TU Delft

Delft University of Technology

Document Version
Final published version

Licence
CC BY

Citation (APA)
Liu, C., Alan, A., Shi, S., & De Schutter, B. (2027). Robust adaptive discrete-time control barrier certificate. *Automatica*, *195*, Article 113310. https://doi.org/10.1016/j.automatica.2026.113310

Important note
To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright
In case the licence states “Dutch Copyright Act (Article 25fa)”, this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse
Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy
Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

This work is downloaded from Delft University of Technology.

Automatica 195 (2027) 113310

Contents lists available at ScienceDirect

Automatica

journal homepage: www.elsevier.com/locate/automatica

Brief paper

Robust adaptive discrete-time control barrier certificate ☆

Changrui Liu *, Anil Alan, Shengling Shi *, Bart De Schutter

Delft Center for Systems and Control, TU Delft, Delft, The Netherlands

ARTICLE INFO

Article history:
Received 9 August 2025
Received in revised form 3 August 2026
Accepted 2 September 2026

Keywords:
Robust adaptive control
Data-driven safety
Discrete-time control barrier function
Online learning

ABSTRACT

This work develops a robust adaptive control strategy for discrete-time systems using Control Barrier Functions (CBFs) to ensure safety under parametric model uncertainty and disturbances. A key contribution of this work is establishing a barrier function certificate in discrete time for general online parameter estimation algorithms. This barrier function certificate guarantees positive invariance of the safe set despite disturbances and parametric uncertainty without access to the true system parameters. In addition, real-time implementation and inherent robustness guarantees are provided. The proposed robust adaptive safe control framework demonstrates that the parameter estimation module can be designed separately from the CBF-based safety filter, simplifying the development of safe adaptive controllers for discrete-time systems. The resulting safe control approach guarantees that the system remains within the safe set while the controller adapts to model uncertainties, making it a promising strategy for discrete-time safety-critical systems.

© 2026 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

1. Introduction

Safety-critical control aims to guarantee that the system state remains within a safe set, i.e., constraint satisfaction, which is crucial in domains where unsafe behaviors can lead to undesirable outcomes (Ames et al., 2016). Among various approaches (e.g., Hamilton–Jacobi reachability analysis (Semnani et al., 2020) and Model Predictive Control (MPC) (Liu and Shi, &, 2026; Rawlings et al., 2017)), Control Barrier Functions (CBFs) have become a powerful and systematic tool for both continuous-time (Ames et al., 2016) and discrete-time systems (Agrawal & Sreenath, 2017). CBFs allow formal encoding of safety as control invariant sets, and they can be integrated on top of traditional control methods (e.g., control Lyapunov functions (Ames et al., 2016)) and learning-based control approaches (e.g., model-based Reinforcement Learning (RL) (Cohen & Belta, 2023)) that may not inherently offer safety guarantees. CBFs have many applications, including control of bipedal robots (Agrawal & Sreenath, 2017), cruise control (Ames et al., 2016), and microgrid energy management (Liu and Shi and Alan et al., 2026), demonstrating their versatility and practical relevance in safety-critical domains.

Despite the recent surge in learning-based approaches for CBFs (Srinivasan et al., 2020), model-based methods remain essential due to their analytical tractability and formal safety guarantees (Ames et al., 2016; Wang & Wen, 2025). Typically, these approaches support modular design through the explicit use of dynamical models (Agrawal & Sreenath, 2017; Ames et al., 2016; Xu et al., 2015). A standard approach is to encode safety constraints through a CBF-based inequality condition on the input, and then to incorporate this condition within a safety filter (Didier & Zeilinger, 2026) that minimizes the adjustment of a given nominal input that may otherwise lead to unsafe behaviors (i.e., constraint violations). However, a fundamental challenge in model-based methods is model uncertainty, which can compromise safety (Wang & Wen, 2025). Accordingly, various robust and adaptive extensions have been developed for CBFs to handle disturbances (Cosner et al., 2024; Jankovic, 2018) or parametric uncertainty (Lopez et al., 2020; Shen et al., 2025; Taylor & Ames, 2020). For disturbances, high-probability safety (Cosner et al., 2024) and guaranteed safety (Jankovic, 2018) have been established for unbounded and bounded disturbances, respectively. There are also results on adaptive safety focusing on parametric uncertainty, employing either Lyapunov-based methods (Lopez et al., 2020; Taylor & Ames, 2020), or batched least-squares methods (Shen et al., 2025). For an in-depth overview, we refer the reader to Wang and Wen (2025).

Although recent progress has been made on robust and adaptive CBFs in continuous time (Lopez et al., 2020; Shen et al., 2025; Taylor & Ames, 2020), the discrete-time case remains largely unexplored despite its relevance to discrete-time control designs (e.g., approximate MPC (Hertneck et al., 2018; Liu and

☆ This paper is part of a project that has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant agreement No. 101018826 - CLariNet). The material in this paper was not presented at any conference. This paper was recommended for publication in revised form by Associate Editor Frederic Mazenc under the direction of Editor Luca Zaccarian.

* Corresponding authors.
E-mail addresses: c.liu-14@tudelft.nl (C. Liu), a.alan@tudelft.nl (A. Alan), s.shi-3@tudelft.nl (S. Shi), b.deschutter@tudelft.nl (B. De Schutter).

https://doi.org/10.1016/j.automatica.2026.113310
0005-1098/© 2026 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

Shi and Alan et al., 2026) and model-based RL (Cohen & Belta, 2023)). Since practical systems rely on sampled sensing and actuation, there is a need to develop adaptive Discrete-Time CBFs (DT-CBFs) (Cosner et al., 2024; Santoyo et al., 2019). The main difficulty for DT-CBFs is that derivative-based arguments are unavailable. In addition, *causal* controllers must compute inputs without future information, a feature overlooked in continuous-time methods (Lopez et al., 2020; Shen et al., 2025). Moreover, Lyapunov adaptation laws lack guarantees on bounded updates or convergence (Lopez et al., 2020), motivating estimation-based approaches that better align with sampled-data settings (Cohen & Belta, 2023). Since standard estimation tools (e.g., least-squares methods) operate on discrete measurements (Ljung, 2010), DT-CBFs provide a more suitable foundation to establish adaptive safety.

This paper extends adaptive safe control to discrete-time systems, particularly those implemented in embedded or sampled-data settings. The proposed method connects to classical adaptive control, where parameter estimation, projection, and set-membership identification are used to ensure robustness (Åström & Wittenmark, 1994; Goodwin & Sin, 2014; Ioannou & Sun, 1996; Milanese et al., 2013). In contrast, this paper reverses the standard perspective: rather than augmenting adaptive laws with safety mechanisms, robust safety is prioritized as the primary objective, and adaptation is subsequently integrated via parameter estimation to reduce conservatism. The contributions of this work are summarized as follows:

- Given a DT-CBF satisfying appropriate conditions, we establish safety guarantees for input-affine discrete-time systems under bounded disturbances and parametric mismatch, thereby embedding robust adaptive DT-CBFs into a practically relevant discrete-time control framework.
- A modular control architecture is proposed in which parameter estimation is decoupled from CBF-based safety filtering, in contrast to existing approaches (Lopez et al., 2020; Taylor & Ames, 2020) where the adaptation law is co-designed with the safety filter. The proposed framework is compatible with any nominal controller and any parameter estimator with a bounded update rate, thereby significantly broadening applicability.
- The *inherent robustness*[1] of the proposed safe control approach is rigorously analyzed, extending results for continuous-time systems without model uncertainty (Xu et al., 2015) to discrete-time uncertain systems.

It is further noted that the primary contribution of this paper lies in establishing a general robust adaptive safe control framework with a given DT-CBF, rather than designing DT-CBFs, developing novel estimation algorithms, or refining the safety filter framework (Didier & Zeilinger, 2026). The rest of the paper is organized as follows. Section 2 provides preliminaries, and Section 3 introduces robust adaptive DT-CBFs. A simple estimation algorithm is given in Section 4. Inherent robustness guarantees are discussed in Section 5. A simulation example is presented in Section 6, and Section 7 concludes the paper and outlines future research directions.

2. Preliminaries

This section establishes the foundation for the main results. We first introduce the system setting, then present robust DT-CBFs, showing how they define safe sets and enforce state safety through input feasibility via positive set invariance. In essence, CBFs constrain admissible controls to keep the state within the safe set or drive it back when it deviates. Finally, we introduce a class of online parameter estimators that provide (i) point estimates, (ii) set estimates, and (iii) uniform bounds on point-estimate increments.

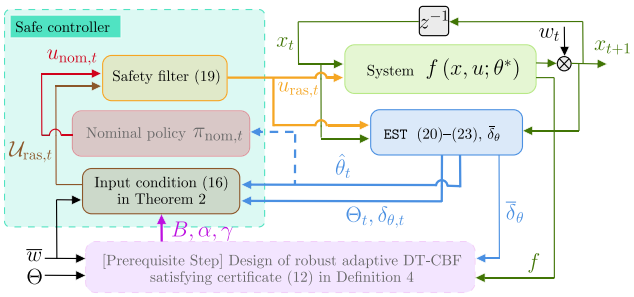


Fig. 1. Block diagram of the proposed robust adaptive safe control framework. Thick arrows (orange and blue) denote information exchange between the safe controller and the system/estimator (EST). The nominal policy and DT-CBF design blocks are rendered grayed out to indicate that they lie outside the scope of the current paper. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

2.1. Notation

Let $\mathbb{R}$ and $\mathbb{N}$ denote the sets of real and natural numbers, respectively, with $\overline{\mathbb{N}} := \mathbb{N} \cup \{+\infty\}$ and $\mathbb{I}_{[a:b]} := \overline{\mathbb{N}} \cap [a, b]$. The matrix or vector p-norm is $\|\cdot\|_p$, where $\|\cdot\|$ defaults to the 2-norm. For a vector x, $x[i]$ is its ith element. The zero vector, one vector, and identity matrix of size n are denoted, respectively, by $\mathbf{0}_n$, $\mathbf{1}_n$, and $\mathbf{I}_n$. A continuous function $\alpha : \mathbb{R} \to \mathbb{R}$ is an extended Class $\mathcal{K}$ function (i.e., $\alpha \in \mathcal{K}_e$) if $\alpha(0) = 0$ and it is strictly increasing. The class of strictly sub-identity functions $\mathcal{C}_{\text{SI}}$ contains $\beta : \mathbb{R} \to \mathbb{R}$ such that $\beta(r) < r$ for all $r > 0$. For a closed set $\mathcal{X} \subseteq \mathbb{R}^n$, define $\Pi_{\mathcal{X}}(\bar{x}) := \arg\min_{x \in \mathcal{X}} \|x - \bar{x}\|$ and $\text{dist}_{\mathcal{X}}(\bar{x}) := \min_{x \in \mathcal{X}} \|x - \bar{x}\|$, respectively. Finally, $\lambda_{\min}(\Gamma)$ denotes the minimum eigenvalue of a symmetric positive-definite matrix Γ.

