Adaptive Safety with Multiple Barrier Functions Using Integral Concurrent Learning

Axton Isaly, Omkar Sudhir Patil, Ricardo G. Sanfelice, Warren E. Dixon

Abstract—**This paper presents an approach to guarantee safety for control systems with uncertain nonlinear dynamics. Constraints on the control input induced by multiple barrier function candidates are developed to ensure forward pre-invariance of a safe set of states despite the uncertainty. Using the adaptive control technique of integral concurrent learning, conservativeness in the input constraints is reduced over time as estimates of the uncertain parameters exponentially converge. The constraints are implemented in a quadratic program that modifies a nominal controller for guaranteed safety. An example is presented showing that the operating region of the system is expanded significantly from the initial size due to the converging estimation error.**

I. Introduction

Safety-critical control based on forward invariance of sets is a useful tool for the synthesis of complex control tasks [1]–[3]. Barrier functions (BF) are a popular technique that allow sets of safe states, encoded by continuously differentiable functions, to be rendered invariant by inducing constraints involving the control input. However, such constraints are generally conservative for uncertain systems as a means to ensure robustness [4], [5]. In invariance-based control, conservativeness takes the form of restricting the state to some subset of the true safe set. This paper explores the idea of leveraging recent advances in data-based adaptive control to reduce the conservativeness that is otherwise present, thereby expanding the operating region of the dynamic system.

A benefit of BF-induced input constraints is that the resulting affine conditions can be implemented in a quadratic program (QP) that modifies the original controller only when necessary to maintain safety (i.e., forward invariance of the safe set) [6]. The problem of ensuring safety while compensating for uncertainty using estimates of the unknown parameters was investigated in [7] and [9]. In [7], an adaptive input constraint was developed based on a strong barrier inequality that required forward invariance of level sets of the BF. The authors in [9] relaxed the input constraint to a zeroing BF-type inequality (cf., [10]), which allowed the state to approach the boundary of the safe set while maintaining invariance. However, both of the aforementioned works defined a composite BF that included the parameter estimation error, which lead to a parameter update law featuring the gradient of the BF. The composite BF restricts the state to a new safe set, a subset of the original set, that is dependent on the estimation error. Because the adaptive safety problem does not require the state to be driven to a point where the gradient of the BF is zero, the estimation error generally does not converge. Therefore, as demonstrated in an example in [8], combining the gradient update law with relaxed zeroing BF conditions can cause large estimation error and introduce limit cycles into the closed-loop dynamics. Moreover, the dependency of the update law in the aforementioned works on the gradient of a particular BF makes the techniques challenging to extend to applications where safe sets are expressed using multiple BF candidates. Expressing safe sets using multiple continuously differentiable functions leads to input constraints that are continuous functions of the state, which is generally not the case for approaches that compose multiple BFs into a single locally Lipschitz function [2]. Using results from [11] and [12], we show that the approach developed in this paper is amenable to problems involving multiple BFs.

To compensate for uncertainty without introducing dependencies on any particular BF, a natural approach is to identify the unknown dynamics. Concurrent learning (CL) techniques are a class of data-driven estimators for systems with linearly parameterized uncertainty [13]–[15]. Integral CL (ICL) is a development that eliminates the need to measure the time derivative of the state trajectory [16]. ICL yields a condition based on stored data that can be used to determine the worst-case estimation error online. Subject to a finite-time excitation condition, the estimation error is exponentially regulated to zero. The idea of using an estimate of the unknown dynamics to reduce conservativeness in a BF-induced input constraint was also explored in [17] and [18], where the authors used Gaussian processes (GP). Although GPs apply to broader classes of systems, they cannot guarantee convergence to the true dynamics, and require measurements of the state derivative. Moreover, GPs are computationally expensive, which the aforementioned works addressed by discretizing either the state space or the dynamics. Finally, the GP technique has not been extended to the case of multiple BFs.

In this paper, we leverage a computable upper bound of the estimation error for ICL to develop an implementable set of input constraints based on multiple BFs. After the

Axton Isaly, Omkar Sudhir Patil, and Warren E. Dixon are with the Department of Mechanical and Aerospace Engineering, University of Florida, Gainesville FL 32611-6250, USA. Email: {axtonisaly1013, patilomkarsudhir, wdixon}@ufl.edu

Ricardo G. Sanfelice is with the Department of Electrical and Computer Engineering, University of California, Santa Cruz, USA. Email: ricardo@ucsc.edu

This research is supported in part by Office of Naval Research Grant N00014-13-1-0151; AFOSR award numbers FA9550-19-1-0169, FA9550-19-1-0053, and FA9550-20-1-0238; NSF Grants no. ECS-1710621, CNS-1544396, and CNS-2039054; and by CITRIS and the Banatao Institute at the University of California. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the sponsoring agency.

finite-time excitation condition is reached, conservativeness in the constraints is significantly reduced because both the estimation error and its upper bound are exponentially convergent to zero. To aid the design, we show that ICL can be implemented in a standalone data-driven estimator where the typical state feedback gradient component is eliminated. A key advantage is that we certify safety based on the original BF candidate, rather than forming a composite BF with the estimation error. The resulting safe set from using a composite BF may be empty. We show that the new technique can be implemented in a QP involving multiple constraints to guarantee forward pre-invariance of the set where all BF candidates are nonpositive. Conditions that a priori guarantee feasibility of the QP are presented for a single BF. The developed technique is demonstrated in an example showing that the trajectory remains within the safe set, while the operating region of the system grows significantly relative to its initial size, which is quantified by a two order of magnitude increase in the value of the BF.

