

CBF-RL: Safety Filtering Reinforcement Learning in Training with Control Barrier Functions

*Lizhi Yang, *Blake Werner, Massimiliano de Sa, Aaron D. Ames

Abstract—Reinforcement learning (RL), while powerful and expressive, can often prioritize performance at the expense of safety. Yet safety violations can lead to catastrophic outcomes in real-world deployments. Control Barrier Functions (CBFs) offer a principled method to enforce dynamic safety—traditionally deployed *online* via safety filters. While the result is safe behavior, the fact that the RL policy does not have knowledge of the CBF can lead to conservative behaviors. This paper proposes CBF-RL, a framework for generating safe behaviors with RL by enforcing CBFs *in training*. CBF-RL has two key attributes: (1) minimally modifying a nominal RL policy to encode safety constraints via a CBF term, and (2) safety filtering of the policy rollouts in training. Theoretically, we prove that continuous-time safety filters can be deployed via closed-form expressions on discrete-time roll-outs. Practically, we demonstrate that CBF-RL internalizes the safety constraints in the learned policy—both enforcing safer actions and biasing towards safer rewards—enabling safe deployment without the need for an online safety filter. We validate our framework through ablation studies on navigation tasks and on the Unitree G1 humanoid robot, where CBF-RL enables safer exploration, faster convergence, and robust performance under uncertainty, enabling the humanoid robot to avoid obstacles and climb stairs safely in real-world settings without a runtime safety filter.

I. INTRODUCTION

Humanoid robots are capable of interacting with environments designed for humans. However, the complex environment, high-dimensional robot dynamics, and noise of the sensors also make them highly vulnerable to unsafe control inputs. One unsafe action could lead to damage to both the robot and its surroundings, and thus ensuring safety is essential. Meanwhile, reinforcement learning (RL) has emerged as a powerful tool for humanoid robots to achieve diverse skills, but focuses mostly on performance [1]–[3] and expressiveness [4]–[7]. In this paper, we propose integrating formal safety mechanisms with the powerful exploration and exploitation abilities of RL so that learned policies can reduce or prevent catastrophic behaviors. To achieve this, we turn to *Control Barrier Functions (CBFs)* [8] for a principled way to encode state-based safety constraints as forward-invariant sets. The CBF conditions are often enforced using *safety filters* [9]: quadratic programs satisfying the safety constraint by minimally modifying a proposed control input.

There are two key approaches to instantiating safety filters in RL. The first approach is *safety filtering* the RL-proposed action and projects it into the safe set before execution [10]–[13] or performing constrained updates of the gradient [14],

* denotes equal contribution. All authors affiliated with Caltech MCE. This research is supported in part by the Technology Innovation Institute (TII), BP p.l.c., and by The Dow Chemical Company project #227027AW.




Fig. 1. A humanoid robot trained to climb stairs with the CBF-RL framework. Safety is injected into training by both filtering the policy-proposed actions and also provide safety rewards in addition to task and regularization rewards. During deployment the CBF policy retains safe behavior without a runtime filter.

[15]. This guarantees safety at runtime, but the filter must remain in the loop at deployment, and the learned policy may never internalize the constraint. This prevents a high-dimensional agent, like a humanoid robot, from discovering novel or efficient behaviors since the exploration space is pruned too aggressively. Both also require solving an optimization program at every control step, which may be computationally expensive. The second approach is *reward shaping* where a residual augments the reward term to penalize states that approach or violate constraint boundaries [16]–[21] and encourages the agent towards safer behaviors without active filtering. This alone does not directly enforce safe actions during training and is often sensitive to the choice of penalty weights, possibly being insufficient in safety-critical applications. This paper proposes a fusion of these two approaches to integrating CBFs into RL.

A. Contributions

In this paper, we show that safety filtering and reward shaping are complementary, proposing *CBF-RL*: a *dual approach* that applies both a closed-form CBF-based safety filter and a barrier-inspired reward term during training. This safety filters a nominal RL policy, enabling it to learn safe

behaviors. Catastrophic unsafe actions are prevented by the active filter, and the reward term biases the policy toward avoiding safety interventions. Therefore, the policy has direct corrective supervision—it observes what it would have done, how the filter corrects it, and how the reward changes. It also learns to propose actions that directly satisfy the barrier condition. This enables the policy to show safe behaviors at deployment time without an active filter.

We evaluate CBF-RL, and its dual approach to safe RL, with ablations using a 2D navigation task with dynamics randomness, analyzing task completion rates and robustness tests. We also train humanoid locomotion policies in IsaacLab [22] with full-order dynamics and domain randomization for obstacle avoidance and stair climbing tasks. The approach is validated on hardware using a Unitree G1 humanoid with zero-shot sim-to-real policies to exhibit the effectiveness of the proposed framework in high-dimensional, complex systems. With the dual-trained policy, the robot can successfully navigate around obstacles and climb stairs under command sequences that would lead to failure with nominal policies that don’t leverage the CBF-RL framework.

Our contributions are as follows:

- *Conceptually*: We propose a dual CBF-RL training framework that uses both CBF-based active filtering and barrier-inspired rewards during training, and can be deployed without a filter.
- *Theoretically*: We provide a continuous-to-discrete-relationship analysis of CBF-RL and a closed-form solution for light-weight integration.
- *Practically*: We empirically demonstrate across simulated, and hardware experiments that policies trained using the dual approach can internalize safety and reduce unsafe actions at deployment.

B. Related Work

Due to the ease of modifying rewards for training, a number of works incorporate safety value functions into the reward structure to guide policies toward safety. [16], [17] use a fixed penalty, [18], [19], [21] utilize correction-proportional penalties, and [20] proposes an explicit barrier-inspired reward shaping mechanism to reduce unsafe exploration. These methods encourage safe behavior but do not explicitly direct the policy to exhibit safe behaviors during learning, instead solely relying on the policy to discover safer actions on its own, leading to slower training.

To address this, runtime CBF-based safety filters during training minimally modify the actions of the policy such that the system stays in the user-defined forward-invariant safe set typically by solving an optimization program, be it discrete-time due to the nature of RL training [10]–[13], [15], [23], or continuous-time [24]–[26] which empirically shows improved safety. For humanoid locomotion, these methods are not ideal as humanoid robots have tight real-time and computation power constraints and inaccurate state estimation from sensor noise. We instead filter only in simulation during training, and show that the resulting policy retains safety even without a runtime filter. We further

provide theoretical verification that under certain conditions, continuous-time CBF can be used as conditions for forward invariance for RL simulations that are discrete in nature, and thus we can use the closed-form solution to the CBF-QP to accelerate each step in training.

Many papers utilize model-based approaches [11], [15], [24], [27], which require access to accurate dynamics models. While theoretically appealing, they are less practical for high-dimensional humanoids where dynamics are complex and uncertain. Some works also do constrained updates of the gradient to ensure that the policy remains safe [14], [15]; however, they also require solving optimization programs at each gradient update step and limit the exploration of the policy. Our framework is model-free, requiring only derivatives of the reduced-order model (e.g. for the kinematics of a humanoid robot, its Jacobian J), and emphasizes lightweight integration with standard policy-gradient RL, in our case proximal policy optimization (PPO) [28]. This also leads to the benefit of the policy being able to venture closer to the constraint boundaries, ensuring rich exploration.

