

Safe Multi-Agent Interaction through Robust Control Barrier Functions with Learned Uncertainties

Richard Cheng, Mohammad Javad Khojasteh, Aaron D. Ames, and Joel W. Burdick

Abstract—Robots operating in real world settings must navigate and maintain safety while interacting with many heterogeneous agents and obstacles. Multi-Agent Control Barrier Functions (CBF) have emerged as a computationally efficient tool to guarantee safety in multi-agent environments, but they assume perfect knowledge of both the robot dynamics and other agents’ dynamics. While knowledge of the robot’s dynamics might be reasonably well known, the heterogeneity of agents in real-world environments means there will *always* be considerable uncertainty in our prediction of other agents’ dynamics. This work aims to learn high-confidence bounds for these dynamic uncertainties using Matrix-Variate Gaussian Process models, and incorporates them into a robust multi-agent CBF framework. We transform the resulting min-max robust CBF into a quadratic program, which can be efficiently solved in real time. We verify via simulation results that the nominal multi-agent CBF is often violated during agent interactions, whereas our robust formulation maintains safety with a much higher probability and adapts to learned uncertainties.

I. INTRODUCTION

Collision-free robot navigation in natural multi-agent environments is vital for a myriad of robotic applications, such as self-driving cars, navigation in crowds, etc. However, placing robots in rapidly evolving, uncertain environments introduces many challenges in guaranteeing safety [1]–[5]. Uncertainty in the prediction of other agents’ trajectories is inevitable, and robots should learn and account for this uncertainty to ensure safe operation.

In this work, we utilize the Multi-Agent Control Barrier Function (CBF) proposed in [6] to generate low-level controllers that guarantee collision-free behavior. The Multi-Agent CBF uses an optimization-based controller to prevent the robot from entering unsafe sets (i.e. states leading to inevitable collision) in a *minimally invasive* fashion. However, previous work either assumes that the robot dynamics and the other agents’ dynamics are perfectly known, or considers highly conservative worst-case bounds [7], [8]. It is clear though that all natural environments (e.g. settings with humans) are fraught with varying levels of uncertainty. For example, human trajectories remain notoriously difficult to predict and are highly stochastic [9]. Thus, both capturing and accounting for uncertainty is crucial for safe navigation.

Statement of contributions: The goal of this paper is to (1) learn uncertainty bounds *online* from agents’ observed trajectories, and (2) incorporate those uncertainty bounds into

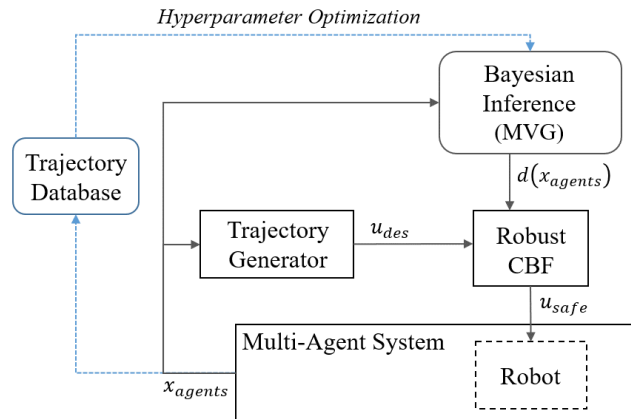


Fig. 1: Diagram overviewing the control structure. Our approach guarantees safety by utilizing a Bayesian Inference Module to learn dynamic uncertainties, and handles them with our proposed Robust CBF module.

the Multi-Agent CBF while maintaining computational efficiency of the underlying controller (i.e. a quadratic program).

Our approach focuses first on learning high-confidence polytopic bounds on the, *possibly coupled*, uncertainties in both the robot dynamics and other agents’ dynamics. To achieve this, we utilize Matrix-Variate Gaussian Processes (MVG) and optimize their hyperparameters *offline* from interaction data; this allows us to predict ellipsoidal uncertainty in our dynamics *online*, which we convert to an uncertainty polytope given a desired confidence level. Using these polytopic bounds, we formulate a robust CBF as a min-max optimization problem over the robot controls and the potential uncertainties, respectively. We then transform this min-max problem into a quadratic program that can be efficiently solved to find a safe control action that is robust with respect to our estimated uncertainty. See Fig. 1 for an overview of our approach.