Notation: Let $\mathbb{R}_{\geq 0} \triangleq [0,\infty)$, $\mathbb{R}_{>0} \triangleq (0,\infty)$, $\mathbb{Z}^{+}$ denote the set of positive integers, $\mathbb{R}^{n\times m}$ be the space of $n\times m$ dimensional matrices, $0_{n\times m}$ be an $n\times m$ matrix of zeros, and $I_{n\times n}$ be an $n\times n$ identity matrix. The notation $i\in[k]$ is shorthand for $i\in\{1,2,\cdots,k\}$. For a vector function $B:\mathbb{R}^{n}\to\mathbb{R}^{d}$, each component is indexed so that $B(x)\triangleq[B_1(x),B_2(x),\ldots,B_d(x)]^T$. The inequality $B(x)\leq 0$ means that $B_i(x)\leq 0$ for all $i\in[d]$. The Euclidean norm is denoted by $\|\cdot\|$. For a set $A\subset\mathbb{R}^{n}$, ∂A denotes the boundary of A and $U(A)$ denotes some open neighborhood around A.

II. Dynamic Model and Problem Formulation

A. Dynamic Model

Consider the nonlinear dynamic system

$$\dot{x}=Y(x,t)\,\theta+g(x)\,u, \tag{1}$$

with the state $x\in\mathbb{R}^{n}$, control input $u\in\mathbb{R}^{m}$, control effectiveness matrix $g:\mathbb{R}^{n}\to\mathbb{R}^{n\times m}$, a known regression matrix $Y:\mathbb{R}^{n}\times\mathbb{R}_{\geq 0}\to\mathbb{R}^{n\times p}$, and a vector of constant, unknown system parameters $\theta\in\mathbb{R}^{p}$. It is assumed that g is locally Lipschitz, and $Y(x,t)$ is locally Lipschitz in x and continuous in t. The following assumption is imposed on the unknown parameter vector to ensure compensation for the initial estimation error.

Assumption 1. The parameter vector θ takes values from a compact set $\Theta\subset\mathbb{R}^{p}$. There exist real numbers $\bar{\theta}$ and $\bar{\vartheta}$ such that, for every $\theta_1,\theta_2\in\Theta$, $\|\theta_1\|\leq\bar{\theta}$ and $\|\theta_1-\theta_2\|\leq\bar{\vartheta}$.

Given a controller $\kappa:\mathbb{R}^{n}\times\mathbb{R}_{\geq 0}\to\mathbb{R}^{m}$, the closed-loop dynamic system defined by κ is given by

$$\dot{x}=Y(x,t)\,\theta+g(x)\,\kappa(x,t)\,. \tag{2}$$

A solution $\phi:\operatorname{dom}\phi\to\mathbb{R}^{n}$ to the closed-loop system, starting from $\phi_0\in\mathbb{R}^{n}$ at $t=0$, is a locally absolutely continuous function $t\mapsto\phi(t)$ that satisfies (2) for almost all $t\in\operatorname{dom}\phi$, where $\operatorname{dom}\phi\subseteq\mathbb{R}_{\geq 0}$. A solution is said to be *complete* if $\operatorname{dom}\phi$ is unbounded, and it is *maximal* if there is no solution ϕ' such that $\phi(t)=\phi'(t)$ for all $t\in\operatorname{dom}\phi$ with $\operatorname{dom}\phi$ a proper subset of $\operatorname{dom}\phi'$ [12].

B. Problem Formulation

We consider the problem of designing a controller that ensures forward invariance of a set of safe states. We assume that the safe set $\mathcal{S}\subset\mathbb{R}^{n}$ can be defined by the zero-sublevel set of multiple, scalar-valued, continuously differentiable functions, which form a vector function $B:\mathbb{R}^{n}\to\mathbb{R}^{d}$. The function B is called a *BF candidate* [12] if it defines $\mathcal{S}$ as

$$\mathcal{S}\triangleq\{x\in\mathbb{R}^{n}:B(x)\leq 0\}\,, \tag{3}$$

where $B(x)\triangleq[B_1(x),B_2(x),\ldots,B_d(x)]^T$. We also define $\mathcal{S}_i\triangleq\{x\in\mathbb{R}^{n}:B_i(x)\leq 0\}$ for each $i\in[d]$. Note that $\mathcal{S}$ is necessarily closed if B is continuous. The following notions of forward invariance are defined for a set $\mathcal{S}\subset\mathbb{R}^{n}$ according to [12].

Definition 1. The set $\mathcal{S}$ is *forward pre-invariant* for a closed-loop dynamic system if, for each $\phi_0\in\mathcal{S}$ and each maximal solution ϕ starting from ϕ_0, $\phi(t)\in\mathcal{S}$ for all $t\in\operatorname{dom}\phi$.

