

# Safe Navigation under State Uncertainty: Online Adaptation for Robust Control Barrier Functions

Ersin Daş<sup>1</sup>, Rahal Nanayakkara<sup>2</sup>, Xiao Tan<sup>1</sup>, Ryan M. Bena<sup>1</sup>,  
Joel W. Burdick<sup>1</sup>, Paulo Tabuada<sup>2</sup>, and Aaron D. Ames<sup>1</sup>

**Abstract**—Measurements and state estimates are often imperfect in control practice, posing challenges for safety-critical applications, where safety guarantees rely on accurate state information. In the presence of estimation errors, several prior robust control barrier function (R-CBF) formulations have imposed strict conditions on the input. These methods can be overly conservative and can introduce issues such as infeasibility, high control effort, etc. This work proposes a systematic method to improve R-CBFs, and demonstrates its advantages on a tracked vehicle that navigates among multiple obstacles. A primary contribution is a new optimization-based online parameter adaptation scheme that reduces the conservativeness of existing R-CBFs. In order to reduce the complexity of the parameter optimization, we merge several safety constraints into one unified numerical CBF via Poisson’s equation. We further address the dual relative degree issue that typically causes difficulty in vehicle tracking. Experimental trials demonstrate the overall performance improvement of our approach over existing formulations.

**Index Terms**—Safe navigation, constrained control, control barrier functions, state uncertainty, robust nonlinear control

## I. INTRODUCTION

THE increasing use of autonomous robots in complex environments emphasizes the need for safety-critical control [1]. Control barrier functions (CBFs) [2] provide conditions, by leveraging system dynamics, that render a prescribed safe set forward invariant and, in practice, lead to safety filters that minimally modify the nominal input only when necessary to maintain safety. However, model uncertainties, including external disturbances, time delays, and state estimation errors, can lead to safety violations in real-world control applications [3]. For instance, modern mobile robots and uncrewed aerial vehicles often rely upon visual-inertial odometry (VIO) for localization and navigation. Under VIO, the uncertainty in the vehicle’s location estimate grows between global positioning updates, and is an ever-present source of state estimation error that can render safety guarantees meaningless unless they are robustified against these persistent errors.

Early robust CBF methods addressed unmodeled dynamics through input-to-state safety margin [4]. Later formulations

Manuscript received: August, 26, 2025; Revised November, 21, 2025; Accepted December, 18, 2025.

This paper was recommended for publication by Editor Lucia Pallottino upon evaluation of the Associate Editor and Reviewers’ comments.

This work was supported in part by TII under project #A6847 and in part by DARPA under the LINC Program. (Corresponding author: Ersin Daş.)

<sup>1</sup>E. Daş, X. Tan, R. M. Bena, J. W. Burdick, and A. D. Ames are with the Department of Mechanical and Civil Engineering, California Institute of Technology, Pasadena, CA 91125, USA, {ersindas,xiaotan,ryanbena,jburdick,ames}@caltech.edu

<sup>2</sup>R. Nanayakkara and P. Tabuada are with the Electrical and Computer Engineering Department, University of California at Los Angeles, Los Angeles, CA 90095, USA, rahaln@ucla.edu, tabuada@ee.ucla.edu

Digital Object Identifier (DOI) 10.1109/LRA.2026.3653366

©2026 IEEE

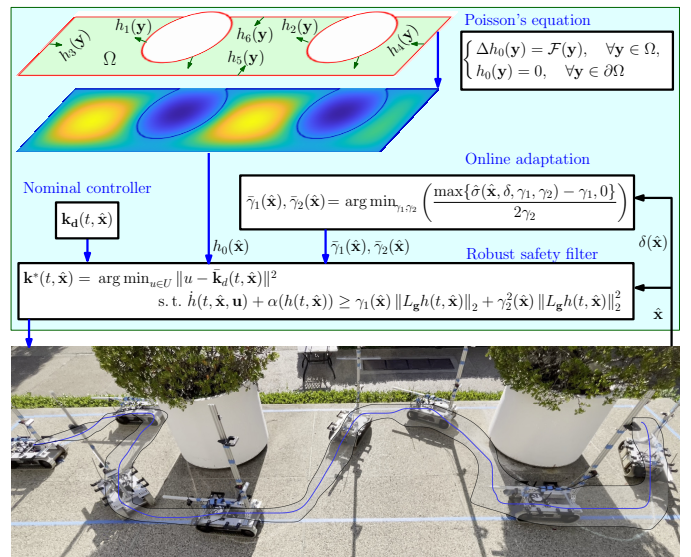


Fig. 1. The proposed robust CBF (R-CBF) method in an outdoor safety-critical navigation scenario with a tracked mobile robot. Our R-CBF-QP-based safety filter maintains safety under state estimation uncertainty.

ensure safety by introducing adaptive robustness margins to reduce conservatism [3], [5]. One recent approach [6] learns a predictive robustness term that is updated online to tighten the CBF inequality whenever the full-order robot dynamics model deviates from its reduced-order model, assuming accurate state or output measurements. To enhance robustness in safety constraints, [7] presents a learning-based framework that enables online adaptation of class- $\mathcal{K}_\infty$  functions for systems under input constraints.

Beyond model uncertainty, robustness to measurement errors is notoriously challenging because CBF conditions are evaluated on the estimated state. This problem has been approached by using more conservative conditions [8], [9], by considering specific types of state estimators [10], [11], or by adopting stochastic approaches [12]. Although effective in specific control problems, each of these methods imposes specific requirements on the state estimator or uncertainty bounds, which may not be applicable to some real-world applications. To overcome these limitations, a concept for robust CBFs (R-CBFs) is introduced in [13] to ensure safety in the presence of state uncertainty, without requiring prior knowledge of the uncertainty magnitude. Unlike the approach proposed in [10], [11], R-CBFs do not depend on specific classes of state observers. Nevertheless, R-CBFs account for estimation errors through conditions with constant robustness parameters, which may lead to overcompensation for uncertainty, ultimately reducing system performance in an undesirable manner.

**IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.**

In this paper, we present a safety-critical control framework that is robust to state estimation uncertainty, with a focus on practical implementation. We build upon the R-CBF formulation in [13] and propose an optimization-based procedure for the online tuning of the robustifying parameters in the R-CBF condition. This strategy enables the controller to adaptively guarantee robustness to state uncertainty by optimizing the R-CBF parameters using perturbations sampled locally about the current state estimate.

In order to account for multiple safety specifications—which commonly arise in real-world autonomy problems [14]–[17]—and reduce the computational complexity of the online parameter optimization, we consolidate them into a single CBF using Poisson safety functions [18], [19]. This also facilitates the generation of safe sets from perception data. Rather than composing safety constraints with approximations, the Poisson safety function provides a smooth scalar field by solving Poisson’s equation with Dirichlet boundary conditions defined by the unsafe and safe set boundaries.

We demonstrate the effectiveness of our approach through hardware experiments on a tracked mobile robot, where state estimation is provided by VIO. We model the robot as a unicycle system. Our objective is to provide collision safety even under the time-varying uncertainty introduced by the VIO system. The linear and angular velocity inputs to this vehicle affect the safety function at two different orders of differentiation, resulting in *dual relative degree* behavior. Inspired by [20], we first design a CBF for the single integrator subsystem, then we leverage a Lyapunov function to synthesize a CBF that has a relative degree of one with respect to both inputs for the nonlinear unicycle system. The synthesized CBF explicitly incorporates the effect of angular velocity into the safety condition, overcoming a limitation of conventional CBFs for the unicycle system, where angular velocity does not contribute to the satisfaction of the safety specifications.

We organize this paper as follows. Section II introduces preliminaries. Section III presents an online adaptation method for the robustifying parameters of R-CBFs, aimed at reducing conservativeness. Section IV presents simulations and experiments, while Section V concludes the paper.

## II. PRELIMINARIES

### A. Control Barrier Functions

We consider a nonlinear control-affine system:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{u}, \quad \mathbf{x} \in \mathcal{X} \subseteq \mathbb{R}^n, \quad \mathbf{u} \in \mathcal{U} \subseteq \mathbb{R}^m, \quad (1)$$

where the drift dynamics  $\mathbf{f}: \mathcal{X} \rightarrow \mathbb{R}^n$  and the actuation matrix  $\mathbf{g}: \mathcal{X} \rightarrow \mathbb{R}^{n \times m}$  are locally Lipschitz continuous functions. We assume the set of admissible control inputs  $\mathcal{U}$  is an  $m$ -dimensional convex polytope. For a given time-varying feedback controller  $\mathbf{k}: \mathbb{R}_{\geq 0} \times \mathcal{X} \rightarrow \mathcal{U}$ , which is piecewise continuous in  $t$  and locally Lipschitz continuous in  $\mathbf{x}$ , the closed-loop system,  $\mathbf{f}_{cl}: \mathbb{R}_{\geq 0} \times \mathcal{X} \rightarrow \mathbb{R}^n$ :

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{k}(t, \mathbf{x}) \triangleq \mathbf{f}_{cl}(t, \mathbf{x}), \quad (2)$$

has a unique solution:  $\mathbf{x}(t) = \mathbf{x}_0 + \int_{t_0}^t \mathbf{f}_{cl}(\tau, \mathbf{x}(\tau))d\tau$ ,  $t > t_0$ , for any initial condition  $\mathbf{x}(t_0) = \mathbf{x}_0 \in \mathcal{X}$ . In this paper, we assume that the system is forward complete.