2.2. System description

Consider a class of input-affine discrete-time systems:

$$\begin{aligned} x_{t+1} &= f(x_t, u_t; \theta) + w_t \\ &= f_{\text{d}}(x_t) - [\boldsymbol{\phi}(x_t)]^\top \theta + g(x_t) u_t + w_t, \end{aligned} \tag{1}$$

where $x_t \in \mathcal{X} \subset \mathbb{R}^n$ and $u_t \in \mathcal{U} \subseteq \mathbb{R}^m$ are the state and input at time step t. The set $\mathcal{X}$ is compact, and $\mathcal{U}$ is closed to ensure the existence of solutions for the safety filter optimization problem (Ames et al., 2016), Lopez et al. (2020). These settings are practically valid for most engineering systems, where physical constraints or design specifications naturally confine the state and input variables to a bounded operational regime (Agrawal & Sreenath, 2017; Ames et al., 2016; Lopez et al., 2020). The model f_{d} captures known dynamics (Cosner et al., 2024) or discretization effects (Mania et al., 2022), while g characterizes the input coupling. The term $[\boldsymbol{\phi}(x_t)]^\top \theta$ represents the uncertain drift, which models unknown damping or state-coupled disturbances (Luo et al., 2022; Taylor & Ames, 2020) through a known continuous kernel $\boldsymbol{\phi}$ and an unknown constant parameter $\theta \in \Theta$. Accurate values for f_{d}, $\boldsymbol{\phi}$, and g can be obtained via full state measurement. In addition, the disturbance w_t lies in a known polytope[2] given by

$$\mathcal{W} = \{w \in \mathbb{R}^n \mid H_w w \leq \mathbf{1}_{n_w}\}, \tag{2}$$

[1] In this paper, inherent robustness refers to steering the state towards safe sets when the current state is unsafe (Xu et al., 2015, Section 2.2), which should not be confused with robustness against model uncertainty (Didier & Zeilinger, 2026; Lopez et al., 2020).

[2] A polytope is a *bounded* polyhedron (i.e., an intersection of a finite set of closed half-spaces) (Borrelli et al., 2017, Chapter 2.1).

where $H_w \in \mathbb{R}^{n_w \times n}$ and $n_w \geq n$ is the number of linear inequalities. It is clear that $\mathbf{0}_n$ lies in the interior of $\mathcal{W}$, and we further define $\overline{w} := \max_{w \in \mathcal{W}} \|w\|$. The true parameter $\theta^* \in \Theta$ is a *constant* and the prior known parameter set Θ is a polytope given by

$$\Theta = \{\theta \in \mathbb{R}^q \mid H_\theta \theta \leq h_\theta\}, \tag{3}$$

where $H_\theta \in \mathbb{R}^{n_\theta \times q}$ with $n_\theta \geq q$, and $h_\theta \in \mathbb{R}^{n_\theta}$. Both the disturbance and the model parameter are considered to lie in polytopes, which is motivated by practical considerations, as many disturbances and parameters are naturally bounded and can be modeled by box-constrained sets (Ackermann et al., 1993; Ames et al., 2016; Lopez et al., 2020). Theoretically, polytopes can also inner-approximate arbitrary compact sets, making this a flexible and effective modeling choice. Given an estimate $\hat{\theta} \in \Theta$, its associated estimation error is defined as

$$e_\theta := \hat{\theta} - \theta^*. \tag{4}$$

The estimation error e_θ cannot be computed exactly since θ^* is unknown. However, for a specific $\hat{\theta}$, an upper bound on its associated estimation error can be computed by solving the following optimization problem:

$$\varepsilon_\theta(p) := \max_{\theta \in \Theta} \|\theta - \hat{\theta}\|_p, \tag{5}$$

and it holds that[3] $\|e_\theta\| \leq \varepsilon_\theta(2) \leq \varepsilon_\theta(1)$. Following (1), the *true* system under disturbances is given by

$$x_{t+1} = f(x_t, u_t; \theta^*) + w_t. \tag{6}$$

The system state evolves according to (6); however, the controller does not have access to the true parameter θ^*.

Remark 1. The controller design for (1) does not require smooth kernels; with non-smooth kernels, the model can represent PieceWise-Affine (PWA) dynamics (Mania et al., 2022). The input-affine structure and linear parameter dependence are versatile, standard modeling choices (Köhler et al., 2021; Nijmeijer & van der Schaft, 1990) that do not trivialize the safe control problem. This framework is widely used in application examples such as inverted pendulums (Luo et al., 2022) and vehicle cruise control (Taylor & Ames, 2020).

2.3. Robust safe control

This subsection introduces robust DT-CBFs, which constitute a natural and modest extension of continuous-time CBFs to discrete-time systems with bounded disturbances, building on prior work on continuous-time systems with bounded disturbances (Jankovic, 2018) and discrete-time systems without disturbances (Agrawal & Sreenath, 2017).

Definition 1 (*Robust Safety*)**.** A set $\mathcal{X}_s \subseteq \mathcal{X}$ is positive invariant for (6) if $x_0 \in \mathcal{X}_s$ implies that $\forall \{w_t\}_{t=0}^{\infty} \in \mathcal{W}^\infty$, $x_t \in \mathcal{X}_s$ for all $t \in \mathbb{N}$. In this case, the system (6) is robustly safe w.r.t. $\mathcal{X}_s$, and $\mathcal{X}_s$ is called a robust safe set.

Definition 1 addresses only robustness against disturbances. It applies to the true parameter θ^* but need not hold for all parameters $\theta \in \Theta$. Characterizing safe sets using CBFs is useful for safe controller design (Ames et al., 2016). For discrete-time systems, DT-CBFs are powerful tools to achieve this objective (Agrawal & Sreenath, 2017; Cosner et al., 2024). Specifically, for a continuous barrier function $B : \mathcal{X} \times \Theta \to \mathbb{R}$, its induced safe set is defined using the 0-superlevel set of $B(\cdot, \theta) : \mathcal{X} \to \mathbb{R}$ for all $\theta \in \Theta$ as

$$\mathcal{S}_{B,\theta} := \{x \in \mathcal{X} \mid B(x, \theta) \geq 0\}, \tag{7}$$

and its boundary $\partial \mathcal{S}_{B,\theta}$ and interior $\text{int}(\mathcal{S}_{B,\theta})$ are defined, respectively, by $\partial \mathcal{S}_{B,\theta} := \{x \in \mathcal{X} \mid B(x, \theta) = 0\}$ and $\text{int}(\mathcal{S}_{B,\theta}) := \{x \in \mathcal{X} \mid B(x, \theta) > 0\}$.

Remark 2. This work adopts barrier functions that depend on the model parameter θ as in Lopez et al. (2020) and Taylor and Ames (2020) since most barrier functions are model-based (Ames et al., 2016). Nonetheless, CBFs that are derived mainly from the safety requirements without using θ are also commonly used in safe control applications, e.g., adaptive cruise control (Taylor & Ames, 2020) and aircraft pitch control (Lopez et al., 2020).

Definition 2 (*Lipschitz Continuity*)**.** The function $B : \mathcal{X} \times \Theta \to \mathbb{R}$ is locally Lipschitz continuous on the joint space $\mathcal{X} \times \Theta$ if there exist constants $L_{B,x}, L_{B,\theta} \geq 0$ such that $\forall x', x'' \in \mathcal{X}$, $\theta', \theta'' \in \Theta$,

$$|B(x', \theta') - B(x'', \theta'')| \leq L_{B,x}\|x' - x''\| + L_{B,\theta}\|\theta' - \theta''\|.$$

Definition 3 (*Robust DT-CBF*)**.** The function $B_{\theta^*} := B(\cdot, \theta^*)$ is a robust DT-CBF for (6) on $\mathcal{S}_{B,\theta^*}$ if it is Lipschitz continuous and there exists $\alpha \in \mathcal{K}_e \cap \mathcal{C}_{SI}$ such that $\forall x \in \mathcal{X}$, $\exists u \in \mathcal{U}$ satisfying

$$B_{\theta^*}(f(x, u; \theta^*)) - B_{\theta^*}(x) - L_{B,x}\overline{w} \geq -\alpha\left(B_{\theta^*}(x)\right), \tag{8}$$

and (8) is called the robust control barrier certificate against disturbances.

In discrete time, Lipschitz continuity is adopted to quantify the perturbation of B. The barrier certificate (8) also defines the set of robust safe control inputs for $\mathcal{S}_{B,\theta^*}$ as

$$\mathcal{U}_{\text{rs},\theta^*}(x) = \{u \in \mathcal{U} \mid (8) \text{ holds}\},$$

where the subscript θ^* emphasizes that characterizing $\mathcal{U}_{\text{rs},\theta^*}(x)$ requires knowing the true parameter. The following theorem is a standard result on CBF-based safety guarantees, and the proof is omitted for brevity since it follows the same procedure as in Agrawal and Sreenath (2017).

Theorem 1 (*Robust Safety*)**.** *If the continuous function $B_{\theta^*} : \mathcal{X} \to \mathbb{R}$ is a robust DT-CBF for (6) on $\mathcal{S}_{B,\theta^*}$, then $\mathcal{U}_{\text{rs},\theta^*}(x)$ is non-empty for all $x \in \mathcal{X}$. In addition, if a state-feedback control policy $\pi : \mathcal{X} \to \mathbb{R}^m$ satisfies $\pi(x) \in \mathcal{U}_{\text{rs},\theta^*}(x)$, $\forall x \in \mathcal{X}$, then the closed-loop system $x_{t+1} = f(x_t, \pi(x_t); \theta^*) + w_t$ is robustly safe w.r.t. $\mathcal{S}_{B,\theta^*}$ against disturbances w_t.*

Since knowing θ^* is often unrealistic, the certificate (8) cannot be directly used for safe control when θ^* is unknown. A common remedy is to robustify against worst-case model mismatch by considering a constant nominal $\hat{\theta}$ and adding terms that scale with the estimation error, yielding a worst-case safety guarantee similar to Theorem 1. However, such designs lead to conservative inputs and may have worse performance. Because parametric uncertainty is *epistemic* and can be reduced using data (Wang & Wen, 2025), incorporating adaptive safety mechanisms with online parameter estimation is more desirable.