In response to the issue of stochastic systems, robust extensions address uncertainty in dynamics or sensing through disturbance observers, Gaussian Process models, or robustified CBFs [24]–[26], [29]. These methods improve reliability under uncertainty but add significant complexity and computational burden. Here we show that by relying on domain randomization during training, dual-trained policies remain safer under uncertainty of the system without explicit models.

Another line of work learns barrier-like certificates or relates value functions to barrier properties [30], [31]. These methods aim to automate the design of barrier functions or embed them in differentiable layers. Our approach assumes analytic barrier functions and focuses on pragmatic integration with RL training rather than barrier discovery.

The above works have been applied to domains such as spacecraft inspection [12], drone control [13], autonomous driving [23], and driver assistance [24]. Much of the prior work captures part of the safety challenge, but none directly combines filtering and shaping in a way that allows a policy to *internalize safety* during training and then act autonomously without a filter at deployment, especially not on a high-dimensional system such as a humanoid robot. This distinction defines the novelty of our contribution.

II. BACKGROUND

Reinforcement Learning We consider the standard RL formulation of a Markov decision process (MDP) $(\mathcal{X}, \mathcal{U}, P, r, \gamma)$. At each timestep k , an agent selects an action $\mathbf{u}_k \in \mathcal{U}$, according to a policy $\pi_\theta(\mathbf{u}_k|\mathbf{x}_k)$ based on an observed state $\mathbf{x}_k \in \mathcal{X}$, and receives a reward $r(\mathbf{x}_k, \mathbf{u}_k)$. The environment follows the transition dynamics $\mathbf{x}_{k+1} \sim P(\cdot|\mathbf{x}_k, \mathbf{u}_k)$. The goal is then to maximize the expected discounted return $\mathbb{E}[\sum_{k=0}^{\infty} \gamma^k r(\mathbf{x}_k, \mathbf{u}_k)]$. During deployment, actions are selected as the conditional expectation $\mathbf{u}_k = E[\pi_\theta(\mathbf{u}_k|\mathbf{x}_k)]$. In this work, we specifically use model-free policy-gradient actor-critic methods such as PPO [28], though our approach is agnostic to the specific RL algorithm.

RL also often suffers from reward sparsity when unsafe events like obstacle collisions are rare. Thus the policy seldom experiences consequences of unsafe actions, leading to vanishing gradients and unstable, slow training. As such, one part of our method also does reward densification to mitigate this by designing $r(\mathbf{x}_k, \mathbf{u}_k)$ to give informative nonterminal signals related to safety.

Reduced-Order Models Consider a continuous-time system with dynamics $\dot{\mathbf{x}} = \phi(\mathbf{x}, \mathbf{u})$, where $\mathbf{x} \in \mathbb{R}^n$ is the state and $\mathbf{u} \in \mathbb{R}^m$ is the control input. In our case the state \mathbf{x} may be very high-dimensional, e.g. joint positions and velocities of a robot. Thus we consider a reduced-order state $\mathbf{q} \in \mathbb{R}^{n_q}$, $n_q < n$, representing key lower-dimensional features such as the robot's center of mass position. We define a projection $\mathbf{p} : \mathbb{R}^n \rightarrow \mathbb{R}^{n_q}$ that projects the full-order state onto the reduced-order state. Given a locally Lipschitz continuous feedback control law $\mathbf{v} = \mathbf{k}(\mathbf{q})$, the reduced-order state \mathbf{q} is

$$\dot{\mathbf{q}} = \frac{\partial \mathbf{p}}{\partial \mathbf{x}} \phi(\mathbf{x}, \psi(\mathbf{x}, \mathbf{v})) \approx \mathbf{f}(\mathbf{q}) + \mathbf{g}(\mathbf{q})\mathbf{v} = \mathbf{f}(\mathbf{q}) + \mathbf{g}(\mathbf{q})\mathbf{k}(\mathbf{q}), \quad (1)$$

where $\psi(\mathbf{x}, \mathbf{v})$ is a control interface lifting the reduced-order input \mathbf{v} to a full-order input \mathbf{u} . See [30], [32] for the connections between reduced and full order models.

Control Barrier Function Safety Filters In the control barrier function framework, a set of ‘‘safe states’’ for the system, $\mathcal{S} \subseteq \mathbb{R}^{n_q}$, is encoded as the zero superlevel set of a continuously differentiable function $h : \mathbb{R}^{n_q} \rightarrow \mathbb{R}$,

$$\mathcal{S} := \{\mathbf{q} \in \mathbb{R}^{n_q} \mid h(\mathbf{q}) \geq 0\}, \quad (2)$$

$$\partial \mathcal{S} := \{\mathbf{q} \in \mathbb{R}^{n_q} \mid h(\mathbf{q}) = 0\}, \quad (3)$$

$$\text{int}(\mathcal{S}) := \{\mathbf{q} \in \mathbb{R}^{n_q} \mid h(\mathbf{q}) > 0\}. \quad (4)$$

The aim of safety-critical control is to design a feedback control law $\mathbf{k}(\mathbf{q})$ that renders \mathcal{S} forward invariant.

Definition 1. A set $\mathcal{S} \subseteq \mathbb{R}^{n_q}$ is *forward invariant* for (1) if, for every initial state $\mathbf{q}(t_0) \in \mathcal{S}$, the resulting state trajectory $\mathbf{q} : I \subseteq \mathbb{R} \rightarrow \mathbb{R}^{n_q}$ remains in \mathcal{S} for all $t \in I \cap \mathbb{R}_{\geq t_0}$.

For control-affine systems as in (1), the forward invariance of \mathcal{S} can be enforced using CBFs [8], [9].

Definition 2. Let $\mathcal{S} \subset \mathbb{R}^{n_q}$ be a set as in (2), with h satisfying $\nabla h|_{\partial \mathcal{S}} \neq 0$. Then, h is a *control barrier function* on \mathbb{R}^{n_q} if there exists¹ a $\gamma \in \mathcal{K}_\infty^e$ such that for all $\mathbf{q} \in \mathbb{R}^{n_q}$,

$$\sup_{\mathbf{v} \in \mathbb{R}^m} \left\{ \dot{h}(\mathbf{q}, \mathbf{v}) = L_{\mathbf{f}}h(\mathbf{q}) + L_{\mathbf{g}}h(\mathbf{q})\mathbf{v} \right\} > -\gamma(h(\mathbf{q})). \quad (5)$$

Given a CBF h and function $\gamma \in \mathcal{K}_\infty^e$, the set of control inputs satisfying (5) at \mathbf{q} is given by

$$\mathcal{U}_{\text{CBF}}(\mathbf{q}) = \left\{ \mathbf{v} \in \mathbb{R}^m \mid \dot{h}(\mathbf{q}, \mathbf{v}) \geq -\gamma(h(\mathbf{q})) \right\}. \quad (6)$$

Any locally Lipschitz controller \mathbf{k} for (1) for which $\mathbf{k}(\mathbf{q}) \in \mathcal{U}_{\text{CBF}}(\mathbf{q}) \forall \mathbf{q} \in \mathcal{S}$ enforces forward invariance of \mathcal{S} [8].