Definition 2. The set $\mathcal{S}$ is *forward invariant* for a closed-loop dynamic system if it is forward pre-invariant and for each $\phi_0\in\mathcal{S}$, each maximal solution ϕ starting from ϕ_0 is complete.

Due to parametric uncertainty in (1), guaranteeing forward invariance of the safe set typically requires restricting the system to operate in some subset of $\mathcal{S}$. When the true dynamics are known, the state trajectory can be permitted to explore the entire safe set. Consequently, there is motivation to expand the operating region of the system by developing an estimate of the unknown parameter vector θ. Because the estimation error is unknown, an estimator alone cannot guarantee forward invariance. As an alternative to adaptive update laws featuring the gradient of the BF, we seek an estimator that provides some computable indication of the estimation error.

III. Estimator Design

ICL is an adaptive control technique that yields a finite-time excitation condition to verify exponential convergence to the true parameter vector θ [16]. The computable eigenvalues of a regression matrix based on stored data determine the rate of exponential convergence and provide an indication of the worst-case estimation error. Moreover, the ICL strategy does not depend on the gradient of a specific BF candidate. ICL is therefore well suited as an estimator in the problem described in Section II-B.

Let the controller $(x,t)\mapsto\kappa(x,t)$ be continuous in x and t. Integrating (2) along a solution ϕ to the closed-loop dynamic system in (2) yields

$$\phi(t)-\phi(t-\Delta t)=\mathcal{Y}(t)\,\theta+\mathcal{K}(t)\,, \tag{4}$$

for all $t \in \mathrm{dom}\,\phi$ such that $t > \Delta t$, where $\Delta t \in \mathbb{R}_{>0}$ is a user-defined constant determining the size of the window of integration. In (4),

$$\mathcal{Y}(t) \triangleq \int_{t-\Delta t}^{t} Y(\phi(\tau), \tau)\, \mathrm{d}\tau, \tag{5}$$

$$\mathcal{K}(t) \triangleq \int_{t-\Delta t}^{t} g(\phi(\tau))\, \kappa(\phi(\tau), \tau)\, \mathrm{d}\tau. \tag{6}$$

Data in ICL is gathered by sampling $\mathcal{Y}$ and $\mathcal{K}$ at discrete time instants. Consider a sequence of sampling times $\{t_j\}_{j=1}^{N(t)}$ such that $\Delta t < t_1 < \cdots < t_{N(t)} \le t$, where $N(t)$ indicates the time-dependent number of samples. Using the notation $\mathcal{Y}_j \triangleq \mathcal{Y}(t_j)$ and $\mathcal{K}_j \triangleq \mathcal{K}(t_j)$, an ICL update law is defined as

$$\dot{\hat{\theta}} \triangleq k_{CL} \sum_{j=1}^{N(t)} \mathcal{Y}_j^T \left(\phi(t_j) - \phi(t_j - \Delta t) - \mathcal{K}_j - \mathcal{Y}_j \hat{\theta} \right), \tag{7}$$

starting from $\hat{\theta}_0 \in \Theta$, where $k_{CL} \in \mathbb{R}_{>0}$ is the adaptation gain. From (4) and (7), the estimation error evolves according to

$$\dot{\tilde{\theta}} = -k_{CL} \sum_{j=1}^{N(t)} \mathcal{Y}_j^T \mathcal{Y}_j \tilde{\theta}, \tag{8}$$

where $\tilde{\theta} \triangleq \theta - \hat{\theta}$.

Provided the collected data is sufficiently rich, the update law in (7) leads to exponential convergence of the estimation error. Given a solution ϕ to (2), define the function $\lambda_{min} : \mathrm{dom}\,\phi \to \mathbb{R}$ such that $\lambda_{min}(t)$ denotes the minimum eigenvalue of the matrix $\sum_{j=1}^{N(t)} \mathcal{Y}_j^T \mathcal{Y}_j$ at time t. The function λ_{min} is piecewise constant between the sampling times t_j and t_{j+1}, and non-negative, where the latter fact is because a matrix product of the form $\mathcal{Y}_j^T \mathcal{Y}_j$ is symmetric and at least positive semi-definite [19]. The following lemma provides an upper bound of the estimation error that depends on λ_{min}. The minimum eigenvalue can be monitored online by numerically computing $\mathcal{Y}(t)$ using measurements of $\phi(t)$. See [15], [16], [20] for implementations of ICL that ensure λ_{min} is non-decreasing.

Assumption 2. The closed-loop dynamic system is sufficiently excited over a finite duration of time, namely, there exist $\underline{\lambda} \in \mathbb{R}_{>0}$ and $T > \Delta t$ such that, for any complete solution to (2), $\lambda_{min}(t) \ge \underline{\lambda}$ for all $t \ge T$ [16], [21].