2.4. Online parameter estimation

Given Θ in (3), an estimator EST provides a point estimate $\hat{\theta}_t$ and a polytopic set $\Theta_t \subseteq \Theta$ at each time step t such that $\theta^* \in \Theta_t$. These sets are initialized as $\Theta_0 = \Theta$ and recursively refined, for instance, via set-membership identification. Leveraging (4) and (5), the instantaneous estimation error and its associated bound are defined as

[3] Given $x \in \mathbb{R}^n$, $\|x\|_2 \leq \|x\|_1$ (Horn & Johnson, 2012).

$$e_{\theta,t} := \hat{\theta}_t - \theta^*, \tag{9a}$$

$$\varepsilon_{\theta,t}(p) := \max_{\theta\in\Theta_t} \|\theta - \hat{\theta}_t\|_p. \tag{9b}$$

The module EST performs recursive updates according to the following abstracted law:

$$\hat{\theta}_{t+1} = \mathcal{T}_{\mathtt{ADP},t}(\hat{\theta}_t), \tag{10}$$

where the time-varying operator $\mathcal{T}_{\mathtt{ADP},t}$ depends on online closed-loop data and the set estimate Θ_t. The estimate increment is defined as

$$\delta_{\theta,t} := \hat{\theta}_{t+1} - \hat{\theta}_t = \mathcal{T}_{\mathtt{ADP},t}(\hat{\theta}_t) - \hat{\theta}_t. \tag{11}$$

Assumption 1. There exists a uniform upper bound $\bar{\delta}_\theta$ on the estimate increments $\delta_{\theta,t}$, i.e., $\|\delta_{\theta,t}\| \le \bar{\delta}_\theta$, $\forall t \in \mathbb{N}$.

Since Θ is compact, Assumption 1 holds trivially since $\|\delta_{\theta,t}\| \le \max_{\theta',\theta''\in\Theta} \|\theta' - \theta''\|$, i.e., the diameter of Θ.

3. Robust adaptive safe control

This section introduces robust adaptive DT-CBFs, providing a sufficient condition for their certification. Building on this, we derive an online input condition to compute safe controls within a safety filter framework. Finally, we discuss some practical implementation issues.

3.1. Problem formulation

Given a candidate function $B : \mathcal{X} \times \Theta \to \mathbb{R}$ defining the safe set in (7), we first provide the conditions required for B to qualify as a *robust adaptive DT-CBF*.

Definition 4 (*Robust Adaptive DT-CBF*)**.** Given an estimator EST satisfying Assumption 1, a barrier function $B : \mathcal{X} \times \Theta \to \mathbb{R}$ is a robust adaptive DT-CBF for (6) and EST if it is Lipschitz continuous and there exists $\alpha \in \mathcal{K}_\infty \cap \mathcal{C}_{\mathrm{SI}}$ such that, for all $x \in \mathcal{X}$, there exists $u \in \mathcal{U}$ satisfying

$$\begin{aligned}&\Delta B^\star(x,u) - L_{B,x}\overline{w} - E^\star_\theta(x)\\ &\quad\ge -\alpha\left(B^\star(x) - \frac{(\bar{\varepsilon}_\theta(p))^2}{2\gamma}\right),\end{aligned} \tag{12}$$

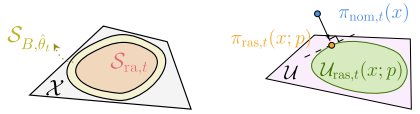
where $\gamma > 0$, $B^\star(x) := \min_{\theta\in\Theta} B(x,\theta)$, $\Delta B^\star(x,u) := \min_{\theta\in\Theta} \{B(f(x,u;\theta),\theta) - B(x,\theta)\}$, $\bar{\varepsilon}_\theta(p)$ is the worst-case error p-norm ($p \in \{1,2\}$) defined as

$$\bar{\varepsilon}_\theta(p) := \max_{\theta,\theta'\in\Theta} \|\theta - \theta'\|_p, \tag{13}$$

and $E^\star_\theta(x) = \left(L_{B,x}\|\boldsymbol{\phi}(x)\| + \frac{\bar{\delta}_\theta}{\gamma}\right)\bar{\varepsilon}_\theta(p) + \frac{\bar{\delta}_\theta^2}{2\gamma} + L_{B,\theta}\bar{\delta}_\theta$, with $L_{B,x}$ and $L_{B,\theta}$ given in Definition 2 and $\bar{\delta}_\theta$ characterized by EST as in Assumption 1. In this case, the condition (12) is called the robust adaptive Control Barrier Certificate (CBC) against disturbances and parametric model mismatch.

Assumption 2. For the model (1) with sets $\mathcal{W}$ and Θ defined in (2) and (3), respectively, and an estimator EST satisfying Assumption 1, there exists a known robust adaptive DT-CBF $B : \mathcal{X} \times \Theta \to \mathbb{R}$ available for controller synthesis.

The condition (12) provides a discrete-time certificate without using θ^*, extending the framework in Lopez et al. (2020). Assumption 2 is a standard feasibility condition (Agrawal & Sreenath, 2017; Lopez et al., 2020), ensuring that a safety certificate exists and that the underlying optimization problem is well-posed. In practice, such functions are typically obtained by parameterizing B and solving for its coefficients via numerical optimization (Agrawal & Sreenath, 2017; Ames et al., 2016). Assuming a valid robust adaptive DT-CBF is available, the objective is (i) to design a less conservative condition to certify safe inputs with parameter estimation and (ii) to use this condition to modify a potentially unsafe nominal policy $\pi_{\mathrm{nom},t} : \mathcal{X} \to \mathbb{R}^m$ such that the constraints encoded by the DT-CBF are satisfied.



(a) Robustification (b) Safety filter

Fig. 2. Sketch of (a) adaptively robustified safe sets, and (b) the safety filter projection when the nominal input is unsafe.

Remark 3. Although the Lipschitz-continuity requirement restricts the search space, it is standard in the literature because commonly used CBFs are continuously differentiable (Lopez et al., 2020; Taylor & Ames, 2020), implying the required Lipschitz continuity. In practice, this assumption is not restrictive; in fact, Lipschitz-continuous DT-CBFs are widely used in several discrete-time safe-control applications (Agrawal & Sreenath, 2017; Cosner et al., 2024; Luo et al., 2022).

Remark 4. The considered nominal controller may either ignore safety entirely during design (e.g., basic PID controllers) or incorporate safety constraints without formal guarantees. Examples of the latter include model-based RL with safety-based rewards or approximate MPC that does not inherently preserve the safety properties of the original MPC controller (Hertneck et al., 2018).

3.2. Robust adaptive safe control

To address parametric uncertainty, we define the *robust adaptive* counterpart of a given B with estimate $\hat{\theta}_t$ by

$$B_{\mathrm{ra},t}(x) := B(x,\hat{\theta}_t) - \frac{1}{2}e_{\theta,t}^\top \Gamma^{-1} e_{\theta,t}, \tag{14}$$

where Γ is a symmetric positive definite matrix such that $\lambda_{\min}(\Gamma) = \gamma$ with γ given in (12). Accordingly, the *robust adaptive* safe set $\mathcal{S}_{\mathrm{ra},t}$ is defined using $B_{\mathrm{ra},t}(\cdot)$ as

$$\mathcal{S}_{\mathrm{ra},t} := \{x \in \mathcal{X} \mid B_{\mathrm{ra},t}(x) \ge 0\}. \tag{15}$$

It is noted that $\mathcal{S}_{\mathrm{ra},t} \subseteq \mathcal{S}_{B,\hat{\theta}_t}$ (see Fig. 2(a)), and $\mathcal{S}_{\mathrm{ra},t} = \mathcal{S}_{B,\theta^*}$ if and only if $\hat{\theta}_t = \theta^*$.

Definition 5 (*Timed Robust Safety*)**.** Given $T \in \overline{\mathbb{N}}$, $\{\mathcal{X}_{\mathrm{s},t}\}_{t=0}^T$ satisfying $x_0 \in \mathcal{X}_{\mathrm{s},0}$ and $\forall t \in \mathbb{I}_{[0:T]}$, $\mathcal{X}_{\mathrm{s},t} \subseteq \mathcal{X}$ is sequentially positively invariant for (6) if $\forall t \in \mathbb{I}_{[0:T-1]}$, $x_t \in \mathcal{X}_{\mathrm{s},t}$ implies that $x_{t+1} \in \mathcal{X}_{\mathrm{s},t+1}$ for all $w_t \in \mathcal{W}$. In this case, the system (6) is said to be robustly safe w.r.t. $\{\mathcal{X}_{\mathrm{s},t}\}_{t=0}^T$. In addition, if $\mathcal{X}_{\mathrm{s},t} \subseteq \mathcal{X}_{\mathrm{s}}^*$, $\forall t \in \mathbb{I}_{[0:T]}$, the system (6) is said to be robustly safe w.r.t. $\mathcal{X}_{\mathrm{s}}^*$ over $\mathbb{I}_{[0:T]}$.

Finite-horizon safety has been widely studied in CBF-based control (Cosner et al., 2024). The timed safety notion in Definition 5 is both more practical and more general than its infinite-horizon counterpart, as safety is typically required only over a finite operational period (Santoyo et al., 2019); setting $T = +\infty$ recovers infinite-horizon safety.

Theorem 2 (*Robust Adaptive Safety*). *Given $x_0 \in \mathcal{S}_{\mathrm{ra},0}$, a horizon $T \in \overline{\mathbb{N}}$, an estimator* EST *satisfying Assumption 1, and a robust adaptive DT-CBF $B : \mathcal{X} \times \Theta \to \mathbb{R}$ with $\gamma > 0$ and function $\alpha \in \mathcal{K}_\infty \cap \mathcal{C}_{\mathrm{SI}}$, for the robust adaptive online input condition*

$$B\big(f(x,u;\hat{\theta}_t),\hat{\theta}_t\big) - B(x,\hat{\theta}_t) - L_{B,x}\overline{w} - E_{\theta,t}(x) \geq -\alpha\left(B(x,\hat{\theta}_t) - \frac{(\varepsilon_{\theta,t}(p))^2}{2\gamma}\right), \tag{16}$$

the admissible input set $\mathcal{U}_{\mathrm{ras},t}(x;p)$ defined as

$$\mathcal{U}_{\mathrm{ras},t}(x;p) := \{u \in \mathcal{U} \mid (16) \text{ holds}\} \tag{17}$$

is non-empty, $\forall x \in \mathcal{X}$ and $t \in \mathbb{I}_{[0,T-1]}$. The term $E_{\theta,t}(x) = \left(L_{B,x}\|\boldsymbol{\phi}(x)\| + \frac{\|\delta_{\theta,t}\|}{\gamma}\right)\varepsilon_{\theta,t}(p) + L_{B,\theta}\|\delta_{\theta,t}\| + \frac{\|\delta_{\theta,t}\|^2}{2\gamma}$ with $\delta_{\theta,t}$ defined in (11). In addition, if a time-varying control policy $\pi_t : \mathcal{X} \to \mathbb{R}^m$ satisfies $\pi_t(x) \in \mathcal{U}_{\mathrm{ras},t}(x;p)$ for all $x \in \mathcal{X}$ and $t \in \mathbb{I}_{[0,T-1]}$, then the time-varying closed-loop system

$$x_{t+1} = f(x_t, \pi_t(x_t); \theta^*) + w_t \tag{18}$$

is robustly safe w.r.t. $\{\mathcal{S}_{\mathrm{ra},t}\}_{t=0}^{T}$ and thus robustly safe w.r.t. $\{\mathcal{S}_{B,\hat{\theta}_t}\}_{t=0}^{T}$ (cf. Definition 5) against both parametric model mismatch and disturbances.

