

One Filter to Deploy Them All: Robust Safety for Quadrupedal Navigation in Unknown Environments

Albert Lin , Shuang Peng , and Somil Bansal , *Member, IEEE*

Abstract—As learning-based methods for legged robots rapidly grow in popularity, it is important that we can provide safety assurances efficiently across different controllers and environments. Existing works either rely on *a priori* knowledge of the environment and safety constraints to ensure system safety or provide assurances for a specific locomotion policy. To address these limitations, we propose an observation-conditioned reachability-based (OCR) safety-filter framework. Our key idea is to use an OCR value network (OCR-VN) that predicts the optimal control-theoretic safety value function for new failure regions and dynamic uncertainty during deployment time. Specifically, the OCR-VN facilitates rapid safety adaptation through two key components: a LiDAR-based input that allows the dynamic construction of safe regions in light of new obstacles and a disturbance estimation module that accounts for dynamics uncertainty in the wild. The predicted safety value function is used to construct an adaptive safety filter that overrides the nominal quadruped controller when necessary to maintain safety. Through simulation studies and hardware experiments on a Unitree Go1 quadruped in agile planar navigation tasks, we demonstrate that the proposed framework can automatically safeguard a wide range of hierarchical quadruped controllers, adapts to novel environments, and is robust to unmodeled dynamics *without a priori access to the controllers or environments*—hence, “One Filter to Deploy Them All.”

Index Terms—Adaptive safety, Hamilton–Jacobi reachability analysis, robust verification, safe legged locomotion, safety filtering.

I. INTRODUCTION

LEGGED robots hold immense potential across diverse real-world applications, such as hazardous inspections [1], [2], search and rescue missions [3], [4], entertainment [5], [6], and public safety [7]. A fundamental requirement in these scenarios is the ability to operate reliably in cluttered and *a priori unknown* environments. However, achieving this reliability

poses a significant challenge, as legged locomotion controllers must balance high performance with safety (e.g., avoiding collisions) during deployment. This work focuses on designing quadrupedal controllers that enable safe, collision-free planar navigation while maintaining agility despite uncertainties in obstacles and dynamics.

Existing approaches to designing (safe) controllers for legged locomotion can be broadly categorized into model-based and reinforcement learning (RL)-based methods. Model-based methods provide provable safety guarantees using frameworks such as model predictive control (MPC), barrier functions, and reachability analysis [8], [9], [10], [11], [12]. However, model mismatches and online computational burden limit the applicability of these approaches in the wild. In contrast, RL-based controllers have demonstrated impressive agility in complex terrains and unstructured environments [13], [14], [15], [16], [17], [18], [19], [19], [20], [21], [22]. However, these controllers typically prioritize agility, treating collision avoidance as a soft constraint during training, which can result in unsafe behaviors in cluttered or unseen environments.

To address these shortcomings, recent works have explored safety critics and backup policies to safeguard RL-based controllers [23], [24], [25], [26], [27], [28], [29], [30]. These approaches precompute or learn safety critics that indicate when the nominal policy is deemed unsafe. While promising, existing methods often fail to ensure safety beyond the distribution of dynamics, environments, or locomotion policies encountered during training. In addition, learning reliable backup policies for unknown environments remain a persistent challenge.

In this work, we propose an observation-conditioned reachability (OCR) safety-filter framework, designed to integrate the agility of a nominal robot policy with robust safety for planar navigation tasks in environments with unknown obstacles and dynamical uncertainties. Our key idea is to use an OCR Value Network (OCR-VN), which predicts the optimal control-theoretic safety value function for new failure regions and dynamical uncertainties encountered during deployment. Specifically, the proposed framework achieves this adaptability through two key components: first, it leverages an observation-based exteroceptive input (a LiDAR scan in this work) to dynamically adapt the safety value function directly from raw sensory inputs, enabling robust collision avoidance in different scenarios with onboard sensing and computation. Second, it employs a disturbance estimation module to compute bounds on dynamics uncertainty (e.g., due to slippage, modeling inaccuracies, or

Received 22 May 2025; revised 26 September 2025; accepted 27 October 2025. Date of publication 16 December 2025; date of current version 16 January 2026. This work was supported in part by a NASA Space Technology Graduate Research Opportunity, in part by the NSF CAREER Program under award 2240163, and in part by the DARPA ANSR program. This article was recommended for publication by Associate Editor K. Akbari Hamed and Editor S. Behnke upon evaluation of the reviewers’ comments. (*Corresponding author: Albert Lin.*)

Albert Lin is with the University of Southern California, Los Angeles, CA 90089 USA, and also with Stanford University, Stanford, CA 94305 USA (e-mail: albertkl@stanford.edu).

Shuang Peng is with the University of Southern California, Los Angeles, CA 90089 USA (e-mail: shuangpe@usc.edu).

Somil Bansal is with Stanford University, Stanford, CA 94305 USA (e-mail: somil@stanford.edu).

The experiment videos can be found on the project website here: https://sia-lab-git.github.io/One_Filter_to_Deploy_Them_All. Project Website: https://sia-lab-git.github.io/One_Filter_to_Deploy_Them_All

Digital Object Identifier 10.1109/TRO.2025.3644957

low-level tracking errors) using recent state-action histories and adapts the robustness of the safety value function based on these bounds.

The OCR framework predicts when a nominal policy might violate safety and provides corrective control commands for the robot if necessary. A key strength of the proposed framework is its generality—it can be deployed with a wide variety of nominal legged locomotion policies without requiring any retraining or policy-specific tuning. In addition, we propose a Hamilton–Jacobi reachability-based method to train the OCR-VN, ensuring robust and efficient safety filtering. We validate our approach through extensive simulations and real-world experiments on a Unitree Go1 quadruped, demonstrating that the OCR framework provides a reliable safety layer across multiple existing legged locomotion policies (both model-based and learning-based) and a variety of environments, without requiring prior knowledge of the specific policy or environment. In summary, our key contributions are:

- 1) A reachability-based safety-filtering framework that ensures safety across diverse quadruped controllers and environments, without a priori access to the controller or the environment;
- 2) An online adaptation mechanism that dynamically adapts the system safety to real-world environment variations and modeling uncertainties;
- 3) Simulation and hardware experiments demonstrating the superior efficacy and robustness of the proposed approach in ensuring safe legged locomotion.

II. RELATED WORKS

A. Safe Legged Locomotion

1) *Model-based safety*: Traditional approaches for obstacle avoidance use collision-free motion planning techniques in the configuration space [31], [32], [33]. They satisfy kinematic safety constraints but do not consider the dynamics of the system, limiting motions to slow quasistatic trajectories. However, recent advances in agile locomotion and its applications have resulted in the need to consider dynamics.

Model-based approaches, such as MPC, use hand-designed or learned dynamics models to compute optimal maneuvers that are dynamically feasible and satisfy safety constraints [26], [34], [35], [36], [37], [38], [39]. Despite their impressive performance, such model-based approaches are generally computationally intensive for online settings and can run into safety feasibility issues, especially in cluttered obstacle environments. In addition, although they often perform well in settings that are captured accurately by their models, the safety guarantees are not robust to model mismatches that a robot might encounter in the wild [13], [40].

2) *RL-based safety*: Given the challenges associated with existing model-based approaches for agile locomotion, model-free RL-based approaches have emerged as popular alternatives. RL-based approaches have found remarkable success in synthesizing efficient and robust locomotion in the real world, especially as the availability of high-fidelity simulators has increased [41], [42]. They are well-suited to handle complex high-dimensional

systems, multimodal feedback signals, and difficult-to-specify task objectives [43].

Previous studies have optimized locomotion policies for specific skills such as agility [44], [45], resilience [46], [47], [48], and difficult terrain traversal [13], [20], [21], [48], [49], [50], [51], [52], [53], [54]. However, these works typically focus on maximizing agility without regard to safe navigation. These methods can be combined with high-level collision-free planners; however, they suffer from the aforementioned limitations of model-based controllers and restricted mobility [29], [38], [55], [56]. Other works consider safe navigation during the learning process by including a large collision penalty in the reward function to incentivize collision-avoidance [14], [55], [57], [58], [59], [60], [61], [62], [63], [64]. Unfortunately, there are no formal guarantees of safety, and the synthesized locomotion policies can degrade in safety when transferred to the real world due to a distribution shift away from the environments seen during training.

Recently, constrained reinforcement learning (CRL) has emerged as a prominent framework for incorporating safety considerations into policy learning by formulating decision-making problems as constrained Markov decision processes (CMDPs) [65]. CRL methods aim to optimize performance while satisfying constraints, often expressed as cost or risk thresholds, either through Lagrangian-based optimization, reward shaping, or primal-dual approaches [66], [67], [68], [69]. While CRL is attractive for its ability to learn directly from environment interactions, it typically provides probabilistic or empirical safety assurances within the distribution of training. Thus, the resulting policies are not automatically robust to new environments with unseen dynamical uncertainties. This can lead to a loss of safety assurances in practice.

3) *Certificate-Based Safety*: In order to provide rigorous safety assurances, many works have proposed certificate-based safety methods within both model-based [26], [34], [35], [37], [39] and RL-based [23], [24], [25], [27], [28], [29], [30], [70] frameworks, most often using control barrier functions or reachability-based value functions. These methods are typically reliant on the offline availability of a certificate function or dynamics, which limits their applicability to complex real-world systems.

Some recent works learn adaptive safety certificates and recovery policies to ensure safety at runtime [23], [24], [25], [27], [28], [29], [70]. He et al. [29] proposed Agile But Safe (ABS), an approach that co-designs performance and control-theoretic safety controllers via RL and switches between them as necessary to maintain safety. Since the computed safety assurances are policy-dependent, they can be suboptimal and overly conservative, depending on the quality of the nominal policy. Moreover, the safety controller needs to be retrained as the policy is fine-tuned, which can be cumbersome and time-consuming. In addition, since approaches like ABS lack active mechanisms to account for unmodeled dynamics during deployment (e.g., slippage), the computed safety assurances are valid only within the distribution of environments and locomotion policies seen during training. These limitations are especially relevant due to the rising popularity of diverse learning-based policies, whose safety assurances must be updated as they are fine-tuned.

Our work most closely aligns with certificate-based safety methods like ABS, but we address several existing limitations. Instead of computing the policy-conditioned safety function, we compute the optimal control-theoretic safety controller via Hamilton–Jacobi (HJ) reachability analysis. This enables us to construct a safety filter for any nominal controller without needing to recompute the safety assurances. We discard the need for high-fidelity simulators by using a reduced-order system and robustly handle the model gap as an adversarial disturbance in the dynamics. By estimating and adapting to disturbance bounds during deployment, our proposed framework is able to ensure safety more robustly across a range of settings and policies compared to previous works.

B. Reachability-Based Safety Filters

This work extends the class of Hamilton–Jacobi (HJ) reachability-based filters [71], [72], [73], which ensure safety at deployment by overriding a nominal controller when necessary to preserve safety. Due to the computational burden, traditional reachability-based safety filters are typically constructed offline for systems assumed to be known a priori [74]. In the case of quadrupedal navigation, a quadruped using a specific locomotion policy is often abstracted as a reduced-order system for computational tractability. The dynamics of the system depend on the locomotion policy; thus, reachability-based filters will not readily safeguard different locomotion policies. In addition, it is challenging to adapt safety assurances to different obstacle configuration and terrain properties due to the computational burden [75], [76]. This work overcomes these limitations by distilling the safety solutions for a wide range of system settings into an observation-conditioned reachability-based value network (OCR-VN). During deployment, the robot adapts by querying the trained OCR-VN with the current state, control, and observation history.

1) *Connections to control barrier functions*: As explained in Section IV-B, we employ reachability-based safety filters. Readers may notice similarities between this type of filter and the popular CBF-QP formulation [77]. This connection is not coincidental: the infinite-time value functions computed via HJ reachability satisfy the conditions of valid CBFs [78]. However, our approach inherits several important advantages. First, HJ reachability provides a constructive framework for synthesizing optimal certificates for general nonlinear systems, whereas traditional CBF methods lack such a procedure [74]. Second, HJ reachability is designed to account for worst-case disturbances, which is crucial for safety-critical settings. Finally, the QP induced by the HJ value function is guaranteed to remain feasible, unlike standard CBF-QPs that may suffer infeasibility under control constraints [71]. These distinctions motivate our choice to build upon reachability-based safety filters in the proposed OCR framework.