Organization: Sec. II goes over related work in the safe collision avoidance literature, Sec. III provides background information on CBFs used to guarantee safety, and Matrix-Variate Gaussian Processes used to estimate correlated uncertainties. Sect. IV introduces the robust multi-agent CBF and shows how it can be efficiently solved under given uncertainty bounds. Sec. V looks at how to learn these confidence bounds over the uncertainty. Finally, Sec. VI presents simulation results illustrating the benefits of our algorithm and verifying the safety of our controller.

R. Cheng, A. D. Ames, and J. W. Burdick are with the Department of Mechanical and Civil Engineering, California Institute of Technology. M. J. Khojasteh is with the Department of Electrical Engineering, California Institute of Technology. {rcheng,mjkhogas,ames,jwb}@caltech.edu




Fig. 2: Sample path of a multi-agent system based on the nominal CBF (cf. [6]) and our proposed Robust CBF. The robot (blue) tries to navigate from a start position to random goal position while avoiding collisions with other agents (red). Approximately half of the other agents blindly travel towards their own randomly chosen goal, while the rest exhibit varying degrees of collision-avoidance behavior (the robot does not know their behavior a priori). For more details and results of simulations, see Sec. VI-A. (a) Initial robot/environment configuration, (b) Intermediate configuration, (c) Intermediate configuration showing that the nominal CBF controller experiences collision (*top*), while the robust CBF avoids collision (*bottom*). (d) Final configuration before robot reaches its goal position (star). See <https://youtu.be/hXg5kZ086Lw> for the simulation videos.

II. RELATED WORK

Multi-agent collision avoidance has been a long-studied problem with different approaches proposed for enabling safe control in varying situations. Velocity obstacles is a popular approach that involves limiting control actions to a set of “safe” actions, though its constant velocity with linear dynamics assumption is limiting [10], [11]. Related works in this direction have loosened these assumptions, but require significant sampling of the action space and do not incorporate dynamic uncertainty [12], [13]. More recently, Buffered Voronoi Cells (BVC) have been proposed as a tool to provide safety guarantees with only positional information, though safety guarantees are provided only under linear dynamics and without uncertainty [14], [15]. Also, [16]–[18] provide safety guarantees, under worst case disturbances, by solving the Hamilton-Jacobi-Isaacs equation to obtain a minimally invasive control law. However, the heavy computational expense prohibits applicability to large-scale multi-agent systems.

Reinforcement learning methods have emerged recently, which directly learn actions in multi-agent settings in response to the observed environment [19], [20]. However, these methods provide no formal guarantees of safety, and as such are prone to collision in novel environments. Furthermore, they have not been shown to scale to settings with many agents. To ensure safety in the reinforcement learning settings, the work [21] combined a safe backup policy with the learned policies. [22] incorporated safety constraints into the reinforcement learning framework, although this work considered only *decoupled* uncertainties and was limited to polytopic safety constraints. Bayesian inference was utilized in [23]–[25] to learn system dynamics in an online manner

while ensuring safety (with high probability) using Control Barrier Functions for a single agent.

Control Barrier Functions (CBF) are a tool for enforcing set invariance of dynamical systems [26]–[28], and they have been used to guarantee collision avoidance in multiagent settings by projecting desired actions, to the closest (in least-squares sense) safe actions according to a CBF condition [29]–[31]. However, defining a valid CBF for general systems remains a challenge, especially in cases with significant uncertainty. Recent works have looked at learning CBFs for general systems, either implicitly or explicitly [32], [33]. A multi-agent CBF was defined explicitly for multi-agent systems in the case of continuous-time linear dynamics [6], [7]. However, this multi-agent CBF does not incorporate uncertainty or nonlinearity in the dynamics. Other work on robust CBFs deals only with highly conservative worst-case bounds [8]. We seek to address this gap in this paper.

