

Modular Adaptive Safety-Critical Control

Max H. Cohen and Calin Belta

Abstract—This paper presents an adaptive control approach for uncertain nonlinear systems subject to safety constraints that allows for modularity in the selection of the parameter estimation algorithm. Such modularity is achieved by unifying the concepts of input-to-state stability (ISS) and input-to-state safety (ISSf) via control Lyapunov functions (CLFs) and control barrier functions (CBFs), respectively. In particular, we propose a class of exponential ISS-CLFs and ISSf high order CBFs that can be combined with a general class of parameter estimation algorithms akin to those found in the literature on concurrent learning adaptive control. We demonstrate that the unification of ISS and ISSf in an adaptive control setting allows for maintaining a single set of parameter estimates for both the CLF and CBF that can be generated by a class of update laws satisfying a few general properties. The modularity of our approach is demonstrated via numerical examples by comparing performance in terms of stability and safety across different parameter estimation algorithms.

I. INTRODUCTION

Adaptive control theory [1], [2] is concerned with simultaneous learning and control of uncertain dynamical systems. In traditional adaptive control, learning often manifests itself as the estimation of uncertain parameters associated with the underlying dynamical system, whereas control is synonymous with stabilization to a set point or tracking of a desired reference trajectory. In the context of adaptive control of nonlinear systems, designs are often classified as either *Lyapunov-based* [3] or *modular* [4]. Lyapunov-based designs [1, Ch. 3-4] typically rely on the design of a control Lyapunov function (CLF) [5] for a modified dynamical system, where potentially destabilizing parameter estimation errors are eliminated through the use of a Lyapunov-based parameter update law. Such designs have the benefit of guaranteeing asymptotic stability (or even exponential stability [6]–[10]), but restrict the design of the parameter estimation algorithm since it is tightly coupled to the associated CLF. On the other hand, modular designs [1, Ch. 5-6] decouple the design of controller and update law: a controller and parameter estimator satisfying a few general properties can be combined to enforce weaker forms of stability, such as input-to-state stability (ISS) [11]. Allowing modularity in the estimation algorithm is motivated by the fact that various estimation procedures may provide certain benefits compared to Lyapunov-based update laws. For example, least-squares based estimation algorithms generally exhibit

faster convergence than gradient descent based estimation algorithms employed in Lyapunov-based approaches.

Motivated by the need for certifiably correct behavior of modern autonomous systems, adaptive control techniques have recently been applied to more complicated control problems, such as guaranteeing safety [12]–[18], often formalized using set theoretic notions [19], and enforcing more general temporal logic specifications [20]–[22]. In [12], classical Lyapunov-based adaptive control designs are extended to a safety-critical setting, where a control barrier function (CBF) [23] is designed for a modified dynamical system and a CBF-based adaptive update law is leveraged to eliminate parameter estimation errors that could otherwise lead to safety violation. Although [12] provides a foundation for extending traditional nonlinear adaptive control designs to enforcing safety using CBFs, the particular approach taken therein is conservative in the sense that the proposed method restricts the system to the nonnegative superlevel sets of the safe set, rather than only the zero superlevel set as is common when using CBFs [13]. By leveraging known bounds on the system’s uncertain parameters, works such as [13]–[17] reduce such conservatism by taking a “robust adaptive” approach whereby safety is guaranteed by accounting for the worst-case parameter estimation error, which is reduced online as more data about the system is collected. In [18] such robust adaptive approaches are extended to CBFs with high relative degree [24]–[27] in which the control input may not directly influence the derivative of the CBF candidate.

The aforementioned approaches to enforcing safety using adaptive control techniques have demonstrated success on a wide variety of problems; however, they suffer from a combination of the following limitations: 1) they restrict the update laws/estimation algorithms that can be used to guarantee safety [12]–[18]; 2) they require precise knowledge of the bounds on the system parameters/uncertainties [13]–[18]; 3) they require redundant parameter estimation in the sense that multiple estimates of the same parameters are needed if the adaptive safety controller is combined with an adaptive stabilizing controller to achieve a performance objective [12], [13], [15], [18]. The main objective of this paper is to develop a framework for modular adaptive safety-critical control that addresses the previous limitations by 1) allowing freedom in the selection of the parameter update law/estimation algorithm, 2) not requiring precise knowledge of bounds on the system parameters/uncertainties, and 3) allowing for the use of a single set of parameter estimates that are shared between the safety and performance (stabilizing) controller. We accomplish this objective by unifying the concept of ISS with that of *input-to-state safety* (ISSf) [28]–[30], extending ideas

The authors are with the Department of Mechanical Engineering, Boston University, Boston, MA, USA {maxcohen, cbelta}@bu.edu. This work is supported by the NSF under grants DGE-1840990 and IIS-2024606. Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF.

from traditional modular nonlinear adaptive control [4] to a safety-critical setting. In particular, we show how a general class of parameter estimation algorithms can be combined with a particular class of ISS-CLFs [31] and a class of ISSf high order CBFs to simultaneously guarantee ISS and ISSf of the underlying system in a modular fashion. The drawback of this modularity is that we establish safety using an ISSf framework, which studies the invariance of inflated safe sets parameterized by the magnitude of a disturbance (parameter estimation error) perturbing the nominal system dynamics.