III. PROBLEM SETUP

See Table I for notation. In this work, we are interested in ensuring the safety of a quadruped robot for agile planar navigation tasks in an *a priori unknown* environment $e \in \mathcal{E}$.

TABLE I
NOMENCLATURE

Symbol	Definition.
t, T	Time, time horizon.
$e \in \mathcal{E}$	Environment.
$x \in \mathcal{X}$	State.
$u \in \mathcal{U}$	Control input.
$d \in \mathcal{D}$	Disturbance input.
f	System dynamics.
$\pi, \pi^{\text{high}}, \pi^{\text{low}}$	Hierarchical, high-level, and low-level policies.
$\xi_{x,t}^{\pi}(\tau)$	State achieved at time τ by starting at initial state x at time t and applying policy π over $[t, \tau]$.
\mathcal{F}, l	Failure set and function.
$o \in \mathcal{O}$	Observation.
$x_r \in \mathcal{X}_r \subseteq \mathcal{X}, f_r$	Reduced-order state and dynamics.
$\omega = (v, w)$	Twist ω consisting of velocity v and yaw rate w .
$\bar{d}_{p_x, p_y}, \bar{d}_{p_\theta}$	Disturbance bound in p_x, p_y and in p_θ .
$\bar{d}_r = [\bar{d}_{p_x, p_y}, \bar{d}_{p_\theta}]$	Disturbance bound for reduced-order system.
V, V_ψ	Ground-truth and learned value functions.
Abbreviation	Definition.
OCR	Observation-conditioned reachability.
ABS	Agile But Safe [29].
WTW	Walk-These-Ways [21].
MPC	Model predictive control policy by Unitree [81].
PS	Predictive sampling-based planner.
NVE	Naive planner.
HMN	Human-teleoperated planner.

Here, e contains all the information needed to inform the effects of the environment on dynamics, as well as failure regions. For example, e can include the terrain geometry, friction coefficients, and obstacle locations.

We model the quadruped as a nonlinear dynamical system with state $x \in \mathcal{X}$, control $u \in \mathcal{U}$, and dynamics $\dot{x} = f(x, u; e)$ governing how x evolves over time until a final time horizon T . The dynamics are also affected by the environment e , e.g., by the effects of a slippery floor. We denote the robot observations from proprioception and/or exteroception as $o_x^e = h(x; e) \in \mathcal{O}$. For exteroception, we primarily deal with LiDAR scans in this work, though other sensors can also be used. We denote the set of failure states as $\mathcal{F}^e \subseteq \mathcal{X}$ (e.g., collision states) which the robot is not allowed to enter. The failure set can be represented by the zero-sublevel set of a Lipschitz-continuous function $l^e : \mathcal{X} \rightarrow \mathbb{R}$, i.e., $x \in \mathcal{F}^e \Leftrightarrow l^e(x) \leq 0$. Note that \mathcal{F}^e, l^e are also functions of e . We assume that \mathcal{F}^e is static and leave it to future work to handle dynamic obstacles.

Let π_{nom} denote a hierarchical nominal policy for the quadruped that takes the system history and outputs the robot control. We assume that π_{nom} consists of a high-level planner, $\pi_{\text{nom}}^{\text{high}}$, that provides twist commands $\omega := (v, w)$ consisting of a forward velocity v and a yaw rate w . These twist commands are tracked by a low-level locomotion policy, $\pi_{\text{nom}}^{\text{low}}$, e.g., an RL-based policy [21], [29], [44] or an MPC-based policy [32], [79], that ultimately provides control inputs u for the robot. This architecture is popular in the legged robotics literature where $\pi_{\text{nom}}^{\text{high}}$ is typically designed for collision avoidance and navigation, and $\pi_{\text{nom}}^{\text{low}}$ could be an agile locomotion policy that can handle different terrains that the robot might encounter.

Let $\xi_{x,e,t}^{\pi}(\tau)$ denote the state achieved at time $\tau \in [t, T]$ by starting at initial state x at time t and applying the control policy

π over $[t, \tau]$ in environment e . Our goal is to compute a safe policy π_{safe} that ensures that the quadruped remains outside of the failure set at all times, i.e., $\pi_{\text{safe}} : \forall \tau \in [0, T], \xi_{x,e,0}^{\pi_{\text{safe}}}(\tau) \notin \mathcal{F}^e$, while preserving the underlying performance of π_{nom} to the extent possible. The key challenge in designing such a policy is that the robot environment (and hence the failure set and the robot dynamics) is not known beforehand, necessitating a real-time update of the safety policy with the environment. A second challenge stems from the fact that we want the safety framework to be agnostic to different nominal policies.

IV. BACKGROUND

A. Hamilton–Jacobi Reachability

We build upon Hamilton–Jacobi (HJ) reachability analysis, a popular formal verification tool for computing safety guarantees for general nonlinear dynamical systems [80], [81]. For reachability analysis, we will consider a more general form of dynamics $\dot{x} = f(x, u, d)$, where $d \in \mathcal{D}$ represents the disturbance. Later on, we will use d to model potential uncertainty in the system dynamics model. We also omit the dependence on environment for now for brevity purposes.

HJ reachability analysis is concerned with computing the system’s initial-time Backward Reachable Tube, which we denote as BRT. We define BRT as the set of all initial states $x \in \mathcal{X}$ starting from which, for all control signals $\mathbf{u}(\cdot)$, there exists a disturbance signal $\mathbf{d}(\cdot)$ such that the system will inevitably enter the failure set \mathcal{F} within the time horizon $[0, T]$:

$$\text{BRT} := \{x \in \mathcal{X} : \forall \mathbf{u}(\cdot), \exists \mathbf{d}(\cdot), \exists \tau \in [0, T], \xi_{x,0}^{\mathbf{u}(\cdot), \mathbf{d}(\cdot)}(\tau) \in \mathcal{F}\}. \quad (1)$$

By ensuring that the system remains outside of BRT, we guarantee system safety for the time horizon T .

In HJ reachability, computing BRT is formulated as a robust optimal control problem. First, we implicitly represent the failure set \mathcal{F} by a failure function $l(x)$ whose zero-sublevel set yields \mathcal{F} : $\mathcal{F} = \{x \in \mathcal{X} : l(x) \leq 0\}$. $l(x)$ is commonly the signed distance function to \mathcal{F} . Next, we define the cost function corresponding to a control signal $\mathbf{u}(\cdot)$ and disturbance signal $\mathbf{d}(\cdot)$ to be the minimum of $l(x)$ over the trajectory starting from state x and time t :

$$J_{\mathbf{u}(\cdot), \mathbf{d}(\cdot)}(x, t) := \min_{\tau \in [t, T]} l\left(\xi_{x,t}^{\mathbf{u}(\cdot), \mathbf{d}(\cdot)}(\tau)\right). \quad (2)$$

Since the control aims to avoid \mathcal{F} under worst-case disturbance, the value function corresponding to this robust optimal control problem is

$$V(x, t) := \max_{\mathbf{u}(\cdot)} \min_{\mathbf{d}(\cdot)} J_{\mathbf{u}(\cdot), \mathbf{d}(\cdot)}(x, t). \quad (3)$$

By defining our optimal control problem in this way, we can easily recover BRT using the value function. The value function being nonpositive implies that the failure function is nonpositive somewhere along the optimal trajectory, or in other words, that the system will inevitably enter \mathcal{F} . Conversely, the value function being positive implies that there exists a control signal that will prevent the system from entering \mathcal{F} even under the worst-case disturbance signal. Thus, BRT is computed as the

zero-sublevel set of the initial-time value function:

$$\text{BRT} = \{x \in \mathcal{X} : V(x, 0) \leq 0\}. \quad (4)$$

The value function in (3) can be computed using dynamic programming, resulting in the following final value Hamilton–Jacobi–Isaacs variational inequality (HJI-VI):

$$\begin{aligned} \min\{D_t V(x, t) + H(x, t, \nabla V(x, t)), l(x) - V(x, t)\} &= 0 \\ V(x, T) &= l(x) \quad \forall t \in [0, T]. \end{aligned} \quad (5)$$

$D_t V(x, t)$ and $\nabla V(x, t)$ represent the temporal derivative and spatial gradient of the value function $V(x, t)$, respectively. The Hamiltonian $H(x, t, \nabla V(x, t))$ encodes how the control and disturbance interact with the system dynamics:

$$H(x, t, \nabla V(x, t)) := \max_{u \in \mathcal{U}} \min_{d \in \mathcal{D}} \nabla V(x, t) \cdot f(x, u, d). \quad (6)$$

Under mild assumptions, the value function computed as the solution to (5) is bounded, Lipschitz continuous, and differentiable almost everywhere [81]. The value function also induces the optimal safety controller:

$$\mathbf{u}^*(x, t) := \arg \max_{u \in \mathcal{U}} \min_{d \in \mathcal{D}} \nabla V(x, t) \cdot f(x, u, d). \quad (7)$$

Intuitively, the optimal safety controller aligns the system dynamics in the direction of the value function’s gradients, thus steering the system towards higher-value states, i.e., away from \mathcal{F} . An important result of HJ reachability theory is that safety is guaranteed despite worst-case disturbances if the system starts outside of BRT and applies the control in (7) at the BRT boundary [71].

B. HJ Reachability-Based Safety Filtering

In this work, we will use HJ reachability analysis to maintain the safety of the quadruped robot during deployment via HJ reachability-based safety filtering, where the computed safety value function is used to construct a safety filter that guarantees system safety. While there are many different types of safety filters that can be constructed, we use the smooth least-restrictive safety filter, which aims to maximally preserve the underlying performance of a given nominal policy π_{nom} by intervening as seldomly and as lightly as possible [71]. When outside of the system BRT, the smooth least-restrictive filter π_{safe} outputs the nominal control. At the boundary of the BRT, π_{safe} outputs a safe control as close as possible to the nominal control by solving what amounts to a quadratic program (QP) for control-affine systems

$$\pi_{\text{safe}}(x, t) = \begin{cases} \pi_{\text{nom}}(x, t), & V(x, t) > 0 \\ \pi_{\text{QP}}(x, t), & V(x, t) = 0 \end{cases} \quad (8)$$

where $\pi_{\text{QP}}(x, t)$ is obtained by solving

$$\begin{aligned} \arg \min_{u \in \mathcal{U}} \|u - \pi_{\text{nom}}(x, t)\|_2^2 \\ \text{s.t. } D_t V(x, t) + \min_{d \in \mathcal{D}} \nabla V(x, t) \cdot f(x, u, d) &= 0. \end{aligned} \quad (9)$$

Intuitively, the constraint on Line (9) enforces safe control of the system at the BRT boundary, where the system safety could



Fig. 1. Our proposed observation-conditioned reachability-based (OCR) safety-filter framework automatically safeguards different controllers in diverse settings for agile planar navigation tasks *without a priori access to the controllers or environments*. A trained OCR value network governs the switch between nominal and filtered control using an onboard LiDAR sensor. The framework successfully safeguards a variety of high-level planners, including (a) learning-based, (c, f, k) model-based, (b, d, g, h, i, j) human teleoperated, and (e) naive planners, on top of different low-level locomotion policies, including (a, f, i, j, k) learning-based and (b, c, d, e, g, h) model-based policies. Safety is maintained despite (a, b, c) narrow corridors, (d, i, j) rough terrains, (e, k) dynamic obstacles, (f) external disturbances, and (h) collision-seeking human teleoperation.