1: A function $\alpha : \mathbb{R} \rightarrow \mathbb{R}$ belongs to \mathcal{K}_∞^e if it is increasing, continuous, and satisfies $\alpha(0) = 0$, $\lim_{r \rightarrow \pm\infty} \alpha(r) = \pm\infty$.

For real-world robotic systems with zero-order hold, sampled-data implementations, it is generally impossible to choose continuous control actions that create the closed-loop system (1). As such, for the remainder of this work, we focus on the discrete-time analogues of these systems,

$$\mathbf{q}_{k+1} = \mathbf{F}(\mathbf{q}_k) + \mathbf{G}(\mathbf{q}_k)\mathbf{v}_k, \quad \forall k \in \mathbb{Z}_{\geq 0}, \quad (7)$$

where $\mathbf{F} : \mathbb{R}^{n_q} \rightarrow \mathbb{R}^{n_q}$ and $\mathbf{G} : \mathbb{R}^{n_q} \rightarrow \mathbb{R}^{n_q \times m}$ are the discretization of (1) over a time interval $\Delta t > 0$ for a constant input \mathbf{v} . Here, we focus on enforcing forward invariance at sample times² $k\Delta t$. This can be achieved using a discrete-time CBF [34], [35].

Definition 3. Let $\mathcal{S} \subseteq \mathbb{R}^{n_q}$ be as in (2) and $\rho \in [0, 1]$. The function h is a *discrete-time control barrier function (DTCBF)* for (7) if $\forall \mathbf{q} \in \mathcal{S}$ there exists a $\mathbf{v} \in \mathbb{R}^m$ for which

$$h(\mathbf{F}(\mathbf{q}) + \mathbf{G}(\mathbf{q})\mathbf{v}) \geq \rho h(\mathbf{q}). \quad (8)$$

Similar to CBFs and continuous-time systems, DTCBFs keep the discrete-time system (7) safe for each $k \in \mathbb{Z}_{\geq 0}$. Here the value of h is lower bounded by a geometrically decaying curve, $h(\mathbf{q}_k) \geq \rho^k h(\mathbf{q}_0)$ [34]. Thus, they can be used to generate safe control actions through an optimization program wherein a desired but potentially unsafe input $\mathbf{v}_k^{\text{des}} \in \mathbb{R}^m$ is minimally modified to produce a safe input:

$$\mathbf{v}_k^{\text{safe}} = \arg \min_{\mathbf{v}_k \in \mathbb{R}^m} \|\mathbf{v}_k - \mathbf{v}_k^{\text{des}}(\mathbf{q}_k)\|^2 \quad (9)$$

$$\text{s.t. } h(\mathbf{F}(\mathbf{q}_k) + \mathbf{G}(\mathbf{q}_k)\mathbf{v}_k) \geq \rho h(\mathbf{q}_k). \quad (10)$$

III. DUAL APPROACH TO CBF-RL

In order to apply CBFs to RL pipelines, we characterize the relationship between continuous-time CBFs and discrete updates of the RL environment. With this relationship, we can utilize the closed-form solution of the continuous-time CBF-QP to understand its effect on the RL environment.

Lemma 1. Suppose $h : \mathbb{R}^{n_q} \rightarrow \mathbb{R}$ is a C^1 CBF for the continuous-time single integrator $\dot{q} = v$, where $q, v \in \mathbb{R}^{n_q}$. Let $k_s : \mathbb{R}^{n_q} \rightarrow \mathbb{R}^{n_q}$ be any safe, locally Lipschitz controller for the continuous-time integrator, satisfying

$$\nabla h(\mathbf{q})^\top k_s(\mathbf{q}) \geq -\alpha h(\mathbf{q}), \quad \forall \mathbf{q} \in \mathbb{R}^{n_q}, \quad (11)$$

for some $\alpha > 0$. Let $\mathbf{f}_{\Delta t} : \mathbb{R}^{n_q} \times \mathbb{R}^{n_q} \rightarrow \mathbb{R}^{n_q}$, $(\mathbf{q}, \mathbf{v}) \mapsto \mathbf{q} + \Delta t \mathbf{v}$ be the Euler discretization of the single integrator with time step $\Delta t > 0$. There exists a continuous function $R : \mathbb{R}^{n_q} \times \mathbb{R}^{n_q} \rightarrow \mathbb{R}$ such that $\forall \mathbf{q} \in \mathbb{R}^{n_q}$, $\lim_{\|\mathbf{w}\| \rightarrow 0} \frac{R(\mathbf{q}, \mathbf{w})}{\|\mathbf{w}\|} = 0$ and

$$h(\mathbf{f}_{\Delta t}(\mathbf{q}, k_s(\mathbf{q}))) \geq (1 - \Delta t \alpha)h(\mathbf{q}) - |R(\mathbf{q}, \Delta t k_s(\mathbf{q}))|. \quad (12)$$

Proof. By [36, Chapter 5.2 (2)], there exists a continuous function $R : \mathbb{R}^{n_q} \times \mathbb{R}^{n_q} \rightarrow \mathbb{R}$ satisfying $\lim_{\|\mathbf{w}\| \rightarrow 0} \frac{R(\mathbf{q}, \mathbf{w})}{\|\mathbf{w}\|} = 0$ and $h(\mathbf{q} + \mathbf{w}) = h(\mathbf{q}) + \nabla h(\mathbf{q})^\top \mathbf{w} + R(\mathbf{q}, \mathbf{w}) \forall \mathbf{q}, \mathbf{w} \in \mathbb{R}^{n_q}$. Taking $\mathbf{w} = \Delta t k_s(\mathbf{q})$, we calculate $h(\mathbf{q} + \Delta t k_s(\mathbf{q}))$ as

$$= h(\mathbf{q}) + \Delta t \cdot [\nabla h(\mathbf{q})^\top k_s(\mathbf{q})] + R(\mathbf{q}, \Delta t k_s(\mathbf{q})) \quad (13)$$

$$\geq h(\mathbf{q}) - \Delta t \cdot \alpha h(\mathbf{q}) - |R(\mathbf{q}, \Delta t k_s(\mathbf{q}))|. \quad (14)$$

2: See [33] for a discussion of zero-order-hold, intersample safety.




Fig. 2. CBF-RL framework: For one given task, the user defines the safety barrier function $h(\mathbf{q})$ and the accompanying $\nabla h(\mathbf{q})$. During training, the RL policy proposes action $\mathbf{v}^{\text{policy}}$; the CBF safety filter then calculates the closed-form solution of the CBF-QP $\mathbf{v}^{\text{filtered}}$ and the safety reward r_{cbf} based on proposed action $\mathbf{v}^{\text{policy}}$ and agent configuration \mathbf{q} . The RL agent executes $\mathbf{v}^{\text{filtered}}$ in the massively-parallel discretized environment, and the policy is updated with the combination of task, regularization, and safety rewards $r = r_{\text{nominal}} + r_{\text{cbf}}$. During deployment, the policy is able to output safe actions $\mathbf{v}^{\text{policy}}$ without needing an explicit runtime filter. Sample code for the single integrator example with CBF-RL reward core construction can be found at <https://github.com/lzyang2000/cbf-rl-navigation-demo>.