The contributions of this paper are threefold. First, we present a modular approach to nonlinear adaptive stabilization using a class of exponential ISS-CLFs [31]. We combine such CLFs with a class of parameter estimators characteristic of those found in *concurrent learning* adaptive control [6]–[10] and show that our approach guarantees ISS and, under suitable excitation conditions outlined in [6]–[10], asymptotic stability of the origin. Second, we generalize the class of ISSf-CBFs introduced in [28]–[30] to high relative degree safety constraints [24]–[27], which allows for the systematic construction of a candidate safe set from a given safety constraint on the lower order system dynamics. Finally, we demonstrate the versatility of our approach via numerical examples in which the uncertain parameters of a system are learned online using different update laws while guaranteeing ISS and ISSf.

The remainder of this paper is organized as follows. Sec. II covers preliminaries on ISS and ISSf. Sec. III presents our approach to modular adaptive stabilization. Sec. IV introduces the notion of an ISSf high order CBF. Sec. V presents simulations and Sec. VI contains concluding remarks. Due to space limitations, all proofs and some simulation examples are omitted and are available in the preprint of this paper.

II. PRELIMINARIES AND PROBLEM FORMULATION

Consider the uncertain nonlinear control affine system

$$\dot{x} = f(x) + F(x)\theta + g(x)u, \quad (1)$$

where $x \in \mathbb{R}^n$ is the system state, $u \in \mathbb{R}^m$ is the control input, and $\theta \in \mathbb{R}^p$ is a vector of uncertain parameters. The vector field $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ captures the system drift dynamics, the columns of $g : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times m}$ represent vector fields describing the control directions, and the matrix-valued function $F : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times p}$ is a known basis of nonlinear features for the uncertain parameters. We assume f , g , and F , are locally Lipschitz continuous, and that $f(0) = 0$ and $F(0) = 0$ so that the origin is an equilibrium point of (1) with $u = 0$. Our main objective is to learn the uncertain parameters in (1) completely online, while guaranteeing stability and safety. To this end, let $\hat{\theta} \in \mathbb{R}^p$ be an estimate of θ and define $\tilde{\theta} := \theta - \hat{\theta}$ as the parameter estimation error. Using $\tilde{\theta}$, system (1) is equivalent to

$$\dot{x} = f(x) + F(x)\hat{\theta} + g(x)u + F(x)\tilde{\theta}, \quad (2)$$

and given a controller $k : \mathbb{R}^n \times \mathbb{R}^p \rightarrow \mathbb{R}^m$, locally Lipschitz in both its arguments, we obtain the closed-loop system by

fixing $u = k(x, \hat{\theta})$ as

$$\dot{x} = f(x) + F(x)\hat{\theta} + g(x)k(x, \hat{\theta}) + F(x)\tilde{\theta}. \quad (3)$$

When $\tilde{\theta}$ is viewed as a disturbance input to the nominal closed-loop dynamics $f(x) + F(x)\hat{\theta} + g(x)k(x, \hat{\theta})$, a tool for studying the stability of (3) is *input-to-state stability*.

Definition 1 ([1]). System (3) is said to be *input-to-state stable* (ISS) if for each $x(0) \in \mathbb{R}^n$ and each $\tilde{\theta}(\cdot) \in \mathcal{L}_\infty$ the trajectory of the closed-loop system (3) satisfies

$$\|x(t)\| \leq \beta(\|x(0)\|, t) + \iota(\|\tilde{\theta}\|_\infty), \quad \forall t \in \mathbb{R}_{\geq 0}, \quad (4)$$

for some $\beta \in \mathcal{KL}$ and $\iota \in \mathcal{K}_\infty$. If $\beta(r, s) = cre^{-\lambda s}$ for some positive constants $c, \lambda \in \mathbb{R}_{>0}$, then (2) is said to be *exponentially ISS* (eISS).

In addition to stability, we are interested in studying the safety properties of (1), which is often associated with the concept of set invariance [19]. Given a locally Lipschitz feedback controller for (1), a set

$$\mathcal{C} = \{x \in \mathbb{R}^n \mid h(x) \geq 0\}, \quad (5)$$

where $h : \mathbb{R}^n \rightarrow \mathbb{R}$ is continuously differentiable, is said to be *forward invariant* if the resulting solution $x : I(x(0)) \rightarrow \mathbb{R}^n$ satisfies $x(t) \in \mathcal{C}$ for all $t \in I(x(0))$, where $I(x(0)) \subseteq \mathbb{R}_{\geq 0}$ is the solution's maximal interval of existence from an initial condition of $x(0) \in \mathbb{R}^n$. When studying the disturbed system (2), the notion of *input-to-state safety* (ISSf), introduced in [28]–[30], provides a methodology to study the impact of uncertainties on safety. In particular, the ISSf framework is concerned with establishing the forward invariance of an inflated version of (5) defined as

$$\mathcal{C}_\delta = \{x \in \mathbb{R}^n \mid h(x) + \gamma(\delta) \geq 0\}, \quad \gamma \in \mathcal{K}_\infty. \quad (6)$$

Definition 2 ([30]). System (3) is said to be *input-to-state safe* (ISSf) on a set \mathcal{C} as in (5) if there exists a γ as in (6) such that for all $\delta \in \mathbb{R}_{\geq 0}$ and all $\tilde{\theta}(\cdot) \in \mathcal{L}_\infty$ satisfying $\|\tilde{\theta}\|_\infty \leq \delta$, the set \mathcal{C}_δ defined as in (6) is forward invariant.