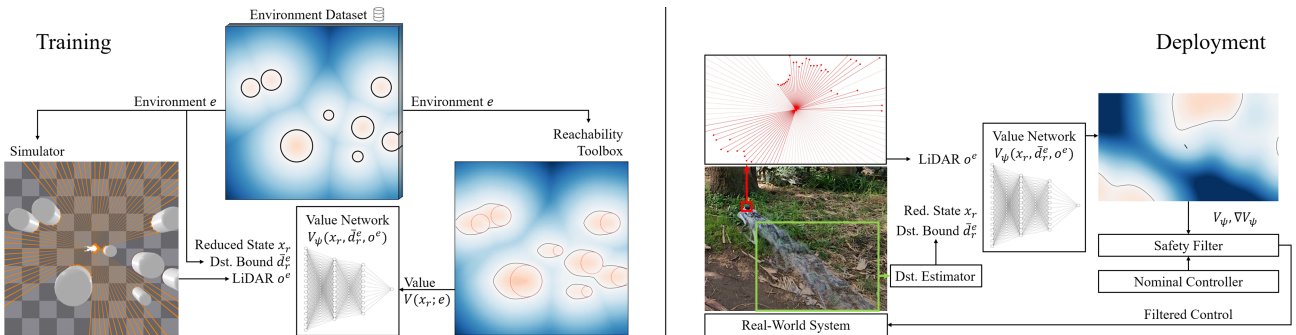


Fig. 2. OCR framework. (Left) During training, we generate environments with random obstacles and disturbance bounds. The OCR-VN is trained to predict the value function (visualized over a grid) using the disturbance bound, the LiDAR reading, and the state. (Right) During deployment, the OCR-VN is queried with the observed LiDAR reading, the disturbance bound estimated using the most recent state and action history, and the current state estimate to construct an adaptive safety filter.

be in jeopardy. It has been shown that the filter defined by (8) is guaranteed to maintain system safety under worst-case disturbances as long as the system starts outside of the BRT [71]. In the next section, we propose to construct an adaptive version of this safety filter by using an OCR-VN that predicts the safety value function for new failure regions and dynamic uncertainties encountered during deployment.

V. APPROACH

The main difficulty in directly applying traditional HJ reachability methods in Section IV to the problem in Section III is that the system dynamics $f(x, u; e)$ and the failure set \mathcal{F}^e are functions of the unknown environment e . Thus, a safety controller must be able to adapt to both *dynamics uncertainty* and *environment uncertainty*. To overcome this challenge, we

propose the observation-conditioned reachability (OCR) safety-filter framework, which uses an OCR Value Network (OCR-VN) to predict a safety value function conditioned on the recent system and observation history. The online adaptation is enabled by two key components: a LiDAR-based exteroceptive input that encodes new obstacle configurations (see Section V-A3) and a disturbance estimation module (see Section V-B1). The resulting safety value function is then used to filter the nominal policy to ensure safety while maintaining the agility of the underlying policy. Ultimately, our safety framework provides high-level twist commands, which, when combined with the low-level nominal policy, maintains the robot’s safety.

The OCR framework can be divided into two distinct phases: training and deployment. These are illustrated in Fig. 2. During the training phase, we collect a diverse dataset of raw LiDAR observations and uncertainty bounds to train the OCR-VN on.

During the deployment phase, we query the OCR-VN with onboard observations and estimated uncertainty bounds to automatically ensure safety across different locomotion policies and environments. These two phases are described in more detail next.

A. Training Phase of the OCR Framework

1) *System dynamics model*: During the training phase, we aim to distill the ground-truth value functions for a diverse range of environments e into an OCR-VN. The first difficulty that we encounter is how to model the full quadruped dynamics $f(x, u; e)$, which is high-dimensional and complex. Our key insight is that many quadruped control schemes are hierarchically composed of a high-level navigation planner and a low-level locomotion policy. From the perspective of the high-level planner, the quadruped, along with its locomotion policy, form a system with a *reduced-order* dynamics model. Thus, we propose to use a reduced-order dynamics model $f_r(x_r, \omega)$, where $x_r \in \mathcal{X}_r \subseteq \mathcal{X}$ is the reduced state and ω is the control input of the reduced-order model. We capture any possible modeling errors of $f_r(x_r, \omega)$ as disturbances in the system. We propose to estimate bounds on these disturbances during deployment for robust adaptive safety guarantees.

We set $f_r(x_r, \omega)$ to the dynamics of a 3-D Dubins car system with state $x_r = (p_x, p_y, p_\theta)$, where (p_x, p_y) is the quadruped's 2-D location, and p_θ is the quadruped's heading. We model the error in $f_r(x_r, \omega)$ as an unknown additive disturbance d_r^e , which is a function of the underlying environment e . Thus, the reduced-order dynamics $f_r(x_r, \omega, d_r^e)$ is given as

$$\dot{p}_x = v \cos p_\theta + d_{p_x}^e, \quad \dot{p}_y = v \sin p_\theta + d_{p_y}^e, \quad \dot{p}_\theta = w + d_{p_\theta}^e \quad (10)$$

with the control input $\omega = (v, w)$ being a commanded twist that includes forward velocity $v \in [v_{\min}, v_{\max}]$ and yaw rate $w \in [w_{\min}, w_{\max}]$. The disturbance input $d_r^e = (d_{p_x}^e, d_{p_y}^e, d_{p_\theta}^e)$ consists of a bounded additive disturbance in position $\| [d_{p_x}^e, d_{p_y}^e] \| \leq \bar{d}_{p_x, p_y}^e$ and a bounded additive disturbance in heading $|d_{p_\theta}^e| \leq \bar{d}_{p_\theta}^e$. We use \bar{d}_r^e to denote the disturbance bound tuple $\bar{d}_r^e = (\bar{d}_{p_x, p_y}^e, \bar{d}_{p_\theta}^e)$, which is a function of the underlying environment e . Since the disturbances are additive in all state variables, the disturbance bounds can always be chosen large enough to contain any possible modeling errors, although at the cost of model conservatism. In order to be robust to the modeling error, we assume that the disturbances are adversarial in nature.

Remark: We adopt a Dubins car model as the reduced-order dynamics to balance model fidelity with tractability, since grid-based reachability methods scale poorly to high-dimensional systems. This simplification is sufficient for capturing the safety-relevant aspects of agile planar navigation. Importantly, the OCR framework is agnostic to this choice. Other reduced-order models and disturbance assumptions may be employed depending on the desired tradeoff between fidelity, conservatism, and computational complexity.

2) *Data generation*: Based on our modeling choices in Section V-A1, we let the underlying environment $e = (\mathcal{F}^e, \bar{d}_r^e)$ consist of the failure set \mathcal{F}^e and the disturbance bound \bar{d}_r^e

described in Section V-A1. We generate each environment e by randomly spawning 2-D obstacles at various locations and sampling a disturbance bound. For each of the generated environments, we compute the initial-time ground-truth value function $V(x_r, 0; e)$ using the `hj_reachability` Python toolbox [82], which we denote as $V(x_r; e)$ for brevity. We compute the converged value function, that is, when we take the time horizon $T \rightarrow \infty$. This ground-truth value function is then used to generate data for training the OCR-VN.

Specifically, for each training environment, we generate training pairs $((x_r, \bar{d}_r^e, o^e), V(x_r; e))$ by setting the system origin to different states in the environment and rendering the corresponding observation o^e at the origin. For each system origin and its corresponding observation o^e , we sample the value function $V(x_r; e)$ and its spatial gradients $\nabla V(x_r; e)$ at uniformly random states x_r (specified in the frame of the new system origin) that are not occluded by obstacles. The collected data are then used to train the OCR-VN via supervised learning as described in the next section. Details on the specific parameters of our data generation are provided next.

To generate an environment, we spawn n circular 2-D obstacles, where n is uniformly sampled from $\{1, \dots, 10\}$. The radius r of each obstacle is uniformly sampled from $[0.1, 1]$ m. The location (p_x, p_y) of each obstacle is uniformly sampled from $[-5, 5]$ m \times $[-5, 5]$ m. Disturbance bounds \bar{d}_{p_x, p_y}^e and $\bar{d}_{p_\theta}^e$ are uniformly sampled from $[0, 1]$ m/s and $[0, 2]$ rad/s, respectively. In this work, we use LiDAR observations which consists of 100 evenly spaced angles from $[-\pi, \pi]$ rad and clipped to within $[0.2, 10]$ m, determined by hardware limits.

Remark: For training, we model obstacles as circles to balance computational simplicity with sufficient expressiveness for LiDAR-based perception. This choice enables efficient dataset generation while exposing the OCR-VN to diverse observation patterns. Although circular obstacles are used in training, the framework generalizes in practice to more complex geometries (e.g., walls, irregular shapes). Users may extend the data generation process as needed to suit higher dimensional observations or more specialized scenarios.

For obtaining the ground-truth value function, we use a grid of shape $(100, 100, 60)$ spanning $[-5, 5]$ m \times $[-5, 5]$ m \times $[-\pi, \pi]$ rad and a time horizon of 2 s, when the value function approximately converges. The control bounds for the system are $v \in [0, 2]$ m/s and $|w| \leq 2$ rad/s. We generate a total of 1000 training environments and 100 validation environments. Generating the datasets takes roughly 35 min on an NVIDIA 3090Ti GPU. The ground-truth value function $V(x_r; e)$ and the corresponding LiDAR observation o^e for a validation environment are shown in Fig. 3.

We generate a training batch by first sampling 10 training environments. For each environment, we set the system origin to 10 different states sampled outside of the obstacle set and capture the corresponding LiDAR observation, resulting in ten different egocentric observations. This is done to increase the diversity of LiDAR observations seen during training without substantially increasing the computational effort, which is important because the observations are a high-dimensional network input. For each sampled system origin, we query the ground-truth value

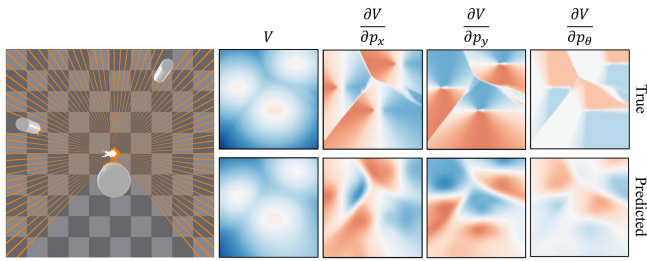


Fig. 3. (Left) A LiDAR observation o^e in a validation environment where $\bar{d}_{p_x, p_y}^e = 0.82$ m/s, $\bar{d}_{p_\theta}^e = 0.56$ rad/s. (Right top-row) The ground-truth value function and its spatial gradients. (Right bottom-row) OCR-VN predictions using o^e and \bar{d}_r^e . As shown above, the OCR-VN predictions for the value function and its spatial gradients are highly accurate.

function and its spatial gradients at 500 random states that are not occluded by obstacles. This ultimately results in $N = 50\,000$ samples per training batch.

3) *OCR-VN architecture and training*: We train an OCR-VN to predict $V(x_r; e)$ from the reduced system state x_r , disturbance bound \bar{d}_r^e , and the LiDAR observation o^e captured at the system origin. Let us denote the value predicted by the OCR-VN as $V_\psi(x_r, \bar{d}_r^e, o^e)$, where ψ denotes the learnable weights of the neural network. We choose V_ψ to be a multilayer perceptron with 4 hidden layers of 512 neurons each. We desire V_ψ to accurately model the spatial gradients of V as well, since they will be used in the safety filter construction. Thus, we use sinusoidal activations, which have been shown to lead to more effectively gradient modeling than ReLU activations [83].

For each training batch as described in Section V-A2, we minimize the MSE training loss

$$\frac{1}{N} \sum_{i=1}^N (\|V_\psi^i(x_r, \bar{d}_r^e, o^e) - V^i(x_r; e)\|_2^2 + \| \nabla_x V_\psi^i(x_r, \bar{d}_r^e, o^e) - \nabla_x V^i(x_r; e) \|_2^2) \quad (11)$$

which includes a loss term for both the value and its spatial gradients. We use the Adam optimizer with a learning rate of 10^{-5} . Training converges in roughly 8 hours on an NVIDIA 3090Ti GPU. As shown in Fig. 3, the OCR-VN can accurately model the value function and its spatial gradients for a validation environment using a single LiDAR observation o^e . Please see the appendix of the extended version of this article for a study on the importance of the sinusoidal activations and the spatial gradient loss term.¹

4) *OCR-VN Calibration*: The OCR-VN will inevitably contain learning errors that can be critical for safety. To safeguard against these errors, we calibrate the OCR-VN by shifting its output by a probabilistic error bound δ computed on the validation dataset. We will refer to δ as the ‘‘calibration level’’ and the shifted network output as the ‘‘calibrated output’’. We employ conformal prediction, a popular uncertainty quantification tool in the machine learning literature, to compute δ up to a desired confidence β and violation rate ϵ [84], [85].