Our work builds upon the current literature by computing the minimally invasive action necessary to maintain safety in a computationally efficient manner *for discrete-time systems while incorporating nonlinear dynamics and learned uncertainty*.

III. BACKGROUND

Our robotic system is represented by nonlinear control-affine dynamics in discrete-time:

$$x_{t+1} = \begin{bmatrix} p_{t+1} \\ v_{t+1} \\ z_{t+1} \end{bmatrix} = \underbrace{\begin{bmatrix} f_p(x_t) \\ f_v(x_t) \\ f_z(x_t) \end{bmatrix}}_{f(x_t)} + \underbrace{\begin{bmatrix} g_p(x_t) \\ g_v(x_t) \\ g_z(x_t) \end{bmatrix}}_{g(x_t)} u + \underbrace{\begin{bmatrix} d_p(x_t) \\ d_v(x_t) \\ d_z(x_t) \end{bmatrix}}_{d(x_t)}, \quad (1)$$

where $p \in \mathbb{R}^2$, $v \in \mathbb{R}^2$, and $z \in \mathbb{R}^{n-4}$ denote position, velocity and other states, respectively. Here, f_j , g_j , and d_j are real-valued functions, for $j \in \{p, v, z\}$, and $u \in \mathcal{U}$ where $\mathcal{U} := \{u \in \mathbb{R}^m : \|u\|_2 \leq u_{max}\}$. The functions $f(x)$ and $g(x)$ are assumed to be *known*, whereas $d(x)$ represents *unknown* uncertainty in the dynamics, which we model with a Gaussian process [cf. Sec. III-B]. We assume that this system has relative degree 2 with respect to the positional output p ; in discrete time, this directly implies that $g_p(x) = \mathbf{0}_{2 \times 2}$. Similarly, let us represent the other agents within our multi-agent system with dynamics,

$$x_{t+1}^{(i)} = \begin{bmatrix} p_{t+1}^{(i)} \\ v_{t+1}^{(i)} \\ z_{t+1}^{(i)} \end{bmatrix} = \underbrace{\begin{bmatrix} f_p^{(i)}(x_t) \\ f_v^{(i)}(x_t) \\ f_z^{(i)}(x_t) \end{bmatrix}}_{f^{(i)}(x_t)} + \underbrace{\begin{bmatrix} d_p^{(i)}(x_t) \\ d_v^{(i)}(x_t) \\ d_z^{(i)}(x_t) \end{bmatrix}}_{d^{(i)}(x_t)}, \quad (2)$$

where $i \in \mathbb{N}$ indexes each of the other agents in our system. We assume the control input for other agents are a (uncertain) function of their state at the given time, so we do not show control inputs explicitly in (2).

As our robot is interacting with other unknown agents, it will be important for us to account for the uncertainties, $d, d^{(i)}$, when considering safety via CBFs.

A. Control Barrier Functions

Consider a safe set, \mathcal{C} , defined by the super-level set of a continuously differentiable function $h : \mathbb{R}^n \rightarrow \mathbb{R}$:

$$\mathcal{C} := \{x \in \mathbb{R}^n : h(x) \geq 0\}. \quad (3)$$

To maintain safety during the learning process, the system state must always remain within the safe set \mathcal{C} (i.e. the set \mathcal{C} is *forward invariant* with respect to the system dynamics). A set \mathcal{C} is forward invariant if for every $x_0 \in \mathcal{C}$, $x(t) \in \mathcal{C}$ for all $t \geq 0$. In our multi-agent setting, this set could include all states where collision can be avoided given the robot's input bounds. *Control barrier functions* utilize a Lyapunov-like argument to provide a sufficient condition for ensuring forward invariance of the safe set \mathcal{C} under controlled dynamics.