III. MODULAR ADAPTIVE ISS

In this section we introduce a modular adaptive control approach to stabilization by exploiting a class of exponential ISS control Lyapunov functions (eISS-CLF) [31]. Although our development leverages ISS-CLFs, we illustrate that, under suitable assumptions, the controller derived from this ISS-CLF guarantees asymptotic stability of the closed-loop system, rather than ISS. In principle our approach can be used with any parameter estimation algorithm that guarantees boundedness of the estimation error; however, we specialize our results to a particular class of parameter estimators whose characteristics are outlined in the following lemma.

Lemma 1. Consider a parameter update law $\dot{\hat{\theta}} = \tau(\hat{\theta}, t)$, with τ locally Lipschitz in its first argument and piecewise continuous in its second, and a Lyapunov-like function $V_\theta : \mathbb{R}^p \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$, continuously differentiable in both its arguments, satisfying

$$\eta_1 \|\tilde{\theta}\|^2 \leq V_\theta(\hat{\theta}, t) \leq \eta_2 \|\tilde{\theta}\|^2, \quad \forall (\tilde{\theta}, t) \in \mathbb{R}^p \times \mathbb{R}_{\geq 0}, \quad (7)$$

for some $\eta_1, \eta_2 \in \mathbb{R}_{>0}$. Provided

$$\dot{V}_{\tilde{\theta}}(\tilde{\theta}, t) \leq 0, \quad \forall (\tilde{\theta}, t) \in \mathbb{R}^p \times \mathbb{R}_{\geq 0}, \quad (8)$$

then $\tilde{\theta}(\cdot) \in \mathcal{L}_{\infty}$. Furthermore, if there exists a pair $(\eta_3, T) \in \mathbb{R}_{>0} \times \mathbb{R}_{\geq 0}$ such that

$$\dot{V}_{\tilde{\theta}}(\tilde{\theta}, t) \leq -\eta_3 \|\tilde{\theta}\|^2, \quad \forall (\tilde{\theta}, t) \in \mathbb{R}^p \times \mathbb{R}_{\geq T}, \quad (9)$$

then $\tilde{\theta}(\cdot) \in \mathcal{L}_2 \cap \mathcal{L}_{\infty}$ and

$$\|\tilde{\theta}(t)\| \leq \frac{\eta_2}{\eta_1} \|\tilde{\theta}(0)\| e^{-\frac{\eta_3}{2\eta_2}(t-T)}, \quad \forall t \in \mathbb{R}_{\geq 0}. \quad (10)$$

The condition in (8) requires that the origin of the parameter estimation error dynamics is stable in the sense of Lyapunov – a property satisfied by standard estimation algorithms (cf. [4]). The condition in (9) is more restrictive. It asks, after a certain time period, for the parameter estimates to exponentially converge to their true values. Traditionally, this is only guaranteed under prohibitive persistence of excitation (PE) conditions that generally cannot be verified in practice [6]. Over the past decade, however, a suite of tools termed *concurrent learning adaptive control* [6] have emerged that relax the PE condition by maintaining a sufficiently rich “history stack” of input-output data [7]. In this regard, the interval $[0, T)$ in Lemma 1 corresponds to an initial input-output data collection phase; once a sufficiently rich dataset has been collected such data can be exploited to ensure exponential convergence of the parameter estimates (see [7] for data collection strategies). We refer the interested reader to [6]–[8] for a more thorough introduction and to [9], [10], [32], [33] for specific instances of update laws satisfying the conditions of Lemma 1. The class of parameter estimators outlined in the preceding lemma will be combined with the notion of an eISS-CLF to establish stability of (2) in the presence of uncertain parameters.

Definition 3. A continuously differentiable positive definite function $V : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0}$ is said to be an *exponential input-to-state stable control Lyapunov function* (eISS-CLF) for (2) if there exist positive constants $c_1, c_2, c_3, \varepsilon \in \mathbb{R}_{>0}$ such that

$$c_1 \|x\|^2 \leq V(x) \leq c_2 \|x\|^2, \quad \forall x \in \mathbb{R}^n, \quad (11a)$$

$$\inf_{u \in \mathbb{R}^m} \dot{V}(x, \hat{\theta}, u) < -c_3 V(x) - \frac{1}{\varepsilon} \|L_F V(x)\|^2, \quad (11b)$$

for all $(x, \hat{\theta}) \in (\mathbb{R}^n \setminus \{0\}) \times \mathbb{R}^p$ where

$$\dot{V}(x, \hat{\theta}, u) := L_f V(x) + L_F V(x) \hat{\theta} + L_g V(x) u.$$

When the uncertain parameters in (2) are *matched* the construction of an eISS-CLF can be performed by constructing a CLF for the nominal dynamics $\dot{x} = f(x) + g(x)u$. We show in the following theorem that combining a parameter estimator satisfying the conditions of Lemma 1 with a controller satisfying the conditions in (11) renders the closed-loop system ISS and, under additional conditions, renders the origin asymptotically stable.

Theorem 1. If V is an eISS-CLF for (2) and the conditions of Lemma 1 hold, then any controller $u = k(x, \hat{\theta})$ locally Lipschitz on $(\mathbb{R}^n \setminus \{0\}) \times \mathbb{R}^p$ satisfying

$$\dot{V}(x, \hat{\theta}, k(x, \hat{\theta})) \leq -c_3 V(x) - \frac{1}{\varepsilon} \|L_F V(x)\|^2,$$

for all $(x, \hat{\theta})$ renders (3) eISS and $\lim_{t \rightarrow \infty} x(t) = 0$.