Our procedure is as follows. Select a desired confidence $\beta \in (0, 1)$ and violation rate $\epsilon \in (0, 1)$. Sample N calibration points according to a distribution \mathbb{P} over the validation dataset

TABLE II
CALIBRATION LEVELS FOR THE OCR-VN WITH $N = 10^7, \beta = 10^{-12}$

Violation Rate ϵ	10^{-1}	10^{-2}	10^{-3}
Calibration Level δ (m)	0.22	0.49	0.99

Δ . Compute the conformal scores $\{s_i\}_{i=1}^N$ as the prediction errors $s_i = V_\psi^i(x_r, \bar{d}_r^e, o^e) - V^i(x_r; e)$. The following theorem provides a probabilistic bound on the OCR-VN error.

Theorem 1 (Conformal OCR-VN Calibration): Compute the number of ‘‘outliers’’ k as

$$\arg \max_k k : \sum_{i=0}^k \binom{N}{i} \epsilon^i (1 - \epsilon)^{N-i} \leq \beta \quad (12)$$

where β, ϵ , and N are as defined above. Compute the calibration level δ as the $\frac{N-k}{N}$ quantile of the conformal scores $\{s_i\}_{i=1}^N$. Then, with probability at least $1 - \beta$ over the draws of the calibration samples, the following holds:

$$\mathbb{P}_{(x_r, e, \bar{d}_r^e, o^e) \in \Delta} (V_\psi(x_r, \bar{d}_r^e, o^e) - V(x_r; e) > \delta) \leq \epsilon. \quad (13)$$

The proof of Theorem 1 is presented in the appendix of the extended version of this article.¹ Ignoring the confidence parameter β for a moment, Inequality (13) tells us that the volume of the validation dataset for which the OCR-VN overestimates the safety value by more than δ is bounded by the violation parameter ϵ . This enables us to be sure, up to a desired ϵ , that the true safety value is at least as large as that predicted by the OCR-VN, once adjusted by δ . However, this holds only for the distribution from which the calibration samples are drawn, which can lead to safety issues in practice. Care should be taken to use a validation dataset that will reflect the expected deployment settings. Although we do not explore it in this work, we suggest future work can improve the robustness of this calibration step using, e.g., out-of-distribution detection methods during deployment.

To interpret the confidence parameter β , note that the calibration level δ is a random variable that depends on the randomly sampled calibration set. It might be the case that we happen to draw an unrepresentative calibration set, in which case the ϵ bound does not hold. β controls the probability of this adverse event, which regards the correctness of the probabilistic guarantee given by Inequality (13). Fortunately, β goes to 0 exponentially with N , so β can be chosen sufficiently small, such as 10^{-12} , when we sample large N . $1 - \beta$ will then be so close to 1 that it does not have any practical importance.

We set $N = 10^7, \beta = 10^{-12}$ and compute δ corresponding to various ϵ in Table II. We select $\delta = 0.49$ m associated with $\epsilon = 10^{-2}$ as a comfortable tradeoff between conservatism and performance for the tasks in this work.

B. Deployment Phase of the OCR Framework

In this section, we describe the steps involved in deploying the OCR framework onto a quadruped robot. During deployment,

¹https://sia-lab-git.github.io/One_Filter_to_Deploy_Them_All/resources/manuscript.pdf

the quadruped periodically receives a LiDAR observation o^e from an onboard sensor as a function of the environment e . Using a LiDAR-based state localization algorithm, the quadruped maintains an estimate of its reduced state x_r in the relative frame defined by o^e . The quadruped then estimates the disturbance bound \bar{d}_r^e as a function of the most recent state and action history $(x_r^{i-k:i}, \omega^{i-k:i-1})$ using an online disturbance bound estimation scheme. It queries the OCR-VN with the input (x_r, \bar{d}_r^e, o^e) to construct an adaptive safety filter. We describe the details of the online disturbance bound estimation scheme and the adaptive safety filter next.

1) *Online disturbance bound estimation*: We propose to estimate the disturbance bound \bar{d}_r^e using the most recent state and action history, to enable rapid adaptation to dynamical uncertainty. Let $x_r^{i-k:i}$ and $\omega^{i-k:i-1}$ denote a discrete history of estimated states and twist commands ending at time τ . To efficiently roll out the state trajectory in discrete time for online computation, we use Euler's method. Since disturbance is additive in $f_r(x_r, \omega, d_r^e)$ in (10), we have

$$\begin{aligned} x_r^i &\approx x_r^{i-k} + \sum_{j=i-k}^{i-1} \eta \cdot \left(f_r(x_r^j, \omega^j) + (d_r^e)^j \right) \\ &\approx x_r^{i-k} + \sum_{j=i-k}^{i-1} \eta \cdot f_r(x_r^j, \omega^j) + \eta \sum_{j=i-k}^{i-1} (d_r^e)^j \end{aligned}$$

where η is the discrete time step. To ensure that we capture systemic disturbances in the dynamics and ignore transient noise in the state estimation, we assume that the disturbance input takes the form of a low-frequency signal and remains a constant d_r^e throughout the short history

$$x_r^i \approx x_r^{i-k} + \sum_{j=i-k}^{i-1} \eta \cdot f_r(x_r^j, \omega^j) + \eta \cdot k \cdot d_r^e.$$

Thus, we estimate the disturbance d_r^e as

$$d_r^e \approx \frac{x_r^i - \hat{x}_r^i}{\eta \cdot k} \quad (14)$$

where $\hat{x}_r^i := x_r^{i-k} + \sum_{j=i-k}^{i-1} \eta \cdot f_r(x_r^j, \omega^j)$ is the state predicted by disturbance-free dynamics. Intuitively, d_r^e is the error in the state prediction normalized by the prediction time horizon $\eta \cdot k$. We set $\eta \cdot k = 2$ s in practice.

By computing d_r^e over the most recent φ discrete time steps ending at the current time t , i.e., for $\tau = t - j \cdot \eta, \forall j \in \{0, 1, \dots, \varphi - 1\}$, we construct a sliding window of estimated disturbances $d_r^{1:\varphi}$. To ensure good coverage of the estimated disturbances, we compute the disturbance bound \bar{d}_r^e as

$$\bar{d}_r^e = |\mu((d_r^{1:\varphi})_c)| \pm b \cdot \sigma((d_r^{1:\varphi})_c) \quad (15)$$

where the coverage parameter $c \in [0, 1]$ determines the middle fraction of the sorted window of disturbances $d_r^{1:\varphi}$ to be considered. This makes sure that the safety value function does not become overly conservative due to outlier disturbances. $\mu(\cdot)$ is

the sample mean, and $\sigma(\cdot)$ is the sample standard deviation. $b \geq 0$ is the spread of disturbances, in standard deviation units, that we compute the disturbance bound for, modulating the degree of robustness of the disturbance bound. A larger choice of b results in a more robust disturbance bound. We set $c = 0.8$, $b = 2$, and $\eta \cdot \varphi = 2$ s in practice. We compute two separate disturbance bounds: \bar{d}_{p_x, p_y}^e bounding the uncertainty in the position dynamics and $\bar{d}_{p_\theta}^e$ bounding the uncertainty in the yaw dynamics. For computing a bound on the norm of d_{p_x, p_y}^e , we use (15) on a window of disturbance norms instead of raw values.

Remark: Although we adopt a simple sliding window-based estimator for its simplicity, rapid adaptability, and computational efficiency, the OCR framework is agnostic to this choice. More advanced estimation approaches, such as Kalman filters or Gaussian Process regression, could be used depending on the desired balance between accuracy, efficiency, and generalization capabilities [24], [86]. See the appendix of the extended version of this article for a comparison with an extended Kalman filter-based estimator.¹

2) *Adaptive safety filter*: We use the OCR-VN to construct a smooth least-restrictive safety filter (see Section IV-B) safeguarding a potentially unsafe nominal controller that is given to us. Specifically, we compute a filtered high-level policy $\pi_{\text{safe}}^{\text{high}}$ which, when combined with the underlying low-level policy $\pi_{\text{nom}}^{\text{low}}$, maintains the overall system safety.

We compute the hierarchical filtered policy $\pi_{\text{safe}}(x, \bar{d}_r^e, o^e)$ as follows:

$$\begin{aligned} \pi_{\text{safe}}(x, \bar{d}_r^e, o^e) &= \begin{cases} \pi_{\text{nom}}(x), & V_\psi(x_r, \bar{d}_r^e, o^e) > \delta \\ \pi_{\text{nom}}^{\text{low}} \circ \pi_{\text{QP}}^{\text{high}}(x_r, \bar{d}_r^e, o^e), & V_\psi(x_r, \bar{d}_r^e, o^e) \leq \delta \end{cases} \quad (16) \end{aligned}$$

where $\pi_{\text{QP}}^{\text{high}}(x_r, \bar{d}_r^e, o^e)$ is obtained by solving

$$\begin{aligned} \min_{\omega \in \Omega, s \geq 0} & \left(\|\omega - \pi_{\text{nom}}^{\text{high}}(x_r)\|_2^2 + \lambda s^2 \right) \\ \text{s.t.} \min_{d_r \in \mathcal{D}_r^e} & \nabla V_\psi(x_r, \bar{d}_r^e, o^e) \cdot f_r(x_r, \omega, d_r) \geq -s \quad (17) \end{aligned}$$

where the slack variable s ensures that there always exists a feasible solution, and λ (set to 10^3) is the relative weight of s in the objective. \mathcal{D}_r^e is the set of disturbances respecting the disturbance bound \bar{d}_r^e . Intuitively, (17) minimally adjusts the nominal policy so that the resulting twist commands ensure system safety in the deployment environment. Note that the OCR-VN is differentiable everywhere, and its gradients can be computed efficiently via automatic differentiation. Furthermore, the introduction of the slack variable ensures that the QP problem in (17) is always feasible, especially when the disturbance bounds are overly conservative. The disturbance optimization in the constraint in (17) is independent of the optimization of ω , since the disturbance enters additively into f_r in (10). Since f_r is affine in control and disturbance, (17) is a quadratic program (QP) that we can solve efficiently online, making the proposed framework amenable for real-world robotic systems.

We remark that the design of the filter in (16) is grounded in reachability theory. If V_ψ perfectly models the underlying value

function and the system starts outside of the BRT, then (16) is guaranteed to maintain safety [71].

Remark: When the quadruped is controlled by an *end-to-end* nominal controller, we can use the OCR-VN framework to maintain safety with a minor adjustment. While $V_\psi > 0$, we permit the nominal controller to execute unhindered. Otherwise, we must provide our own high-level planner $\pi_{\text{nom}}^{\text{high}}(x)$ for solving (17) and a low-level policy $\pi_{\text{nom}}^{\text{low}}$ for execution.

VI. SIMULATION EXPERIMENTS

We test the proposed framework on a Unitree Go1 quadruped in Isaac Sim (this section) and on a hardware testbed (Section VII) in a variety of unseen obstacle settings and terrains, as well as on different nominal controllers. The goal of our simulation studies and experiments is to answer the following questions about the proposed framework:

- 1) Can the safety filter dynamically adapt to unknown obstacles to maintain safety?
- 2) Can the safety filter adapt to system and environment uncertainty (e.g., uneven terrains, slippery surfaces, etc.)?
- 3) Can the safety filter maintain safety under different nominal policies?

A. Nominal Controllers

We test our framework on several different hierarchical nominal policies. A nominal policy is generated by selecting one of the three high-level planners and one of the three low-level locomotion policies described below. These policies are selected to contain both RL-based policies as well as MPC-based policies. We test all different combinations of high-level and low-level policies, but present the results for only a few of them here for brevity purposes.

In addition, we include an end-to-end nominal policy from ABS [29] (called ABS-Agile here on), as an illustration of how our method can generalize to end-to-end nominal policies. To implement our method on top of the end-to-end ABS-Agile policy, we compute the safety twist commands using the proposed filter framework in (16) with the predictive sampling-based (PS) high-level planner (see Section VI-A1a) and use the ABS-recovery policy (see Section VI-A2c) to track the filtered twist commands. Whenever the system safety is not at risk, we use the nominal ABS-Agile policy.