Combining h terms, the result follows. \square

Using Lemma 1, we bound the evolution of the barrier function along trajectories of the Euler-discretized integrator.

Theorem 1 (Continuous to Discrete Safety). *Consider the setting of Lemma 1. Suppose there exists a compact, forward-invariant set $K \subseteq \mathbb{R}^{n_q}$ for the discrete-time integrator dynamics $\mathbf{q}_{k+1} = \mathbf{f}_{\Delta t}(\mathbf{q}_k, k_s(\mathbf{q}_k))$. Then, $\forall \Delta t > 0$, $\mu(\Delta t) = \sup_{\mathbf{q} \in K} |R(\mathbf{q}, \Delta t k_s(\mathbf{q}))| < +\infty$ and $\lim_{\Delta t \rightarrow 0} \mu(\Delta t)/\Delta t = 0$. Further, provided $(1 - \Delta t \alpha) \in [0, 1)$, for $\mathbf{q}_0 \in K$,*

$$h(\mathbf{q}_k) \geq (1 - \Delta t \alpha)^k h(\mathbf{q}_0) - \frac{\mu(\Delta t)}{\Delta t \alpha}, \quad \forall k \geq 0. \quad (15)$$

Remark 1. As a consequence of the limiting behavior of μ , as we take the discretization step Δt to zero, the standard DTCBF bound, $h(\mathbf{q}_k) \geq (1 - \Delta t \alpha)^k h(\mathbf{q}_0)$, is recovered.

Proof. Fix $\Delta t > 0$. Since R and k_s are continuous, compactness of K implies $\mu(\Delta t)$ is finite $\forall \Delta t > 0$. To establish the limit $\lim_{\Delta t \rightarrow 0} \mu(\Delta t)/\Delta t = 0$, we note that $\lim_{\Delta t \rightarrow 0} \frac{|R(\mathbf{q}, \Delta t k_s(\mathbf{q}))|}{\Delta t} = 0 \quad \forall \mathbf{q} \in K$ implies

$$\sup_{\mathbf{q} \in K} \lim_{\Delta t \rightarrow 0} \frac{|R(\mathbf{q}, \Delta t k_s(\mathbf{q}))|}{\Delta t} = 0. \quad (16)$$

Compactness of K and joint continuity of $|R(\mathbf{q}, \Delta t k_s(\mathbf{q}))|$ in \mathbf{q} and Δt implies uniform continuity of $|R(\mathbf{q}, \Delta t k_s(\mathbf{q}))|$; this lets us interchange limit and supremum. We conclude $\lim_{\Delta t \rightarrow 0} \sup_{\mathbf{q} \in K} \frac{|R(\mathbf{q}, \Delta t k_s(\mathbf{q}))|}{\Delta t} = 0 \Rightarrow \lim_{\Delta t \rightarrow 0} \frac{\mu(\Delta t)}{\Delta t} = 0$.

Now, we establish the bound using a comparison system. If $\mathbf{y}_{k+1} = \rho \mathbf{y}_k - |\mathbf{d}_k|$ for all $k \geq 0$ and $|\mathbf{d}_k| \leq \mu$, a geometric series argument establishes $\mathbf{y}_k \geq \rho^k \mathbf{y}_0 - (\frac{1}{1-\rho})\mu \quad \forall k \geq 0$. Fix $\mathbf{q}_0 \in K$. Since K is forward invariant for the closed-loop system, $|R(\mathbf{q}_k, \Delta t k_s(\mathbf{q}_k))| \leq \mu(\Delta t) \quad \forall k \geq 0$. Using Lemma 1, we take $\rho = 1 - \Delta t \alpha$, $\mu(\Delta t) = \sup_{\mathbf{q} \in K} |R(\mathbf{q}, \Delta t k_s(\mathbf{q}))|$, and conclude by comparison with \mathbf{y}_k that the bound

$$h(\mathbf{q}_k) \geq (1 - \Delta t \alpha)^k h(\mathbf{q}_0) - \frac{\mu(\Delta t)}{\Delta t \alpha}, \quad (17)$$

does indeed hold for all $\mathbf{q}_0 \in K$. \square

The above means that with small enough Δt (typically $\Delta t \leq 0.01$ s, governed by the stability limits and solve time of the physics engine), continuous-time CBF tools can be directly applied to discrete-time RL environments. We provide the two core parts of CBF-RL shown in Fig. 2:

Safety-filtering during training At training time, the policy would generate desired actions that are not necessarily safe $\mathbf{v}_k^{\text{policy}}$ at step k . A safety filter is then applied to enforce safe behaviors and guide the system to 'learn' the safety filter as part of the natural closed-loop dynamics. As we have proven the relationship between the continuous-time and discrete-time CBF conditions, we replace the typically nonlinear DTCBF constraint with the continuous-time CBF first-order inequality. This reduces the safety filtering problem to a single linear-constraint quadratic program (QP) [9]:

$$\mathbf{v}_k^{\text{safe}} = \arg \min_{\mathbf{v}_k \in \mathbb{R}^{n_q}} \frac{1}{2} \|\mathbf{v}_k - \mathbf{v}_k^{\text{policy}}\|^2 \quad (18)$$

$$\text{s.t. } \nabla h(\mathbf{q}_k)^\top \mathbf{v}_k \geq -\alpha h(\mathbf{q}_k). \quad (19)$$

While this QP provides a safe control input, solving an optimization program numerically at every step of RL training is computationally undesirable, especially in massively parallel environments such as IsaacLab [22]. Fortunately, since the QP has a single linear constraint, it can be solved analytically in closed-form:

$$\mathbf{v}_k^{\text{safe}} = \begin{cases} \mathbf{v}_k^{\text{policy}}, & \text{if } \mathbf{a}_k^\top \mathbf{v}_k^{\text{policy}} \geq b_k \\ \mathbf{v}_k^{\text{policy}} + \frac{(b_k - \mathbf{a}_k^\top \mathbf{v}_k^{\text{policy}}) \mathbf{a}_k}{\|\mathbf{a}_k\|^2}, & \text{o.w.} \end{cases} \quad (20)$$

$$\mathbf{a}_k := \nabla h(\mathbf{q}_k), \quad b_k := -\alpha h(\mathbf{q}_k). \quad (21)$$

Penalizing unsafe behavior In addition to filtering the policy actions at training time, we also quantify how safe the environment is to inform training. To this end, we modify the rewards to include r_{CBF} defined as

$$r_{\text{cbf}}(\mathbf{q}_k, \mathbf{v}_k) = \min\left(\mathbf{a}_k^\top \mathbf{v}_k^{\text{policy}} - b_k, 0\right) \quad (22)$$

$$+ \left(\exp\left(-\frac{\|\mathbf{v}_k^{\text{policy}} - \mathbf{v}_k^{\text{safe}}\|^2}{\sigma^2}\right) - 1\right) \quad (23)$$

and the whole reward is thus $r = r_{\text{nominal}} + r_{\text{cbf}}$.

Intuitively, this reward penalizes actions whenever the safety filter is activated, and also incentivizes the model to take actions as close to the safe actions as possible to reduce the intervention of the filter. Because the penalty term $\exp(\dots) - 1$ is strictly lower-bounded by -1 , it provides a smooth learning signal without causing unbounded instability with respect to the primary task rewards. The whole algorithm is as shown in Alg. 1.