The proof of Theorem 2 is provided in Appendix A. Notably, condition (12) guarantees that the admissible set $\mathcal{U}_{\mathrm{ras},t}(x;p)$ is non-empty, while the online safe control condition (16) depends solely on estimated parameters. To ensure safety with minimal intervention, we employ a safety filter that projects the nominal input onto the admissible set (see Fig. 2(b)). Specifically, for a nominal policy $\pi_{\mathrm{nom},t} : \mathcal{X} \to \mathbb{R}^m$, the robust adaptive safe policy $\pi_{\mathrm{ras},t} : \mathcal{X} \to \mathcal{U}$ is implicitly defined by the following constrained optimization:

$$\pi_{\mathrm{ras},t}(x;p) = \underset{\nu \in \mathcal{U}_{\mathrm{ras},t}(x;p)}{\arg\min} \ \|\nu - \pi_{\mathrm{nom},t}(x)\|. \tag{19}$$

This policy is time-varying as $\mathcal{U}_{\mathrm{ras},t}$ evolves to capture historical state and input data via the estimator EST.

Remark 5. The safe initial state condition ($x_0 \in \mathcal{S}_{\mathrm{ra},0}$), verified via the error bound $\varepsilon_{\theta,0}(p)$, does not restrict the argument of $\alpha(\cdot)$ in (12) and (16) to positive values. For instance, condition (12) remains valid whenever a feasible input $u \in \mathcal{U}$ yields a sufficiently large $\Delta B^\star(x,u)$, regardless of the sign of the argument of $\alpha(\cdot)$.

Remark 6. The online input condition (16) is less conservative than the CBC (12), which is reserved for offline verification. In contrast to Lopez et al. (2020), the robustified function (14) employs the true error $e_{\theta,t}$ instead of a set-induced bound on $e_{\theta,t}$. Notably, the invariance property for robustified CBFs using bounds on $e_{\theta,t}$ is valid in discrete time but not in continuous time; see Appendix B.

3.3. Practical implementation

Deriving a closed-form policy $\pi_{\mathrm{ras},t}(\cdot;p)$ is generally intractable. In practice, the safe control input $\pi_{\mathrm{ras},t}(x_t;p)$ is computed by solving (19) for the current state x_t. This requires characterizing $\mathcal{U}_{\mathrm{ras},t}(x_t;p)$ via condition (16), which depends on $\varepsilon_{\theta,t}(p)$. While computing $\varepsilon_{\theta,t}(2)$ typically involves a non-convex program, using $\varepsilon_{\theta,t}(1)$ reduces the calculation to an efficient linear program. Since $\varepsilon_{\theta,t}(2) \leq \varepsilon_{\theta,t}(1)$, the 1-norm-based condition is more conservative but computationally advantageous.

While (19) is generally non-convex, the input-affine dynamics ensure f is linear in u. If the barrier candidate B is concave in x, then $\mathcal{U}_{\mathrm{ras},t}(x;p)$ is a convex set, and computing $\pi_{\mathrm{ras},t}(x_t;p)$ becomes a convex optimization problem with formal convergence and optimality guarantees (Boyd & Vandenberghe, 2004). Although state-concavity is standard in DT-CBF research (Agrawal & Sreenath, 2017; Zeng et al., 2021), designing such functions is application-specific and remains outside the scope of this work.

Remark 7. Employing $\varepsilon_{\theta,t}(1)$ in (16) introduces conservatism, which increases the suboptimality of $u_{\mathrm{ras},t}(1)$ compared to $u_{\mathrm{ras},t}(2)$. This is particularly relevant when the nominal control policy is derived from optimal control frameworks such as discrete-time CLFs (Bof et al., 2018), model-based RL (Cohen & Belta, 2023), or approximate MPC (Hertneck et al., 2018; Liu and Shi and Alan et al., 2026). An alternative approach involves uniformly adopting the ∞-norm for all error bounds, including Eqs. (5), (9b), and (13), as well as the bounds in Definition 2 and Assumption 1.

4. Online estimation and causal synthesis

Following the parameter estimator in Lorenzen et al. (2019), we instantiate the estimator EST using a clipped recursive-least-squares (RLS) update with set-membership projection (Milanese et al., 2013). We further present a *causal* integration of the safety filter and estimator, ensuring that the computation of $u_{\mathrm{ras},t}(p)$ relies solely on information available up to time step t.

For $t \geq 1$, both $\hat{\theta}_t$ and Θ_t are determined by the information flow $\mathcal{F}_{t-1} := (x_0, u_0, x_1, u_1, \ldots, x_{t-1})$, and a causal synthesis means computing u_t based on $\mathcal{F}_t$. The initial estimate satisfies $\hat{\theta}_0 \in \Theta_0$ with $\Theta_0 = \Theta$. Given x_t, the residual is defined by $r_t := x_t - f_{\mathrm{d}}(x_{t-1}) - g(x_{t-1})u_{t-1}$, and it holds that $r_t + [\boldsymbol{\phi}(x_{t-1})]^\top\theta^* \in \mathcal{W}$, i.e., $H_w(r_t + [\boldsymbol{\phi}(x_{t-1})]^\top\theta^*) \leq \mathbf{1}_{n_w}$ by (2). Equivalently, $\theta^* \in \Delta_{\theta,t}$ with the non-falsified polytopic set $\Delta_{\theta,t}$ given by

$$\Delta_{\theta,t} = \left\{\theta \in \mathbb{R}^q \mid H_w[\boldsymbol{\phi}(x_{t-1})]^\top\theta \leq \mathbf{1}_{n_w} - H_w r_t\right\}. \tag{20}$$

Given the current set estimate $\Theta_t = \{\theta \in \mathbb{R}^q \mid H_{\theta,t}\theta \leq h_{\theta,t}\}$, the updated set estimate is computed as

$$\Theta_{t+1} = \Theta_t \cap \Delta_{\theta,t} \tag{21}$$

which is a polytope described by an augmented set of hyperplanes[4] inherited from Θ_t and $\Delta_{\theta,t}$. For the point estimate, a *prior* is first computed based on the clipped RLS parameter update as

$$\hat{\theta}_{t+1}^{[\mathrm{p}]} = \hat{\theta}_t - \mu_t \underbrace{\boldsymbol{\phi}(x_{t-1})\left(x_t - f(x_{t-1}, u_{t-1}; \hat{\theta}_t)\right)}_{:=\psi_{\theta,t}}, \tag{22}$$

where $\mu_t = \min\left\{(\|\psi_{\theta,t}\| + \epsilon)^{-1}, \|\boldsymbol{\phi}(x_{t-1})\|^{-2}\right\}$ is the clipped learning rate with $\epsilon > 0$ being a small constant. Note that (22) recovers a gradient descent update (Lorenzen et al., 2019), while the proposed clipped RLS update aligns with gradient clipping techniques (Kingma & Ba, 2014). The final estimate $\hat{\theta}_{t+1}$ is obtained via a simple point-to-set projection:

$$\hat{\theta}_{t+1} = \Pi_{\Theta_{t+1}}(\hat{\theta}_{t+1}^{[\mathrm{p}]}). \tag{23}$$

Due to $\|\mu_t\psi_{\theta,t}\| \leq 1$ and the triangle inequality, $\|\delta_{\theta,t}\| \leq \min\{1, \max_{\theta',\theta'' \in \Theta}\|\theta' - \theta''\|\}$, consistent with Assumption 1. As u_t depends on $\mathcal{F}_t$, causality is ensured by aligning $\hat{\theta}_t$ and Θ_t with the delayed information flow $\mathcal{F}_{t-1}$.

Algorithm 1 details the core procedure. At each step t, $\hat{\theta}_t$ and $\delta_{\theta,t}$ are used to define $\mathcal{U}_{\mathrm{ras},t}(x_t;p)$. A safe input u_t is then computed via the safety filter, and the system state evolves to x_{t+1}. A summary of the proposed safe control approach is provided in Fig. 1. Note that the nominal controller may also use the estimates $\hat{\theta}_t$ via certainty-equivalence model-based approaches (Liu et al., 2024, 2025).

[4] Unbounded growth in the number of hyperplanes can be avoided by overapproximating the intersection with a hyperrectangle (or hyperbox) of fixed dimension (Lorenzen et al., 2019).

Remark 8. The adaptation bound in Assumption 1 is reminiscent of the averaging methods in classical adaptive control (Åström & Wittenmark, 1994). However, its role in the proposed method is not to perform averaging, but to directly limit the adaptation rate, which is necessary to prevent excessive growth of B and thus to maintain the input condition (16).

Remark 9. Convergence of $\hat{\theta}_t$ to θ^* typically requires persistent excitation (Ljung, 2010), often achieved through active exploration (Mania et al., 2022). In contrast, the proposed approach ensures safety by robustifying against the uncertainty set Θ_t containing θ^*. While convergence minimizes conservatism as $E_{\theta,t}(x)$ in (16) approaches zero, our safety guarantees do not rely on achieving it. This work aims to ensure safety for a broad class of estimators; designing specific estimators or ensuring persistent excitation is beyond the scope of the current paper.

Algorithm 1 Robust Adaptive Safe Control

Require: B, α, x_0, $\hat{\theta}_0$, Θ, $\overline{w}$, $L_{B,x}$, $L_{B,\theta}$, Γ
  $\Theta_0 \leftarrow \Theta$, $\mu_0 \leftarrow 0$, $\Delta_{\theta,0} \leftarrow \mathbb{R}^q$
  for $t = 0$ to $T-1$ **do**
    Obtain nominal policy $\pi_{\text{nom},t}$
    Compute $\varepsilon_{\theta,t}(p)$ ▷ (9b)
    Compute $\hat{\theta}_{t+1}$ and $\delta_{\theta,t}$ ▷ (20)–(23), (11)
    Determine $\mathcal{U}_{\text{ras},t}(x;p)$ ▷ (16)–(17)
    Compute u_t using the safety filter ▷ (19)
    Obtain x_{t+1} by applying u_t
  end for

5. Inherent robustness

Following the analysis in Xu et al. (2015), this section examines the inherent robustness of (16), extending the analysis to discrete-time uncertain systems. While Section 3 focuses on safety guarantees, analyzing system behavior under unsafe initial conditions (i.e., $x_0 \notin \mathcal{S}_{B,\hat{\theta}_0}$) provides another important aspect of the proposed safe control approach. In this regime, the safety filter (19) provides a restorative action that steers trajectories toward $\mathcal{S}_{\text{ra},t}$. Rather than guaranteeing safety, this section establishes asymptotic convergence of x_t to $\mathcal{S}_{\text{ra},t}$.