We characterize safety through the notion of forward invariance. Let the time-varying safe set  $\mathcal{C}(t)$  be the 0-superlevel set of a continuously differentiable function  $h: \mathbb{R}_{\geq 0} \times \mathcal{X} \rightarrow \mathbb{R}$ :

$$\mathcal{C}(t) \triangleq \{\mathbf{x} \in \mathcal{X} \subseteq \mathbb{R}^n : h(t, \mathbf{x}) \geq 0\}.$$

This set is said to be *forward invariant* if, for any initial condition  $\mathbf{x}(t_0) \in \mathcal{C}(t_0)$ , the solution  $\mathbf{x}(t)$  stays in  $\mathcal{C}(t)$ ,  $\forall t \geq t_0$ . The closed-loop system (2) is *safe* with respect to the *safe set*  $\mathcal{C}(t)$  if  $\mathcal{C}(t)$  is forward invariant. To design controllers that guarantee such forward invariance properties, CBFs [2] provide a systematic framework.

**Definition 1** (Control Barrier Function, [2]). Let  $\mathcal{C}(t) \subseteq \mathcal{X}$  be the 0-superlevel set of a continuously differentiable function  $h: \mathbb{R}_{\geq 0} \times \mathcal{X} \rightarrow \mathbb{R}$  with 0 as a regular value. The function  $h$  is a *control barrier function* for (1) on  $\mathcal{C}(t)$  if there exists an extended class- $\mathcal{K}_\infty$  function\*  $\alpha \in \mathcal{K}_\infty^e$  such that  $\forall \mathbf{x} \in \mathcal{C}(t)$ :

$$\sup_{\mathbf{u} \in \mathcal{U}} [L_{\mathbf{f}}h(t, \mathbf{x}) + L_{\mathbf{g}}h(t, \mathbf{x})\mathbf{u} + \partial_t h(t, \mathbf{x})] \geq -\alpha(h(t, \mathbf{x})),$$

where  $L_{\mathbf{f}}h = \frac{\partial h(t, \mathbf{x})}{\partial \mathbf{x}} \mathbf{f}(\mathbf{x})$ ,  $L_{\mathbf{g}}h = \frac{\partial h(t, \mathbf{x})}{\partial \mathbf{x}} \mathbf{g}(\mathbf{x})$ ,  $\partial_t h = \frac{\partial h(t, \mathbf{x})}{\partial t}$ .

We can utilize CBFs to design a *safety filter*. In particular, consider a baseline *nominal controller*  $\mathbf{k}_d: \mathbb{R}_{\geq 0} \times \mathcal{X} \rightarrow \mathcal{U}$  that is locally Lipschitz continuous but potentially unsafe, along with a CBF  $h$  with a corresponding  $\alpha \in \mathcal{K}_\infty^e$  for system (1). We enforce safety constraints using the CBF-Quadratic Program (CBF-QP) [2]:

$$\begin{aligned} \mathbf{k}^*(t, \mathbf{x}) = \arg \min_{\mathbf{u} \in \mathcal{U}} \quad & \|\mathbf{u} - \mathbf{k}_d(t, \mathbf{x})\|^2 \\ \text{s.t.} \quad & \dot{h}(t, \mathbf{x}, \mathbf{u}) \geq -\alpha(h(t, \mathbf{x})). \end{aligned}$$

## III. ROBUST CONTROL BARRIER FUNCTIONS WITH ONLINE PARAMETER ADAPTATION

We first define the concept of a time-varying robust CBF (R-CBF) to account for state uncertainty. Then, we introduce an optimization method for online adaptation of the R-CBF tuning parameters.

### A. Robust Control Barrier Functions

In practice, feedback control systems rely on an estimate,  $\hat{\mathbf{x}} \in \mathcal{X} \subseteq \mathbb{R}^n$  of the true state  $\mathbf{x} \in \mathcal{X} \subseteq \mathbb{R}^n$  for the calculation of the control input. This estimate is typically obtained through an observer or filter, and is subject to bounded estimation uncertainties satisfying  $\|\mathbf{x} - \hat{\mathbf{x}}\| \leq \delta(\mathbf{x})$ , for some  $\delta: \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$ , where  $\mathbf{e} \triangleq \mathbf{x} - \hat{\mathbf{x}}$  is the state estimation error.

**Assumption 1.** Every non-positive real number is a regular value of  $h$ ,  $\forall t \geq 0$ , i.e.,  $\nabla_{\mathbf{x}} h(t, \mathbf{x}) \neq 0$  when  $h(t, \mathbf{x}) \leq 0$ ,  $\forall t \geq 0$ .

**Assumption 2.** Every non-positive super-level set of  $h$  is compact and uniformly bounded, i.e.,  $\mathcal{C}_\beta(t) \triangleq \{\mathbf{x} \in \mathcal{X} : h(t, \mathbf{x}) \geq -\beta\}$  is compact for all  $\beta \in \mathbb{R}_{\geq 0}$ ,  $t \geq 0$ , and  $\mathcal{C}_\beta(t) \subseteq S$  for all  $t \geq 0$ , where  $S$  is some compact set.

\*A continuous function  $\alpha: [0, a) \rightarrow \mathbb{R}_{\geq 0}$ , where  $a > 0$ , belongs to class- $\mathcal{K}$  ( $\alpha \in \mathcal{K}$ ) if it is strictly monotonically increasing and  $\alpha(0) = 0$ . A continuous function  $\alpha: \mathbb{R} \rightarrow \mathbb{R}$  belongs to the set of extended class- $\mathcal{K}_\infty$  functions ( $\alpha \in \mathcal{K}_\infty^e$ ) if it is strictly monotonically increasing,  $\alpha(0) = 0$ ,  $\lim_{r \rightarrow \infty} \alpha(r) = \infty$  and  $\lim_{r \rightarrow -\infty} \alpha(r) = -\infty$ . A continuous function  $\beta: [0, a) \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$  belongs to class- $\mathcal{KL}$  ( $\beta \in \mathcal{KL}$ ), if for every  $s \in \mathbb{R}_{\geq 0}$ ,  $\beta(\cdot, s)$  is a class- $\mathcal{K}$  function and for every  $r \in [0, a)$ ,  $\beta(r, \cdot)$  is decreasing and  $\lim_{s \rightarrow \infty} \beta(r, s) = 0$ .

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

**Definition 2.** A time-varying compact set  $\mathcal{S}(t) \subseteq \mathbb{R}^n$  is asymptotically stable if there exists  $\rho \in \mathcal{KL}$  such that

$$\|\mathbf{x}(t)\|_{\mathcal{S}(t)} \leq \rho(\|\mathbf{x}(t_0)\|_{\mathcal{S}(t_0)}, t - t_0), \quad \forall t \geq t_0,$$

where  $\|\mathbf{x}(t)\|_{\mathcal{S}(t)} \triangleq \inf_{y \in \mathcal{S}(t)} \|\mathbf{x}(t) - y\|$  is the distance between the point  $\mathbf{x}$  and the set  $\mathcal{S}$  at time  $t$ .

The notion of a robust control barrier function [13] has been proposed to synthesize safety-critical controllers under state estimation uncertainty.

**Definition 3** (Robust Control Barrier Function, [13]). A continuously differentiable function  $h: \mathbb{R}_{\geq 0} \times \mathcal{X} \rightarrow \mathbb{R}$ , is a *robust control barrier function* parameterized by state-dependent functions  $\gamma_1, \gamma_2: \mathcal{X} \rightarrow \mathbb{R}_{> 0}$  for system (1) with respect to the 0-superlevel set of  $h$ ,  $\mathcal{C}(t) \subseteq \mathcal{X}$ , if there exists  $\alpha \in \mathcal{K}_{\infty}^e$  such that:

$$\begin{aligned} & \sup_{\mathbf{u} \in \mathcal{U}} [L_f h(t, \hat{\mathbf{x}}) + L_g h(t, \hat{\mathbf{x}}) \mathbf{u} + \partial_t h(t, \hat{\mathbf{x}}) + \alpha(h(t, \hat{\mathbf{x}}))] \\ & \geq \gamma_1(\hat{\mathbf{x}}) \|L_g h(t, \hat{\mathbf{x}})\|_2 + \gamma_2^2(\hat{\mathbf{x}}) \|L_g h(t, \hat{\mathbf{x}})\|_2^2, \end{aligned} \quad (3)$$

for all  $\hat{\mathbf{x}} \in \hat{\mathcal{X}}$ , where  $\hat{\mathbf{x}} = \mathbf{x} - \mathbf{e}$ .