1) High-Level Planners:

a) *Predictive sampling-based planner (PS):* PS is a sampling-based high-level MPC planner that optimizes the robot trajectory to reach the goal while avoiding obstacles. The cost function for a discrete control sequence $\omega = \{\omega^j\}_{j=1}^{\lfloor T/\eta \rfloor}$ evaluated at the reduced state x_r is given by

$$J_\omega(x_r) = \sum_{j=1}^{\lfloor T/\eta \rfloor} \left(\|\hat{x}_r^j - g\| + c \cdot \mathbb{1}\{\hat{x}_r^j \in \hat{\mathcal{L}}^e\} \right)$$

where T is the prediction horizon and η is the discrete time step. \hat{x}_r^j is the future state at time $j \cdot \eta$ predicted by the reduced-order dynamics f_r in (10) after applying ω to the system. g is the goal state, c is a collision penalty, and $\hat{\mathcal{L}}^e$ is the estimated obstacle map computed online using obtained LiDAR scans

via *tinySLAM* [87]. Intuitively, the cost function is minimized by control sequences that bring the system to the goal fastest without collisions. To compute the optimal high-level commands, PS samples N control sequences $\omega_i, \forall i = 1, \dots, N$ from a Gaussian centered at ω_{seed} with standard deviation σ . The cost $J_{\omega_i}(x_r)$ for each sequence is computed and finally the best control sequence ω_{best} is selected which minimizes the cost. PS executes the first twist command in ω_{best} and sets $\omega_{\text{seed}} = \omega_{\text{best}}$ for the next control iteration. We set $N = 1,000, \sigma = 0.5, T = 4$ s, $\eta = 0.2$ s, $c = 10^9$, and the initial ω_{seed} to the control range center.

b) *Naive planner (NVE):* NVE performs basic goal-seeking without obstacle avoidance. Let $p_{\theta_{\text{goal}}}$ be the angle to the goal relative to the robot's heading. NVE computes

$$v = \begin{cases} v_{\text{max}}, & |p_{\theta_{\text{goal}}}| \leq \frac{\pi}{2} \\ v_{\text{min}}, & |p_{\theta_{\text{goal}}}| > \frac{\pi}{2} \end{cases}$$

$$w = \begin{cases} w_{\text{min}} \cdot \min \left\{ \left| \frac{p_{\theta_{\text{goal}}}}{p_{\theta_{\text{max}}}} \right|, 1 \right\}, & p_{\theta_{\text{goal}}} \leq 0 \\ w_{\text{max}} \cdot \min \left\{ \left| \frac{p_{\theta_{\text{goal}}}}{p_{\theta_{\text{max}}}} \right|, 1 \right\}, & p_{\theta_{\text{goal}}} > 0 \end{cases}$$

where $p_{\theta_{\text{max}}}$ is the heading difference beyond which yaw rate is maximized. We set $p_{\theta_{\text{max}}} = \frac{\pi}{4}$ rad. Intuitively, NVE turns greedily toward the goal and commands maximum/minimum forward velocity when the goal is ahead/behind.

c) *Human teleoperation (HMN):* a human teleoperator provides the high-level twist commands.

2) Low-Level Policies:

a) *Walk-these-ways (WTW):* An RL-based policy that encodes a structured family of locomotion strategies to enable diverse task generalization such as crouching, hopping, running, and others [21]. During training, the policy is rewarded for accurate twist tracking and gait following. During deployment, different locomotion strategies can be used by changing gait parameters. We use the gait parameters corresponding to a 3 Hz trot for its stability and agility.

b) *Legged MPC by Unitree:* A model-based policy provided by Unitree Robotics as a ROS package for the Go1 quadruped robot [79]. The policy tracks twist commands using an internal MPC-based algorithm that computes motor torque controls. It is widely used by consumers of the Go1 quadruped hardware.

c) *ABS recovery policy:* An RL-based policy used in the ABS framework to track twist commands as fast as possible in order to serve as a backup shielding policy [29]. During training, the policy is rewarded for accurate twist tracking and maintaining a stable posture. Similar to its role in the ABS framework, we will use the ABS-Recovery policy as the low-level backup shielding policy for the ABS-Agile policy.

B. Baseline and Ablations

We evaluate our framework against a range of nominal controllers and the end-to-end safety framework, ABS, proposed in [29]. ABS integrates an RL-based agile policy with a safety-oriented recovery policy. A value function predictor determines when to switch to the recovery policy, which then generates




Fig. 4. OCR framework (green/red: nominal/filtered) safeguards different nominal controllers navigating to a goal (cyan) in an environment with a payload of -0.6 kg and a friction of 0.7 . By themselves (white), the (a) PS + WTW, (b) NVE + WTW, and (c) ABS-Agile controllers fail to maintain safety due to dynamical uncertainty caused by low payload and friction.

safe twist commands to be executed by the recovery policy. We refer the interested readers to [29] for more details. Compared to ABS, the key innovations of this work include the consideration of worst-case dynamical uncertainties in the online setting for robust safety, as well as the computation of a policy-agnostic optimal value function for automatic safety filtering across different nominal controllers.

We also evaluate the OCR framework without the disturbance estimation module (OCR \ DE) and calibration step (OCR \ C), as well as if the safety filter is not used at all (no filter) to understand the importance of each of these modules.

Finally, we include comparisons with a robust CBF-based baseline, a high-level MPC-based baseline, and a low-level MPC-based baseline in the appendix of the extended version of this article to illustrate the difficulty of the safety problem for existing traditional methods.¹

C. Metrics

Each experiment trial ends in either a success, a collision, or a timeout after 60 s. Across successful runs, we report the average velocity \bar{v} of the quadruped along the trajectory, the average rate of safety filter activation \bar{r} , and the average minimum distance to the obstacle set \bar{q} .

D. Simulation Setups

We test our framework as well as baselines on randomly generated environments. Each environment contains four circular obstacles, where each obstacle is spawned at some location drawn uniformly randomly from $[-2, 2] \text{ m} \times [-2, 2] \text{ m}$. Each obstacle has a radius drawn uniformly randomly from $[0.1, 1] \text{ m}$. Each environment also draws an additional payload uniformly randomly from $[-1, -0.5] \cup [0.5, 1] \text{ kg}$ and a ground friction coefficient $[0.5, 0.75] \cup [1.25, 1.5]$ for dynamical variation. We remark that these payloads and frictions are contained within the range used to train the nominal policies and baselines so that the comparison is fair, but they represent the outer limits to effectively test robustness. The quadruped must navigate from $(-5, 0) \text{ m}$ to $(5, 0) \text{ m}$ without colliding (see Fig. 4).

In addition, we present hand-designed environments that stress-test and highlight the abilities of our framework.

E. Simulation Results

Recall that the goal of our simulation studies and experiments is to answer the following questions: can the safety filter adapt to unknown obstacles in the environment, uncertainty in the system and environment, and different nominal policies?




Fig. 5. (In color) Safety rates across settings (1000 trials).

1) *Safeguarding different nominal controllers*: We first evaluate the efficacy of the OCR framework and its ablations for safeguarding various nominal controllers in different obstacle settings. Numerical results are listed in Table III. The OCR framework achieves high success rates, low collision rates, and comfortable distances to the obstacle set across all nominal controllers, highlighting its ability to automatically safeguard different nominal controllers in different environments without *a priori* knowledge. Its success rates are consistently high regardless of whether the underlying nominal control has built-in obstacle avoidance (ABS-Agile and PS + WTW) or not (NVE + WTW). Fig. 4 illustrates the robot trajectories under the OCR framework and the framework's ability to ensure safety when the nominal controller would cause collision.

The results in Table III also clearly show the importance of the disturbance estimation module and the calibration step for enabling the OCR framework to be robust to modeling and learning errors, respectively. Interestingly, even though the ABS framework is particularly designed to safeguard the ABS-Agile policy, our method also outperforms the ABS framework on ABS-Agile. A key reason for this is the robustness of the proposed framework not only to environment uncertainty (i.e., unknown obstacles) but also to dynamics uncertainty, as we discuss next.

2) *Robustness to dynamical uncertainty*: Results in Table III demonstrate the ability of the OCR framework to handle variation in the system dynamics. Recall that each environment has its own payload and friction parameters, which significantly influence the dynamics of the system. In addition, different low-level policies have different abilities to track velocity and yaw rate commands. Despite these challenges, the OCR framework consistently achieves a high success rate without *a priori* access to the underlying nominal policy or simulation environments.

To study robustness more thoroughly, we plot safety rates across environment settings for different frameworks using the ABS-Agile policy as the nominal policy in Fig. 5. The OCR

TABLE III
SIMULATION RESULTS FOR OCR, ABLATIONS, AND ABS ACROSS NOMINAL CONTROLLERS (100 TRIALS)

Controller	Filter	Success Rate \uparrow	Collision Rate \downarrow	Timeout Rate \downarrow	\bar{v} (m/s) \uparrow	\bar{r} \downarrow	\bar{q} (m) \uparrow
ABS-Agile	No Filter	0.75	0.25	0.00	2.10	0.00	0.41
	ABS	0.80	0.20	0.00	2.03	0.03	0.41
	OCR \ DE	0.40	0.35	0.25	0.98	0.62	0.41
	OCR \ C	0.81	0.19	0.00	1.70	0.28	0.41
	OCR (ours)	0.91	0.09	0.00	1.22	0.59	0.58
PS + WTW	No Filter	0.72	0.28	0.00	1.36	0.00	0.43
	OCR \ DE	0.78	0.22	0.00	0.84	0.61	0.42
	OCR \ C	0.90	0.09	0.01	1.21	0.18	0.44
	OCR (ours)	1.00	0.00	0.00	0.97	0.42	0.66
	NVE + WTW	No Filter	0.20	0.80	0.00	1.70	0.00
OCR \ DE		0.21	0.79	0.00	0.93	0.61	0.37
OCR \ C		0.45	0.55	0.00	1.32	0.28	0.39
OCR (ours)		0.91	0.08	0.01	1.04	0.47	0.62

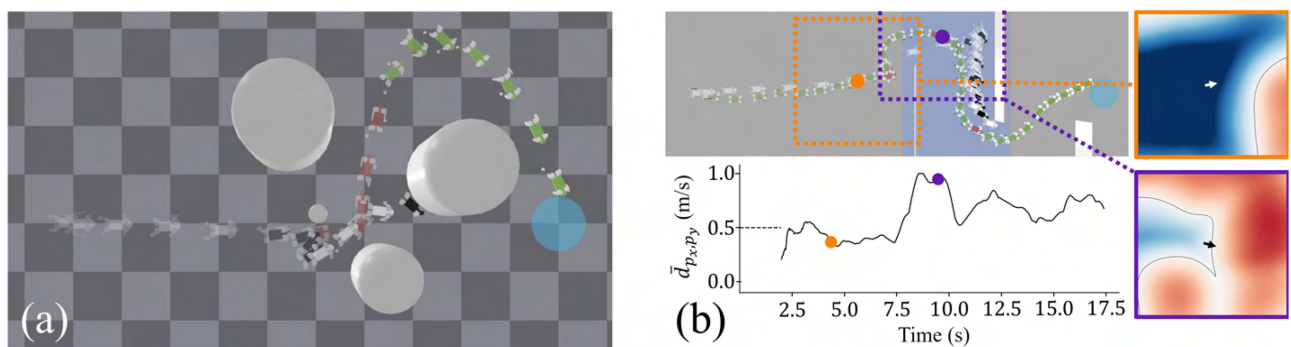


Fig. 6. OCR (green/red: nominal/filtered) and ABS (white/black: nominal/filtered) frameworks in (a) a validation environment with a payload of -0.9 kg and a friction of 0.5 and (b) a hand-designed obstacle configuration with a region of low friction (blue). In (b), we plot the evolution of the estimated disturbance bound in position \bar{d}_{p_x, p_y}^e , as well as the OCR-VN predictions at two different states (orange, purple).

framework achieves high safety rates across all settings, surpassing the ABS baseline particularly in low-friction environments, but only when using disturbance estimation. This leads us to conclude that disturbance estimation plays a key role in the robustness of the proposed framework. For all numerical results, see the appendix of the extended version of this article.¹

Fig. 6(a) visualizes the behavior of the OCR and ABS frameworks in an environment with high perturbations in dynamics (a payload of -0.9 kg and a friction of 0.5). The OCR framework estimates large disturbances online and accordingly intervenes to steer the system away from collision. The adaptation facilitated by the disturbance estimation module is visualized in Fig. 6(b) in a hand-designed environment with a sudden friction change in the blue region. Correspondingly, the proposed framework automatically estimates larger disturbance bounds in this region of low friction (the purple point), which subsequently, leads to a more aggressive safety filter to ensure system safety. The ABS framework, in contrast, intervenes later and allows the quadruped to approach dangerously close to the wall obstacle, eventually leading to collision. For videos illustrating the framework in action, including the evolution of estimated disturbance bounds, the value function predictions, and the safety filter, please see the project website.²

²https://sia-lab-git.github.io/One_Filter_to_Deploy_Them_All

3) *Safety-Performance tradeoff*: Although the results in Table III indicate that the OCR framework achieves superior success rates and collision rates, they also suggest that this comes at a tradeoff with performance, as shown by the decrease in average forward velocity \bar{v} after OCR-VN calibration. To investigate the quality of this safety-performance tradeoff, we collect results across various safety margins for both the OCR framework and the ABS baseline. The results are listed in Table IV.