Definition 6. Let $\{\mathcal{X}_t\}_{t=0}^{\infty}$ with $\mathcal{X}_t \subseteq \mathcal{X}^*$ be a sequentially positively invariant sequence (cf. Definition 5) for a general system $x_{t+1} = h_t(x_t) + w_t$, where $h_t : \mathcal{X} \to \mathbb{R}^n$ is continuous for all $t \in \overline{\mathbb{N}}$, and $w_t \in \mathcal{W}$. The state x_t is said to converge asymptotically to[5] $\mathcal{X}_{\text{sup}} := \limsup_{t\to\infty} \mathcal{X}_t$ on $\mathcal{X}^*$ if $\lim_{t\to\infty} \text{dist}_{\mathcal{X}_{\text{sup}}}(x_t) = 0$ for all $x_0 \in \mathcal{X}^*$.

Following (15), we define the corresponding *time-varying* continuous energy functions on $\mathcal{X}$ as follows:

$$V_t(x) = \begin{cases} 0 & \text{if } x \in \mathcal{S}_{\text{ra},t} \\ -B_{\text{ra},t}(x) & \text{if } x \in \mathcal{X} \setminus \mathcal{S}_{\text{ra},t} \end{cases}, \quad \forall t \in \overline{\mathbb{N}}. \tag{24}$$

It is obvious that $\forall t \in \overline{\mathbb{N}}$, $V_t(x) > 0$ for $x \in \mathcal{X} \setminus \mathcal{S}_{\text{ra},t}$. Given that $\mathcal{S}_{\text{ra},t}$ is sequentially positively invariant (see Appendix A), V_t can be used to show that x_t converges to the time-varying set $\mathcal{S}_{\text{ra},t}$ asymptotically if the input u_t satisfies (12), which is stated in the following theorem.

Theorem 3 (*Inherent Robustness*). *Given an estimator satisfying Assumption 1 and a robust adaptive DT-CBF $B : \mathcal{X} \times \Theta \to \mathbb{R}$ (cf. Definition 4), then x_t converges asymptotically to $\limsup_{t\to\infty} \mathcal{S}_{\text{ra},t}$ on $\mathcal{X}$ (cf. Definition 6) for the system (6) with $u_t \in \mathcal{U}_{\text{ras},t}(x_t;p)$, where $\mathcal{U}_{\text{ras},t}(x_t;p)$ is constructed as in (17) via (16).*

The proof of Theorem 3 is provided in Appendix C. This result implies that control inputs certified by (16) guarantee $x_t \in \mathcal{S}_{\text{ra},t}$ as $t \to \infty$. Although $\mathcal{S}_{\text{ra},t}$ does not generally converge to $\mathcal{S}$, asymptotic convergence to the time-varying set $\mathcal{S}_{\text{ra},t}$ still holds. This property also applies when using the maximized error-bound CBF design (Lopez et al., 2020) discussed in Remark 6. Furthermore, if $\alpha(\cdot)$ is linear or admits a linear lower bound (i.e., $\alpha(r) \geq kr$ for some $k > 0$), exponential convergence is achieved (Bof et al., 2018).

6. Numerical simulations

In this section, two simulation examples are provided. The first one involves vehicle cruise control, a standard benchmark in safe control using CBFs (Taylor & Ames, 2020; Xiao et al., 2021). The second one considers speed control of Permanent-Magnet Synchronous Motors (PMSMs). For both examples, the RLS-based estimator of Section 4 is adopted. The simulations[6] are implemented in Python 3.13.11.

6.1. Vehicle cruise control

Applying forward Euler discretization ($\Delta t = 0.1$ s) to the continuous-time vehicle model in Taylor and Ames (2020) yields

$$\begin{bmatrix} v_{t+1} \\ d_{t+1} \end{bmatrix} = \begin{bmatrix} v_t - \frac{\Delta t}{M} F_{\text{r}}(v_t) \\ d_t + \Delta t(v_{\text{f}} - v_t) \end{bmatrix} + \begin{bmatrix} \frac{\Delta t}{M} \\ 0 \end{bmatrix} u_t + \begin{bmatrix} w_{v,t}\Delta t \\ w_{d,t}\Delta t \end{bmatrix}.$$

where v, M, and d denote velocity, vehicle mass, and distance to the leading vehicle, respectively. The leading vehicle's nominal velocity $v_{\text{f}} \in [v_{\text{lb}}, v_{\text{ub}}]$ is unknown, and the input $u \in [-u_{\max}, u_{\max}]$ is the ego vehicle's traction force. The resistance force $F_{\text{r}}(v) = F_{\text{roll}} + \mu_{\text{vis}} v + \mu_{\text{aero}} v^2$ models drag effects for $v > 0$ (Taylor & Ames, 2020). While F_{roll} and μ_{vis} are determined offline, $\mu_{\text{aero}} \in [\mu^-, \mu^+]$ is identified online (Xiao et al., 2021). Disturbances $w_{v,t} \in [-\overline{w}_v, \overline{w}_v]$ and $w_{d,t} \in [-\overline{w}_d, \overline{w}_d]$ account for wind effects and variations in v_{f}, respectively. Nominal parameters and known bounds are summarized in Table 1. The model follows the general form (1) with box-constrained disturbance $w_t = [w_{v,t}\Delta t, w_{d,t}\Delta t]^\top$ and parameter $\theta = [\mu_{\text{aero}}, v_f]^\top$ satisfying (2)–(3). We enforce safety via the time-headway constraint $d \geq 1.8v$ (Taylor & Ames, 2020) using a constructed parameter-free DT-CBF $B(v,d) = d - 1.8v - 0.5$ with $\alpha(x) = (1 - 10^{-4})x$. Numerical verification confirms that (12) holds for $p = 2$ and $\Gamma = 10^5 \mathbf{I}_2$. For numerical illustration, a myopic (i.e., one-step prediction) adaptive MPC (aMPC) controller serves as the nominal controller, tracking $v_{\text{ref}} = 30$ m/s and utilizing the estimator from Section 4 to obtain estimates $\hat{\mu}_{\text{aero}}$ and $\hat{v}_f$. We compare adaptive MPC with a robust adaptive safety filter (**aMPC-raCBF**), MPC with a robust safety filter (**MPC-rCBF**), adaptive MPC (**aMPC**), and myopic robust adaptive MPC (**raMPC**) proposed in Köhler et al. (2021). Fig. 3 shows the performance of the estimator, demonstrating accurate parameter recovery over time. Safety results are summarized in Fig. 4, where the middle subplot shows that the proposed robust adaptive CBF-based safety filter maintains constraint satisfaction despite disturbances and model mismatch. This is achieved by actively modifying the control inputs as safety limits are approached (Fig. 4, bottom). Notably, the unconstrained

[5] Given a sequence of sets $\{\mathcal{X}_t\}_{t=0}^{\infty}$, $\limsup_{t\to\infty} \mathcal{X}_t := \bigcap_{\tau=0}^{\infty} \left(\bigcup_{t\geq\tau} \mathcal{X}_t\right)$ (Folland, 1999).

[6] The full codebase and specific implementations are available at: https://github.com/lcrekko/dt_ra_cbf.

Table 1
Parameters and bounds used in cruise control simulation.

Nominal parameters			Bounds		
Symbol	Value	Unit	Symbol	Value	Unit
M	1650	kg	$[v_{\text{lb}}, v_{\text{ub}}]$	[20, 32]	m/s
F_{roll}	125	N	$[\mu^-, \mu^+]$	[0.1, 0.6]	kg/m
μ_{vis}	1.2	kg/s	$\overline{w}_v$	2	m/s^2
μ^*_{aero}	0.55	kg/m	$\overline{w}_d$	5	m/s
v^*_{f}	22	m/s	$u_{\max}$	10 000	N

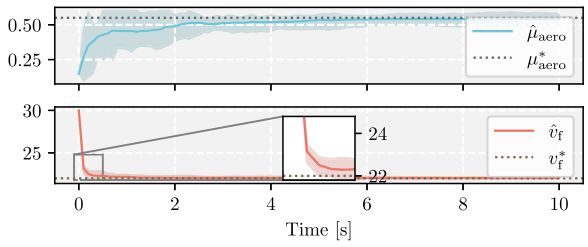


Fig. 3. Parameter estimation of vehicle example using unconstrained adaptive MPC with robust adaptive safety filter. The shaded area is the min–max envelope encompassing different disturbance realizations.

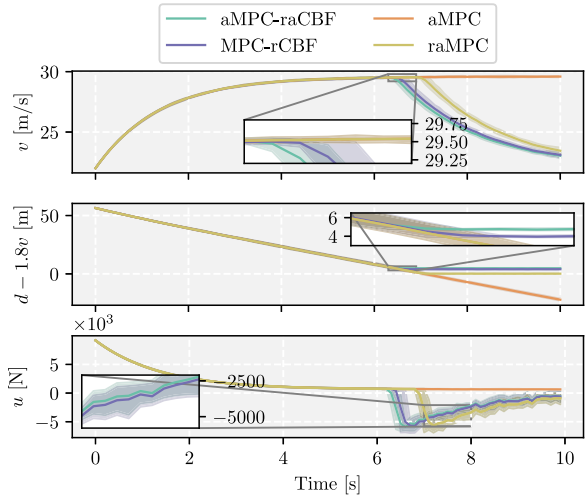


Fig. 4. Comparison of adaptive MPC with a robust adaptive safety filter (aMPC-raCBF), MPC with a robust safety filter (MPC-rCBF), adaptive MPC (aMPC), and robust adaptive MPC (raMPC) (Köhler et al., 2021). The shaded area denotes the min–max envelope across various disturbance realizations, modeled as uniformly distributed noise.

aMPC baseline fails to enforce safety; in this myopic case, it is equivalent to a control Lyapunov function (CLF) approach with an adaptive model. While constrained MPC is a valid alternative for safe control, our purpose is not to outperform constrained MPC but to demonstrate the safety guarantee of CBFs. Indeed, CBFs can be integrated into MPC (Zeng et al., 2021) or used to safeguard controllers that lack inherent guarantees, such as PID control, reinforcement learning (Cohen & Belta, 2023), or approximate MPC (Hertneck et al., 2018; Liu and Shi and Alan et al., 2026). Although robust adaptive MPC with constraint tightening (Köhler et al., 2021) also ensures safety, it reacts more slowly to the time-headway safety requirement in our simulation setting, necessitating more conservative braking beginning at around $t = 7$ s.