Compared to the robust CBF formulation in [13], which introduces a general robustness function, our R-CBF definition uses explicit state-dependent robustness terms  $\gamma_1$  and  $\gamma_2$ . This structured choice is a specialization of the framework in [13] and is tailored to simplify the online adaptation procedure. Given an R-CBF  $h$  and a corresponding  $\alpha$ , the pointwise set of all control values that satisfy (3) is given by

$$\begin{aligned} K_{R\text{-CBF}}(t, \hat{\mathbf{x}}) & \triangleq \{ \mathbf{u} \in \mathcal{U} \mid \dot{h}(t, \hat{\mathbf{x}}, \mathbf{u}) + \alpha(h(t, \hat{\mathbf{x}})) \\ & \geq \gamma_1(\hat{\mathbf{x}}) \|L_g h(t, \hat{\mathbf{x}})\|_2 + \gamma_2^2(\hat{\mathbf{x}}) \|L_g h(t, \hat{\mathbf{x}})\|_2^2 \}. \end{aligned} \quad (4)$$

We can establish formal robust safety guarantees based on Definition 3 with the help of the following theorem:

**Theorem 1.** Let  $h: \mathbb{R}_{\geq 0} \times \mathcal{X} \rightarrow \mathbb{R}$  be an R-CBF for (1), on its 0-superlevel set  $\mathcal{C}(t)$  with an  $\alpha \in \mathcal{K}_{\infty}^e$ . Assume that  $\|\mathbf{x} - \hat{\mathbf{x}}\| \leq \delta(\mathbf{x})$ , for all  $\mathbf{x} \in \mathcal{X}$ , for some  $\delta: \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$ , and that Assumption 1 and 2 hold. Then, any locally Lipschitz feedback controller  $\mathbf{k}(t, \hat{\mathbf{x}})$  with Lipschitz constant  $\mathcal{L}_{\mathbf{k}}$ , satisfying  $\mathbf{k}(t, \hat{\mathbf{x}}) \in K_{R\text{-CBF}}(t, \hat{\mathbf{x}})$  for all  $(t, \hat{\mathbf{x}}) \in \mathbb{R}_{\geq 0} \times \hat{\mathcal{X}}$ , renders the original set  $\mathcal{C}(t)$  safe and asymptotically stable for (1) if:

$$\delta(\mathbf{x}) \leq \frac{\gamma_1(\mathbf{x})}{\mathcal{L}_{\mathbf{k}}}, \quad \forall \mathbf{x} \in \mathcal{C}(t). \quad (5)$$

Otherwise, for any compact superlevel set  $\mathcal{C}_{\beta}(t)$ , defined in Assumption 2, the same controller ensures safety and asymptotic stability of the inflated set given by

$$\tilde{\mathcal{C}}(t) \triangleq \{ \mathbf{x} \in \mathcal{X} : h(t, \mathbf{x}) \geq \alpha^{-1}(-\bar{\phi}^2(t, \gamma_1(\mathbf{x}), \gamma_2(\mathbf{x}), \delta(\mathbf{x}))) \},$$

where

$$\begin{aligned} \bar{\phi}(t, \gamma_1(\mathbf{x}), \gamma_2(\mathbf{x}), \delta(\mathbf{x})) & = \frac{\sigma_{\beta}(t, \delta(\mathbf{x})) - \gamma_1(\mathbf{x})}{2\gamma_2(\mathbf{x})}, \\ \sigma_{\beta}(t, \delta(\mathbf{x})) & \triangleq \sup_{\hat{\mathbf{x}} \in \mathcal{C}_{\beta}(t)} \left[ \sup_{\|\mathbf{e}\| \leq \delta(\mathbf{x})} [\|\mathbf{k}(t, \hat{\mathbf{x}}) - \mathbf{k}(t, \hat{\mathbf{x}} + \mathbf{e})\|] \right]. \end{aligned} \quad (6)$$

*Proof.* The proof follows from Theorem 2 and Corollary 3 in [13], except for the explicit time dependence of  $h$ , and state-dependent parameters  $\gamma_1, \gamma_2$ . The explicit time dependence of  $h$  is handled by our assumption that the superlevel sets of  $h$  are uniformly bounded in time.  $\square$

*Remark 1.* Definition 3 and Theorem 1 extend the R-CBF framework from [13] in two key ways: it employs state-dependent parameters  $\gamma_1(\hat{\mathbf{x}}), \gamma_2(\hat{\mathbf{x}})$  rather than constant parameters  $\gamma_1, \gamma_2 \in \mathbb{R}_{> 0}$ , and it characterizes robust safety for time-varying CBFs. Thus, this definition generalizes *tunable* R-CBFs [4], and Corollary 2 in [13] becomes a special case when the state-dependent parameters take the form  $\gamma_1(\hat{\mathbf{x}}) = \frac{\gamma_1}{\varepsilon_1(h(t, \hat{\mathbf{x}}))}$  and  $\gamma_2^2(\hat{\mathbf{x}}) = \frac{\gamma_2^2}{\varepsilon_2(h(t, \hat{\mathbf{x}}))}$ , where  $\varepsilon_1, \varepsilon_2: \mathbb{R} \rightarrow \mathbb{R}_{> 0}$  are continuous functions that satisfy  $\varepsilon_i(r) = 1$  when  $r \leq 0$ , and  $\frac{d\varepsilon_i}{dr} > 0$  when  $r > 0$ , for  $i = 1, 2$ .

## B. Online Parameter Adaptation

In this subsection, an online adaptation method for the robustifying parameters in the R-CBF inequality is presented. We omit the dependence on  $\mathbf{x}$  from  $\gamma_1, \gamma_2$  and  $\delta$  for clarity in our notation.

Recall that a controller satisfying (4) guarantees that the following set will be forward invariant and asymptotically stable:

$$\tilde{\mathcal{C}}(t) \triangleq \{ \mathbf{x} \in \mathcal{X} : h(t, \mathbf{x}) \geq \alpha^{-1}(-\phi^2(t, \gamma_1, \gamma_2, \delta)) \},$$

where  $\phi: \mathbb{R}_{\geq 0} \times \mathbb{R}_{> 0} \times \mathbb{R}_{> 0} \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$  is a measure of the set inflation, given by

$$\phi(t, \gamma_1, \gamma_2, \delta) = \frac{\max\{\sigma_{\beta}(t, \delta) - \gamma_1, 0\}}{2\gamma_2}. \quad (7)$$

Note that the max function in (7) combines the two cases discussed in Theorem 1. The function  $\sigma_{\beta}$ , which is defined in (6), with some superlevel set of  $h$ ,  $\mathcal{C}_{\beta}$ , upper bounds the maximum possible change in the feedback input for a given bound on the state estimation error. We also note that this function  $\sigma_{\beta}$  will be a function of  $\gamma_1, \gamma_2$ , since the closed-loop controller  $\mathbf{k}$  in (6) will be a function of  $\gamma_1, \gamma_2$ .

Equation (7) gives an explicit expression about the safe set inflation for an uncertainty bound  $\delta$  and the robustifying parameters  $\gamma_1$  and  $\gamma_2$ . The uncertainty bound  $\delta$  information, albeit time-varying and possibly conservative, is usually available, for example, provided by VIO. In the following, we then aim to minimize this set inflation by choosing proper parameters  $\gamma_1$  and  $\gamma_2$ .

To select the parameters online, we first introduce a local version of  $\sigma_{\beta}(t, \delta(\mathbf{x}))$  in (6) around the current state estimate  $\hat{\mathbf{x}}$ , since CBF conditions are evaluated pointwise at the current state estimate:

$$\bar{\sigma}(t, \hat{\mathbf{x}}, \delta, \gamma_1, \gamma_2) = \sup_{\|\mathbf{e}\| \leq \delta} (\|\mathbf{k}(t, \hat{\mathbf{x}}) - \mathbf{k}(t, \hat{\mathbf{x}} + \mathbf{e})\|), \quad (8)$$

which upper bounds the variation of the feedback input due to state estimation error locally at the given state estimate. It is still not possible to compute (8) explicitly, as it requires evaluating the supremum around the current estimate. We now approximate this by first sampling  $N \in \mathbb{N}$  points,  $\{\tilde{\mathbf{x}}_1, \dots, \tilde{\mathbf{x}}_N\}$ , from the closed ball of radius  $\delta$  centered at  $\hat{\mathbf{x}}$  in  $\mathbb{R}^n$ ,  $B_{\delta}(\hat{\mathbf{x}})$ , and then evaluating:

$$\hat{\sigma}(t, \hat{\mathbf{x}}, \delta, \gamma_1, \gamma_2) = \max_{\mathbf{x} \in \{\tilde{\mathbf{x}}_1, \dots, \tilde{\mathbf{x}}_N\}} \|\mathbf{k}(t, \mathbf{x}) - \mathbf{k}(t, \hat{\mathbf{x}})\|. \quad (9)$$