Definition 1: Given a set $\mathcal{C} \in \mathbb{R}^n$ defined by (3), the continuously differentiable function $h : \mathbb{R}^n \rightarrow \mathbb{R}$ is a *discrete-time control barrier function* (CBF) for dynamical system (1) if there exists $\eta \in [0, 1]$ such that for all $x_t \in \mathcal{C}$,

$$\sup_{u_t \in \mathcal{U}} \text{CBC}(x_t, u_t) \geq 0, \quad (4)$$

for $\text{CBC}(x_t, u_t) := h(x_{t+1}(u_t)) + (\eta - 1)h_t(x_t)$, where CBC stands for control barrier condition.

If a function $h(x)$ is a CBF, then there exists a controller such that the set \mathcal{C} is forward invariant [27], [34]. In other words, system safety is guaranteed by ensuring satisfaction of condition (4). Our goal is to ensure safety, by computing a minimally invasive control action that satisfies (4), in an online fashion. In particular, we utilize the following multi-agent CBF inspired by [6]:

$$h(x) := \frac{\Delta p^T \Delta v}{\|\Delta p\|} + \sqrt{a_{max}(\|\Delta p\| - D_s)}, \quad (5)$$

where a_{max} represents our robot's max acceleration in the collision direction, D_s is the collision margin, $\Delta p = p - p^{(i)}$ is the positional difference between the agents, and $\Delta v = v - v^{(i)}$ is the velocity difference between the agents.

The work [6] introduced a CBF similar to (5) for *continuous-time linear* systems, such that a_{max} can be determined easily. In Sec. IV, under an appropriate assumption, we show that (5) is a valid CBF for the discrete-time nonlinear dynamics (1) and (2).

B. Matrix-Variate Gaussian Process Regression

Here we illustrate how Bayesian learning can be used to acquire a distribution over the uncertainty in the dynamics. Since we are estimating a multivariate uncertainty $d(x)$, we must consider potential correlations in its components. Thus, we use the Matrix-Variate Gaussian Process (MVG) model to learn the system dynamics and uncertainty from data. By learning $\mu_d(x)$ and $\Sigma_d(x)$ in tandem with the controller, we can obtain high probability confidence intervals on the unknown dynamics, which adapt/shrink as we obtain more information (i.e. measurements) on the system. We first start by defining the MVG distribution [24], [35]–[37] as follows.

Definition 2: We say the random matrix $\mathbf{X} \in \mathbb{R}^{n \times m}$ is distributed according to a MVG distribution when its probability density function is defined as:

$$p(\mathbf{X}; \mathbf{M}, \Sigma, \Omega) := \frac{\exp\left(-\frac{1}{2} \text{tr}\left[\Omega^{-1}(\mathbf{X} - \mathbf{M})^T \Sigma^{-1}(\mathbf{X} - \mathbf{M})\right]\right)}{(2\pi)^{nm/2} \det(\Sigma)^{m/2} \det(\Omega)^{n/2}}, \quad (6)$$

where $M \in \mathbb{R}^{n \times m}$ denotes the mean, and $\Sigma \in \mathbb{R}^{n \times n}$ encodes the covariance matrix of the rows, and $\Omega \in \mathbb{R}^{m \times m}$ encodes the covariance matrix of the columns. In this case, we write $\mathbf{X} \sim \mathcal{M}\mathcal{N}(\mathbf{M}, \Sigma, \Omega)$, and we have $\text{vec}(X) \sim \mathcal{N}(\text{vec}(\mathbf{M}), \Omega \otimes \Sigma)$, where $\text{vec}(X) \in \mathbb{R}^{nm}$ is the vectorization of X , obtained by stacking the columns of X , and \otimes is the Kronecker product.