Lemma 1. *Let the controller $(x,t) \mapsto \kappa(x,t)$ be continuous, and consider a solution ϕ to the closed-loop system in (2). Suppose that Assumption 1 holds, and let $\hat{\theta}$ be updated according to (7) with $\hat{\theta}_0 \in \Theta$. Then*

$$\left\| \tilde{\theta}(t) \right\| \le \tilde{\theta}_{UB}(t) \tag{9}$$

for all $t \in \mathrm{dom}\,\phi$, where

$$\tilde{\theta}_{UB}(t) \triangleq \bar{\vartheta}\, exp\left(-\int_0^t k_{CL} \lambda_{min}(\tau)\, d\tau \right). \tag{10}$$

Additionally, $\left\|\tilde{\theta}(t)\right\| \le \left\|\tilde{\theta}(0)\right\|$ for all $t \in \mathrm{dom}\,\phi$. Furthermore, if Assumption 2 holds and ϕ is complete, then $\tilde{\theta}$ is exponentially regulated in the sense that $\left\|\tilde{\theta}(t)\right\| \le \left\|\tilde{\theta}(0)\right\| exp\left(-\underline{\lambda}(t-T)\right)$ for all $t \in \mathrm{dom}\,\phi$.

IV. Design of Safety Constraints

We now consider the problem of designing a controller to ensure forward invariance of $\mathcal{S}$ in (3). The goal is to use the bound in (9) to expand the operating region of the system as uncertainty is reduced. To guide the control design, consider the following corollary, which is obtained by specializing Theorem 1 in [12] to the non-hybrid dynamics in (1), and treating time as a state. The corollary provides infinitesimal conditions for forward pre-invariance[1].

Corollary 1. *Consider a continuously differentiable BF candidate $B : \mathbb{R}^n \to \mathbb{R}^d$, and let the controller $(x,t) \mapsto \kappa(x,t)$ be continuous. The closed set $\mathcal{S}$ in (3) is forward pre-invariant for the closed-loop dynamic system in (2) if, for all $i \in [d]$,*

$$\nabla B_i^T(x)\left(Y(x,t)\theta + g(x)\kappa(x,t)\right) \le 0 \tag{11}$$

for all $x \in U(M_i) \backslash \mathcal{S}_i$ and $t \in \mathbb{R}_{\ge 0}$, where $U(M_i)$ is any neighborhood around the set $M_i \triangleq \{x \in \partial\mathcal{S} : B_i(x) = 0\}$.

We will enforce (11) by designing a constraint on the control input. Such constraints are commonly enforced using QPs. The constraint is developed with a continuous function that satisfies the following condition.

(C1) The open set $\mathcal{D} \subset \mathbb{R}^n$ and continuous function $\gamma : \mathbb{R}^n \to \mathbb{R}^d$ are such that $\mathcal{S} \subset \mathcal{D}$ and, for each $i \in [d]$, $\gamma_i(x) \ge 0$ for all $x \in U(M_i) \backslash \mathcal{S}_i$.

Note that choosing $\gamma_i(x) \ge 0$ for all $x \in \mathcal{D} \backslash \mathcal{S}_i$ is sufficient for the condition on γ in (C1) to hold. Given γ and $\mathcal{D}$ satisfying (C1), the inequality in (11) can be enforced by constraining the control input. For any $x \in \mathcal{D}$ and $t \in \mathbb{R}_{\ge 0}$, a safe input $u \in \mathbb{R}^m$ must satisfy

$$\nabla B_i^T(x)\left(Y(x,t)\theta + g(x)u\right) \le -\gamma_i(x), \tag{12}$$

for all $i \in [d]$.

Remark 1. A common choice for the function γ that satisfies (C1) on any open set $\mathcal{D} \supset \mathcal{S}$ is $\gamma_i(x) \triangleq \alpha(B_i(x))$, where α is an extended class $\mathcal{K}$ function [10]. The aforementioned zeroing BF selection was shown in [10] to yield asymptotic stability of the set $\mathcal{S}$ for the case of a scalar BF, but is in some cases stronger than required for forward invariance [12].

Since θ is unknown, (12) cannot be implemented directly in a QP. To account for uncertainty, an upper bound for the unknown term $\nabla B_i^T(x) Y(x,t)\theta$ will be used to ensure that (12) holds. As an illustrative example, we first consider a robust control approach to the problem. Using Holder's

[1] See Proposition 2 in [12] for conditions under which a forward pre-invariant set is forward invariant.

inequality, satisfaction of the following constraint implies that (12) is also satisfied:

$$\left\|\nabla B_i^T(x) Y(x,t)\right\| \bar{\theta} + \nabla B_i^T(x) g(x) u \leq -\gamma_i(x), \quad (13)$$

where $\bar{\theta}$ was defined in Assumption 1. Constraint (13) depends on known information and can therefore be implemented in a QP. However, the robust approach introduces conservativeness in the sense that the set of control inputs that satisfy (13) is generally a subset of those that satisfy (12).