Algorithm 1 RL Training with Discrete-Time CBF Safety

- 1: Initialize policy parameters θ , initial configuration q_0 , safety function h
 - 2: **for** step = 1 to N_{steps} **do**
 - 3: Initialize q_0 , observation \mathbf{o}_0
 - 4: **for** $k = 0$ to $T - 1$ **do**
 - 5: $\mathbf{v}_k^{\text{policy}} \leftarrow \pi_\theta(\mathbf{o}_k)$
 - 6: $\mathbf{q}_{k+1}^{\text{policy}} \leftarrow \mathbf{q}_k + \Delta t \mathbf{v}_k^{\text{policy}}$
 - 7: $\mathbf{a}_k = \nabla h(\mathbf{q}_k)$, $b_k = -\alpha h(\mathbf{q}_k)$
 - 8: Compute CBF condition: $c = \mathbf{a}_k^\top \mathbf{v}_k^{\text{policy}} - b_k$
 - 9: **if** $c \geq 0$ **then** $\mathbf{v}_k^{\text{safe}} \leftarrow \mathbf{v}_k^{\text{policy}}$,
 - 10: **else** $\mathbf{v}_k^{\text{safe}} \leftarrow \mathbf{v}_k^{\text{policy}} + \frac{(b_k - \mathbf{a}_k^\top \mathbf{v}_k^{\text{policy}}) \mathbf{a}_k^\top}{\|\mathbf{a}_k\|^2}$
 - 11: $\mathbf{q}_{k+1}^{\text{safe}} \leftarrow \mathbf{q}_k + \Delta t \mathbf{v}_k^{\text{safe}}$
 - 12: $\mathbf{q}_{k+1}^{\text{env}}, \mathbf{o}_{k+1} \leftarrow \text{ENVIRONMENT_UPDATE}(\mathbf{q}_{k+1}^{\text{safe}})$
 - 13: $r \leftarrow w \cdot [\min(\mathbf{a}_k^\top \mathbf{v}_k^{\text{policy}} - b_k, 0) + \exp\left(-\frac{\|\mathbf{v}_k^{\text{policy}} - \mathbf{v}_k^{\text{safe}}\|^2}{\sigma^2}\right) - 1] + R(\mathbf{q}_{k+1}^{\text{env}})$
 - 14: Store transition: $(\mathbf{q}_k, \mathbf{o}_k, \mathbf{q}_{k+1}^{\text{policy}}, \mathbf{q}_{k+1}^{\text{safe}}, \mathbf{q}_{k+1}^{\text{env}}, r)$
 - 15: **end for**
 - 16: Update policy parameters $\theta \leftarrow \eta \nabla_\theta \mathcal{L}(\theta, r)$
 - 17: **end for**
-

Single Integrator Example To demonstrate our proposed approach and analyze the effect of each component of CBF-RL, we perform extensive ablation studies on a single integrator navigation task. The agent is placed into a 2D world with obstacles and tasked with going to a specified goal. The starting, obstacle, and goal locations are all randomized, and extra care is taken such that the agent always initializes in a safe state following [13]. The single integrator has dynamics $\mathbf{q}_{k+1} = \mathbf{q}_k + \mathbf{v}_k \Delta t$ where $\mathbf{q} = (x, y)$ is the robot position, \mathbf{v}_k is the velocity of the agent, and Δt is the timestep size. The safety barrier function h is defined as

$$h(\mathbf{q}) = \min \left\{ \min_j (\|\mathbf{q} - \mathbf{p}_j\| - (r_{\text{agent}} + r_j)), \quad (24)$$

$$x - r_{\text{agent}}, (L - x) - r_{\text{agent}}, \\ y - r_{\text{agent}}, (L - y) - r_{\text{agent}} \right\}$$

$$\nabla h(\mathbf{q}) = \begin{cases} \frac{\mathbf{q} - \mathbf{p}_{j^*}}{\|\mathbf{q} - \mathbf{p}_{j^*}\|}, & j^* \in \arg \min_j h_j(\mathbf{q}), \\ \pm e_x, & \text{if left/right wall active,} \\ \pm e_y, & \text{if bottom/top wall active,} \end{cases} \quad (25)$$

where \mathbf{p}_j is the position of the j th obstacle with radius r_j , the agent also has radius r_{agent} and the size of the world is

L . Thus we can formulate the training-time safety filter with (20) and define our reward modification associated with the safety with (23). Full reward terms are shown in Table II.

Ablation To validate our method, we train 4 variants (Dual, Reward-only, Filter-only, Nominal) and test 12 variants following Table III for 1500 steps with 4096 parallel environments. We also evaluate policies trained with filtered action then deployed without a runtime filter (rt. filt.). The training progress can be seen in Fig. 3, where we observe that the Dual and Filter-only approaches achieve rapid convergence while remaining safe throughout training. Furthermore, as illustrated by the trajectory comparisons in Fig. 4 as an example from over 1000 random test environments, the Dual approach is able to reach the goal in both domain randomized and non-domain randomized settings, even without a runtime filter, while the other methods fail to do so. Notably, the Filter Only approach performs well only with an active safety filter and degrades markedly without it.

Robustness To further investigate the robustness of the policy induced by domain randomization (DR), we train the dual approach with noise on the dynamics model, i.e. $\mathbf{q}_{k+1} = \mathbf{q}_k + (\mathbf{v}_k + \mathbf{d})\Delta t$ where \mathbf{d} follows the standard normal distribution scaled by 20% to the maximum velocity. It is observed that the policy trained with the dual method overall suffers least from the dynamics disturbance as can be seen in Table I.

TABLE I. Success rates over 1000 random test environments for different methods, with and without DR. The dual policy consistently performs well and suffers less degradation due to dynamics uncertainty.

| | Dual | Dual (w/o rt. filt.) | Reward Only |
|-------|--------------------|-----------------------------|--------------|
| No DR | 99.0% | 92.7% | 91.9% |
| DR | 99.0% (-0%) | 91.7% (-1%) | 87.6%(-4.3%) |
| | Filter Only | Filter Only (w/o rt. filt.) | Nominal |
| No DR | 98.8% | 38.7% | 51.4% |
| DR | 96.7%(-2.1%) | 36.8%(-1.9%) | 55.0%(+3.6%) |

IV. CBF-RL FOR SAFE HUMANOID LOCOMOTION

In the following section, we present two different use cases of CBF-RL in humanoid locomotion to show the generality and performance of our method. Each uses a different set of user-specified reduced-order coordinates and a valid CBF for the resulting reduced-order dynamics. The nominal rewards and observations (history length 5) follow [37]. Note that for blind stair climbing, we employ the asymmetric actor-critic method [38] and provide a height scan of size $1\text{m} \times 1.5\text{m}$ with a resolution of 0.1m to the critic observation with a history length of 1. We use an MLP policy with 3 hidden layers of [512, 256, 128] neurons running at 50Hz, outputting the joint position setpoints of the 12-DoF lower body for the Unitree G1. We train in IsaacLab with 4096 environments with $\Delta t = 0.005\text{s}$. Each episode lasts for a maximum of 20,000 steps on the NVIDIA RTX 4090 GPU and perform zero-shot hardware transfer experiments to verify our approach.