Now, at any given  $\hat{\mathbf{x}}$  and  $\delta$ , we substitute  $\hat{\sigma}(t, \hat{\mathbf{x}}, \delta, \gamma_1, \gamma_2)$  as an approximation of  $\sigma_{\beta}(t, \delta(\mathbf{x}))$  in (7), and select the optimal

**IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.**

constants  $\gamma_1, \gamma_2$  by solving:

$$\bar{\gamma}_1(t, \hat{\mathbf{x}}, \delta), \bar{\gamma}_2(t, \hat{\mathbf{x}}, \delta) = \arg \min_{\gamma_1, \gamma_2 \in \mathbb{R}_{>0}} \left( \frac{\max\{\hat{\sigma}(t, \hat{\mathbf{x}}, \delta, \gamma_1, \gamma_2) - \gamma_1, 0\}}{2\gamma_2} \right). \quad (10)$$

Since it is complex to explicitly compute  $\frac{d\hat{\sigma}}{d\gamma_1}, \frac{d\hat{\sigma}}{d\gamma_2}$ , we resort to using derivative-free optimization methods. We remark that  $\hat{\sigma}$  in (9) is a practical, sampling-based approximation of  $\bar{\sigma}$ .

*Remark 2.* To obtain a robust safe control  $\mathbf{k}(t, \mathbf{x})$ , we may replace the nominal CBF condition in the CBF-QP with an R-CBF condition to obtain an R-CBF-QP. When optimizing the parameters  $\gamma_1$  and  $\gamma_2$  via (10), this R-CBF-QP needs to be solved for each sampling point  $\hat{\mathbf{x}}_i$ . When  $\mathcal{U} = \mathbb{R}^m$ , closed-form solutions for a QP with a single constraint can facilitate fast computation.

If we utilize multiple CBFs, resulting from multiple independent safety specifications, directly in the same QP, then we need to optimize two different robustness parameters for each CBF, which consequently increases the dimensionality of the parameter optimization problem. In contrast, by using a single unified CBF constraint in the R-CBF-QP, we only need to optimize a single set of parameters  $\bar{\gamma}_1, \bar{\gamma}_2$ , thereby simplifying the optimization process. This motivates the need for a general CBF unification method, as will be described in the next section.

### C. Unifying CBFs via Poisson Safety Functions

A safety function  $h_0: \mathbf{y} \mapsto h_0(\mathbf{y}) \in \mathbb{R}$  is a continuous function that indicates the value of safety for a particular output coordinate  $\mathbf{y} \in \mathbb{R}^p$ . Crucially, a coordinate is considered safe when  $h_0(\mathbf{y}) \geq 0$  and unsafe when  $h_0(\mathbf{y}) < 0$ . A safety function  $h_0$  differs from a CBF  $h$  in that  $h_0$  characterizes the desired output-space safe set without considering its relationship with the underlying system dynamics. To generate  $h_0$ , we employ the Poisson safety function method introduced in [18], [19].

First, we must construct a bounded domain  $\Omega$  corresponding to the interior of some desired safe set. For the case of multiple safety functions  $h_i$  (these can be CBFs) associated with 0-superlevel sets  $\mathcal{C}_i$  for  $i \in \{1, \dots, N_h\}$ , the new desired safe set can be defined as the intersection of the 0-superlevel sets of these functions:  $\mathcal{C} = \bigcap_{i=1}^{N_h} \mathcal{C}_i$ , since the safety task is to satisfy all constraints simultaneously. Equivalently, this intersection can be expressed via the min function [21]:

$$\Omega \triangleq \bigcap_{i=1}^{N_h} \mathcal{C}_i = \left\{ \mathbf{y} \in \mathbb{R}^p : \min_{i \in \{1, \dots, N_h\}} [h_i(\mathbf{y})] \geq 0 \right\}.$$

Given this domain  $\Omega$ , we construct a unified safety function  $h_0$  on  $\Omega$  by solving a Dirichlet problem (boundary value problem) for Poisson's equation [18]:

$$\begin{cases} \Delta h_0(\mathbf{y}) = \mathcal{F}(\mathbf{y}), & \forall \mathbf{y} \in \Omega, \\ h_0(\mathbf{y}) = 0, & \forall \mathbf{y} \in \partial\Omega, \end{cases} \quad (11)$$

where  $\Delta = \frac{\partial}{\partial \mathbf{y}} \cdot \frac{\partial}{\partial \mathbf{y}}$  is the Laplacian operator,  $\mathcal{F}: \Omega \rightarrow \mathbb{R}_{<0}$  is a forcing function that yields  $h_0(\mathbf{y}) \geq 0$  in safe regions, and the boundary  $\partial\Omega$  corresponds to the desired zero level set. Similarly, we can derive  $h_0$  in unsafe regions  $\hat{\Omega}$  by modifying the forcing function such that  $\mathcal{F}: \hat{\Omega} \rightarrow \mathbb{R}_{>0}$ .

The resultant Poisson safety function  $h_0$  can be used directly as a CBF for systems whose inputs have relative degree one<sup>†</sup>

<sup>†</sup>This is true for systems characterized by single integrator dynamics.




Fig. 2. Poisson safety function. The function  $h_0$  describes the spatial safety specification, and it can be used to generate a CBF, as in (16).

with respect to  $h_0$ . More generally,  $h_0$  can be dynamically extended to generate a CBF.

## IV. AN APPLICATION TO SAFE TRACKING FOR TRACKED ROBOTS

### A. Safety Filter Construction

This section experimentally demonstrates the effectiveness of the proposed safe controller on a tracked mobile robot, which is tasked to follow a reference trajectory while avoiding collisions.

Suppose that the robot is governed by unicycle dynamics:

$$\begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{\theta} \end{bmatrix} = \begin{bmatrix} \cos \theta & 0 \\ \sin \theta & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} v \\ \omega \end{bmatrix}, \quad (12)$$

where  $\mathbf{x} = [x \ y \ \theta]^\top \in \mathcal{X} = \mathbb{R}^2 \times \mathbb{S}^1$  represents the planar positions and heading angle with respect to an inertial frame, and  $\mathbf{u} = [v \ \omega]^\top \in \mathbb{R}^2$  denotes the linear and angular velocities as control inputs.

We define a reference-tracking control law:

$$\mathbf{k}_d(t, \hat{\mathbf{x}}) = \begin{bmatrix} K_v \|\mathbf{r}_e(t, \hat{\mathbf{x}})\| \\ K_\omega (\theta_d(t, \hat{\mathbf{x}}) - \theta) \end{bmatrix}, \quad (13)$$

where  $K_v, K_\omega \geq 0$  are controller gains for the linear and angular velocities,  $\theta_d: \mathbb{R}_{\geq 0} \times \mathcal{X} \rightarrow \mathbb{S}^1$  is the desired robot heading, and  $\mathbf{r}_e(t, \mathbf{x}) \triangleq [x_d(t) \ y_d(t)]^\top - [x \ y]^\top$  is the position-error vector with respect to the desired position  $(x_d, y_d)$  in  $\mathbb{R}^2$ . The desired heading angle is given by

$$\theta_d(t, \mathbf{x}) = \text{atan}(y_d(t) - y, x_d(t) - x). \quad (14)$$

For our particular safety-critical scenario (as pictured in Fig. 1), we consider a reference trajectory parameterized by:

$$x_d(t) = v_{\text{ref}} t, \quad y_d(t) = a_d \sin(c_d x_d(t) + \varphi_d), \quad (15)$$

with  $a_d, c_d$ , and  $\varphi_d$  being the amplitude, spatial frequency, and phase shift, respectively, and  $v_{\text{ref}}$  is the desired linear velocity in the  $x$ -axis.