Using ICL, conservativeness can be reduced as data about the system is collected. An adaptive control approach is developed by first noticing that, using the definition of $\tilde{\theta}$,

$$\nabla B_i^T(x) Y(x,t)\theta = \nabla B_i^T(x)\left(Y(x,t)\tilde{\theta} + Y(x,t)\hat{\theta}\right). \quad (14)$$

Again using Holder's inequality with the bound from Lemma 1, the following holds whenever $\tilde{\theta}_{UB}$ is defined along a solution ϕ to (2),

$$\nabla B_i^T(x) Y(x,t)\tilde{\theta}(t) \leq \left\|\nabla B_i^T(x) Y(x,t)\right\| \tilde{\theta}_{UB}(t), \quad (15)$$

for all $x \in \mathbb{R}^n$ and $t \in \operatorname{dom}\phi$. Based on (14) and (15), and recalling that $\tilde{\theta}_{UB}$ can be computed online from known information, an implementable safety constraint is designed for each $i \in [d]$ as

$$\theta_{con,i}(x,t) + \nabla B_i^T(x) g(x) u \leq -\gamma_i(x), \quad (16)$$

where

$$\theta_{con,i}(x,t) \triangleq \min\Big\{\left\|\nabla B_i^T(x) Y(x,t)\right\| \bar{\theta}, \quad (17)$$
$$\nabla B_i^T(x) Y(x,t)\hat{\theta}(t) + \left\|\nabla B_i^T(x) Y(x,t)\right\| \tilde{\theta}_{UB}(t)\Big\}.$$

Because $\nabla B_i^T(x) Y(x,t)\theta \leq \theta_{con,i}(x,t)$ for all $x \in \mathbb{R}^n$ and $t \in \operatorname{dom}\phi$, any input $u \in \mathbb{R}^m$ that satisfies (16) also satisfies (12). As we show in Section V, enforcing the constraints defined in (16) ensures that $\mathcal{S}$ is forward pre-invariant. Moreover, under the assumptions of Lemma 1, the upper bound $\tilde{\theta}_{UB}(t)$ converges to zero exponentially, and $\hat{\theta}(t)$ converges to θ as $t \to \infty$. The result is that, in contrast to the robust control approach, the upper bound $\theta_{con,i}(x,t)$ becomes a close approximation of the unknown term $\nabla B_i^T(x) Y(x,t)\theta$ over time. Therefore, the conservativeness that is introduced into $\theta_{con,i}(x,t)$ to ensure safety is significantly reduced and the control input can be selected from a less restrictive set.

Remark 2. The robust upper bound from (13) is used in (17) since that bound may be smaller at the initial time, until sufficient data has been collected to reduce the upper bound $\tilde{\theta}_{UB}(t)$. As noted in ([7], Remark 4), a semi-robust approach, in addition to assuming that the parameter vector θ is bounded (Assumption 1), appears to be necessary in the context of invariance control since invariance must be guaranteed for a specific set.

A. Feasibility

The constraint-based approach described above is contingent on the constraints being feasible. Guaranteeing feasibility a priori for problems involving multiple BF candidates is challenging in general and beyond the scope of this work. The authors in [22] explored continuity of nonlinear programs, but required the assumption that a certain set-valued mapping describing the feasible set of control inputs was nonempty. In other works, feasibility is ensured by assuming the existence of a control BF [3]. The following condition will ensure that the developed controller is feasible regardless of the initial state.

(C2) (Feasibility) For each $x \in \mathcal{D}$ and $t \in \mathbb{R}_{\geq 0}$, there exists $u \in \mathbb{R}^m$ such that, for all $i \in [d]$,

$$\nabla B_i^T(x) g(x) u \leq -\gamma_i(x) - \left\|\nabla B_i^T(x) Y(x,t)\right\| \bar{\theta}. \quad (18)$$

In the simpler case of a scalar BF candidate, it is possible to guarantee condition (C2) a priori by checking only the points where the gradient (with respect to u) of the constraint has rank zero. The following proposition shows this fact and, that by requiring a strict inequity to hold, one can guarantee the control input $u = 0_{m\times 1}$ satisfies (18) in some neighborhood of any rank-zero point, implying that a continuous and bounded control input can be selected in that neighborhood.

Proposition 1. *Let $b : \mathbb{R}^n \to \mathbb{R}$ be a continuously differentiable, scalar BF candidate, and let the function γ and the set $\mathcal{D}$ satisfy (C1). Define the set $G \triangleq \{x \in \mathbb{R}^n : \nabla b^T(x) g(x) = 0_{1\times m}\}$. Suppose that $x^* \in G$ implies that*

$$\gamma(x^*) + \left\|\nabla b^T(x^*) Y(x^*,t)\right\| \bar{\theta} < 0, \quad (19)$$

for all $t \in \mathbb{R}_{\geq 0}$. Then (C2) is satisfied with b in place of B_i. In particular, for any $x^* \in G$ and $t^* \in \mathbb{R}_{\geq 0}$, there exists a neighborhood U around (x^*, t^*) such that the vector $u = 0_{m\times 1}$ satisfies (18) for all $(x,t) \in U$.

V. Stability Analysis

The following theorem shows that a closed-loop controller subjected to the constraints designed in Section IV enforces forward pre-invariance of the safe set $\mathcal{S} = \{x \in \mathbb{R}^n : B(x) \leq 0\}$, given that the constraints are feasible in a neighborhood of the safe set. It is notable that the finite excitation condition of Assumption 2 is not necessary to prove the theorem. Meeting the finite excitation condition does, however, lead to reduced conservativeness because the upper bound $\theta_{con,i}(x,t)$ becomes a close approximation of $\nabla B_i^T(x) Y(x,t)\theta$ when $\tilde{\theta}_{UB}(t)$ is small.