Given an eISS-CLF as in Def. 3, inputs satisfying the conditions of Theorem 1 can be computed for any $(x, \hat{\theta})$ by solving the quadratic program (QP)

$$\min_{u \in \mathbb{R}^m} \frac{1}{2} u^\top u \quad (12)$$

$$\text{subject to } \dot{V}(x, \hat{\theta}, u) \leq -c_3 V(x) - \frac{1}{\varepsilon} \|L_F V(x)\|^2.$$

As noted in [30], [34], the strict inequality in (11) helps to establish Lipschitz continuity of the QP-based controller away from the origin and is independent of the non-strict inequality that a particular controller must satisfy to guarantee stability. This observation also applies to the CBFs outlined in the next section. Continuity at the origin can be ensured provided the eISS-CLF V satisfies the *small control property* (see [34] for further details).

IV. MODULAR ADAPTIVE ISSF

In this section we shift our attention to the problem of establishing safety of (2) in the presence of uncertain parameters. Importantly, we aim to establish such safety guarantees without having to commit to any particular parameter update law as in [12]–[18]. We accomplish this by unifying the ISSf framework [28]–[30] with the high order CBF (HOCBF) framework [24]–[27], which allows one to recursively compute a candidate safe set from a user-defined high relative degree safety constraint. We note that efforts towards this unification have been explored by [35] in the context of safety *verification* of interconnected systems. Here, we present a formulation better suited for control synthesis in the context of uncertain systems.

Definition 4 ([27]). A function $h : \mathbb{R}^n \rightarrow \mathbb{R}$ is said to have relative degree $r \in \mathbb{N}$ for (2) with respect to u on a domain $\mathcal{D} \subseteq \mathbb{R}^n$ if

- 1) h is r -times differentiable;
- 2) for all $x \in \mathbb{R}^n$ and all $i \in \{0, 1, \dots, r-2\}$, we have $L_g L_f^i h(x) = 0$;
- 3) $L_g L_f^{r-1} h(x) \neq 0$ for all $x \in \mathcal{D}$.

The relative degree of a function h for (2) with respect to $\tilde{\theta}$ is defined similarly by replacing g with F in Def. 4. Now consider a function $h : \mathbb{R}^n \rightarrow \mathbb{R}$ of relative degree $r \in \mathbb{N}$ for (2) with respect to u and define

$$\mathcal{S} := \{x \in \mathbb{R}^n \mid h(x) \geq 0\}, \quad (13)$$

as the *constraint set* that we desire to render invariant.

Assumption 1 ([18]). The constraint function h in (13) has relative degree $r \in \mathbb{N}$ for (2) with respect to both u and $\tilde{\theta}$.

The implication of Assumption 1 is that *both* the control input and uncertain parameters only appear in the r -th total

derivative of h , which will facilitate the development of affine constraints on the control input that are sufficient to establish ISSf, and is not restrictive provided the uncertain parameters are matched. Given a constraint function h of relative degree $r \in \mathbb{N}$ as in (13) and a collection of sufficiently smooth $\alpha_i \in \mathcal{K}_\infty^e$, $i \in \{1, \dots, r\}$, we define the sequence of functions

$$\begin{aligned}\psi_0(x) &:= h(x) \\ \psi_i(x) &:= \dot{\psi}_{i-1}(x) + \alpha_i(\psi_{i-1}(x)), \quad \forall i \in \{1, \dots, r-1\}, \\ \psi_r(x, u) &:= \dot{\psi}_{r-1}(x, u) + \alpha_r(\psi_{r-1}(x)).\end{aligned}\quad (14)$$

We associate to each $i \in \{1, \dots, r\}$ a set $\mathcal{C}^i \subset \mathbb{R}^n$ defined as the zero-superlevel set of ψ_{i-1} as

$$\mathcal{C}^i := \{x \in \mathbb{R}^n \mid \psi_{i-1}(x) \geq 0\}, \quad (15)$$

and define a candidate safe set as

$$\mathcal{C} := \bigcap_{i=1}^r \mathcal{C}^i. \quad (16)$$

We now aim to develop a control strategy that renders \mathcal{C} from (16) ISSf with respect to the parameter estimation error. To this end, we define the sequence of functions

$$\rho_i(x, \delta) := \psi_{i-1}(x) + \gamma_i(\delta) \quad (17)$$

for all $i \in \{1, \dots, r\}$, where ψ_{i-1} is defined as in (14), and each $\gamma_i \in \mathcal{K}_\infty$. Similar to (15), we associate to each ρ_i a set $\mathcal{C}_\delta^i \subset \mathbb{R}^n$ as

$$\mathcal{C}_\delta^i := \{x \in \mathbb{R}^n \mid \rho_i(x, \delta) \geq 0\}, \quad (18)$$

and define an inflated version of (16) as

$$\mathcal{C}_\delta := \bigcap_{i=1}^r \mathcal{C}_\delta^i, \quad (19)$$

whose forward invariance will be established using the notion of an *ISSf high order CBF* (ISSf-HOCBF).