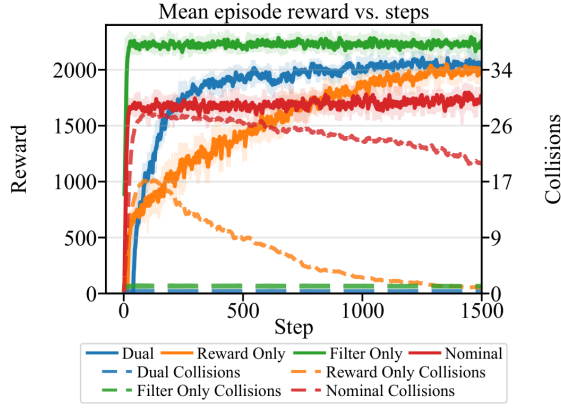


Fig. 3. Training progress with the Dual, Reward Only, Filter Only and Nominal methods. Dual and Filter Only achieve faster convergence and avoid training-time safety violations.

TABLE II. Reward terms for single integrator navigation. Agent is rewarded for staying alive and closing the distance to the goal and penalized for hitting the obstacles / walls, exceeding the set time to complete the task, and violating the CBF conditions or not proposing safer actions

| Reward | Formula |
|-----------------------|---|
| r_{goal} | $1.0 \cdot \mathbf{1}(\text{goal reached})$ |
| r_{obstacle} | $-1.0 \cdot \mathbf{1}(\text{obstacle collision})$ |
| r_{wall} | $-1.0 \cdot \mathbf{1}(\text{wall collision})$ |
| r_{progress} | $20.0 \cdot \frac{\ \mathbf{p}_{t-1} - \mathbf{g}\ - \ \mathbf{p}_t - \mathbf{g}\ }{v_{\text{max}} \Delta t} \cdot \mathbf{1}(\text{active})$ |
| r_{alive} | $0.01 \cdot \mathbf{1}(\text{active})$ |
| r_{cbf} | $100 \cdot \left(\min(\nabla h(\mathbf{q})^\top \mathbf{v} + \alpha h(\mathbf{q}), 0) + \exp\left(-\frac{\ \mathbf{v}_{\text{policy}} - \mathbf{v}_{\text{safe}}\ ^2}{0.5^2}\right) - 1 \right) \cdot \mathbf{1}(\text{active})$ |
| r_{timeout} | $-10.0 \cdot \mathbf{1}(\text{time exceeded})$ |

A. Task Definition

Planar Obstacle Avoidance First, we consider the task where a humanoid robot needs to avoid obstacles during locomotion without intervention, even when the velocity command is to collide with the obstacle. Thus, we can simplify the safety problem as a single integrator problem where the policy modulates the robot’s planar velocities $\mathbf{v}_{\text{planar}}^{\text{base}} = [v_x, v_y]$ to maintain safe distances from the closest obstacle, a cylinder centered at \mathbf{p}_o^r in the robot frame. We define the safety function $h(\mathbf{p}) = \|\mathbf{p}_o^r\| - R_r - R_o$ where R_r and R_o are the radii of the robot and the obstacle respectively. We then train our robot with the CBF reward:

$$r_{\text{obstacle cbf}} = \min\left(\frac{\mathbf{p}_o^r}{\|\mathbf{p}_o^r\|} \cdot \mathbf{v}_{\text{planar}} + \alpha h(\mathbf{p}), 0\right) + \exp\left(-\frac{\|\mathbf{v}_{\text{planar}}^{\text{base}} - \mathbf{v}_{\text{planar}}^{\text{safe}}\|^2}{\sigma^2}\right) - 1. \quad (26)$$

Stair Climbing Second, we consider the task of humanoid locomotion on stairs. For this task, we consider the kinematic model of the foot as our reduced-order model $\mathbf{q}_{k+1}^{\text{sw}} = \mathbf{q}_k^{\text{sw}} + \Delta t \mathbf{J}^{\text{sw}}(\mathbf{q}_k^{\text{sw}}) \mathbf{v}_k^{\text{sw}}$, where $\mathbf{q}^{\text{sw}} = [p_x, p_y]^T$ is the swing foot’s position in the body frame, \mathbf{J}^{sw} comprises the rows of the robot’s body Jacobian associated with the foot’s position, and \mathbf{v}^{sw} is robot’s joint velocities. In climbing stairs, one

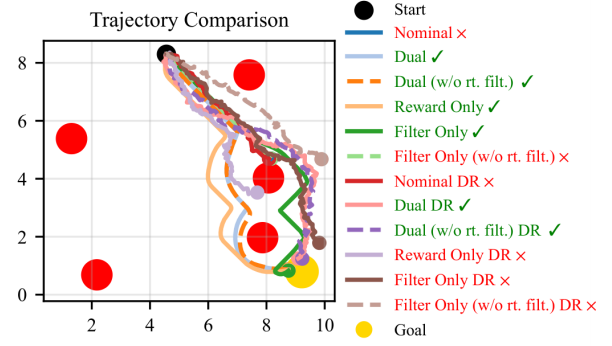


Fig. 4. Trajectory comparisons of one example simulation. The black dot is the start and the yellow circle is the goal. Success = reaching the goal; failure = collision with obstacles or wall.

TABLE III. List of method configurations for single integrator navigation. We ablate all permutations of the two main components of CBF-RL.

| Method | Training | Deployment | DR |
|--------------------------------|---------------|-------------------|-----|
| Nominal | Nominal | No Runtime Filter | No |
| Dual | Reward+Filter | Runtime Filter | No |
| Dual (w/o rt. filt.) | Reward+Filter | No Runtime Filter | No |
| Reward Only | Reward | No Runtime Filter | No |
| Filter Only | Filter | Runtime Filter | No |
| Filter Only (w/o rt. filt.) | Filter | No Runtime Filter | No |
| Nominal DR | Nominal | No Runtime Filter | Yes |
| Dual DR | Reward+Filter | Runtime Filter | Yes |
| Dual (w/o rt. filt.) DR | Reward+Filter | No Runtime Filter | Yes |
| Reward Only DR | Reward | No Runtime Filter | Yes |
| Filter Only DR | Filter | Runtime Filter | Yes |
| Filter Only (w/o rt. filt.) DR | Filter | No Runtime Filter | Yes |

problem is the robot hitting its toe against the next stair riser. We design the barrier as the distance to a hyperplane tangent to the stair after the one it is currently stepping on: $h(\mathbf{q}) = p_x^{\text{stair}} - p_x$, where p_x^{stair} is the x position of the hyperplane in the body reference frame. Thus the CBF reward is:

$$r_{\text{next stair cbf}} = \min\left(\left(-\mathbf{J}_x^{\text{sw}}(\mathbf{q})\mathbf{v} + \alpha h(\mathbf{q}), 0\right) + \exp\left(-\frac{\|\mathbf{q} - \mathbf{q}_{\text{safe}}\|^2}{\sigma^2}\right) - 1. \quad (27)$$

added to the nominal rewards including modifications to the feet clearance reward where the reference feet height now is dependent on the stair at the front of the robot and also a penalty on the swing foot force.