Theorem 1. *Let $B : \mathbb{R}^n \to \mathbb{R}^d$ be a continuously differentiable BF candidate, and suppose that Assumption 1 holds. Let the function γ and the set $\mathcal{D}$ satisfy (C1) and (C2). Along any solution to (2), let the parameter estimate $\hat{\theta}$ be updated* *according to (7) with $\hat{\theta}_0 \in \Theta$, define $\tilde{\theta}_{UB}$ according to (10), and let κ^* be a control law generated by the following QP:*

$$\kappa^*(x,t) \triangleq \underset{u\in\mathbb{R}^m}{\arg\min} \ \|u-\kappa_{nom}(x,t)\|^2 \tag{20}$$
$$s.t.\ \nabla B_i^T(x)\, g(x) u \le -\gamma_i(x) - \theta_{con,i}(x,t),\ \forall i \in [d],$$

where $\theta_{con,i}$ is defined for each $i \in [d]$ as in (17), and $\kappa_{nom}: \mathbb{R}^n \times \mathbb{R}_{\ge 0} \to \mathbb{R}^m$ is a nominal controller. If $(x,t) \mapsto \kappa^*(x,t)$ is continuous, then the set $\mathcal{S}$ is forward pre-invariant for the closed-loop dynamics defined by κ^*.

Proof: Let ϕ be a solution to the closed-loop system with $\phi_0 \in \mathcal{S}$. It will be shown that $\phi(t) \in \mathcal{S}$ for all $t \in \operatorname{dom}\phi$. Condition (C2) is sufficient to conclude that the QP in (20) has a solution for all $x \in \mathcal{D}$ since $\theta_{con,i}(x,t) \le \left\|\nabla B_i^T(x)\, Y(x,t)\right\| \bar{\theta}$. Because the controller κ^* satisfies the constraints in (20), it follows that for each $i \in [d]$, the controller ensures that[2]

$$\dot{B}_i(\phi) \le \nabla B_i^T(\phi)\, Y(\phi,t)\,\theta - \gamma_i(\phi) - \theta_{con,i}(\phi,t) \tag{21}$$

for all $t \in \operatorname{dom}\phi$ such that $\phi(t) \in \mathcal{D}$. From (13)-(15), $\nabla B_i^T(\phi)\, Y(\phi,t)\,\theta \le \theta_{con,i}(\phi,t)$. Thus,

$$\dot{B}_i(\phi) \le -\gamma_i(\phi). \tag{22}$$

Because $\mathcal{D}$ is open and contains the closed set $\mathcal{S}$, $\mathcal{D}$ contains a neighborhood of $\partial\mathcal{S}$. Therefore, from the definition of M_i, there are neighborhoods such that $U(M_i) \subset U(\partial\mathcal{S}) \subset \mathcal{D}$. By condition (C1), it follows that, for any $i \in [d]$, if $\phi(t) \in U(M_i) \setminus \mathcal{S}_i$, then $\gamma(\phi) \ge 0$, and

$$\dot{B}_i(\phi) \le 0. \tag{23}$$

The proof of Theorem 1 in [12] shows that a contradiction with (23) is obtained by assuming that there exists $t_2 \in \operatorname{dom}\phi$ such that $\phi(t_2) \notin \mathcal{S}$. We conclude that $\phi(t) \in \mathcal{S}$ for all $t \in \operatorname{dom}\phi$. Since the aforementioned properties hold for any solution starting from $\mathcal{S}$, we conclude that $\mathcal{S}$ is forward pre-invariant. ■

VI. Example

The following example system, taken from [16], will be used to illustrate the reduced conservativeness provided by the ICL estimator in (7) and the controller κ^* in (20).

$$\dot{x} = \begin{bmatrix} x_1^2 & \sin(x_2) & 0 & 0 \\ 0 & x_2 \sin(t) & x_1 & x_1 x_2 \end{bmatrix} \theta + u, \tag{24}$$

where $x \in \mathbb{R}^2$ and $u \in \mathbb{R}^2$. The unknown parameters were selected as $\theta = \begin{bmatrix} 5 & 10 & 15 & 20 \end{bmatrix}^T$. The controller developed in [16] was used as a nominal tracking controller,

$$\kappa_{nom}(x,t) = \dot{x}_d(t) - Y(x,t)\,\theta - Ke, \tag{25}$$

where $e \triangleq x - x_d(t)$. The desired trajectory is $x_d(t) = r(t)\,[2\cos(t), \sin(t)]^T$, where $r(t) = 0.27t$. To better demonstrate how the adaptive safety constraints in the QP in (20) affect tracking performance, the nominal controller

[2]When $\phi(t)$ appears as the argument of a function, we subsequently omit the dependency on t, for brevity.