Definition 5. A function $h : \mathbb{R}^n \rightarrow \mathbb{R}$ of relative degree r with respect to u is said to be an *input-to-state safe high order control barrier function* for (2) on \mathcal{C}_δ defined as in (19) if $\nabla \psi_{i-1}(x) \neq 0$ for all $x \in \partial \mathcal{C}^i$ for all $i \in \{1, \dots, r\}$ and there exist sufficiently smooth $\alpha_i \in \mathcal{K}_\infty^e$, $i \in \{1, \dots, r\}$, and a positive constant $\varepsilon \in \mathbb{R}_{>0}$ such that for all $x \in \mathbb{R}^n$ and all $\hat{\theta} \in \mathbb{R}^p$

$$\sup_{u \in \mathbb{R}^m} \dot{\hat{\psi}}(x, \hat{\theta}, u) > -\alpha_r(\psi_{r-1}(x)) + \frac{\|L_F \psi_{r-1}(x)\|^2}{\varepsilon}, \quad (20)$$

where

$$\dot{\hat{\psi}}(x, \hat{\theta}, u) := L_f \psi_{r-1}(x) + L_F \psi_{r-1}(x) \hat{\theta} + L_g \psi_{r-1}(x) u,$$

and ψ_{r-1} is defined by h and the choice of α_i from (14).

Given the above definition, we define the pointwise set of all control values satisfying the condition in (20) as

$$\underbrace{\left\{ u \in \mathbb{R}^m \mid \dot{\hat{\psi}}(x, \hat{\theta}, u) \geq -\alpha_r(\psi_{r-1}(x)) + \frac{\|L_F \psi_{r-1}(x)\|^2}{\varepsilon} \right\}}_{=: K_h(x, \hat{\theta})}$$

and show in the following theorem that any locally Lipschitz controller belonging to K_h renders \mathcal{C}_δ forward invariant.

Theorem 2. Let h be an ISSf-HOCBF for (2) on \mathcal{C}_δ and suppose the conditions of Lemma 1 and Assumption 1 hold. Let $k : \mathbb{R}^n \times \mathbb{R}^p \rightarrow \mathbb{R}^m$ be a locally Lipschitz control policy satisfying $k(x, \hat{\theta}) \in K_h(x, \hat{\theta})$ for all $(x, \hat{\theta}) \in \mathbb{R}^n \times \mathbb{R}^p$ and suppose $\|\hat{\theta}\|_\infty \leq \delta$. Then, the control policy $u = k(x, \hat{\theta})$ renders \mathcal{C}_δ forward invariant for the closed-loop system with each γ_i defined as

$$\begin{aligned}\gamma_r(\delta) &:= -\alpha_r^{-1}\left(-\frac{\varepsilon \delta^2}{4}\right), \\ \gamma_i(\delta) &:= -\alpha_i^{-1}(-\gamma_{i+1}(\delta)), \quad \forall i \in \{1, \dots, r-1\}.\end{aligned}\quad (21)$$

Given an eISS-CLF and ISSf-HOCBF one can compute inputs satisfying the conditions of Theorem 2 and (relaxed) conditions of Theorem 1 by solving the QP

$$\begin{aligned}\min_{u \in \mathbb{R}^m} \quad & \frac{1}{2} \|u - k(x, \hat{\theta})\|^2 \\ \text{s.t.} \quad & \dot{\hat{\psi}}(x, \hat{\theta}, u) \geq -\alpha_r(\psi_{r-1}(x)) + \frac{\|L_F \psi_{r-1}(x)\|^2}{\varepsilon},\end{aligned}\quad (22)$$

where $k(x, \hat{\theta})$ satisfies the conditions of Theorem 1. Rather than combining the CLF and CBF in a single QP as in [23], [34], here we filter the solution of (12) through (22). This obviates the need to select an appropriate penalty on the relaxation of the CLF constraint, which can lead to controllers with large Lipschitz constants if chosen improperly [34].

V. NUMERICAL EXAMPLES

We consider a simple obstacle avoidance scenario for a planar mobile robot modeled as a double integrator with nonlinear drag effects of the form [36]

$$\ddot{q} = -D\dot{q}\|\dot{q}\| + u, \quad (23)$$

where $q = [q_1 \ q_2]^\top \in \mathbb{R}^2$ denotes the robot's position, $u \in \mathbb{R}^2$ its commanded acceleration, and $D \in \mathbb{R}^{2 \times 2}$ a diagonal matrix of damping coefficients. Defining $x := [q^\top \ \dot{q}^\top]^\top \in \mathbb{R}^4$ allows (23) to be expressed as in (1) with

$$\dot{x} = \underbrace{\begin{bmatrix} \dot{q} \\ 0 \end{bmatrix}}_{f(x)} + \underbrace{\begin{bmatrix} 0_{2 \times 2} \\ \text{diag}(\dot{q}\|\dot{q}\|) \end{bmatrix}}_{F(x)} \underbrace{\begin{bmatrix} D_1 \\ D_2 \end{bmatrix}}_{\theta} + \underbrace{\begin{bmatrix} 0_{2 \times 2} \\ I_{2 \times 2} \end{bmatrix}}_{g(x)} u, \quad (24)$$