Without calibration, the OCR framework (OCR \ C) achieves roughly the same agility and success rate as the ABS baseline (ABS w/ Threshold -0.05). As the safety margin increases (by decreasing the threshold at which ABS triggers and decreasing the calibration level of the OCR-VN), both frameworks experience a comparable decrease in the agility of the system. However, in contrast to the ABS framework, the OCR framework provides substantially better success rates and collision rates as the conservatism increases. This is due to the difference in *optimality* of the learned value functions between the two frameworks, as discussed further in the next section. In addition, the speeds achieved remain in the agile 1 to 2 m/s range across safety margins. Based on these results, we conclude that the OCR framework achieves a reasonable safety-performance tradeoff.

4) *Further comparisons with ABS*: We now further compare the proposed OCR framework with ABS, as it leverages a similar safety filtering framework. One of the key advantages of the OCR framework is that it grants the user the freedom to

TABLE IV
SAFETY-PERFORMANCE TRADEOFF FOR OCR AND ABS ACROSS DIFFERENT SAFETY MARGINS (100 TRIALS)

Filter	Success Rate	Collision Rate	Timeout Rate	\bar{v} (m/s)	\bar{r}	\bar{q} (m)
ABS w/ Threshold -0.05	0.80	0.20	0.00	2.03	0.03	0.41
ABS w/ Threshold -0.10	0.78	0.22	0.00	2.00	0.03	0.41
ABS w/ Threshold -0.20	0.83	0.17	0.00	1.94	0.05	0.41
ABS w/ Threshold -0.40	0.77	0.17	0.06	1.85	0.08	0.42
ABS w/ Threshold -0.60	0.67	0.25	0.08	1.47	0.19	0.44
OCR \ C	0.81	0.19	0.00	1.70	0.28	0.41
OCR w/ C -0.1	0.79	0.21	0.00	1.72	0.24	0.41
OCR w/ C -0.2	0.78	0.22	0.00	1.42	0.32	0.42
OCR w/ C -0.3	0.86	0.14	0.00	1.30	0.40	0.43
OCR w/ C -0.4	0.90	0.10	0.00	1.14	0.47	0.46
OCR w/ C -0.5 (ours)	0.91	0.09	0.00	1.22	0.59	0.58




Fig. 7. OCR (green/red : nominal/filtered) and ABS (white/black : nominal/filtered) frameworks in environments with hand-designed obstacle configurations. In (b), the OCR framework is used without calibration due to the highly cluttered setting.

choose the most appropriate nominal controller for the task at hand without needing to recompute safety assurances, whereas the ABS framework is specific to a particular nominal policy. We demonstrate this utility in Fig. 7(a), where the robot must navigate two sharp turns. The OCR framework successfully completes the task using the PS + WTW nominal controller. On the other, the ABS baseline gets stuck, because the learning-based ABS-Agile policy is not suited for the required navigation style.

The policy-independence of the OCR framework is a consequence of the *optimality* of the underlying value function synthesized during the training phase described in Section V-A. We demonstrate the advantages of using the optimal value function in Fig. 7(b) in a hand-designed “dead-end” environment. Since the setting is highly cluttered, we use OCR \ C to reduce conservatism. We remark that this remains a fair evaluation with respect to the ABS baseline because OCR \ C retains a comparable safety rate to ABS in Table III.

Using a control search guided by a policy-conditioned value function, the ABS baseline suffers from unnecessary interventions and suboptimality related to the quality of the ABS-Agile policy. On the other hand, OCR framework reasons about the optimal safety behavior (regardless of the nominal policy), reducing unnecessary interventions and permitting task progress.

VII. HARDWARE EXPERIMENTS

We deploy the OCR framework with nominal controllers from Section VI-A on a Unitree Go1 quadruped robot equipped with an onboard Slamtec RPLiDAR A2 sensor. The quadruped maintains an estimate of its state in the global frame via *tinyslam* [87], a LiDAR-based simultaneous localization and mapping

(SLAM) algorithm, which is subsequently used for disturbance bound estimation.

A. Hardware Setups

We first quantitatively and qualitatively compare the proposed framework and all baselines in two different experiment settings. In the first setting, the robot needs to navigate through an obstacle maze consisting of walls and circular obstacles. The second setting is the same as the first, except that a rectangular region directly preceding the first obstacle has a very low friction due to an oil-soaked tarp, to test the dynamic safety adaptivity of different methods. The setups are shown in Fig. 8. In our experiments, we also include results on a few additional nominal policies to stress-test our system.

In addition, we qualitatively test our framework in a diverse set of real-world experiments, including cluttered environments, rough terrains, external disturbances, and adversarial human teleoperation, with videos on the project website.²

B. Hardware Results

The hardware results listed in Table V reaffirm the findings in simulation. Namely, the OCR framework is effective at safeguarding a diverse set of controllers. The framework also displays a significant robustness to changes in the system dynamics, as evidenced by the high success rates even in the slippery condition, where the ABS baseline fails. Fig. 8 illustrates the difference between the OCR and ABS frameworks in the slippery condition. Whereas the ABS baseline intervenes too late to maintain safety, the OCR framework slows down the robot and turns in time to avoid collision.




Fig. 8. Hardware experiments with a slippery region outlined in white. (Left) The ABS baseline collides due to drifting caused by the slippery floor. (Right) The OCR framework (green/red : nominal/filtered) stops and turns in time to prevent a collision, demonstrating superior robustness to changes in the system dynamics.

TABLE V
HARDWARE RESULTS FOR OCR AND ABS ACROSS NOMINAL CONTROLLERS (10 TRIALS)

Controller	Filter	Normal Condition			Slippery Condition		
		# Successes \uparrow	# Collisions \downarrow	# Timeouts \downarrow	# Successes \uparrow	# Collisions \downarrow	# Timeouts \downarrow
ABS-Agile	No Filter	9	1	0	1	9	0
	ABS	9	1	0	2	8	0
	OCR (ours)	10	0	0	8	1	1
PS + WTW	No Filter	5	5	0	1	9	0
	OCR (ours)	9	1	0	8	2	0
NVE + WTW	No Filter	0	10	0	0	10	0
	OCR (ours)	9	0	1	8	2	0
PS + MPC	No Filter	5	5	0	0	10	0
	OCR (ours)	10	0	0	9	1	0
NVE + MPC	No Filter	0	10	0	0	10	0
	OCR (ours)	9	0	1	7	3	0

We further demonstrate the capabilities of the framework in a variety of real-world scenarios. Videos of these demonstrations can be found on the project website.² Many of these scenarios are illustrated in Fig. 1.

1) *Safeguarding different nominal controllers*: Fig. 1 illustrates the ability of the OCR framework to automatically safeguard a variety of high-level planners, including (a) learning-based, (c, f, k) model-based, (b, d, g, h, i, j) human teleoperated, and (e) obstacle-agnostic planners, on top of different low-level locomotion policies, including (a, f, i, j, k) learning-based and (b, c, d, e, g, h) model-based policies.

2) *External disturbances*: Subplot (f) in Fig. 1 demonstrates the framework’s ability to preserve safety even under forceful disturbances. After receiving a kick while on a slippery floor, the framework guides the robot to turn sharply from the obstacle and move away to maintain safety.

3) *Adversarial human teleoperation*: In subplot (h) of Fig. 1, an adversarial human teleoperator attempts to collide the robot into a pillar multiple times. The filter activates only when necessary and causes the robot to veer to avoid collision while respecting the commanded input as much as possible.

4) *Dynamic obstacles*: The framework’s ability to react to dynamic obstacles in real time is illustrated in subplots (e, k) in Fig. 1. The robot slows down from a velocity of roughly 2 m/s

to provide enough turn radius to avoid the box placed suddenly by a human in front of the robot in subplot (e).

5) *Cluttered indoor environments*: The robot navigates narrow corridors in subplots (a, b, c) in Fig. 1. Due to the conservatism of the OCR framework, a few adjustments are needed for good performance in highly cluttered settings. First, we discard LiDAR readings outside of a front-facing cone spanning $\pi/2$ rad. We believe this is necessary because of a shift in distribution of LiDAR readings from what is seen during training, which has sparse obstacles and no walls. Second, we use the uncalibrated output of the OCR-VN. We theorize that the conservatism of the OCR framework is an inherent result of using a worst-case analysis for disturbance, which makes the safety problem especially difficult in crowded settings. We defer addressing these limitations to future work.

6) *Rough Terrain*: The robot avoids collisions on rough outdoor terrains in subplots (d, i, j) of Fig. 1, further highlighting the ability of the proposed framework to adapt to dynamical uncertainties.

VIII. LIMITATIONS

Although our proposed framework achieves high safety rates in both simulation and hardware results in Tables III and V, the

framework still exhibits a nonzero collision rate. We attribute this to several failure modes of the framework that we identify and discuss next.

First, during deployment, the system may experience disturbances that exceed the estimated disturbance bounds which the framework is designed to safeguard against. Furthermore, there is an unavoidable delay before environment changes will reflect in the estimated disturbance bounds. During this delay period, the system can encounter failure before the framework has a chance to adapt.

Second, the OCR-VN can contain learning errors that are critical to safety. While the calibration scheme presented in Section V-A4 can help us better understand the degree of these errors, we cannot directly extrapolate the theoretical guarantees from a validation dataset to the real world, due to potential distribution shifts. In particular, the calibration may not be valid in settings different than the distribution of LiDAR scans seen during training, e.g., when the robot is fully surrounded by obstacles. In addition, the use of only circular obstacles during training can inhibit the generalization abilities of the OCR-VN to settings with irregular geometries, especially if the observation space is higher dimensional, e.g., occupancy grids or camera images. It would be interesting to explore more expressive scenario generation methods, online calibration methods and out-of-distribution detection methods to overcome these challenges.

In addition to these failure modes which affect safety, there are several limitations of the framework that affect its performance. For some low-level locomotion policies and environment settings, the error in the reduced-order dynamics model can be very large, leading to high dynamical uncertainty. Large estimated disturbances can also appear as an artifact of latency and state estimation issues with real-world hardware. The resulting disturbance bounds can produce overly conservative BRTs, which severely inhibit progress by the system. In highly cluttered environments, this can manifest as stalling behavior.

The conservatism of the framework is also partly a consequence of our choice to model the dynamics uncertainty as *adversarial* in nature. Future works can explore approaches to overcome these issues by using proactive methods of modeling and estimating dynamics and disturbances, for example, by using observations of the environment or learning controller-specific dynamics to anticipate future system behavior and reduce uncertainty, as well as considering different characterizations of disturbances.

Finally, while our proposed framework can, in principle, accommodate a wide range of reduced-order models and uncertainty characterizations, the specific implementation and experiments in this work are based on the use of a simplified 3-D Dubins car model. This choice is motivated by the particular application scenario considered. As such, the presented results do not capture the full richness of quadrupedal locomotion dynamics, including contact transitions, body compliance, or high-dimensional control, which may be required for more challenging tasks. A key challenge in extending the framework to more complex models lies in the computational burden of HJ reachability analysis, which does not scale well with the

dimensionality of the state space. However, recent advances, such as decomposition methods [88] and learning-based approximations [83], [89] offer promising directions for integrating more expressive models in the future.