Next, we construct a safe set such that the mobile robot must remain within a safe rectangular region defined by  $x_{\min} \leq x \leq x_{\max}$  and  $y_{\min} \leq y \leq y_{\max}$  while avoiding two circular obstacles centered at  $(x_{o,i}, y_{o,i}) \in \mathbb{R}^2$  with radius  $R_i$ ,  $i \in \{1, 2\}$ . As such, the desired unified safe set is the intersection of the 0-superlevel sets of the following functions:

$$\begin{aligned} h_1(x, y) &= \|[x, y]^\top - [x_{o,1}, y_{o,1}]^\top\| - R_1, \quad i \in \{1, 2\}, \\ h_3(x) &= x - x_{\min}, \quad h_4(x) = x_{\max} - x, \\ h_5(y) &= y - y_{\min}, \quad h_6(y) = y_{\max} - y. \end{aligned}$$

We generate  $h_0$ , depicted in Fig. 2, by numerically solving (11). Because the safe set is static, this solution can be precomputed offline on a uniform discrete mesh. During real-time operation, we evaluate  $h_0$  using bilinear interpolation,

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

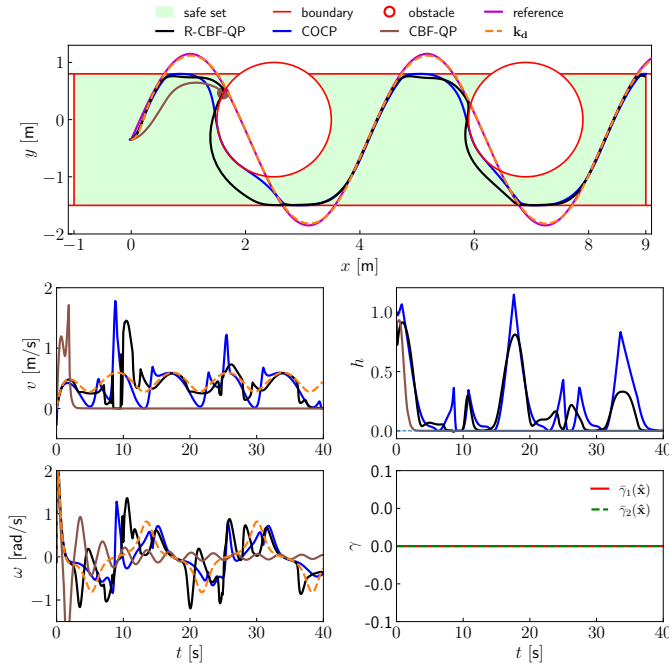


Fig. 3. Nominal (no state uncertainty): COCP, vanilla CBF-QP, and our R-CBF-QP with online adaptation. **Top**: Robot trajectories with the safe set and the reference. The brown dot ( $\bullet$ ) marks a vanilla CBF-QP deadlock: the robot reaches a point where no admissible input both preserves safety and moves toward the reference. The baseline controller  $\mathbf{k}_d$  follows the sinusoidal reference with negligible steady-state error. **Bottom**: Time histories of  $v$ ,  $\omega$ ,  $\gamma_1$ ,  $\gamma_2$ , and  $h$ . Our method maintains  $h \geq 0$  and closely tracks the reference, similar to COCP. The tuned parameters are obtained as  $\bar{\gamma}_1(\hat{\mathbf{x}}) \approx 0$ ,  $\bar{\gamma}_2(\hat{\mathbf{x}}) \approx 0$  with the proposed online tuning method, as  $\mathbf{e} \equiv 0$ .

and we compute its spatial derivatives with various numerical gradient methods.

Considering the dynamics in (12), the Poisson safety function  $h_0$  has relative degree one with respect to the input  $v$  and relative degree two with respect to the input  $\omega$ . This dual relative degree property reduces the utility of employing  $h_0$  directly as a CBF, since the resultant safety filter will only use  $v$  to enforce forward invariance of the desired safe set, completely neglecting the robot's steering capabilities. In accordance with [20], a modified CBF is proposed as

$$h(t, \mathbf{x}) = h_0(x, y) - \frac{1}{\mu} V(t, \mathbf{x}), \quad (16)$$

$$V(t, \mathbf{x}) = 1 - \cos(\theta - \theta_s(t, x, y)), \quad (17)$$

where  $\mu \in \mathbb{R}_{>0}$  is a tuning parameter and  $V: \mathbb{R}_{\geq 0} \times \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$  is a Lyapunov function that characterizes the convergence of  $\theta$  to  $\theta_s(t, x, y)$  under control (13). The function  $\theta_s: \mathbb{R}_{\geq 0} \times \mathbb{R}^2 \rightarrow \mathbb{S}^1$  defines a heading which, when achieved, guarantees safety for the unicycle system. Systematically,  $\theta_s(t, x, y)$  can be chosen as the direction of the safe output vector for an equivalent integrator system. This modified barrier function (16) has been shown to be effective in eliminating deadlocks when avoiding obstacles in the configuration space [22, Example 6] and [20]. Further details on the construction of  $\theta_s$  can be found in [20, Section IV]. It is also worth noting that the  $\lambda$ -superlevel set of  $h$  will be contained within the  $\lambda$ -superlevel set of  $h_0$  for all  $t \geq 0$ , for any  $\lambda \in \mathbb{R}$ , thus preserving compactness of superlevel sets.

In an R-CBF-QP, the nominal controller  $\mathbf{k}_d(t, \hat{\mathbf{x}})$  provided to the QP is directly taken from the reference-tracking con-

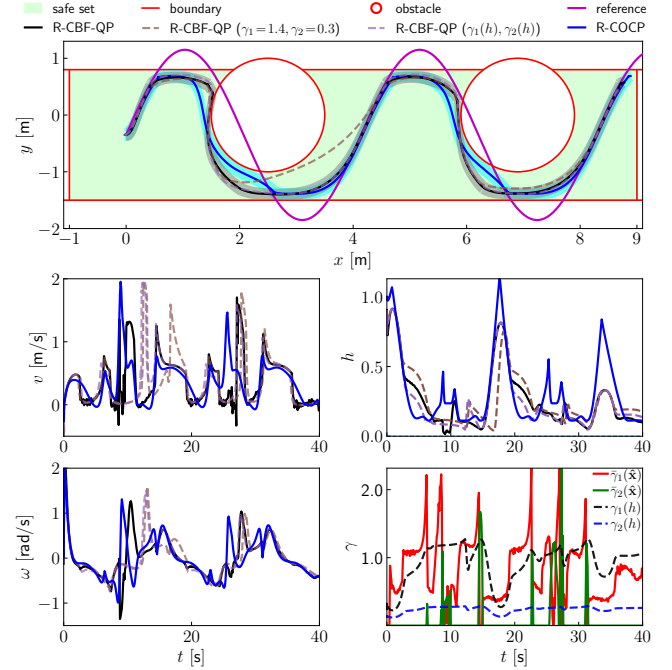


Fig. 4. Robust (bounded state uncertainty): R-COCP, R-CBF-QP with fixed ( $\gamma_1 = 1.4, \gamma_2 = 0.3$ ), and tunable R-CBF-QP with  $\gamma_1(h), \gamma_2(h)$  given in (19), and our R-CBF-QP with online adaptation. **Top**: Trajectories under the measurement uncertainty. Shaded tubes show the uncertainty envelope from bounded measurement error. Our R-CBF-QP and tunable R-CBF-QP track the reference more closely and are comparable to R-COCP near obstacles, while R-CBF-QP with fixed parameters is more conservative. **Bottom**: Time histories of  $v$ ,  $\omega$ ,  $\gamma_1$ ,  $\gamma_2$ , and  $h$ . All methods satisfy  $h \geq 0$ . Our online tuning method increases  $\bar{\gamma}_1, \bar{\gamma}_2$  only when needed, such as near the boundary.

troller,  $\mathbf{k}_d$ . However, near the safety boundary, this nominal input often deviated from the safe direction indicated by the single-integrator velocity  $\hat{v}_s$ . This misalignment can lead to jittery behavior when approaching the boundary. To address this issue, we introduce a practical modification for the baseline control law  $\mathbf{k}_d$ , given in (13), as

$$\bar{\mathbf{k}}_d(t, \hat{\mathbf{x}}) = \begin{bmatrix} |\hat{v}_s(t, \hat{\mathbf{x}})| \\ K_\omega (\theta_s(t, \hat{\mathbf{x}}) - \theta) \end{bmatrix}, \quad (18)$$

which aims to maintain the desired safe linear velocity while aligning the heading angle towards  $\theta_s$ . Practically, this modification projects the potentially unsafe reference trajectory  $[x_d(t) \ y_d(t)]^\top$  onto the safe set, so that the R-CBF-QP only needs to apply small corrections near the boundary, thereby mitigating the jittery behavior.

Finally, robust safety can be ensured by solving the following R-CBF-QP:

$$\begin{aligned} \mathbf{k}^*(t, \hat{\mathbf{x}}) &= \arg \min_{\mathbf{u} \in \mathcal{U}} \|\mathbf{u} - \bar{\mathbf{k}}_d(t, \hat{\mathbf{x}})\|^2 \\ \text{s.t. } &L_f h(t, \hat{\mathbf{x}}) + L_g h(t, \hat{\mathbf{x}}) \mathbf{u} + \partial_t h(t, \hat{\mathbf{x}}) + \alpha(h(t, \hat{\mathbf{x}})) \\ &\geq \bar{\gamma}_1(\hat{\mathbf{x}}) \|L_g h(t, \hat{\mathbf{x}})\|_2 + \bar{\gamma}_2^2(\hat{\mathbf{x}}) \|L_g h(t, \hat{\mathbf{x}})\|_2^2. \end{aligned}$$

## B. Simulations

In this subsection, we numerically demonstrate the effectiveness of our method by comparing it against several benchmarks in both nominal (no uncertainty) and robust (with state uncertainty) scenarios. To establish a baseline for comparison, we generate a safe trajectory by solving a robust

**IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.**

constrained optimal control problem (R-COCP)-based safety filter problem that enforces the safety and input constraints under bounded state estimation errors. For details on this R-COCP, see Appendix.