We continue by modeling $d(x)$ as a MVG on \mathbb{R}^n . Without loss of generality, we assume zero mean for the MVG with positive semi-definite parameter covariance matrix $\Omega \in \mathbb{R}^{m \times m}$, and kernel $\kappa : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$. That is,

$$\text{vec}(d(x_1), \dots, d(x_N)) \sim \mathcal{N}(\mathbf{0}, \Sigma(x) \otimes \Omega),$$

where $\Sigma \in \mathbb{R}^{n \times n}$ with $\Sigma_{i,j} = \kappa(x_i, x_j)$. There are many potential choices for the kernel function $\kappa(x_i, x_j)$ (cf. [38]), though we utilize the simple squared-exponential kernel in this work,

$$\kappa(x_i, x_j) = \sigma^2 \exp\left(\frac{-(x_i - x_j)^2}{2l^2}\right), \quad (7)$$

where σ and l are kernel hyperparameters. Since $d(x)$ is an MVG, the training observations $y_{[N]} := [d(x_1), \dots, d(x_N)]$ at sampling points $x_{[N]} := [x_1, \dots, x_N]$, and the predictive target $d(x_*)$ at query test point, x_* , are jointly Gaussian as

$$\begin{bmatrix} y_N \\ d(x_*) \end{bmatrix} \sim \mathcal{M}\mathcal{N}\left(\mathbf{0}, \begin{bmatrix} K(x_{[N]}, x_{[N]}) & K(x_*, x_{[N]})^T \\ K(x_*, x_{[N]}) & \kappa(x_*, x_*) \end{bmatrix}, \Omega\right),$$

where $K(x_{[N]}, x_{[N]}) \in \mathbb{R}^{n \times n}$ with $[K(x_{[N]}, x_{[N]})]_{i,j} = \kappa(x_i, x_j)$, and $K(x_*, x_{[N]}) \in \mathbb{R}^{1 \times N}$ with $[K(x_*, x_{[N]})]_i = \kappa(x_*, x_i)$. Thus, we can compute the posterior distribution as follows:

$$\begin{aligned} \text{vec}(d(x_*)) &\sim \mathcal{N}(\hat{M}, \hat{\Sigma} \otimes \hat{\Omega}) \\ \hat{M} &= K(X_*, X)^T K(X, X)^{-1} Y \\ \hat{\Sigma} &= \kappa(X_*, X_*) - K(X_*, X)^T K(X, X)^{-1} K(X_*, X) \\ \hat{\Omega} &= \Omega \end{aligned} \quad (8)$$

This allows us to estimate our unknown dynamics and their possibly correlated uncertainties.

IV. ROBUST MULTI-AGENT CBF

In this section, we first show that under certain assumptions, (5) is a multi-agent CBF for our discrete-time nonlinear dynamics. Then, given bounds on the uncertainty in each agents' dynamics, we incorporate robustness to these uncertainties into the CBF while maintaining the computational efficiency of a quadratic program.

Extending Multi-Agent CBF to Discrete-Time, Nonlinear Systems: The multi-agent CBF introduced in [6] was originally designed for *continuous-time linear* systems. However, we prove that under proper assumption, $h(x)$ defined in (5) is a discrete-time CBF for the *discrete-time nonlinear* system (1)/(2). The tradeoff is the additional conservativeness in a_{max} introduced by the following assumption. Intuitively, this assumption ensures that the robot can accelerate in any direction relative to the other agents, as proved in Lemma 1.

Assumption 1: Assume that for all $x \in \mathcal{C}$, $g_v(x)$ is invertible and $\frac{\|\beta_v(x)\|}{\sigma_{min}(g_v(x))u_{max}} < 1$, where $\beta_v(x) = f_v(x) + d_v(x) - f_v^{(i)}(x) - d_v^{(i)}(x) - \Delta v_t$ and $\sigma_{min}(g_v(x))$ is the minimum singular value of $g_v(x)$.