Hardware Deployment Considerations The stochasticity of real-world state estimation hinders the deployment of runtime explicit filters. CBF-RL is able to internalize safety by learning to map noisy observations to inherently safe actions, thus reducing the reliance of such a filter.

B. Hardware Experiments

Obstacle Course The obstacle course comprises 2 parts: the first part is the obstacle avoidance task, where the robot has to prevent itself from colliding with the obstacles, even if the velocity command intends otherwise; the second part comprises stairs constructed out of wooden pallets with a riser height of 0.14m and tread depth of 0.3m. During execution, the robot locates the obstacles approximated as

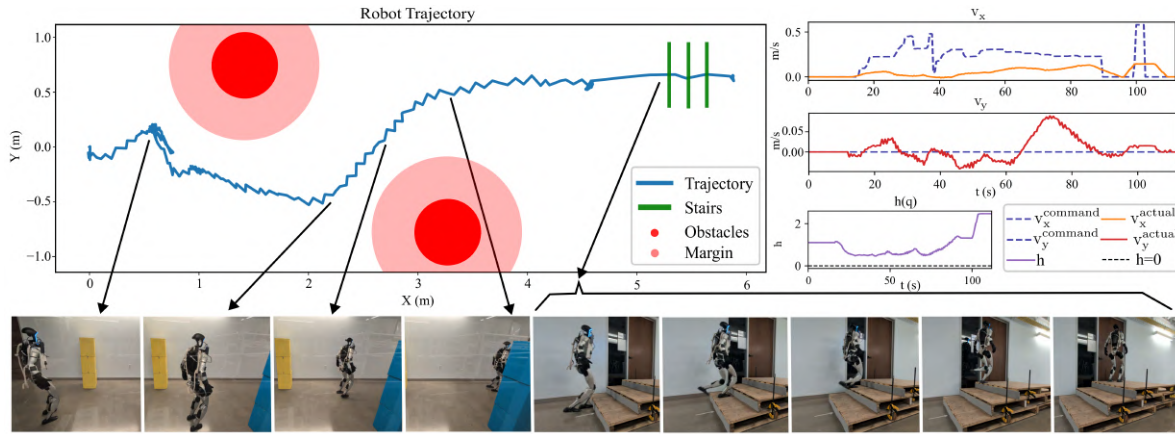


Fig. 5. Robot trajectory, h , and command vs. actual velocity visualization. The robot avoids obstacles approximated as cylinders without a runtime safety filter and climb up stairs. The velocity plots show the robot modulating its own velocities despite the command and the h plot quantifies safety.



Fig. 6. Snapshots of high stairs of riser height 0.3m. The nominal policy clips its feet against the riser and stumbles, as shown with the red CBF violations while the CBF-RL dual trained policy successfully climbs up and down. The yellow star marks the point the robot’s feet collides with the stair riser.




Fig. 7. Snapshots of outdoor experiments. The robot is able to climb up stairs of varying roughness, tread depths and riser heights.

cylinders using the ZED 2 RGB-D camera through point cloud clustering. As seen from the h values in Fig. 5, the robot modulates its own velocity to avoid the obstacle, even when the velocity command prompts it to do so. However, it has no terrain perception and only uses proprioception for stair climbing. Despite this, we observe that the robot uses proprioception to determine when to climb and how high to lift its feet, successfully climbing the wooden stairs.

Stair Climbing Robustness In order to further test the robustness of our method, we experiment both indoors and outdoors. During indoor experiments, we perform continuous runs up and down stairs and also test the performance of the policy on high stairs, with the dual-trained policy able to climb up stairs 0.3m high and the policy trained without the CBF modifications unable to do so, as shown in Fig. 6. Here we note that the robot is able to gage the depth and height of the stairs through proprioception and adjust its footsteps accordingly. For outdoor experiments, we test the dual-trained policy on concrete-poured stairs of different

roughness and sizes, with rougher stairs of riser height 0.14m / tread depth 0.33m and smoother stairs of riser height 0.15m / tread depth 0.4m respectively, as shown in Fig. 7, where the robot could adjust its center of mass by modulating the torso pitch angle to account for deeper and higher stairs.

V. CONCLUSION

This paper introduces CBF-RL, a lightweight dual approach to inject safety into learning by combining training-time CBF safety filtering with reward design, leading to policies that internalize safety and operate without a runtime filter, demonstrating its effectiveness through simulated and real-world experiments. Looking forward, we plan to incorporate automated barrier discovery, perception-based barriers, and extend the application of CBF-RL beyond locomotion to whole-body loco-manipulation, which will address broader humanoid capabilities.

Acknowledgment During the preparation of this work, the authors used AI-assisted capabilities (Claude, ChatGPT)

for grammar and editing enhancements of the text, as well as for debugging software code related to the experiments. The authors have responsibly reviewed and verified all AI-generated content to ensure correctness.