In the nominal case, the benchmarks are CBF-QP and COCP. For the robust case, we compare the proposed R-CBF-QP with online robustness parameter adaptation  $\bar{\gamma}_1(\hat{\mathbf{x}}), \bar{\gamma}_2(\hat{\mathbf{x}})$ , R-CBF-QP with fixed  $\gamma_1, \gamma_2$  parameters, R-CBF-QP with tunable (CBF-dependent)  $\gamma_1(h), \gamma_2(h)$  parameters, and R-COCP. We use two performance metrics  $J_{opt}$  and  $J_t$  (see Appendix) to compare alternative methods. <sup>‡</sup>

To simulate our safety-critical scenario with the real-world setup shown in Fig. 1, we first measure the boundaries of the safe rectangular region as  $x_{\min} = -1.5$  m,  $x_{\max} = 0.8$  m,  $y_{\min} = -2.0$  m,  $y_{\max} = 2.0$  m. And, two circular obstacles are centered at  $(x_{o,1}, y_{o,1}) = (2.5, 0.0)$  m and  $(x_{o,2}, y_{o,2}) = (6.9, 0.0)$  m with radius 0.6 m. We inflate the circular obstacle boundaries by half the robot's length, which is 0.4 m; therefore, we have  $R_1 = R_2 = 1$  m. Using these geometric parameters from the experimental setup, we define the individual CBFs  $h_i$  for  $i \in \{1, \dots, 6\}$ , construct the unified Poisson safety function  $h_0$  as explained in Section III-C, and obtain  $h_0$  shown in Fig. 2.

Next, we choose a sinusoidal reference trajectory tracking task via (15) and (14) such that it yields a safety constraint violation under the nominal controller  $\mathbf{k}_d$  in (13). We set the parameters as  $v_{\text{ref}} = 0.25$  m/s,  $a_d = 1.5$  m,  $c_d = 1.52$  rad/s, and  $\varphi_d = 0$ . We use  $\mathbf{k}_d$  and its safe projection  $\bar{\mathbf{k}}_d$  in (18) with the controller gains  $K_v = 1, K_\omega = 2.5$ . The initial condition is  $\mathbf{x}_0 = [0 \ -0.35 \ 0]^T$ .

For CBF-QP and R-CBF-QPs, the safety-critical control parameters are:  $\alpha = 3$ , and  $\mu = 3.3$ . We pick the fixed robustness parameters for comparison as  $\gamma_1 = 1.4, \gamma_2 = 0.3$ . To modulate the fixed robustness gains in the tunable R-CBF-QP, we utilize an exponential function following [4]:  $\varepsilon_i(h(t, \hat{\mathbf{x}})) = e^{(\eta_i \max\{h(t, \hat{\mathbf{x}}), 0\})}$ , where  $\eta_i > 0$ , and in our simulations, we set  $\eta_1 = 2, \eta_2 = 2$ ; therefore, we have

$$\gamma_1(h) = \frac{\gamma_1}{e^{(\eta_1 \max\{h(t, \hat{\mathbf{x}}), 0\})}}, \gamma_2^2(h) = \frac{\gamma_2^2}{e^{(\eta_2 \max\{h(t, \hat{\mathbf{x}}), 0\})}}. \quad (19)$$

In this study, we employ a grid search method to solve the online tuning problem (10), as proof-of-concept, though other gradient-free optimizers may also be utilized. We sample  $N = 100$  perturbed states, and perform a gradient-free grid search over  $\bar{\gamma}_1, \bar{\gamma}_2$  pairs arranged on a  $400 \times 400$  mesh covering the range  $[0.0001, 4]$  for each axis.

The control input ranges for the robot are  $v \in [-2, 2]$  m/s and  $\omega \in [-2, 2]$  rad/s. And, the control loop operates at 50 Hz. We then solved the QP-based safety filter problems with CVXPY using the OSQP solver. We solve the robust constrained optimal control problem numerically using CasADi [23] software tool using IPOPT solver, warm-starting from a forward rollout obtained with constant linear velocity  $v = v_{\text{ref}}$  and zero angular velocity.

The simulation results for the measurement model  $\hat{\mathbf{x}} \equiv \mathbf{x}$  without uncertainty (nominal) are presented in Fig. 3. Next, we consider a scenario with a bounded state measurement error set (robust) given by  $\mathcal{E} \triangleq [-0.05, 0.05] \times [-0.1, 0.1] \times \{0\}$ , i.e.,  $\mathbf{x} \in \hat{\mathbf{x}} \oplus \mathcal{E}$ , where  $\oplus$  denotes the Minkowski sum. The results for this scenario are shown in Fig. 4.

<sup>‡</sup> The code and additional details of our simulations can be found at <https://github.com/ersindas/robust-CBFs>

TABLE I  
PERFORMANCE METRICS FOR SIMULATION IN THE NOMINAL (NO STATE UNCERTAINTY) AND ROBUST (BOUNDED UNCERTAINTY) SCENARIOS.

Method	Nominal		Robust	
	$J_{opt}$	$J_t$ [s]	$J_{opt}$	$J_t$ [s]
COCP	0.0	16.1	–	–
CBF-QP	22896.4	36.7	–	–
R-COCP	–	–	0.0	21.0
<b>R-CBF-QP</b>	<b>237.1</b>	<b>10.7</b>	<b>198.9</b>	26.9
R-CBF-QP ( $\gamma_1 = 1.4, \gamma_2 = 0.3$ )	–	–	466.1	30.5
R-CBF-QP ( $\gamma_1(h), \gamma_2(h)$ )	–	–	312.5	<b>25.3</b>

TABLE II  
PERFORMANCE METRIC  $J_t$  FOR HARDWARE DEMONSTRATIONS.

Method	$J_t$ [s]
<b>R-CBF-QP</b>	<b>16.6</b>
R-CBF-QP ( $\gamma_{1,2} = 0$ )	26.0
R-CBF-QP ( $\gamma_1(h), \gamma_2(h)$ )	26.3
R-CBF-QP ( $\gamma_1 = 1.0, \gamma_2 = 0.2$ )	35.8
CBF-QP	34.2

**Nominal (no state uncertainty):** The vanilla CBF-QP deadlocks when the robot reaches the first obstacle because the barrier involves input  $v$  at relative degree one, but not  $\omega$ . The filter cannot exploit steering for safety, highlighting the need to address the dual relative degree issue. With no uncertainty, our R-CBF-QP drives the robustness parameters to their lower bounds, realizing  $\bar{\gamma}_1(\hat{\mathbf{x}}) \equiv 0.0001$  and  $\bar{\gamma}_2(\hat{\mathbf{x}}) \equiv 0.0001$ . By incorporating the Lyapunov term in the modified CBF  $h$  (16), our method explicitly integrates vehicle angular velocity into the CBF condition and resolves the deadlock. Moreover, the trajectory of the robot under our R-CBF-QP closely matches COCP while maintaining  $h \geq 0$ , see Fig. 3.

**Robust (bounded state uncertainty):** In the presence of bounded measurement uncertainty, all controllers maintain robust safety, but differ in performance (see Fig. 4). The R-CBF-QP with fixed parameters is noticeably conservative, slowing earlier and deviating farther from the reference, as it uses the same robustness parameters everywhere, while the tunable R-CBF-QP and our online-adapted R-CBF-QP track more closely and are comparable to R-COCP near obstacles. Our online adaptive R-CBF-QP increases  $\bar{\gamma}_1$  only near a boundary to add robustness against measurement uncertainty, while minimizing its value elsewhere. The value of  $\bar{\gamma}_2$  is mostly zero, increases only when safe set inflation is necessary, as explained in Theorem 1.

In simulation, our method achieves the smallest optimal cost,  $J_{opt}$  with  $Q = I, R = I$ , among the CBF-based controllers while also maintaining a competitive deviation period  $J_t$ . In contrast, the tunable R-CBF-QP method yields the lowest deviation period  $J_t$ , but at the expense of a higher optimal cost  $J_{opt}$ . The R-COCP method provides a lower bound for this comparison. For further details, please refer to Table I.