6.2. Speed control of PMSMs

Following the surface-mounted PMSM model in Li and Liu (2009) with the direct-axis (d-axis) current i_{d} set to zero, the

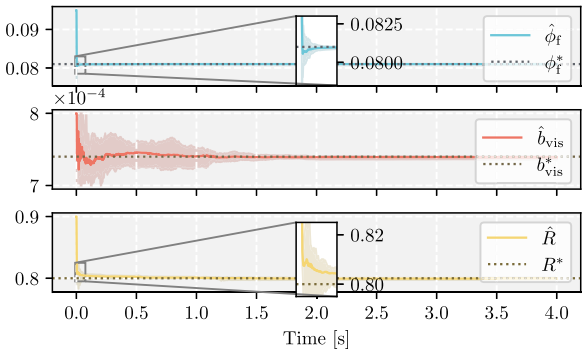


Fig. 5. Parameter estimation of PMSM example using adaptive PID control with robust adaptive safety filter. The shaded area is the min–max envelope across various disturbance realizations, modeled as uniformly distributed noise.

Table 2
Parameters and bounds used in motor control simulation.

Nominal parameters			Bounds		
Symbol	Value	Unit	Symbol	Value	Unit
n_{p}	4	$[-]$	$[\phi_{\text{f}}^-, \phi_{\text{f}}^+]$	[0.065, 0.095]	Wb
J	$2.35 \cdot 10^{-4}$	kg m^2	$[b_{\text{vis}}^-, b_{\text{vis}}^+]$	$[7,\ 8] \cdot 10^{-4}$	N m s
L	$2.9 \cdot 10^{-3}$	H	$[R^-, R^+]$	[0.6, 1]	Ω
ϕ^*_{f}	0.081	Wb	$\overline{v}_\omega$	100	s^{-2}
b^*_{vis}	$7.42 \cdot 10^{-4}$	N m s	$\overline{v}_i$	60	A/s
R^*	0.8	Ω	$I_{\max}$	2.8	A

discrete-time model, under forward Euler discretization with sampling period $\Delta t = 10^{-3}$[s], is given by

$$\begin{bmatrix}\omega_{t+1}\\ i_{\text{q},t+1}\end{bmatrix} = \begin{bmatrix}\left(1 - \frac{b_{\text{vis}}\Delta t}{J}\right) & \frac{n_{\text{p}}\phi_{\text{f}}\Delta t}{J}\\ -\frac{n_{\text{p}}\phi_{\text{f}}\Delta t}{L} & 1 - \frac{R\Delta t}{L}\end{bmatrix}\begin{bmatrix}\omega_t\\ i_{\text{q},t}\end{bmatrix} + \begin{bmatrix}0\\ \frac{\Delta t}{L}\end{bmatrix} u_{\text{q},t} + \begin{bmatrix}\nu_{\omega,t}\Delta t\\ \nu_{i,t}\Delta t\end{bmatrix},$$

where i_{q}, u_{q}, and ω denote the quadrature-axis (q-axis) current, q-axis voltage, and rotor angular velocity, respectively. While the number of pole pairs n_{p}, rotor inertia J, and inductance L are known, the flux linkage $\phi_{\text{f}} \in [\phi_{\text{f}}^-, \phi_{\text{f}}^+]$, viscous friction coefficient $b_{\text{vis}} \in [b_{\text{vis}}^-, b_{\text{vis}}^+]$, and resistance $R \in [R^-, R^+]$ are *unknown* but bounded. Disturbances $\nu_{\omega,t} \in [-\overline{\nu}_\omega, \overline{\nu}_\omega]$ and $\nu_{i,t} \in [-\overline{\nu}_i, \overline{\nu}_i]$ account for load variations and perturbations of i_{d}, respectively. This discretized PMSM model aligns with (1) under the box-constrained disturbance $\nu_t = [\nu_{\omega,t}\Delta t, \nu_{i,t}\Delta t]^\top$ and the parameter $\theta = [\phi_{\text{f}}, b_{\text{vis}}, R]^\top$. The objective is to track ω_{ref} while ensuring the safety constraint $i_{\text{q}} \in [-I_{\max}, I_{\max}]$. To this end, we construct two[7] CBFs $B^+(i_{\text{q}}) = I_{\max} - i_{\text{q}} - 0.05$ and $B^-(i_{\text{q}}) = i_{\text{q}} + I_{\max} - 0.05$ with $\alpha(x) = (1 - 10^{-4})x$. Numerical verification confirms that (12) holds for both CBFs with $p = 1$ and $\Gamma = 10^7\mathbf{I}_3$. The voltage saturation limit (i.e., bound on u_{q}) is 220 V. Nominal parameters and bounds are listed in Table 2. We use an adaptive PID (aPID) controller as the nominal controller and compare aPID control with a robust adaptive safety filter (**aPID-raCBF**), aPID with a robust safety filter (**aPID-rCBF**), PID control with a robust safety filter (**PID-rCBF**), and input-saturated **aPID** control.

Fig. 5 illustrates the high estimation accuracy achieved for all three unknown parameters. The comparative efficacy of the proposed framework is illustrated in Fig. 6. While the baseline **aPID** controller achieves precise tracking, it fails to maintain safety as the q-axis current violates the prescribed bounds. Conversely, the **aPID-rCBF** and **PID-rCBF** variants ensure safety but exhibit extreme conservatism; because the parameter error $\varepsilon_{\theta,t}(p)$ is

[7] For multiple CBFs, the set in (17) is defined by incorporating the input conditions corresponding to each CBF.

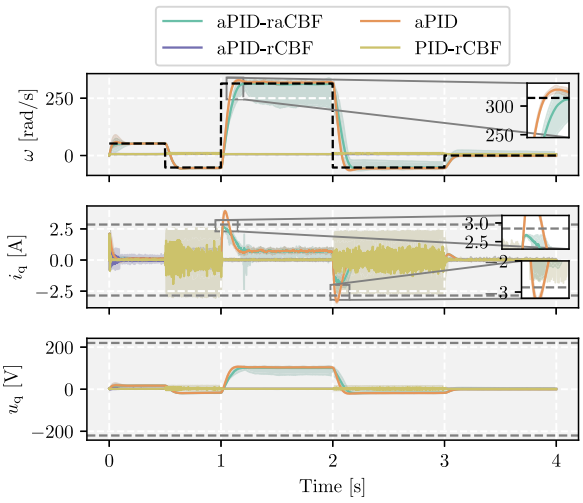


Fig. 6. Comparison of adaptive PID (aPID) control with a robust adaptive safety filter (aPID-raCBF), aPID with a robust safety filter (aPID-rCBF), PID control with a robust safety filter (PID-rCBF), and input-saturated aPID control. The shaded area denotes the min–max envelope across various disturbance realizations, modeled as uniformly distributed noise.

not reduced via estimation, the robustification term $\|\boldsymbol{\phi}(x)\|\varepsilon_{\theta,t}(p)$ severely restricts control effort, resulting in a near-total loss of tracking. In contrast, the proposed robust adaptive safety filter overcomes this issue. By dynamically updating uncertainty sets, it maintains rigorous safety certificates with only marginal tracking degradation, effectively achieving satisfactory tracking performance while maintaining robust safety.

7. Conclusions and future work

In this paper, we have proposed a robust adaptive safe control method for discrete-time systems under parametric uncertainty and disturbances. Unlike traditional continuous-time designs that couple adaptation and control, our framework introduces a modular architecture where the parameter estimator operates independently of the safety filter. Future work will focus on extending this framework to output-feedback systems with imperfect measurements and/or nonlinearly parameterized uncertainty. Additionally, we intend to investigate active exploration for enhanced estimation, predictive safety filter formulations, the formal synthesis of control barrier functions (CBFs) for uncertain systems, and safe control with finite-time robustness guarantees.

Appendix A. Proof of Theorem 2

Proof.

First, due to the definition of $\Delta B^\star(x,u)$ and $B^\star(x)$, the increment bound in Assumption 1, and $\bar{\varepsilon}_\theta(p) \geq \varepsilon_{\theta,t}(p)$, it can be obtained that (12) is a sufficient condition for (16). Therefore, if B is a robust adaptive DT-CBF that satisfies (12), then $\mathcal{U}_{\mathrm{ras},t}(x;p)$ is non-empty. Next, it suffices to prove robust safety through positive invariance. Considering $B_{\mathrm{ra},t+1}(x_{t+1}) - B_{\mathrm{ra},t}(x_t)$, which characterizes the increment of the barrier function, it holds that

$$
\begin{aligned}
&B_{\mathrm{ra},t+1}(x_{t+1}) - B_{\mathrm{ra},t}(x_t)\\
\overset{(11)}{=}&B(x_{t+1},\hat{\theta}_{t+1}) - B(x_t,\hat{\theta}_t) - \frac{1}{2}\delta_{\theta,t}^\top \Gamma^{-1}(\delta_{\theta,t} + 2e_{\theta,t})\\
\geq&B(x_{t+1},\hat{\theta}_{t+1}) - B(x_t,\hat{\theta}_t) - \frac{\|\delta_{\theta,t}\|^2 + 2\|\delta_{\theta,t}\|\varepsilon_{\theta,t}(p)}{2\lambda_{\min}(\Gamma)},
\end{aligned}
$$

where $B(x_{t+1},\hat{\theta}_{t+1}) - B(x_t,\hat{\theta}_t)$ can be further relaxed using the Lipschitz perturbation bound as

$$
\begin{aligned}
&B(x_{t+1},\hat{\theta}_{t+1}) - B(x_t,\hat{\theta}_t) \geq B(f(x_t,\pi_t(x_t);\hat{\theta}_t),\hat{\theta}_t)\\
&- B(x_t,\hat{\theta}_t) - L_{B,x}(\|\boldsymbol{\phi}(x_t)\|\varepsilon_{\theta,t}(p) + \overline{w}) - L_{B,\theta}\|\delta_{\theta,t}\|.
\end{aligned}
$$

Combining the above two inequalities and using the fact that $\pi_t(x_t)$ satisfies (16), we have

$$
\begin{aligned}
&B_{\mathrm{ra},t+1}(x_{t+1}) - B_{\mathrm{ra},t}(x_t)\\
&\geq -\alpha\left(B(x_t,\hat{\theta}_t) - \frac{(\varepsilon_{\theta,t}(p))^2}{2\lambda_{\min}(\Gamma)}\right) \geq -\alpha(B_{\mathrm{ra},t}(x_t)).
\end{aligned}
\tag{A.1}
$$