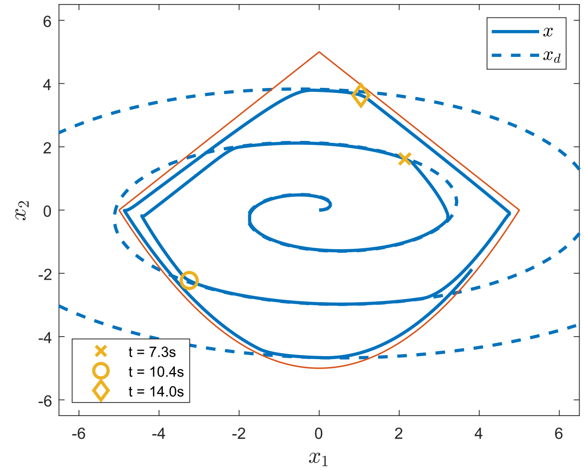

Figure 1. Evolution of the state trajectory for the system in (24). The thin orange line indicates the boundary of the safe set $\mathcal{S}$. The QP in (20) maintains the state inside the safe set although the desired trajectory leaves the set. For visual clarity, the figure shows only the first 18 seconds of the simulation.

is given access to the true parameter values. It is emphasized that the constraints of the QP, namely the functions $\theta_{con,i}$, do not use the true parameter values. The BF is defined as

$$B(x) = \begin{bmatrix} x_1 + x_2 - c \\ -x_1 + x_2 - c \\ \frac{1}{c}x_1^2 - x_2 - c \end{bmatrix}, \tag{26}$$

where $c = 5$. The safe set $\mathcal{S}$ defined by B is displayed in Figure 1 as a thin orange line. The function γ was selected as $\gamma_i(x) \triangleq K_b (B_i(x))^3$, where $K_b \in \mathbb{R}_{>0}$. A cubic function is used to obtain faster growth in the interior of the safe set and slower growth near the boundary. The maximum magnitude and estimation error of the parameters were set to $\bar{\theta} = 43$ and $\bar{\vartheta} = 22$, respectively.

The simulation was implemented in Simulink® using ode45, while the quadprog function was used to implement the controller in (20) with objective function $u^T u - 2\kappa_{nom}^T(x,t)\,u$. Figure 1 shows a phase portrait of the desired and actual state trajectory. Specific times of interest are denoted by symbols in Figure 1 and marked on the horizontal axes of the plots in Figure 2. The forward pre-invariance of $\mathcal{S}$ under κ^* guaranteed by Theorem 1 is validated in Figure 2c by the fact that the BF in (26) remains negative for the entire simulation. There are many instances when the desired trajectory lies outside the safe set, yet the QP maintains the actual trajectory inside the set. Figure 2a shows that deviations from the nominal controller occur to keep the trajectory in the safe set. For example, there are large deviations after $t = 14.0s$ when the desired trajectory leaves the safe set. As predicted by Lemma 1, the estimation error in Figure 2b converges using the modified ICL update law in (7), despite the fact that the tracking error necessarily grows to ensure safety. In contrast, the update law in [16], which used the tracking error in a gradient term, would not achieve parameter convergence for the adaptive safety problem.

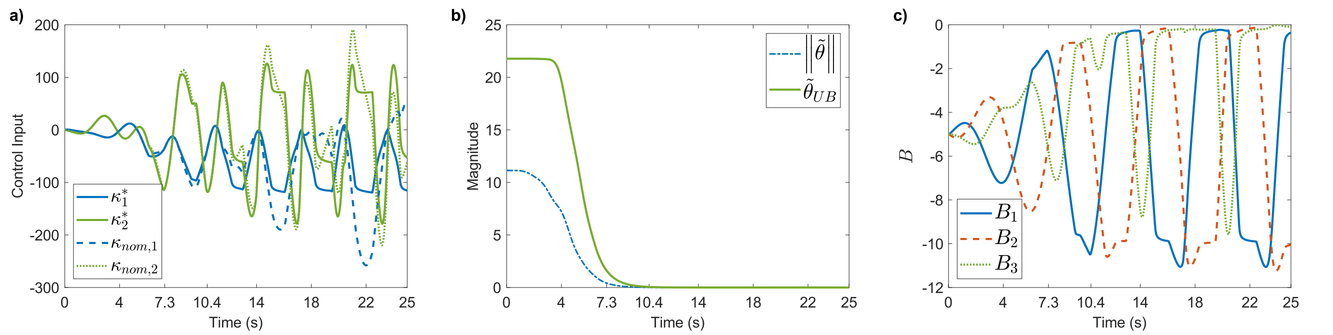

Figure 2. Evolution of the a) QP-modified control input $\kappa^* = \left(\kappa_1^*, \kappa_2^*\right)$ and nominal control input $\kappa_{nom} = (\kappa_{nom,1}, \kappa_{nom,2})$, b) norm of the parameter estimation error vector $\left\|\tilde{\theta}\left(t\right)\right\|$ and its upper bound $\tilde{\theta}_{UB}\left(t\right)$, and c) the BF candidate in (26). Deviations from the nominal control input in a) are necessary to enforce the safety constraint. In b), the estimation error and its upper bound $\tilde{\theta}_{UB}$ both decay exponentially. The barrier function in c) remains negative for the entire simulation, indicating safety.

