at time t_0 . If initial state condition (13) holds, i.e., if $\theta(t_0) < 0.2$ rad and $\dot{\theta}(t_0) < 2 - 10\theta(t_0)$, then as long as $\theta(t) \geq 0.1$ rad, $\dot{\theta}(t) \geq 0$ rad/s, and $(p(t), \dot{p}(t)) \in \mathcal{P}$, we have $\theta(t) < 0.2$ rad and $\dot{\theta}(t) < 2 - 10\theta(t)$. These equations generate the phase portrait of Fig. 2, which successfully reproduces the desired behavior exhibited in Fig. 1.

Our POLICEd controller guarantees that all trajectories entering \mathcal{B} cannot cross its upper bound \bar{b} and hence cannot violate the constraint, whereas some baseline trajectories cross \bar{b} and violate the constraint as shown in Fig. 2.

B. Space Shuttle Landing

We now study the highly-nonlinear dynamics of the space shuttle landing [18]. Original state $x \in \mathbb{R}^3$ is composed of the altitude h of the shuttle, its flight path angle γ , and its velocity v . The dynamics of these state are

$$\dot{h}(t) = v(t) \sin \gamma(t) \quad (22a)$$

$$\dot{\gamma}(t) = \rho(t)v(t)C_L(t)\frac{S}{2m} - \frac{g \cos \gamma(t)}{v(t)} \quad (22b)$$

$$\dot{v}(t) = -\rho(t)v^2(t)C_D(t)\frac{S}{2m} - g \sin \gamma(t), \quad (22c)$$

where the air density satisfies $\rho(t) = \rho_0 e^{-h(t)/H}$ and the lift and drag coefficients take the form

$$\begin{aligned} C_L(t) &= C_{L_0} \sin^2 \alpha(t) \cos \alpha(t) \\ C_D(t) &= C_{D_0} + K C_L^2(t). \end{aligned} \quad (23)$$

The control input is the angle-of-attack α of the shuttle, which makes these dynamics non-affine in control, and hence

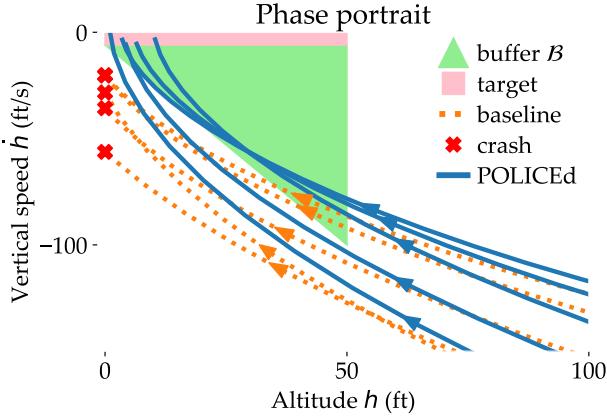


Fig. 3: Phase portrait of the space shuttle landing. POLICEd trajectories (blue) entering buffer \mathcal{B} (green) all converge to a set of target conditions (pink) with small vertical velocity from which landing is feasible. However, the baseline trajectories (dotted orange) reach the ground $h = 0$ with high vertical velocities $\dot{h} \leq -6$ ft/s resulting in a crash of the shuttle (x).

cannot be handled directly by any CBF method¹ [11], [20].

In this scenario, the shuttle starts from a descent configuration with high vertical velocity \dot{h} , which must be drastically reduced to allow for a soft landing. More specifically, we consider initial states typical of a descent phase from a height $h_0 = 500$ ft, velocities $v_0 \in [300, 400]$ ft/s and flight path angles $\gamma_0 \in [-30^\circ, -10^\circ]$. The objective of our controller is to bring the shuttle to a low altitude $h \leq 50$ ft with vertical velocity $\dot{h} \leq 6$ ft/s sufficiently small to allow for a soft landing [18]. We choose an output constraint $y := -h \leq 0$, which has a relative degree 2 for control input α . We build a buffer with $y_{min} = -50$ ft, $y_{max} = 0$ ft, $s_2^{min} = 6$ ft/s and $\dot{y}_{max} = 100$ ft/s.

We introduce state $s := (h, \dot{h}, \gamma)$ and invertible transformation $T(x) = s$ is easily obtainable from pure geometric considerations as it only needs (22a). Thus, Assumption 1 is verified and determining T does not violate our black-box assumption on dynamics (22), (23).

We train two PPO policies [17] to minimize the vertical velocity at touchdown. One of these policies is a standard MLP used as a baseline and the other is our POLICEd version [10]. As seen in Fig. 3 our POLICEd policy successfully enforces the dissipative buffer of Theorem 1 and ensures soft landing of the shuttle contrary to the baseline PPO policy.

V. CONCLUSION AND FUTURE WORK

In this work, we established High Relative Degree POLICEd RL, a novel method to enforce a hard constraint of high relative degree on learned policies, while only using an implicit black-box model of the environment. We built a buffer region where the policy dissipates the generalized inertia of the high relative degree constraint to prevent trajectories from reaching the constraint line. We illustrated our theory on the MuJoCo inverted pendulum and on a space shuttle landing scenario.