where $0_{2 \times 2} \in \mathbb{R}^{2 \times 2}$ is a 2×2 matrix of zeros, $I_{2 \times 2}$ is a 2×2 identity matrix, $\text{diag}(\cdot)$ constructs a diagonal matrix from a vector, and $D_1, D_2 \in \mathbb{R}_{>0}$ are the unknown drag coefficients. Our control objective is to drive (23) to the origin while avoiding an obstacle in the workspace and learning the uncertain parameters online. To estimate the uncertain parameters, we leverage a general class of concurrent learning parameter estimation algorithms [6], [8] based on the method developed in [10]. This method works based on the observation that, along state-control trajectory $(x(\cdot), u(\cdot))$, system (1) can be expressed as

$$\begin{aligned}\int_{t-\Delta t}^t \dot{x}(s) ds &= \int_{t-\Delta t}^t f(x(s)) ds + \int_{t-\Delta t}^t F(x(s)) ds \\ &+ \int_{t-\Delta t}^t g(x(s)) u(s) ds,\end{aligned}$$

for all $t \geq \Delta t$, where $\Delta t \in \mathbb{R}_{>0}$ is the length of an integration window. Defining

$$\mathcal{Y}(t) := \int_{t-\Delta t}^t (\dot{x}(s) - f(x(s)) - g(x(s))u(s))ds,$$

$$\mathcal{F}(t) := \int_{t-\Delta t}^t F(x(s))ds$$

yields the linear relationship for the uncertain parameters

$$\mathcal{Y}(t) = \mathcal{F}(t)\theta. \quad (25)$$

Despite the appearance of \dot{x} in \mathcal{Y} , computing \mathcal{Y} only requires state measurements since $\int_{t-\Delta t}^t \dot{x}(s)ds = x(t) - x(t - \Delta t)$. The parameters can then be recursively estimated online by storing values of \mathcal{Y} and \mathcal{F} at run-time in a history stack¹ $\mathcal{H} = \{(\mathcal{Y}_j, \mathcal{F}_j)\}_{j=1}^N$, which is used to update the parameter estimates to minimize the squared prediction error $E(\hat{\theta}) = \sum_{j=1}^N \|\mathcal{Y}_j - \mathcal{F}_j \hat{\theta}\|^2$. To this end, we consider the following class of update laws

$$\dot{\hat{\theta}} = -\Gamma(t)\nabla E(\hat{\theta})^\top = \Gamma(t) \sum_{j=1}^N \mathcal{F}_j^\top (\mathcal{Y}_j - \mathcal{F}_j \hat{\theta}), \quad (26)$$

which serves as a general template for particular update laws based on the properties of $\Gamma(\cdot)$ as follows:

$$\dot{\Gamma} = 0, \quad (27a)$$

$$\dot{\Gamma} = -\Gamma \left[\sum_{j=1}^N \mathcal{F}_j^\top \mathcal{F}_j \right] \Gamma, \quad (27b)$$

$$\dot{\Gamma} = \beta \Gamma - \Gamma \left[\sum_{j=1}^N \mathcal{F}_j^\top \mathcal{F}_j \right] \Gamma, \quad (27c)$$

$$\dot{\Gamma} = \beta \left[1 - \frac{\|\Gamma\|}{\bar{\Gamma}} \right] \Gamma - \Gamma \left[\sum_{j=1}^N \mathcal{F}_j^\top \mathcal{F}_j \right] \Gamma, \quad (27d)$$

where $\beta \in \mathbb{R}_{>0}$ is a forgetting/discount factor and $\bar{\Gamma} \in \mathbb{R}_{>0}$ is a user-defined constant that bounds $\|\Gamma(t)\|$. With $\dot{\Gamma}$ as in (27), the update law in (26) corresponds to: (27a) gradient descent; (27b) recursive least squares (RLS); (27c) RLS with a forgetting/discount factor; (27d) RLS with a variable forgetting factor. An overview of these online parameter estimation algorithms, including their benefits and drawbacks, can be found in [2, Ch. 8.7]. We emphasize that the purpose of our numerical example is not necessarily to establish superiority of one algorithm over the others; rather, our goal is to demonstrate that, under the assumptions posed in Lemma 1, the stability/safety guarantees of the controller can be decoupled from the design of the parameter estimator, which allows considerable freedom in selecting an estimation algorithm best suited for the problem at hand.

We demonstrate the modularity of our approach (i.e., the ability to decouple the design of the estimator from the controller) by running a set of simulations with randomly

¹This data is stored in the history stack using the singular value maximizing algorithm from [7]

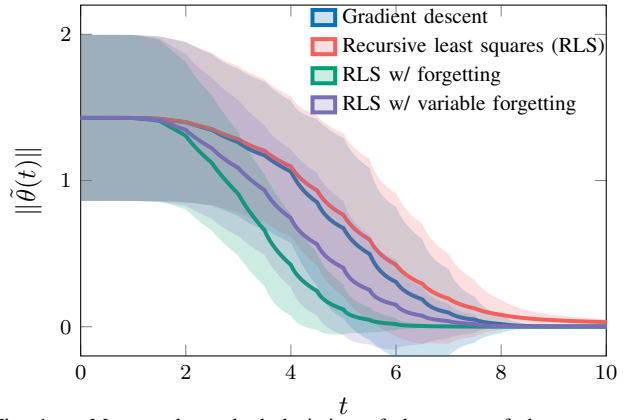


Fig. 1. Mean and standard deviation of the norm of the parameter estimation over time generated by each parameter estimator. The solid lines indicate the average value of $\|\hat{\theta}(t)\|$ across each simulation, and the ribbon surrounding each line corresponds to one standard deviation from the mean.