### C. Experimental Results and Comparisons

We experimentally investigated the effectiveness of the proposed safe controller on a tracked mobile robot. A photo of the robot on an elevated concrete platform during one experiment that involved two static obstacles is shown in

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

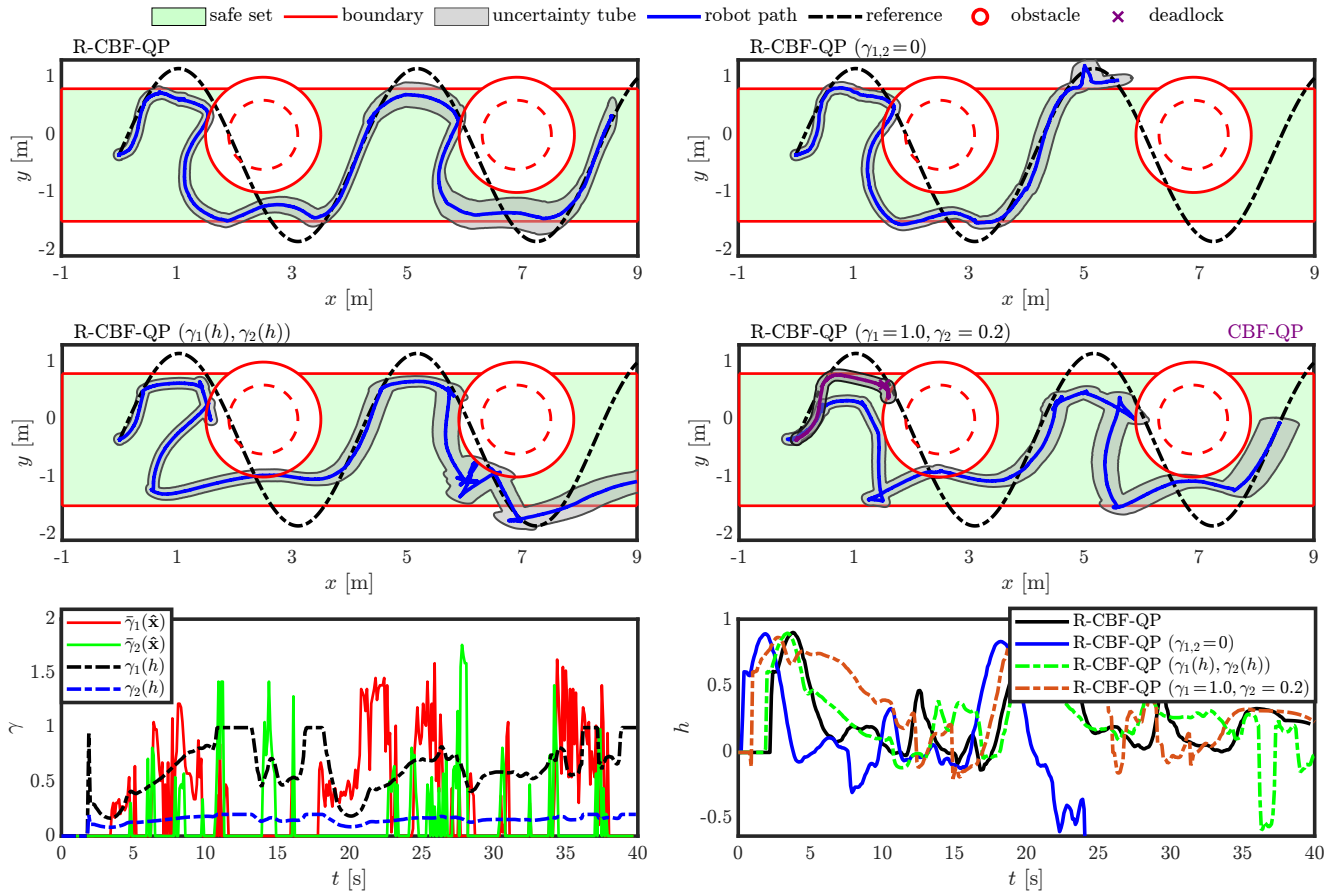


Fig. 5. Experimental comparison on tracked robot. Pose is measured at the robot’s geometric center; therefore, each circular obstacle boundary is inflated by half the robot’s length. Dashed red circles denote the true safe-set boundaries, while solid red circles indicate the inflated boundaries used in the controller synthesis. Gray tubes visualize the 95% confidence envelope of the  $(x, y)$  covariance along the executed paths. **(Left)** (Top) R-CBF-QP with online adaptation of parameters  $\tilde{\gamma}_1, \tilde{\gamma}_2$ , the proposed method. **(Left)** (Middle) R-CBF-QP with tunable robustness parameters, i.e., CBF-dependent  $\gamma_1(h), \gamma_2(h)$ . Tunable R-CBF-QP violated the safety constraint in the presence of time-varying state uncertainty, while its performance was similar to R-CBF-QP with online adaptation in simulation. **(Left)** (Bottom) Robustness parameters for the adaptive and tunable R-CBF-QPs. **(Right)** (Middle) R-CBF-QP without robustness parameters, i.e.,  $\gamma_1=0, \gamma_2=0$ . This non-robust controller violates constraints, and the robot falls from the platform and tips over at 24 sec. The experiment was terminated after this unsafe behavior. This result illustrates the necessity of robustification in practice. The CBF-QP-controlled robot encounters a deadlock at the boundary of the first obstacle, marked by the purple cross ( $\times$ ), since the CBF-QP only uses  $v$ . And, R-CBF-QP with fixed robustness parameters, i.e.,  $\gamma_1=1.0, \gamma_2=0.2$ , this controller is more conservative. **(Right)** (Bottom) The value of CBF  $h$  versus time.

Fig. 1. In this safety-critical navigation scenario, the robot may fall if it crosses the blue-taped rectangle. <sup>§</sup>

Our algorithms operate on a tracked GVR-Bot modified from the iRobot PackBot 510 platform. An NVIDIA Jetson AGX Orin runs our Python and C++ code under the ROS1 framework. The linear and angular body velocity commands are sent from the AGX Orin to an onboard Intel Atom, which converts these commands into individual track speeds, which are regulated through high-rate controllers.

The robot’s perception system uses an Intel RealSense depth camera operating at 30 Hz. A VectorNav VN-100 provides inertial measurements. Our state estimation employs OpenVINS [24], a visual-inertial odometry (VIO) that fuses IMU signals and sparse visual features via an extended Kalman filter. OpenVINS uses first-estimate Jacobians to maintain observability consistency, propagating and updating the full state covariance. The covariance estimates published by OpenVINS provide us with a runtime uncertainty bound. We use a 95% confidence interval.

We evaluated the performance of four controllers in the

scenario shown in Fig. 5, in which the safe set, reference path, and obstacle geometry are the same as in simulation. The state uncertainty bound, provided by VIO, is now time-varying. The hardware study used the same hyperparameters as in Section IV-B except that we employed a coarser  $80 \times 80$  grid for the tunable parameters and set  $\alpha = 1$ . On our platform, solving the R-CBF-QP with this grid takes approximately 13 ms per control step, within the 20 ms budget of the 50 Hz control loop. Our online-adapted R-CBF-QP with  $\tilde{\gamma}_1, \tilde{\gamma}_2$ , consistently preserved safety while staying closest to the reference path. The fixed-gain R-CBF-QP exhibited the expected conservatism (earlier braking and a larger lateral tracking error), while the tunable R-CBF-QP became unsafe in hardware under time-varying state estimation uncertainty. Note that the tunable R-CBF-QP uses the CBF-dependent gains  $\gamma_1(h), \gamma_2(h)$  given in (19), which do not adapt to the time-varying  $\delta(\hat{\mathbf{x}}(t))$ ; this can leave insufficient robustness on hardware and lead to safety violations. The R-CBF-QP without robustness parameters, i.e.,  $\gamma_1=0, \gamma_2=0$  also violated safety and was terminated earlier, see Fig. 5. Quantitative outcomes are presented in Table II: our online adaptive R-CBF-QP achieves the lowest  $J_t$  among safe controllers.

<sup>§</sup>See video at: <https://youtu.be/TNu-Etm0SU>

## V. CONCLUSION

We presented a method to synthesize robust controllers with formal safety guarantees for nonlinear systems that experience state estimation or measurement uncertainties. We tested its performance in a safe trajectory tracking task for a tracked mobile robot. To handle multiple obstacles, we incorporated a unified numerical CBF derived from Poisson's equation. We further addressed the dual relative degree issue that typically causes deadlocks in path tracking. Experimental trials demonstrated the overall performance improvement of our approach over existing formulations.