REFERENCES

- [1] D. Crowley, J. Dao, H. Duan, K. Green, J. Hurst, and A. Fern, "Optimizing bipedal locomotion for the 100m dash with comparison to human running," *arXiv preprint arXiv:2508.03070*, 2025.
- [2] Z. Li, X. B. Peng, P. Abbeel, S. Levine, G. Berseth, and K. Sreenath, "Reinforcement learning for versatile, dynamic, and robust bipedal locomotion control," *The International Journal of Robotics Research*, vol. 44, no. 5, pp. 840–888, 2025.
- [3] T. Peng, L. Bao, and C. Zhou, "Gait-conditioned reinforcement learning with multi-phase curriculum for humanoid locomotion," *arXiv preprint arXiv:2505.20619*, 2025.
- [4] T. He, J. Gao, W. Xiao, Y. Zhang, Z. Wang, J. Wang, Z. Luo, G. He, N. Sobanbabu, C. Pan, Z. Yi, G. Qu, K. Kitani, J. K. Hodgins, L. Fan, Y. Zhu, C. Liu, and G. Shi, "ASAP: Aligning Simulation and Real-World Physics for Learning Agile Humanoid Whole-Body Skills," in *Proceedings of Robotics: Science and Systems*, Los Angeles, CA, USA, June 2025.
- [5] A. Allshire, H. Choi, J. Zhang, D. McAllister, A. Zhang, C. M. Kim, T. Darrell, P. Abbeel, J. Malik, and A. Kanazawa, "Visual imitation enables contextual humanoid control," *arXiv preprint arXiv:2505.03729*, 2025.
- [6] T. E. Truong, Q. Liao, X. Huang, G. Tevet, C. K. Liu, and K. Sreenath, "Beyondmimic: From motion tracking to versatile humanoid control via guided diffusion," *arXiv preprint arXiv:2508.08241*, 2025.
- [7] Z. Su, B. Zhang, N. Rahmani, Y. Gao, Q. Liao, C. Regan, K. Sreenath, and S. S. Sastry, "Hitter: A humanoid table tennis robot via hierarchical planning and learning," *arXiv preprint arXiv:2508.21043*, 2025.
- [8] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs for safety critical systems," *IEEE Transactions on Automatic Control*, vol. 62, no. 8, pp. 3861–3876, 2016.
- [9] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: Theory and applications," in *2019 18th European control conference (ECC)*. Ieee, 2019, pp. 3420–3431.
- [10] R. Cheng, G. Orosz, R. M. Murray, and J. W. Burdick, "End-to-end safe reinforcement learning through barrier functions for safety-critical continuous control tasks," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 33, no. 01, 2019, pp. 3387–3395.
- [11] H. Ma, J. Chen, S. Eben, Z. Lin, Y. Guan, Y. Ren, and S. Zheng, "Model-based constrained reinforcement learning using generalized control barrier function," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 4552–4559.
- [12] D. Van Wijk, K. Dunlap, M. Majji, and K. Hobbs, "Safe Spacecraft Inspection via Deep Reinforcement Learning and Discrete Control Barrier Functions," *Journal of Aerospace Information Systems*, vol. 21, no. 12, pp. 996–1013, Dec. 2024.
- [13] F. P. Bejarano, L. Brunke, and A. P. Schoellig, "Safety Filtering While Training: Improving the Performance and Sample Efficiency of Reinforcement Learning Agents," *IEEE Robotics and Automation Letters*, vol. 10, no. 1, pp. 788–795, Jan. 2025, arXiv:2410.11671 [cs].
- [14] M. Alshiekh, R. Bloem, R. Ehlers, B. Könighofer, S. Niekum, and U. Topcu, "Safe reinforcement learning via shielding," in *32nd AAAI Conference on Artificial Intelligence: AAAI-18*, 2018, pp. 2669–2678.
- [15] H. Zhang, Z. Li, and A. Clark, "Model-based Reinforcement Learning with Provable Safety Guarantees via Control Barrier Functions," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. Xi'an, China: IEEE, May 2021, pp. 792–798.
- [16] H. Krasowski, J. Thumm, M. Müller, L. Schäfer, X. Wang, and M. Althoff, "Provably safe reinforcement learning: Conceptual analysis, survey, and benchmarking," *Transactions on Machine Learning Research*, 2022.
- [17] K. Dunlap, M. Mote, K. Delsing, and K. L. Hobbs, "Run time assured reinforcement learning for safe satellite docking," *Journal of Aerospace Information Systems*, vol. 20, no. 1, pp. 25–36, 2023.
- [18] K. P. Wabersich and M. N. Zeilinger, "A predictive safety filter for learning-based control of constrained nonlinear dynamical systems," *Automatica*, vol. 129, p. 109597, 2021.
- [19] X. Wang, "Ensuring safety of learning-based motion planners using control barrier functions," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 4773–4780, 2022.
- [20] N. Nilaksh, A. Ranjan, S. Agrawal, A. Jain, P. Jagtap, and S. Kolathaya, "Barrier Functions Inspired Reward Shaping for Reinforcement Learning," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, May 2024, pp. 10 807–10 813, arXiv:2403.01410 [cs].
- [21] Z. Wang, T. Ma, Y. Jia, X. Yang, J. Zhou, W. Ouyang, Q. Zhang, and J. Liang, "Omni-perception: Omnidirectional collision avoidance for legged locomotion in dynamic environments," *arXiv preprint arXiv:2505.19214*, 2025.
- [22] M. Mittal, C. Yu, Q. Yu, J. Liu, N. Rudin, D. Hoeller, J. L. Yuan, R. Singh, Y. Guo, H. Mazhar, A. Mandekar, B. Babich, G. State, M. Hutter, and A. Garg, "Orbit: A unified simulation framework for interactive robot learning environments," *IEEE Robotics and Automation Letters*, vol. 8, no. 6, pp. 3740–3747, 2023.
- [23] C. Zhang, L. Dai, H. Zhang, and Z. Wang, "Control Barrier Function-Guided Deep Reinforcement Learning for Decision-Making of Autonomous Vehicle at On-Ramp Merging," *IEEE Transactions on Intelligent Transportation Systems*, vol. 26, no. 6, pp. 8919–8932, June 2025.
- [24] H. Hailemichael, B. Ayalew, L. Kerbel, A. Ivanco, and K. Loiselle, "Safe reinforcement learning for an energy-efficient driver assistance system," *IFAC-PapersOnLine*, vol. 55, no. 37, pp. 615–620, 2022.
- [25] Y. Emam, G. Notomista, P. Glotfelter, Z. Kira, and M. Egerstedt, "Safe reinforcement learning using robust control barrier functions," *IEEE Robotics and Automation Letters*, 2022.
- [26] Y. Cheng, P. Zhao, and N. Hovakimyan, "Safe and efficient reinforcement learning using disturbance-observer-based control barrier functions," in *Learning for Dynamics and Control Conference*. PMLR, 2023, pp. 104–115.
- [27] D. Du, S. Han, N. Qi, H. B. Ammar, J. Wang, and W. Pan, "Reinforcement learning for safe robot control using control lyapunov barrier functions," in *2023 IEEE International Conference on Robotics and Automation, ICRA 2023*. IEEE, 2023, pp. 9442–9448.
- [28] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [29] S. Li and O. Bastani, "Robust model predictive shielding for safe reinforcement learning with stochastic dynamics," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 7166–7172.
- [30] M. H. Cohen and C. Belta, "Safe exploration in model-based reinforcement learning using control barrier functions," *Automatica*, vol. 147, p. 110684, 2023.
- [31] D. C. H. Tan, F. Acero, R. McCarthy, D. Kanoulas, and Z. Li, "Value Functions are Control Barrier Functions: Verification of Safe Policies using Control Theory," Dec. 2023, arXiv:2306.04026 [cs].
- [32] M. H. Cohen, N. Csomay-Shanklin, W. D. Compton, T. G. Molnar, and A. D. Ames, "Safety-critical controller synthesis with reduced-order models," in *2025 American Control Conference (ACC)*. IEEE, 2025, pp. 5216–5221.
- [33] J. Breeden, K. Garg, and D. Panagou, "Control barrier functions in sampled-data systems," *IEEE Control Systems Letters*, vol. 6, pp. 367–372, 2021.
- [34] A. Agrawal and K. Sreenath, "Discrete control barrier functions for safety-critical control of discrete systems with application to bipedal robot navigation," in *Robotics: Science and Systems*, vol. 13. Cambridge, MA, USA, 2017, pp. 1–10.
- [35] M. Ahmadi, A. Singletary, J. W. Burdick, and A. D. Ames, "Safe policy synthesis in multi-agent pomdps via discrete-time barrier functions," in *2019 IEEE 58th Conference on Decision and Control (CDC)*. IEEE, 2019, pp. 4797–4803.
- [36] C. C. Pugh, *Real mathematical analysis*. Springer, 2002.
- [37] K. Zakka, B. Tabanpour, Q. Liao, M. Haiderbhai, S. Holt, J. Y. Luo, A. Allshire, E. Frey, K. Sreenath, L. A. Kahrs, C. Sferrazza, Y. Tassa, and P. Abbeel, "Demonstrating MuJoCo Playground," in *Proceedings of Robotics: Science and Systems*, Los Angeles, CA, USA, June 2025.
- [38] L. Pinto, M. Andrychowicz, P. Welinder, W. Zaremba, and P. Abbeel, "Asymmetric actor critic for image-based robot learning," in *14th Robotics: Science and Systems, RSS 2018*. MIT Press Journals, 2018.