sampled initial conditions for the system state and estimated parameters under each algorithm, and show that, for a given level of uncertainty, the ISSf guarantees are invariant to the particular choice of parameter estimator. For each estimation algorithm we produce 25 different trajectories by uniformly sampling the initial state from $[-1.8, -2.2] \times [1.8, 2.2] \times \{0\} \times \{0\} \subset \mathbb{R}^4$ and the initial parameter estimates from $[0, 3]^2 \subset \mathbb{R}^2$; the true parameters are set to $\theta = [0.8 \ 1.4]^\top$. The hyperparameters for the estimation algorithms are selected as $N = 20$, $\Gamma(0) = 100I_{2 \times 2}$, $\beta = 1$, $\bar{\Gamma} = 1000$. The stabilization objective is achieved by considering the eISS-CLF candidate $V(x) = \frac{1}{2}\|q\|^2 + \frac{1}{2}\|q + \dot{q}\|^2$ with $c_3 = 1$ and $\varepsilon_V = 20$. The safety objective is achieved by considering the constraint function $h(x) = \|q - q_o\|^2 - R_o^2$, where $q_o = [-1 \ 1]^\top$ is the center of the circular obstacle and $R_o = 0.5$ its radius, which has relative degree 2 for (23) with respect to both u and θ as required by Assumption 1. This constraint function is used to construct an ISSf-HOCBF candidate with $\alpha_1(s) = s$, $\alpha_2(s) = \frac{1}{2}s$, and $\varepsilon_h = 1$.

For each simulation, the closed-loop trajectory is generated by the controller in (22), the results of which are provided in Fig. 1 and Fig. 2. As shown in Fig. 2, the trajectories under each update law remain safe and converge to the origin, whereas Fig. 1 illustrates the convergence of the parameter estimation error to zero for each estimation algorithm as predicted by Lemma 1. The curves in Fig. 1 represent the mean and standard deviation of the parameter estimation error over time across all simulations for each estimation algorithm. The results in Fig. 1 illustrate that, on average, the RLS with forgetting factor estimator (27c) produces the fastest convergence of the parameters estimates while also exhibiting low variance across different trajectories. The standard RLS algorithm (27b) produces the slowest convergence, which is expected given that, in general, this algorithm cannot guarantee exponential convergence of the parameter estimates, whereas the others can [2, Ch. 8.7].




Fig. 2. State trajectories generated by each estimation algorithm projected onto the x_1 - x_2 plane. In each plot the gray disk denotes the obstacle. The colors in each plot share the same interpretation as those in Fig. 1.

VI. CONCLUSION

We presented a modular approach to safe adaptive control using CLFs and CBFs. In particular, we unified the concepts of ISS and ISSf to allow for freedom in the estimation algorithm used to learn the uncertain parameters while maintaining ISS and ISSf guarantees. Our hope is that this work facilitates the application of more advanced techniques from the machine learning literature [37] to parameter estimation and learning in a safety-critical setting.

REFERENCES

- [1] M. Krstić, I. Kanellakopoulos, and P. Kokotović, *Nonlinear and adaptive control design*. John Wiley & Sons, 1995.
- [2] J. J. E. Slotine and W. Li, *Applied Nonlinear Control*. Prentice Hall, 1991.
- [3] M. Krstić and P. Kokotović, “Control lyapunov functions for adaptive nonlinear stabilization,” *Syst. Control Lett.*, vol. 26, no. 1, pp. 17–23, 1995.
- [4] M. Krstić and P. Kokotović, “Modular approach to adaptive nonlinear stabilization,” *Automatica*, vol. 32, no. 4, pp. 625–629, 1996.
- [5] E. D. Sontag, “A universal construction of artstein’s theorem on nonlinear stabilization,” *Syst. Control Lett.*, vol. 13, pp. 117–123, 1989.
- [6] G. Chowdhary and E. Johnson, “Concurrent learning for convergence in adaptive control without persistency of excitation,” in *Proc. Conf. Decis. Control*, pp. 3674–3679, 2010.
- [7] G. Chowdhary and E. Johnson, “A singular value maximizing data recording algorithm for concurrent learning,” in *Proc. Amer. Control Conf.*, pp. 3547–3552, 2011.
- [8] G. Chowdhary, T. Yucelen, M. Muhlegg, and E. N. Johnson, “Concurrent learning adaptive control of linear systems with exponentially convergent bounds,” *Int. J. Adapt. Control Signal Process.*, vol. 27, no. 4, pp. 280–301, 2013.
- [9] R. Kamalapurkar, B. Reish, G. Chowdhary, and W. E. Dixon, “Concurrent learning for parameter estimation using dynamic state-derivative estimators,” *IEEE Trans. Autom. Control*, vol. 62, no. 7, pp. 3594–3601, 2017.
- [10] A. Parikh, R. Kamalapurkar, and W. E. Dixon, “Integral concurrent learning: Adaptive control with parameter convergence using finite excitation,” *Int. J. Adapt. Control Signal Process.*, vol. 33, no. 12, pp. 1775–1787, 2019.
- [11] E. D. Sontag, “Smooth stabilization implies coprime factorization,” *IEEE Trans. Autom. Control*, vol. 34, no. 4, pp. 435–443, 1989.
- [12] A. J. Taylor and A. D. Ames, “Adaptive safety with control barrier functions,” in *Proc. Amer. Control Conf.*, pp. 1399–1405, 2020.