¹Adding an integrator $\dot{\alpha}(t) = u(t)$ renders dynamics (22) affine in control at the price of a higher relative degree [19].

Several avenues for future work seem especially interesting. Extending the POLICE algorithm of [12] to enforce multiple affine regions would allow a straightforward extension of this work to guarantee the satisfaction of multiple constraints of high relative degree. Another interesting pursuit would be to investigate how to guarantee constraint satisfaction during the training process of the policy.

REFERENCES

- [1] G. Dulac-Arnold, N. Levine, D. J. Mankowitz, J. Li, C. Paduraru, S. Gowal, and T. Hester, “Challenges of real-world reinforcement learning: Definitions, benchmarks and analysis,” *Machine Learning*, vol. 110, no. 9, pp. 2419 – 2468, 2021.
- [2] S. Gu, L. Yang, Y. Du, G. Chen, F. Walter, J. Wang, Y. Yang, and A. Knoll, “A review of safe reinforcement learning: Methods, theory and applications,” *arXiv preprint arXiv:2205.10330*, 2022.
- [3] E. Altman, *Constrained Markov decision processes*. Routledge, 2021.
- [4] Z. Liu, Z. Guo, Y. Yao, Z. Cen, W. Yu, T. Zhang, and D. Zhao, “Constrained decision transformer for offline safe reinforcement learning,” *arXiv preprint arXiv:2302.07351*, 2023.
- [5] W. Xiao, T.-H. Wang, R. Hasani, M. Chahine, A. Amini, X. Li, and D. Rus, “BarrierNet: Differentiable control barrier functions for learning of safe robot control,” *IEEE Transactions on Robotics*, vol. 39, no. 3, pp. 2289 – 2307, 2023.
- [6] N. Rober, S. M. Katz, C. Sidrane, E. Yel, M. Everett, M. J. Kochenderfer, and J. P. How, “Backward reachability analysis of neural feedback loops: Techniques for linear and nonlinear systems,” *IEEE Open Journal of Control Systems*, vol. 2, pp. 108–124, 2023.
- [7] T.-H. Pham, G. De Magistris, and R. Tachibana, “Optplayer - practical constrained optimization for deep reinforcement learning in the real world,” in *International Conference on Robotics and Automation*. IEEE, 2018, pp. 6236–6243.
- [8] Z. Qin, D. Sun, and C. Fan, “Sablas: Learning safe control for black-box dynamical systems,” *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 1928–1935, 2022.
- [9] Y. Yang, Y. Jiang, Y. Liu, J. Chen, and S. E. Li, “Model-free safe reinforcement learning through neural barrier certificate,” *IEEE Robotics and Automation Letters*, vol. 8, no. 3, pp. 1295–1302, 2023.
- [10] J.-B. Bouvier, K. Nagpal, and N. Mehr, “POLICEd RL: Learning closed-loop robot control policies with provable satisfaction of hard constraints,” in *Robotics: Science and Systems*, 2024.
- [11] J. Breedon and D. Panagou, “High relative degree control barrier functions under input constraints,” in *60th Conference on Decision and Control*. IEEE, 2021, pp. 6119–6124.
- [12] R. Balestrieri and Y. LeCun, “POLICE: Provably optimal linear constraint enforcement for deep neural networks,” in *International Conference on Acoustics, Speech and Signal Processing*. IEEE, 2023, pp. 1–5.
- [13] S. Sastry, *Nonlinear systems: analysis, stability, and control*. Springer Science & Business Media, 1999.
- [14] T. Westenbroek, D. Fridovich-Keil, E. Mazumdar, S. Arora, V. Prabhu, S. S. Sastry, and C. J. Tomlin, “Feedback linearization for uncertain systems via reinforcement learning,” in *International Conference on Robotics and Automation*. IEEE, 2020, pp. 1364 – 1371.
- [15] G. Golub, “Numerical methods for solving linear least squares problems,” *Numerische Mathematik*, vol. 7, pp. 206–216, 1965.
- [16] H. K. Khalil, *Nonlinear Systems*. Prentice Hall, 2002.
- [17] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, “Proximal policy optimization algorithms,” *arXiv preprint arXiv:1707.06347*, 2017.
- [18] N. Harl and S. N. Balakrishnan, “Reentry terminal guidance through sliding mode control,” *Journal of Guidance, Control, and Dynamics*, vol. 33, no. 1, pp. 186 – 199, 2010.
- [19] J. Yao, “Finite time suboptimal control design of nonlinear systems with θ -D technique and implementation to aerospace applications,” Ph.D. dissertation, Missouri University of Science and Technology, 2020.
- [20] Q. Nguyen and K. Sreenath, “Exponential control barrier functions for enforcing high relative-degree safety-critical constraints,” in *American Control Conference*. IEEE, 2016, pp. 322 – 328.