IX. CONCLUSION

We propose the OCR framework, which uses an adaptive safety filter to robustly ensure the safety of a quadruped robot running an a priori unknown locomotion policy in a priori unknown environments using only LiDAR observations. The offline training of the OCR-VN is done *without a priori access to the controllers or a simulator*—hence, “One Filter to Deploy Them All.” In simulation and hardware experiments on a Unitree Go1 quadruped, we demonstrate the superior efficacy of the proposed approach to ensure, in zero-shot fashion, the safety of numerous controllers across diverse environment configurations and perturbed dynamics. Videos demonstrating the efficacy of the proposed approach across various settings can be found on the project website.²

REFERENCES

- [1] S. Halder, K. Afsari, E. Chiou, R. Patrick, and K. A. Hamed, “Construction inspection and monitoring with quadruped robots in future human-robot teaming: A preliminary study,” *J. Building Eng.*, vol. 65, 2023, Art. no. 105814. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2352710222018204>
- [2] C. Gehring et al., “Anymal in the field: Solving industrial inspection of an offshore HVDC platform with a quadrupedal robot,” in *Field and Service Robotics*, G. Ishigami and K. Yoshida, Eds. Singapore: Springer, 2021, pp. 247–260.
- [3] N. Li, J. Cao, and Y. Huang, “Fabrication and testing of the rescue quadruped robot for post-disaster search and rescue operations,” in *Proc. IEEE 3rd Int. Conf. Electron. Technol., Commun. Inf.*, 2023, pp. 723–729.
- [4] C. Cruz Ulloa, J. del Cerro, and A. Barrientos, “Mixed-reality for quadruped-robotic guidance in SAR tasks,” *J. Comput. Des. Eng.*, vol. 10, no. 4, pp. 1479–1489, 2023.
- [5] F. Gao, C. Lei, X. Long, J. Wang, and P. Song, “Design and development of an intelligent pet-type quadruped robot,” in *Proc. IEEE 4th Int. Conf. Multimedia Inf. Process. Retrieval.*, 2021, pp. 366–371.
- [6] H. Lee, “A human–robot interaction entertainment pet robot,” *J. Korean Inst. Intell. Syst.*, vol. 24, no. 2, pp. 179–185, 2014.
- [7] Z. Chen et al., “Autonomous social distancing in urban environments using a quadruped robot,” *IEEE Access*, vol. 9, pp. 8392–8403, 2021.
- [8] A. D. Ames et al., “First steps toward formal controller synthesis for bipedal robots with experimental implementation,” *Nonlinear Anal., Hybrid Syst.*, vol. 25, pp. 155–173, 2017. [Online]. Available: <http://ames.caltech.edu/ames2017first.pdf>
- [9] R. Grandia, A. J. Taylor, A. D. Ames, and M. Hutter, “Multi-layered safety for legged robots via control barrier functions and model predictive control,” in *Proc IEEE Int. Conf. Robot. Automat.*, 2021, pp. 8352–8358.
- [10] H. U. Unlu, V. M. Gonçalves, D. Chaikalis, A. Tzes, and F. Khorrami, “A control barrier function-based motion planning scheme for a quadruped robot,” in *Proc. IEEE Int. Conf. Robot. Automat.*, 2024, pp. 12172–12178.
- [11] M. Tayal and S. Kolathaya, “Safe legged locomotion using collision cone control barrier functions (C3BFS),” 2023, *arXiv:2309.01898*.
- [12] J. Kim, J. Lee, and A. D. Ames, “Safety-critical coordination for cooperative legged locomotion via control barrier functions,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2023, pp. 2368–2375.
- [13] J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, “Learning quadrupedal locomotion over challenging terrain,” *Sci. Robot.*, vol. 5, no. 47, 2020, Art. no. eabc5986. [Online]. Available: <https://www.science.org/doi/abs/10.1126/scirobotics.abc5986>
- [14] D. Hoeller, N. Rudin, D. Sako, and M. Hutter, “Anymal Parkour: Learning agile navigation for quadrupedal robots,” *Sci. Robot.*, vol. 9, no. 88, 2024, Art. no. eadi7566. [Online]. Available: <https://www.science.org/doi/abs/10.1126/scirobotics.adi7566>

- [15] G. Bellegarda and A. Ijspeert, "Visual CPG-RL: Learning central pattern generators for visually-guided quadruped navigation," in *2024 IEEE Int. Conf. Robot. Automat.*, Yokohama, Japan, 2024, pp. 1420–1427, doi: [0.1109/ICRA57147.2024.10611128](https://doi.org/10.1109/ICRA57147.2024.10611128).
- [16] J. Tan et al., "Sim-to-real: Learning agile locomotion for quadruped robots," in *Proc. Robot.: Sci. Syst.*, 2018, doi: [10.15607/RSS.2018.XIV.010](https://doi.org/10.15607/RSS.2018.XIV.010).
- [17] Y.-S. Luo, J. H. Soeseno, T. P.-C. Chen, and W.-C. Chen, "CARL: Controllable agent with reinforcement learning for quadruped locomotion," *ACM Trans. Graph.*, vol. 39, no. 4, pp. 38:1–38:10, Aug. 2020, doi: [10.1145/3386569.3392433](https://doi.org/10.1145/3386569.3392433).
- [18] H. Shi et al., "Terrain-aware quadrupedal locomotion via reinforcement learning," 2023, *arXiv:2310.04675*.
- [19] L. Schneider, J. Frey, T. Miki, and M. Hutter, "Learning risk-aware quadrupedal locomotion using distributional reinforcement learning," in *Proc. 2024 IEEE Int. Conf. Robot. Automat.*, 2024, pp. 11451–11458.
- [20] A. Kumar, Z. Fu, D. Pathak, and J. Malik, "RMA: Rapid motor adaptation for legged robots," in *Proc. Robot.: Sci. Syst.*, 2021, doi: [10.15607/RSS.2021.XVII.011](https://doi.org/10.15607/RSS.2021.XVII.011).
- [21] G. B. Margolis and P. Agrawal, "Walk these ways: Tuning robot control for generalization with multiplicity of behavior," *Proc. Conf. Robot Learn.*, 2022, pp. 22–31, vol. 205.
- [22] S. Gangapurwala, M. Geisert, R. Orsolino, M. Fallon, and I. Havoutis, "Real-time trajectory adaptation for quadrupedal locomotion using deep reinforcement learning," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2021, pp. 5973–5979.
- [23] W. Xiao, T. He, J. Dolan, and G. Shi, "Safe deep policy adaptation," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2024, pp. 17286–17292.
- [24] 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 *Proc. 33rd AAAI Conf. Artif. Intell./31st Innov. Appl. Artif. Intell. Conf. 9th AAAI Symp. Educ. Adv. Artif. Intell.*, 2019, pp. 3387–3395, doi: [10.1609/aaai.v33i01.33013387](https://doi.org/10.1609/aaai.v33i01.33013387).
- [25] W. Zhao, T. He, and C. Liu, "Model-free safe control for zero-violation reinforcement learning," in *Proc. 5th Conf. Robot Learn.*, ser. Proceedings of Machine Learning Research, vol. 164. PMLR, 08–11 Nov. 2022, pp. 784–793. [Online]. Available: <https://proceedings.mlr.press/v164/zhao22a.html>
- [26] Y. K. Nakka, A. Liu, G. Shi, A. Anandkumar, Y. Yue, and S.-J. Chung, "Chance-constrained trajectory optimization for safe exploration and learning of nonlinear systems," *IEEE Robot. Automat. Lett.*, vol. 6, no. 2, pp. 389–396, Apr. 2021.
- [27] B. Thananjeyan et al., "Recovery RL: Safe reinforcement learning with learned recovery zones," *IEEE Robot. Automat. Lett.*, vol. 6, no. 3, pp. 4915–4922, Jul. 2021.
- [28] K.-C. Hsu, A. Z. Ren, D. P. Nguyen, A. Majumdar, and J. F. Fisac, "Sim-to-lab-to-real: Safe reinforcement learning with shielding and generalization guarantees," *Artif. Intell.*, vol. 314, 2023, Art. no. 103811. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0004370222001515>
- [29] T. He, C. Zhang, W. Xiao, G. He, C. Liu, and G. Shi, "Agile but safe: Learning collision-free high-speed legged locomotion," in *Proc. Robot.: Sci. Syst.*, 2024. [Online]. Available: <https://www.roboticsproceedings.org/rss20/p059.html>
- [30] H. Bharadhwaj, A. Kumar, N. Rhinehart, S. Levine, F. Shkurti, and A. Garg, "Conservative safety critics for exploration," in *Proc. Int. Conf. Learn. Representations*, 2021. [Online]. Available: <https://openreview.net/forum?id=iaO86DUuKi>
- [31] D. Kim et al., "Vision aided dynamic exploration of unstructured terrain with a small-scale quadruped robot," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2020, pp. 2464–2470.
- [32] T. Dudzik et al., "Robust autonomous navigation of a small-scale quadruped robot in real-world environments," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2020, pp. 3664–3671.
- [33] R. Buchanan, L. Wellhausen, M. Bjelonic, T. Bandyopadhyay, N. Kottege, and M. Hutter, "Perceptive whole-body planning for multilegged robots in confined spaces," *J. Field Robot.*, vol. 38, no. 1, pp. 68–84, 2021. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/rob.21974>
- [34] M. Gaertner, M. Bjelonic, F. Farshidian, and M. Hutter, "Collision-free mpc for legged robots in static and dynamic scenes," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2021, pp. 8266–8272.
- [35] J.-R. Chiu, J.-P. Sleiman, M. Mittal, F. Farshidian, and M. Hutter, "A collision-free mpc for whole-body dynamic locomotion and manipulation," in *Proc. Int. Conf. Robot. Automat.*, 2022, pp. 4686–4693.
- [36] M. Mattamala, N. Chebrolu, and M. Fallon, "An efficient locally reactive controller for safe navigation in visual teach and repeat missions," *IEEE Robot. Autom. Lett.*, vol. 7, no. 2, pp. 2353–2360, Apr. 2022.
- [37] Q. Liao, Z. Li, A. Thirugnanam, J. Zeng, and K. Sreenath, "Walking in narrow spaces: Safety-critical locomotion control for quadrupedal robots with duality-based optimization," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2023, pp. 2723–2730.
- [38] L. Wellhausen and M. Hutter, "Artplanner: Robust legged robot navigation in the field," *Field Robot.*, vol. 3, pp. 413–434, 2023.
- [39] S. Teng, Y. Gong, J. W. Grizzle, and M. Ghaffari, "Toward safety-aware informative motion planning for legged robots," 2021, *arXiv:2103.14252*.
- [40] F. Jenelten, J. Hwangbo, F. Tresoldi, C. D. Bellicoso, and M. Hutter, "Dynamic locomotion on slippery ground," *IEEE Robot. Automat. Lett.*, vol. 4, no. 4, pp. 4170–4176, Oct. 2019.
- [41] V. Makoviychuk et al., "ISAAC gym: High performance GPU-based physics simulation for robot learning," in *Proc. Neural Inf. Process. Syst. Track Datasets Benchmarks*, 2021, vol. 1.
- [42] N. Rudin, D. Hoeller, P. Reist, and M. Hutter, "Learning to walk in minutes using massively parallel deep reinforcement learning," in *Proc. 5th Annu. Conf. Robot Learn.*, 2021, pp. 91–100, vol. 164. [Online]. Available: <https://openreview.net/forum?id=wK2fDDJ5VcF>
- [43] J. Kober, J. Bagnell, and J. Peters, "Reinforcement learning in robotics: A survey," *Int. J. Robot. Res.*, vol. 32, pp. 1238–1274, 2013.
- [44] G. B. Margolis, G. Yang, K. Paigwar, T. Chen, and P. Agrawal, "Rapid locomotion via reinforcement learning," *Int. J. Robot. Res.*, vol. 43, no. 4, pp. 572–587, 2024, doi: [10.1177/02783649231224053](https://doi.org/10.1177/02783649231224053).
- [45] J. Hwangbo et al., "Learning agile and dynamic motor skills for legged robots," *Sci. Robot.*, vol. 4, no. 26, 2019, Art. no. eaau5872. [Online]. Available: <https://www.science.org/doi/abs/10.1126/scirobotics.aau5872>
- [46] S. Gangapurwala, M. Geisert, R. Orsolino, M. Fallon, and I. Havoutis, "RLOC: Terrain-aware legged locomotion using reinforcement learning and optimal control," *IEEE Trans. Robot.*, vol. 38, no. 5, pp. 2908–2927, Oct. 2022.
- [47] C. Yang, K. Yuan, Q. Zhu, W. Yu, and Z. Li, "Multi-expert learning of adaptive legged locomotion," *Sci. Robot.*, vol. 5, no. 49, 2020, Art. no. eabb2174. [Online]. Available: <https://www.science.org/doi/abs/10.1126/scirobotics.abb2174>
- [48] T. Miki, J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, "Learning robust perceptive locomotion for quadrupedal robots in the wild," *Sci. Robot.*, vol. 7, no. 62, 2022, Art. no. eabk2822. [Online]. Available: <https://www.science.org/doi/abs/10.1126/scirobotics.abk2822>
- [49] Z. Zhuang et al., "Robot Parkour learning," in *Proc. Conf. Robot Learn.*, 2023, pp. 73–92, vol. 229.
- [50] C. Zhang, N. Rudin, D. Hoeller, and M. Hutter, "Learning agile locomotion on risky terrains," in *IEEE/RSJ Int. Conf. Intell. Robot. Syst.*, Abu Dhabi, United Arab Emirates, 2024, pp. 11864–11871, doi: [10.1109/IROS58592.2024.10801909](https://doi.org/10.1109/IROS58592.2024.10801909).
- [51] F. Jenelten, J. He, F. Farshidian, and M. Hutter, "Dtc: Deep tracking control," *Sci. Robot.*, vol. 9, no. 86, 2024, Art. no. eadh5401. [Online]. Available: <https://www.science.org/doi/abs/10.1126/scirobotics.adh5401>
- [52] X. Cheng, K. Shi, A. Agarwal, and D. Pathak, "Extreme parkour with legged robots," *arXiv:2309.14341*, 2023.
- [53] Y. Wang, Z. Jiang, and J. Chen, "Learning robust, agile, natural legged locomotion skills in the wild," in *RoboLetics: Workshop Robot Learn. Athletics, CoRL 2023*, 2023. [Online]. Available: <https://openreview.net/forum?id=b5hiuuX1sm>
- [54] A. Agarwal, A. Kumar, J. Malik, and D. Pathak, "Legged locomotion in challenging terrains using egocentric vision," in *Proc. Conf. Robot Learn.*, 2023, pp. 403–415, vol. 205.
- [55] C. Zhang et al., "Resilient legged local navigation: Learning to traverse with compromised perception end-to-end," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2024, pp. 34–41.
- [56] H. Liu and Q. Yuan, "Safe and robust motion planning for autonomous navigation of quadruped robots in cluttered environments," *IEEE Access*, vol. 12, pp. 69728–69737, 2024.
- [57] D. Hoeller, L. Wellhausen, F. Farshidian, and M. Hutter, "Learning a state representation and navigation in cluttered and dynamic environments," *IEEE Robot. Automat. Lett.*, vol. 6, no. 3, pp. 5081–5088, Jul. 2021.
- [58] R. Yang, M. Zhang, N. Hansen, H. Xu, and X. Wang, "Learning vision-guided quadrupedal locomotion end-to-end with cross-modal transformers," in *Proc. Int. Conf. Learn. Representations*, 2022. [Online]. Available: <https://openreview.net/forum?id=nhn3oo6AB>
- [59] N. Rudin, D. Hoeller, M. Bjelonic, and M. Hutter, "Advanced skills by learning locomotion and local navigation end-to-end," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2022, pp. 2497–2503.