The example demonstrates how the adaptive safety constraint in (20) leads to an expanded operating region for the system. More concretely, the value of the barrier function B can be interpreted as a signed distance from the boundary of the safe set. The increasing trend of the peaks of B in Figure 2c towards zero indicates that the trajectory is allowed closer to the boundary of the safe set. Before $t = 7.3s$, the maximum value of the BF was $B_1 = -1.18$ while the maximum value for $t \geq 10.4s$ was $B_2 = -0.014$. The increase in B is due to the exponential decay of $\tilde{\theta}_{UB}$ in Figure 2b. The function $\tilde{\theta}_{UB}$ introduces conservativeness into the system to ensure that forward pre-invariance is guaranteed despite uncertainty. Figure 2b shows that conservativeness has been largely eliminated from the system after $t = 10.4s$, and additionally validates that $\tilde{\theta}_{UB}\left(t\right)$ upper bounds $\left\|\tilde{\theta}\left(t\right)\right\|$.

VII. Conclusion

A barrier function approach was developed to ensure safety of an uncertain nonlinear system. An ICL adaptation law was used to verify exponential convergence to the true parameter vector. A real-time metric of the estimation performance was used in a set of safety-ensuring input constraints to allow the state to approach the boundary of the safe set. Since the adaptation law is independent of the BF used to describe the safe set, the approach is readily applicable to problems with multiple barrier functions.

References

[1] W. S. Cortez, D. Oetomo, C. Manzie, and P. Choong, "Control barrier functions for mechanical systems: Theory and application to robotic grasping," *IEEE Trans. Control. Syst. Technol.*, 2019.

[2] P. Glotfelter, J. Cortés, and M. Egerstedt, "Nonsmooth barrier functions with applications to multi-robot systems," *IEEE Control Syst. Lett.*, vol. 1, no. 2, pp. 310–315, 2017.

[3] 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, 2016.

[4] T. Gurriet, A. Singletary, J. Reher, L. Ciarletta, E. Feron, and A. Ames, "Towards a framework for realizable safety critical control through active set invariance," in *Proc ACM/IEEE Int. Conf. Cyber Phys. Syst.* IEEE, 2018, pp. 98–106.

[5] Y. Emam, P. Glotfelter, and M. Egerstedt, "Robust barrier functions for a fully autonomous, remotely accessible swarm-robotics testbed," in *Proc. IEEE Conf. Decis. Control*. IEEE, 2019, pp. 3984–3990.

[6] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: Theory and applications," in *Proc. Eur. Control Conf.* IEEE, 2019, pp. 3420–3431.

[7] A. J. Taylor and A. D. Ames, "Adaptive safety with control barrier functions," in *Proc. Am. Control Conf.* IEEE, 2020, pp. 1399–1405.

[8] ——, "Adaptive safety with control barrier functions," *arXiv preprint arXiv:1910.00555*, 2019.

[9] B. T. Lopez, J.-J. E. Slotine, and J. P. How, "Robust adaptive control barrier functions: An adaptive and data-driven approach to safety," *IEEE Control Syst. Lett.*, vol. 5, no. 3, pp. 1031–1036, 2020.

[10] X. Xu, P. Tabuada, J. W. Grizzle, and A. D. Ames, "Robustness of control barrier functions for safety critical control," *IFAC-PapersOnLine*, vol. 48, no. 27, pp. 54–61, 2015.

[11] M. Maghenem and R. G. Sanfelice, "Multiple barrier function certificates for forward invariance in hybrid inclusions," in *Proc. Am. Control Conf.* IEEE, 2019, pp. 2346–2351.

[12] ——, "Sufficient conditions for forward invariance and contractivity in hybrid inclusions using barrier functions," *Automatica*, 2021.

[13] G. Chowdhary and E. Johnson, "Concurrent learning for convergence in adaptive control without persistency of excitation," in *Proc. IEEE Conf. Decis. Control*, 2010, pp. 3674–3679.

[14] G. Chowdhary, T. Yucelen, M. Mühlegg, 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.

[15] 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, July 2017.

[16] 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, Dec. 2019.

[17] L. Wang, E. A. Theodorou, and M. Egerstedt, "Safe learning of quadrotor dynamics using barrier certificates," in *Proc. IEEE Int. Conf. Robot. Autom.* IEEE, 2018, pp. 2460–2465.

[18] R. Cheng, G. Orosz, R. M. Murray, and J. W. Burdick, "End-to-end safe reinforcement learning through barrier functions for safety-critical continuous control tasks," in *Proc. AAAI Conf. Artif. Intell.*, vol. 33, 2019, pp. 3387–3395.

[19] J. E. Gentle, *Matrix Algebra*. Springer, 2007, vol. 10.

[20] G. Chowdhary and E. Johnson, "A singular value maximizing data recording algorithm for concurrent learning," in *Proc. Am. Control Conf.*, 2011, pp. 3547–3552.

[21] A. Parikh, R. Kamalapurkar, and W. E. Dixon, "Target tracking in the presence of intermittent measurements via motion model learning," *IEEE Trans. Robot.*, vol. 34, no. 3, pp. 805–819, 2018.

[22] B. J. Morris, M. J. Powell, and A. D. Ames, "Continuity and smoothness properties of nonlinear optimization-based feedback controllers," in *Proc. IEEE Conf. Decis. Control*. IEEE, 2015, pp. 151–158.