Based on (A.1), if $x_t \in \mathcal{S}_{\mathrm{ra},t}$, i.e., $B_{\mathrm{ra},t}(x_t) \geq 0$, then $B_{\mathrm{ra},t+1}(x_{t+1}) \geq B_{\mathrm{ra},t}(x_t) - \alpha(B_{\mathrm{ra},t}(x_t)) \geq 0$ since $\alpha \in \mathcal{C}_{\mathrm{SL}}$, which implies $x_{t+1} \in \mathcal{S}_{\mathrm{ra},t+1}$. Thus, $\{\mathcal{S}_{\mathrm{ra},t}\}_{t=0}^T$ is positively invariant for the closed-loop system (18) given $\pi_t(x) \in \mathcal{U}_{\mathrm{ras},t}(x;p)$. Moreover, noting that $B_{\mathrm{ra},t}(x) \geq 0 \Longrightarrow B(x,\hat{\theta}_t) \geq 0$, then $\{\mathcal{S}_{B,\hat{\theta}_t}\}_{t=0}^T$ is positively invariant (i.e., robustly safe) for (18). □

Appendix B. Robustified CBFs with error bounds

The true error $e_{\theta,t}$ is used to construct $B_{\mathrm{ra},t}$ in (14) to achieve adaptive safety, which is similar to Taylor and Ames (2020). Although safety can be established, $\{\mathcal{S}_{\mathrm{ra},t}\}_{t=0}^T$ cannot be characterized since $e_{\theta,t}$ is unknown. Considering that the ultimate goal is to achieve safety w.r.t. $\mathcal{S}_{B,\hat{\theta}_t}$, which encodes safety (Lopez et al., 2020; Taylor & Ames, 2020), the proposed method using $B_{\mathrm{ra},t}$ defined in (14) is still valid since $\mathcal{S}_{\mathrm{ra},t} \subseteq \mathcal{S}_{B,\hat{\theta}_t}$.

An alternative is to use the *set-induced* error bound $\eta_{\theta,t}$ satisfying $|e_{\theta,t}| \leq \eta_{\theta,t}$ to construct robustified CBFs (Lopez et al., 2020). Specifically, for all $i \in \mathbb{I}_{[1:q]}$, $\eta_{\theta,t}[i]$ can be computed as

$$
\eta_{\theta,t}[i] = \max_{\theta\in\Theta_t}\theta[i] - \min_{\theta\in\Theta_t}\theta[i]. \tag{B.1}
$$

The bound in (B.1) is conservative as it depends solely on Θ_t. Furthermore, the non-smoothness of the max and min operations precludes an analytical time derivative, forcing continuous-time invariance proofs to rely on the true error rather than the bound used to construct the CBF (Lopez et al., 2020). Consequently, establishing the invariance of safe sets defined by error-bound robustified CBFs is generally infeasible in continuous time[8]. In contrast, the discrete-time setting permits the establishment of such invariance properties as no differentiation is required. The error-bound robustified CBF is defined as:

$$
\overline{B}_{\mathrm{ra},t}(x) := B(x,\hat{\theta}_t) - \frac{1}{2}\eta_{\theta,t}^\top \Gamma^{-1}\eta_{\theta,t}, \tag{B.2}
$$

where Γ is a symmetric positive definite matrix, and the safe set corresponding to $\overline{B}_{\mathrm{ra},t}$ is given by

$$
\overline{\mathcal{S}}_{\mathrm{ra},t} = \{x \in \mathcal{X} \mid \overline{B}_{\mathrm{ra},t}(x) \geq 0\}. \tag{B.3}
$$

Following Definition 4, a modified robust adaptive DT-CBF based on (B.2) and (B.3) can be defined. Moreover, a safety guarantee similar to Theorem 2 follows.

Definition 7. Given a set-based estimator EST, a continuous function $B : \mathcal{X}\times\Theta \to \mathbb{R}$ is an error-bound robust adaptive DT-CBF if it is Lipschitz continuous and there exists $\alpha \in \mathcal{K}_\infty \cap \mathcal{C}_{\mathrm{SI}}$ such that, for all $x \in \mathcal{X}$, there exists $u \in \mathcal{U}$ satisfying

$$
\begin{aligned}
&\Delta B^\star(x,u) - L_{B,x}\overline{w} - E_\eta^\star(x)\\
&\geq -\alpha\left(B^\star(x) - \frac{\|\eta_{\theta,0}\|^2}{2\lambda_{\min}(\Gamma)}\right),
\end{aligned}
\tag{B.4}
$$

[8] In the proof of Lopez et al. (2020, Theorem 2), the conditions $h \geq h_r - \frac{1}{2}\tilde{\vartheta}^\top\Gamma^{-1}\tilde{\vartheta}$ and $h \geq 0$ do not imply $h_r - \frac{1}{2}\tilde{\vartheta}^\top\Gamma^{-1}\tilde{\vartheta} \geq 0$.

where Γ is given in (B.2), $\Delta B^\star(x,u)$ and $B^\star(x)$ are defined in Definition 4, and $E_\eta^\star(x) = \left(L_{B,x}\|\boldsymbol{\phi}(x)\| + L_{B,\theta}\right)\|\eta_{\theta,0}\|$ with $L_{B,x}$ and $L_{B,\theta}$ given in Definition 2. The condition (B.4) is called the error-bound robust adaptive CBC against disturbances and parametric model mismatch.

Theorem 4. *Given $x_0 \in \mathcal{S}_{\mathrm{ra},0}$, a horizon $T \in \mathbb{N}$, a set-based estimator* EST*, and an error-bound robust adaptive DT-CBF with its associated matrix Γ and function $\alpha \in \mathcal{K}_\infty \cap \mathcal{C}_{\mathrm{SI}}$, for the error-bound robust adaptive online input condition*

$$B\big(f(x,u;\hat{\theta}_t),\hat{\theta}_t\big) - B(x,\hat{\theta}_t) - L_{B,x}\overline{w} - E_{\eta,t}(x) \geq -\alpha\left(B(x,\hat{\theta}_t) - \frac{\|\eta_{\theta,t}\|^2}{2\lambda_{\min}(\Gamma)}\right), \tag{B.5}$$

the admissible input set $\overline{\mathcal{U}}_{\mathrm{ras},t}(x)$ defined as

$$\overline{\mathcal{U}}_{\mathrm{ras},t}(x) := \{u \in \mathcal{U} \mid \text{(B.5) holds}\} \tag{B.6}$$

is non-empty, $\forall x \in \mathcal{X}$ and $t \in \mathbb{I}_{[0,T-1]}$. The term $E_{\eta,t}(x)$ in (B.5) equals $L_{B,x}\|\boldsymbol{\phi}(x)\|\|\eta_{\theta,t}\| + L_{B,\theta}\|\delta_{\eta,t}\|$ with $\delta_{\eta,t} = \eta_{\theta,t+1} - \eta_{\theta,t}$. In addition, if a time-varying policy $\overline{\pi}_t : \mathcal{X} \to \mathbb{R}^m$ satisfy $\overline{\pi}_t(x) \in \overline{\mathcal{U}}_{\mathrm{ras},t}(x)$ for all $x \in \mathcal{X}$ and $t \in \mathbb{I}_{[0,T-1]}$, then the closed-loop system

$$x_{t+1} = f(x_t, \overline{\pi}_t(x_t); \theta^*) + w_t \tag{B.7}$$

is robustly safe w.r.t. $\{\overline{\mathcal{S}}_{\mathrm{ra},t}\}_{t=0}^T$ and thus robustly safe w.r.t. $\{\mathcal{S}_{B,\hat{\theta}_t}\}_{t=0}^T$ (cf. Definition 5) against both parametric model mismatch and disturbances.

The proof of Theorem 4 is similar to that of Theorem 2 in Appendix A and is therefore omitted for brevity.

Remark 10. The advantage of using $\eta_{\theta,t}$ as in (B.2) is that $\overline{B}_{\mathrm{ra},t}$ can be explicitly determined. Moreover, using set-membership identification, it holds that $\eta_{\theta,t+1} \leq \eta_{\theta,t}$ (i.e., the error does not increase) since $\Theta_{t+1} \subseteq \Theta_t$, leading to $\overline{\mathcal{S}}_{\mathrm{ra},t} \subseteq \overline{\mathcal{S}}_{\mathrm{ra},t+1}$, i.e., the robust adaptive safe set defined using the error bound is non-shrinking.

Appendix C. More on inherent robustness

We first prove Theorem 3 using V_t defined in (24). The proof is inspired by the proof of Rawlings et al. (2017, Theorem B.13).

Proof. If there exists a $t_0 \in \mathbb{N}$ such that $x_{t_0} \in \mathcal{S}_{\mathrm{ra},t_0}$, then $x_t \in \mathcal{S}_{\mathrm{ra},t}$ holds for all $t \in \mathbb{I}_{[t_0:\infty]}$ by Theorem 2 since $\{B_{\mathrm{ra},t}\}_{t=0}^\infty$ is sequentially positively invariant for the system (6) with $u_t \in \mathcal{U}_{\mathrm{ras},t}(x_t;p)$. Thus, in this case, $\lim_{t\to\infty}\mathrm{dist}_{\mathcal{S}_{\mathrm{ra},t}}(x_t) = 0$. On the other hand, assume $x_t \in \mathcal{X}\setminus\mathcal{S}_{\mathrm{ra},t}$; then we have

$$V_{t+1}(x_{t+1}) - V_t(x_t) \overset{(24)}{=} B_{\mathrm{ra},t}(x_t) - B_{\mathrm{ra},t+1}(x_{t+1}) \overset{(A.1)}{\leq} \alpha(B_{\mathrm{ra},t}(x_t)) < 0, \tag{C.1}$$

where the last inequality is due to $B_{\mathrm{ra},t}(x_t) < 0$ for $x_t \in \mathcal{X}\setminus\mathcal{S}_{\mathrm{ra},t}$ and $\alpha \in \mathcal{K}_{\mathrm{e},\infty}$. Consider the sequence $\{V_t(x_t)\}_{t=0}^\infty$; given that $V_t(x_t) \geq 0$ and (C.1), by the monotone convergence theorem, there exists a $V^* \geq 0$ such that $\lim_{t\to\infty} V_t(x_t) = V^*$. Equivalently, $\forall\epsilon > 0$, $\exists T^*(\epsilon)$ such that $\forall t \geq T^*(\epsilon)$, $V_t(x_t) - V^* < \epsilon$. Next, we prove that $V^* = 0$ by contradiction. Assume $V^* > 0$, and let $\epsilon = -\alpha(-V^*)$. Then, for any $t \geq T^*(-\alpha(-V^*))$, it holds that (i) $V^* \leq V_t(x_t) < V^* - \alpha(-V^*)$ and (ii) $V^* \leq V_{t+1}(x_{t+1}) < V^* - \alpha(-V^*)$, implying $V_{t+1}(x_{t+1}) - V_t(x_t) > \alpha(-V^*)$. On the other hand, by (C.1), we have $V_{t+1}(x_{t+1}) - V_t(x_t) \leq \alpha(-V^*)$, which leads to a contradiction. Therefore, $\lim_{t\to\infty} V_t(x_t) = V^* = 0$, implying $\lim_{t\to\infty}\mathrm{dist}_{\mathcal{S}_{\mathrm{ra},t}}(x_t) = 0$ by the continuity of V_t. □