- [13] B. T. Lopez, J. J. Slotine, and J. P. How, “Robust adaptive control barrier functions: An adaptive and data-driven approach to safety,” *IEEE Contr. Syst. Lett.*, vol. 5, no. 3, pp. 1031–1036, 2021.
- [14] P. Zhao, Y. Mao, C. Tao, N. Hovakimyan, and X. Wang, “Adaptive robust quadratic programs using control lyapunov and barrier functions,” in *Proc. Conf. Decis. Control*, pp. 3353–3358, 2020.
- [15] A. Isaly, O. S. Patil, R. G. Sanfelice, and W. E. Dixon, “Adaptive safety with multiple barrier functions using integral concurrent learning,” in *Proc. Amer. Control Conf.*, pp. 3719 – 3724, 2021.
- [16] M. Black, E. Arabi, and D. Panagou, “A fixed-time stable adaptation law for safety-critical control under parametric uncertainty,” in *Proc. Eur. Control Conf.*, pp. 1328–1333, 2021.
- [17] V. Azimi and S. Hutchinson, “Exponential control lyapunov-barrier function using a filtering-based concurrent learning adaptive approach,” *IEEE Trans. Autom. Control*, 2021.
- [18] M. H. Cohen and C. Belta, “High order robust adaptive control barrier functions and exponentially stabilizing adaptive control lyapunov functions,” in *Proc. Amer. Control Conf.*, pp. 2233–2238, 2022.
- [19] F. Blanchini and S. Miani, *Set-theoretic methods in control*. Springer, 2008.
- [20] S. Sadraddini and C. Belta, “Formal methods for adaptive control of dynamical systems,” in *Proc. Conf. Decis. Control*, pp. 1782–1787, 2017.
- [21] M. H. Cohen, Z. Serlin, K. J. Leahy, and C. Belta, “Temporal logic guided safe model-based reinforcement learning: A hybrid systems approach,” *Nonlinear Analysis: Hybrid Systems*, vol. 47, p. 101295, 2023.
- [22] K. J. Rutledge and N. Ozay, “Correct-by-construction exploration and exploitation for unknown linear systems using bilinear optimization,” in *Proc. Int. Conf. Hybrid Syst. Comp. Control*, 2022.
- [23] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, “Control barrier function based quadratic programs for safety critical systems,” *IEEE Trans. Autom. Control*, vol. 62, no. 8, pp. 3861–3876, 2017.
- [24] Q. Nguyen and K. Sreenath, “Exponential control barrier functions for enforcing high relative-degree safety-critical constraints,” in *Proc. Amer. Control Conf.*, pp. 322–328, 2016.
- [25] W. Xiao and C. Belta, “High order control barrier functions,” *IEEE Trans. Autom. Control*, vol. 67, no. 7, pp. 3655–3662, 2022.
- [26] X. Tan, W. S. Cortez, and D. V. Dimarogonas, “High-order barrier functions: robustness, safety and performance-critical control,” *IEEE Trans. Autom. Control*, vol. 67, no. 6, pp. 3021–3028, 2022.
- [27] J. Breeden and D. Panagou, “High relative degree control barrier functions under input constraints,” in *Proc. Conf. Decis. Control*, pp. 6119–6124, 2021.
- [28] S. Kolathaya and A. D. Ames, “Input-to-state safety with control barrier functions,” *IEEE Contr. Syst. Lett.*, vol. 3, no. 1, pp. 108–113, 2019.
- [29] 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 Contr. Syst. Lett.*, vol. 6, pp. 908–913, 2022.
- [30] A. Alan, A. J. Taylor, C. R. He, A. D. Ames, and G. Orosz, “Control barrier functions and input-to-state safety with application to automated vehicles,” *arXiv preprint arXiv:2206.03568*, 2022.
- [31] S. Kolathaya, J. Reher, A. Hereid, and A. D. Ames, “Input to state stabilizing control lyapunov functions for robust bipedal robotic locomotion,” in *Proc. Amer. Control Conf.*, 2018.
- [32] R. Kamalapurkar, P. Walters, and W. E. Dixon, “Model-based reinforcement learning for approximate optimal regulation,” *Automatica*, vol. 64, pp. 94–104, 2016.
- [33] P. Deptula, Z. I. Bell, F. M. Zegers, R. A. Licitra, and W. E. Dixon, “Approximate optimal influence over an agent through an uncertain interaction dynamic,” *Automatica*, vol. 134, pp. 1–13, 2021.
- [34] M. Jankovic, “Robust control barrier functions for constrained stabilization of nonlinear systems,” *Automatica*, vol. 96, pp. 359–367, 2018.
- [35] Z. Lyu, X. Xu, and Y. Hong, “Small-gain theorem for safety verification under high-relative-degree constraints,” *arXiv preprint arXiv:2204.04376*, 2022.
- [36] B. T. Lopez, J. J. E. Slotine, and J. P. How, “Dynamic tube mpc for nonlinear systems,” in *Proc. Amer. Control Conf.*, pp. 1655–1662, 2019.
- [37] A. C. Wilson, B. Recht, and M. I. Jordan, “A lyapunov analysis of accelerated methods in optimization,” *Journal of Machine Learning Research*, vol. 22, pp. 1–34, 2021.