## REFERENCES

- [1] L. Lindemann, M. Cleaveland, G. Shim, and G. J. Pappas, "Safe planning in dynamic environments using conformal prediction," *IEEE Robot. Autom. Lett.*, vol. 8, no. 8, pp. 5116–5123, 2023.
- [2] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs for safety critical systems," *IEEE Trans. Automatic Control*, vol. 62, no. 8, pp. 3861–3876, 2017.
- [3] E. Daş and J. W. Burdick, "Robust control barrier functions using uncertainty estimation with application to mobile robots," *IEEE Trans. Automatic Control*, vol. 70, no. 7, pp. 4766–4773, 2025.
- [4] A. Alan, A. J. Taylor, C. R. He, G. Orosz, and A. D. Ames, "Safe controller synthesis with tunable input-to-state safe control barrier functions," *IEEE Control Systems Letters*, vol. 6, pp. 908–913, 2021.
- [5] B. T. Lopez, J.-J. E. Slotine, and J. P. How, "Robust adaptive control barrier functions: An adaptive and data-driven approach to safety," *IEEE Control Syst. Lett.*, vol. 5, no. 3, pp. 1031–1036, 2020.
- [6] W. D. Compton, M. H. Cohen, and A. D. Ames, "Learning for layered safety-critical control with predictive control barrier functions," in *Proc. Learn. Dyn. Control*, 2025, pp. 153–165.
- [7] T. Kim, R. I. Kee, and D. Panagou, "Learning to refine input constrained control barrier functions via uncertainty-aware online parameter adaptation," in *IEEE Int. Conf. Robotics and Autom.*, 2025.
- [8] S. Dean, A. J. Taylor, R. K. Cosner, B. Recht, and A. D. Ames, "Guaranteeing safety of learned perception modules via measurement-robust control barrier functions," in *Proc. Conf. Robot Learn.*, 2020, pp. 654–670.
- [9] R. K. Cosner, A. W. Singletary, A. J. Taylor, T. G. Molnar, K. L. Bouman, and A. D. Ames, "Measurement-robust control barrier functions: Certainty in safety with uncertainty in state," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2021, pp. 6286–6291.
- [10] Y. Wang and X. Xu, "Observer-based control barrier functions for safety critical systems," in *Proc. Amer. Control Conf.* IEEE, 2022, pp. 709–714.
- [11] D. R. Agrawal and D. Panagou, "Safe and robust observer-controller synthesis using control barrier functions," *IEEE Control Syst. Lett.*, vol. 7, pp. 127–132, 2022.
- [12] S. Yang, G. J. Pappas, R. Mangharam, and L. Lindemann, "Safe perception-based control under stochastic sensor uncertainty using conformal prediction," in *Proc. IEEE Conf. Decis. Control*, 2023, pp. 6072–6078.
- [13] R. Nanayakkara, A. D. Ames, and P. Tabuada, "Safety under state uncertainty: Robustifying control barrier functions," *arXiv preprint arXiv:2508.17226*, 2025.
- [14] I. Jang and H. J. Kim, "Safe control for navigation in cluttered space using multiple lyapunov-based control barrier functions," *IEEE Robot. Autom. Lett.*, vol. 9, no. 3, pp. 2056–2063, 2024.
- [15] K. Long, C. Qian, J. Cortés, and N. Atanasov, "Learning barrier functions with memory for robust safe navigation," *IEEE Robot. Autom. Lett.*, vol. 6, no. 3, pp. 4931–4938, 2021.
- [16] B. Dai, R. Khorrambakht, P. Krishnamurthy, V. Gonçalves, A. Tzes, and F. Khorrami, "Safe navigation and obstacle avoidance using differentiable optimization based control barrier functions," *IEEE Robot. Autom. Lett.*, vol. 8, no. 9, pp. 5376–5383, 2023.
- [17] L. Lindemann and D. V. Dimarogonas, "Control barrier functions for signal temporal logic tasks," *IEEE Control Syst. Lett.*, vol. 3, no. 1, pp. 96–101, 2018.
- [18] G. Bahati, R. M. Bena, and A. D. Ames, "Dynamic safety in complex environments: Synthesizing safety filters with Poisson's equation," in *Proc. Robotics: Sci. Syst.*, 2025.
- [19] R. M. Bena, G. Bahati, B. Werner, R. K. Cosner, L. Yang, and A. D. Ames, "Geometry-aware predictive safety filters on humanoids: From Poisson safety functions to CBF constrained MPC," in *IEEE-RAS 24th International Conference on Humanoid Robots (Humanoids)*, 2025.
- [20] G. Bahati, R. K. Cosner, M. H. Cohen, R. M. Bena, and A. D. Ames, "Control barrier function synthesis for nonlinear systems with dual relative degree," *arXiv preprint arXiv:2504.00397*, 2025.

- [21] P. Glotfelter, J. Cortés, and M. Egerstedt, "Nonsmooth barrier functions with applications to multi-robot systems," *IEEE Control Syst. Lett.*, vol. 1, no. 2, pp. 310–315, 2017.
- [22] M. H. Cohen, T. G. Molnar, and A. D. Ames, "Safety-critical control for autonomous systems: Control barrier functions via reduced-order models," *Annu. Rev. Control*, vol. 57, p. 100947, 2024.
- [23] J. A. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl, "CasADi: a software framework for nonlinear optimization and optimal control," *Math. Programming Comp.*, vol. 11, pp. 1–36, 2019.
- [24] P. Geneva, K. Eickenhoff, W. Lee, Y. Yang, and G. Huang, "OpenVINS: A research platform for visual-inertial estimation," in *IEEE Int. Conf. Robotics and Automation*, 2020, pp. 4666–4672.
- [25] D. Q. Mayne, M. M. Seron, and S. V. Raković, "Robust model predictive control of constrained linear systems with bounded disturbances," *Automatica*, vol. 41, no. 2, pp. 219–224, 2005.

## APPENDIX

To consider the impact of state errors, we pre-compute tightened safety constraints based on a tube model predictive control (MPC) concept [25]. Specifically, we enlarge each circular obstacle by the maximum expected position error, while the edges of the rectangular boundary are shifted inward by the maximum anticipated error in the corresponding direction. The resulting set is referred to as the tightened set, that is, the 0-superlevel set of  $\bar{h}_i(\mathbf{x})$ , which is defined as the tightened version of our CBFs  $h_i(\mathbf{x})$  for  $i \in \{1, \dots, 6\}$ . Therefore, when  $\bar{h}_i(\mathbf{x}) \geq 0$ , the robot remains within the safe set, even in the presence of estimation errors. Then, we define a robust constrained optimal control problem (R-COCP):

$$\begin{aligned} \mathbf{x}^*(t), \mathbf{u}^*(t) = \arg \min_{\mathbf{x}(\cdot), \mathbf{u}(\cdot)} \int_0^\infty (\|\mathbf{u} - \mathbf{k}_d(t, \mathbf{x})\|^2) dt \\ \text{s.t. } \dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{u}, \mathbf{x}(0) = \hat{\mathbf{x}}(0), \\ \bar{h}_i(\mathbf{x}(t)) \geq 0, \forall t \geq 0, i \in \{1, \dots, 6\}, \\ \mathbf{u}(t) \in \mathcal{U}, \end{aligned}$$

where  $\mathbf{u}^*(t)$  and  $\mathbf{x}^*(t)$  are the respective optimal input and state signals.

In order to quantify the optimality of trajectories with respect to the constrained optimal control problem, we use a standard cost metric:

$$J_{opt} \triangleq \int_0^\infty (\mathbf{x} - \mathbf{x}^*)^\top Q (\mathbf{x} - \mathbf{x}^*) + (\mathbf{u} - \mathbf{u}^*)^\top R (\mathbf{u} - \mathbf{u}^*) dt,$$

where the first term penalizes deviations of the states from the optimal trajectory, while the second term penalizes deviations of the control input from the optimal control input. In this context,  $(\mathbf{x}^*, \mathbf{u}^*)$  is obtained from the solution to R-COCP that incorporates tightened safety constraints.

In addition to the standard cost metric, we introduce another metric that indicates the time optimality for tracking tasks, defined as

$$J_t \triangleq \int_0^\infty \mathbb{1}_{\hat{\mathcal{X}}} (\|\hat{\mathbf{x}}(t) - \hat{\mathbf{x}}_{\mathbf{k}_d}(t)\| \geq \epsilon_x) dt,$$

where  $\epsilon_x \geq 0$ ,  $\hat{\mathbf{x}}_{\mathbf{k}_d}$  is the state under the reference-tracking controller  $\mathbf{k}_d$ , and  $\mathbb{1}_{\hat{\mathcal{X}}}$  is the indicator function such that  $\mathbb{1}_{\hat{\mathcal{X}}} (\|\hat{\mathbf{x}}(t) - \hat{\mathbf{x}}_{\mathbf{k}_d}(t)\| \geq \epsilon_x) = 1$  if  $\|\hat{\mathbf{x}}(t) - \hat{\mathbf{x}}_{\mathbf{k}_d}(t)\| \geq \epsilon_x$ , and 0 otherwise. Therefore,  $J_t$  represents the time during which the error between the resulting states from the safe control and the states derived from  $\mathbf{k}_d$  exceeds  $\epsilon_x$ , and quantifies the practical tracking performance. For our safe control problem, this metric considers only the positional error, replacing the full state vector  $\hat{\mathbf{x}}(t)$  with the robot's position  $[x(t) \ y(t)]^\top$ , and the reference states  $\hat{\mathbf{x}}_{\mathbf{k}_d}$  with the desired position  $[x_d(t) \ y_d(t)]^\top$ .