In fact, following the discussion in Appendix B, by constructing a sequence of energy functions using $\overline{\mathcal{S}}_{\mathrm{ra},t}$, it can be shown that the safe control input characterized by the certificate (B.4) also has the inherent robustness property, which is summarized in the following theorem:

Theorem 5. *Given a set-based estimator and an error-bound robust adaptive DT-CBF $B : \mathcal{X}\times\Theta \to \mathbb{R}$ (cf. Definition 7), x_t converges to $\limsup_{t\to\infty}\overline{\mathcal{S}}_{\mathrm{ra},t}$ asymptotically (cf. Definition 6) on $\mathcal{X}$ for the system (6) with $u_t \in \overline{\mathcal{U}}_{\mathrm{ras},t}(x_t)$, where $\overline{\mathcal{U}}_{\mathrm{ras},t}(x_t)$ is constructed as in (B.6) via (B.5).*

Theorem 5 implies that x_t will be regulated towards the set $\limsup_{t\to\infty}\overline{\mathcal{S}}_{\mathrm{ra},t}$ asymptotically on $\mathcal{X}$ for the closed-loop system (B.7). The proof of Theorem 5 closely follows that of Theorem 3 and is omitted for brevity.

References

Ackermann, J., Bartlett, A., Kaesbauer, D., Sienel, W., & Steinhauser, R. (1993). *Robust control: Systems with uncertain physical parameters*. Springer.

Agrawal, A., & Sreenath, K. (2017). Discrete control barrier functions for safety-critical control of discrete systems with application to bipedal robot navigation. In *Proceedings of robotics: science and systems*. Cambridge, Massachusetts.

Ames, A. D., Xu, X., Grizzle, J. W., & Tabuada, P. (2016). Control barrier function based quadratic programs for safety critical systems. *IEEE Transactions on Automatic Control*, *62*(8), 3861–3876.

Åström, K. J., & Wittenmark, B. (1994). *Adaptive control* (2nd). Inc.: Addison-Wesley Longman Publishing Co..

Bof, N., Carli, R., & Schenato, L. (2018). Lyapunov theory for discrete time systems. arXiv preprint arXiv:1809.05289.

Borrelli, F., Bemporad, A., & Morari, M. (2017). *Predictive control for linear and hybrid systems*. Cambridge University Press.

Boyd, S. P., & Vandenberghe, L. (2004). *Convex optimization*. Cambridge University Press.

Cohen, M. H., & Belta, C. (2023). Safe exploration in model-based reinforcement learning using control barrier functions. *Automatica*, *147*, Article 110684.

Cosner, R. K., Culbertson, P., & Ames, A. D. (2024). Bounding stochastic safety: Leveraging freedman's inequality with discrete-time control barrier functions. *IEEE Control Systems Letters*, *8*, 1937–1942.

Didier, A., & Zeilinger, M. N. (2026). Approximate predictive control barrier function for discrete-time systems. *IEEE Transactions on Automatic Control*, Early Access.

Folland, J. B. (1999). *Real analysis: Modern techniques and their applications*. John Wiley & Sons.

Goodwin, G. C., & Sin, K. S. (2014). *Adaptive filtering prediction and control*. Courier Corporation.

Hertneck, M., Köhler, J., Trimpe, S., & Allgöwer, F. (2018). Learning an approximate model predictive controller with guarantees. *IEEE Control Systems Letters*, *2*(3), 543–548.

Horn, R. A., & Johnson, C. R. (2012). *Matrix analysis* (2nd ed.). Cambridge University Press.

Ioannou, P. A., & Sun, J. (1996). PTR Prentice-Hall Upper Saddle River, NJ.

Jankovic, M. (2018). Robust control barrier functions for constrained stabilization of nonlinear systems. *Automatica*, *96*, 359–367.

Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980.

Köhler, J., Kötting, P., Soloperto, R., Allgöwer, F., & Müller, M. A. (2021). A robust adaptive model predictive control framework for nonlinear uncertain systems. *International Journal of Robust and Nonlinear Control*, *31*(18), 8725–8749.

Li, S., & Liu, Z. (2009). Adaptive speed control for permanent-magnet synchronous motor system with variations of load inertia. *IEEE Transactions on Industrial Electronics*, *56*(8), 3050–3059.

Liu, C., Shi, S., Alan, A., Venayagamoorthy, G. K., & De Schutter, B. (2026). Approximate model predictive control for microgrid energy management via imitation learning. *Engineering Applications of Artificial Intelligence*, *182*, Article 115837.

Liu, C., Shi, S., & De Schutter, B. (2024). Stability and performance analysis of model predictive control of uncertain linear systems. In *2024 IEEE 63rd conference on decision and control* (pp. 7356–7362).

Liu, C., Shi, S., & De Schutter, B. (2025). On the regret of model predictive control with imperfect inputs. *IEEE Control Systems Letters*, *9*, 601–606.

Liu, C., Shi, S., & De Schutter, B. (2026). Certainty-equivalence model predictive control: Stability, performance, and beyond. *IEEE Transactions on Automatic Control*, *71*(7), 4649–4664.

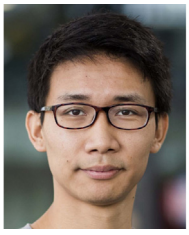
Ljung, L. (2010). Perspectives on system identification. *Annual Reviews in Control*, *34*(1), 1–12.
Lopez, B. T., Slotine, J.-J. E., & How, J. P. (2020). Robust adaptive control barrier functions: An adaptive and data-driven approach to safety. *IEEE Control Systems Letters*, *5*(3), 1031–1036.
Lorenzen, M., Cannon, M., & Allgöwer, F. (2019). Robust MPC with recursive model update. *Automatica*, *103*, 461–471.
Luo, W., Sun, W., & Kapoor, A. (2022). Sample-efficient safe learning for online nonlinear control with control barrier functions. In *International workshop on the algorithmic foundations of robotics* (pp. 419–435). Springer.
Mania, H., Jordan, M. I., & Recht, B. (2022). Active learning for nonlinear system identification with guarantees. *Journal of Machine Learning Research*, *23*(32), 1–30.
Milanese, M., Norton, J., Piet-Lahanier, H., & Walter, É. (2013). *Bounding approaches to system identification*. Springer Science & Business Media.
Nijmeijer, H., & van der Schaft, A. (1990). *Nonlinear dynamical control systems*: *vol. 175*, Springer.
Rawlings, J. B., Mayne, D. Q., & Diehl, M. M. (2017). *Model predictive control: Theory, computation, and design* (2nd ed.). Nob Hill Publishing.
Santoyo, C., Dutreix, M., & Coogan, S. (2019). Verification and control for finite-time safety of stochastic systems via barrier functions. In *2019 IEEE conference on control technology and applications* (pp. 712–717).
Semnani, S. H., Liu, H., Everett, M., Ruiter, A. D., & How, J. P. (2020). Multi-agent motion planning for dense and dynamic environments via deep reinforcement learning. *IEEE Robotics and Automation Letters*, *5*(2), 3221–3226.
Shen, J., Wang, W., Zhou, J., & Lü, J. (2025). Adaptive safety with control barrier functions and triggered batch least-squares identifier. *Automatica*, *173*, Article 112059.
Srinivasan, M., Dabholkar, A., Coogan, S., & Vela, P. A. (2020). Synthesis of control barrier functions using a supervised machine learning approach. In *2020 IEEE/RSJ international conference on intelligent robots and systems* (pp. 7139–7145).
Taylor, A. J., & Ames, A. D. (2020). Adaptive safety with control barrier functions. In *2020 American control conference* (pp. 1399–1405).
Wang, S., & Wen, S. (2025). Safe control against uncertainty: A comprehensive review of control barrier function strategies. *IEEE Systems, Man, and Cybernetics Magazine*, *11*(1), 34–47.
Xiao, W., Belta, C., & Cassandras, C. G. (2021). Adaptive control barrier functions. *IEEE Transactions on Automatic Control*, *67*(5), 2267–2281.
Xu, X., Tabuada, P., Grizzle, J. W., & Ames, A. D. (2015). Robustness of control barrier functions for safety critical control. *IFAC-PapersOnLine*, *48*(27), 54–61.
Zeng, J., Li, Z., & Sreenath, K. (2021). Enhancing feasibility and safety of nonlinear model predictive control with discrete-time control barrier functions. In *2021 60th IEEE conference on decision and control* (pp. 6137–6144).



Changrui Liu received the B.Sc. degree (cum laude & Best 10 Graduates) from the Honors School of Harbin Institute of Technology (HIT), Harbin, China, in 2019, and the M.Sc. degree from TU Delft, Delft, The Netherlands, in 2022. He is currently pursuing the Ph.D. degree at the Delft Center for Systems and Control, TU Delft. His research interests include performance guarantees in machine learning and optimization. He is the recipient of the IFAC Outstanding PhD Research Award (class 2026) at the 23rd IFAC World Congress.



Anil Alan received the B.Sc. degree of Mechanical Engineering at Middle East Technical University, Ankara, Türkiye, in 2012, the M.Sc. degree from Bilkent University, Ankara, in 2017, and the Ph.D. degree from the University of Michigan, Ann Arbor, MI, USA, in 2024. He is currently a Postdoc with the Delft University of Technology, Delft, The Netherlands. His research focuses on safety-critical control.



Shengling Shi is an assistant professor at the Delft Center for Systems and Control (DCSC), TU Delft, and he was a Postdoc at the Massachusetts Institute of Technology. He was also a postdoc at the DCSC, TU Delft. He received the Ph.D. degree from the Eindhoven University of Technology in 2021. His research interests include system identification, model predictive control, and their applications.



Bart De Schutter is a full professor at the Delft Center for Systems and Control of TU Delft. His research interests include integrated learning- and optimization-based control, machine learning for control and decision making, multi-agent control, and control of hybrid systems with applications in intelligent transportation, smart energy, and underwater robotics.