- [60] C. Zhang, N. Rudin, D. Hoeller, and M. Hutter, "Learning agile locomotion on risky terrains," 2024, pp. 11864–11871.
- [61] M. Seo, R. Gupta, Y. Zhu, A. Skoutnev, L. Sentis, and Y. Zhu, "Learning to walk by steering: Perceptive quadrupedal locomotion in dynamic environments," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2023, pp. 5099–5105.
- [62] S. Kareer, N. Yokoyama, D. Batra, S. Ha, and J. Truong, "ViNL: Visual navigation and locomotion over obstacles," in *Proc. 2023 IEEE Int. Conf. Robot. Automat.*, London, U.K., 2023, pp. 2018–2024, doi: [10.1109/ICRA48891.2023.10160612](https://doi.org/10.1109/ICRA48891.2023.10160612).
- [63] J. Truong, M. Rudolph, N. H. Yokoyama, S. Chernova, D. Batra, and A. Rai, "Rethinking sim2real: Lower fidelity simulation leads to higher sim2real transfer in navigation," in *Proc. 6th Annu. Conf. Robot Learn.*, 2022. [Online]. Available: https://openreview.net/forum?id=BxHcg_Zlpxj
- [64] N. Yokoyama et al., "ASC: Adaptive skill coordination for robotic mobile manipulation," *IEEE Robot. Automat. Lett.*, vol. 9, no. 1, pp. 779–786, Jan. 2024.
- [65] E. Altman, *Constrained Markov Decision Processes*. Evanston, IL, USA: Routledge, 2021, pp. 21–55.
- [66] Y. Kim et al., "Not only rewards but also constraints: Applications on legged robot locomotion," *Trans. Robot.*, vol. 40, pp. 2984–3003, Jan. 2024, doi: [10.1109/TRO.2024.3400935](https://doi.org/10.1109/TRO.2024.3400935).
- [67] J. Achiam, D. Held, A. Tamar, and P. Abbeel, "Constrained policy optimization," in *Proc. Int. Conf. Mach. Learn.*, 2017.
- [68] Y. Liu, J. Ding, and X. Liu, "IPO: Interior-point policy optimization under constraints," in *Proc. AAAI Conf. Artif. Intell.*, 2020, vol. 34, pp. 4940–4947.
- [69] T.-Y. Yang, J. Rosca, K. Narasimhan, and P. J. Ramadge, "Projection-based constrained policy optimization," in *Proc. Int. Conf. Learn. Representations*, 2020. [Online]. Available: <https://openreview.net/forum?id=rke3TJrtPS>
- [70] T.-Y. Yang, T. Zhang, L. Luu, S. Ha, J. Tan, and W. Yu, "Safe reinforcement learning for legged locomotion," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2022, pp. 2454–2461.
- [71] J. Borquez, K. Chakraborty, H. Wang, and S. Bansal, "On safety and liveness filtering using Hamilton–Jacobi reachability analysis," *IEEE Trans. Robot.*, vol. 40, pp. 4235–4251, 2024.
- [72] K. P. Wabersich et al., "Data-driven safety filters: Hamilton–Jacobi reachability, control barrier functions, and predictive methods for uncertain systems," *IEEE Control Syst. Mag.*, vol. 43, no. 5, pp. 137–177, Oct. 2023.
- [73] K.-C. Hsu, H. Hu, and J. F. Fisac, "The safety filter: A unified view of safety-critical control in autonomous systems," *Annu. Rev. Control Robot., Auton. Syst.*, vol. 7, 2023.
- [74] S. Bansal, M. Chen, S. Herbert, and C. J. Tomlin, "Hamilton–Jacobi reachability: A brief overview and recent advances," in *Proc. IEEE Conf. Decis. Control*, 2017, pp. 2242–2253.
- [75] S. Bansal, V. Tolani, S. Gupta, J. Malik, and C. Tomlin, "Combining optimal control and learning for visual navigation in novel environments," in *Proc. Conf. Robot Learn.*, 2019, pp. 420–429.
- [76] J. Borquez, K. Nakamura, and S. Bansal, "Parameter-conditioned reachable sets for updating safety assurances online," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2023, pp. 10553–10559.
- [77] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: Theory and applications," in *Proc. 18th Eur. Control Conf.*, 2019, pp. 3420–3431.
- [78] J. J. Choi, D. Lee, K. Sreenath, C. J. Tomlin, and S. L. Herbert, "Robust control barrier–value functions for safety-critical control," in *Proc. 60th IEEE Conf. Decis. Control.*, 2021, pp. 6814–6821, doi: [10.1109/CDC45484.2021.9683085](https://doi.org/10.1109/CDC45484.2021.9683085).
- [79] Unitree Robotics. `unitree_ros`. [Online]. Available: https://github.com/unitreerobotics/unitree_ros
- [80] J. Lygeros, "On reachability and minimum cost optimal control," *Automatica*, vol. 40, no. 6, pp. 917–927, 2004.
- [81] I. Mitchell, A. Bayen, and C. J. Tomlin, "A time-dependent Hamilton–Jacobi formulation of reachable sets for continuous dynamic games," *IEEE Trans. Automat. Control*, vol. 50, no. 7, pp. 947–957, Jul. 2005.
- [82] Edward Schmerling, "hj_reachability: Hamilton–Jacobi reachability analysis in jax," [Online]. Available: https://github.com/StanfordASL/hj_reachability
- [83] S. Bansal and C. J. Tomlin, "DeepReach: A deep learning approach to high-dimensional reachability," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2021, pp. 1817–1824.
- [84] A. N. Angelopoulos and S. Bates, "Conformal prediction: A gentle introduction," *Found. Trends Mach. Learn.*, vol. 16, no. 4, pp. 494–591, 2023, doi: [10.1561/2200000101](https://doi.org/10.1561/2200000101).
- [85] A. Lin and S. Bansal, "Verification of neural reachable tubes via scenario optimization and conformal prediction," in *Proc. 6th Annu. Learn. Dyn. Control Conf.*, (Proceedings of Machine Learning Research Series), A. M. Abate, K. Cannon Margellos, and A. Papachristodoulou, Eds., vol. 242. PMLR, Jul. 15–17, 2024, pp. 719–731. [Online]. Available: <https://proceedings.mlr.press/v242/lin24a.html>
- [86] A. K. Akametalu, J. F. Fisac, J. H. Gillula, S. Kaynama, M. N. Zeilinger, and C. J. Tomlin, "Reachability-based safe learning with Gaussian processes," in *Proc. 53rd IEEE Conf. Decis. Control*, 2014, pp. 1424–1431.
- [87] B. Steux and O. E. Hamzaoui, "tinySLAM: A SLAM algorithm in less than 200 lines C-language program," in *Proc. 11th Int. Conf. Control Automat. Robot. Vis.*, 2010, pp. 1975–1979.
- [88] M. Chen, S. Herbert, and C. J. Tomlin, "Exact and efficient Hamilton–Jacobi guaranteed safety analysis via system decomposition," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2017, pp. 87–92, doi: [10.1109/ICRA.2017.7989015](https://doi.org/10.1109/ICRA.2017.7989015).
- [89] Z. Feng, L. Qiu, and S. Bansal, "Bridging model predictive control and deep learning for scalable reachability analysis," 2025, *arXiv:2505.03830*.



Albert Lin received the B.S.E. degree in computer science from Princeton University, Princeton, NJ, USA, in 2023. He is currently working toward the Ph.D. degree in aeronautics and astronautics with Stanford University, Stanford, CA, where he is advised by Prof. Somil Bansal.

He is a member of the Safe and Intelligent Autonomy Lab. His research interests are at the intersection of machine learning and cognitive science, especially as it relates to robot control and safety.



Shuang Peng received the B.S. degree in robotics engineering from Southeast University, Nanjing, China, in 2022. He is currently working toward the Ph.D. degree in electrical and computer engineering with the University of Southern California, Los Angeles, CA, USA.

He is a member of the Safe and Intelligent Autonomy Lab. His research interests include exploring control and safety assurances for legged robots.



Somil Bansal (Member, IEEE) received the B.Tech. degree in electrical engineering from the Indian Institute of Technology, Kanpur, India, and the M.S. degree in electrical engineering and computer sciences from UC Berkeley, Berkeley, USA, in 2012 and 2014, respectively. He received the Ph.D. degree in electrical engineering and computer sciences from the University of California, Berkeley, in 2020.

He is an Assistant Professor with the Aeronautics and Astronautics Department, Stanford University, where he leads the Safe and Intelligent Autonomy Lab. Previously, he was an Assistant Professor with the Electrical and Computer Engineering Department, University of Southern California. His research focuses on understanding how machine learning methods can be combined with classical, model-based planning and control methods to enable intelligent and safe decision-making.