Lemma 1: Under Assumption 1, the expression (5), defining set \mathcal{C} , represents a discrete-time CBF for system (1), with

$$a_{max} = \min_x \left[\sigma_{min}(g_v(x))u_{max} - \|\beta_v(x)\| \right] > 0. \quad (9)$$

Proof: First, we must show that set \mathcal{C} defined by expression (5) is control invariant for the dynamics (1), given that the robot has acceleration authority in any direction of at least a_{max} for all $x \in \mathcal{C}$. For this, we rely on the same proof structure in [6], with the main difference being that we have discrete-time (rather than continuous-time) dynamics. Let $\Delta \hat{v}(x_t)$ denote the component of velocity $v(x_t)$ in the direction of collision.

$$\Delta \hat{v}(x_t) = \frac{\Delta p^T \Delta v}{\|\Delta p\|} \quad (10)$$

We know that collision can be avoided if we can match the other agent's velocity (i.e. $\Delta v = 0$) by the time we reach them. If we assume that we can accelerate at least a_{max} in any direction, we are guaranteed that we can achieve $\Delta v = 0$

within time $T_c = \frac{-\Delta \hat{v}(x_t)}{a_{max}}$. In our discrete-time formulation, the following condition implies collision avoidance:

$$\begin{aligned} \Delta \hat{v}(x_t) T_c + \|\Delta p\| &\geq D_s, \\ \left(\frac{\Delta p^T \Delta v}{\|\Delta p\|} \right)^2 &\leq a_{max} (\|\Delta p\| - D_s). \end{aligned} \quad (11)$$

Note that this constraint is only active when two agents are moving closer to each other ($\Delta \hat{v} < 0$), and no constraint is needed when two agents are moving away from each other ($\Delta \hat{v} \geq 0$). Therefore, collision can always be avoided under the following condition,

$$-\frac{\Delta p^T \Delta v}{\|\Delta p\|} \leq \sqrt{a_{max} (\|\Delta p\| - D_s)}. \quad (12)$$

Based on our geometric argument, we know that the set $\mathcal{C} := \{x \in \mathbb{R}^n : h(x) \geq 0\}$ defined by $h := \frac{\Delta p^T \Delta v}{\|\Delta p\|} + \sqrt{a_{max} (\|\Delta p\| - D_s)}$ is control invariant. This implies that $h(x)$ is a discrete-time CBF [27], given that the robot can accelerate by at least a_{max} in any direction.

Therefore, our second step is to show that for all $x \in \mathcal{C}$ and any unit vector \hat{e} , it holds that $\sup_{u \in \mathcal{U}} \|((v_{t+1}(x, u) - v_{t+1}^{(i)}(x)) - (v_t - v_t^{(i)})^T \hat{e})\| \geq a_{max} > 0$.

$$\begin{aligned} \sup_{u \in \mathcal{U}} \|((v_{t+1}(x, u) - v_{t+1}^{(i)}(x)) - (v_t - v_t^{(i)})^T \hat{e})\| & \\ &= \sup_{u \in \mathcal{U}} \|(\beta_v(x) + g_v(x)u)^T \hat{e}\| \\ &= \sup_{u \in \mathcal{U}} \|\hat{e}^T \beta_v(x) + \hat{e}^T g_v(x)u\| \\ &\geq \sup_{u \in \mathcal{U}} \|\hat{e}^T g_v(x)u\| - \|\hat{e}^T \beta_v(x)\| \\ &\geq \sigma_{min}(g_v(x))u_{max} - \|\beta_v(x)\| \\ &\geq \min_x \left[\sigma_{min}(g_v(x))u_{max} - \|\beta_v(x)\| \right] = a_{max} \\ &> 0, \end{aligned} \quad (13)$$

where the last inequality follows directly from Assumption 1. Therefore, we are guaranteed that the robot can accelerate by at least $a_{max} > 0$ in any direction. Combined with the first part of the proof, this shows that the set \mathcal{C} defined by (5) is a discrete-time CBF. ■

Incorporating Robustness into CBF: While uncertainty in robot/environmental dynamics can be directly incorporated into the Control Barrier Condition (CBC) for simple systems/constraints (e.g. linear CBFs) [22], this is not the case for the multi-agent CBF with discrete-time dynamics. Unfortunately, uncertainty cannot be directly incorporated into the CBC while maintaining a quadratic program.

Consider our CBF (5) and the dynamics defined in (1) and (2). Based on these, we can compute the following CBC:

$$\begin{aligned}
CBC(x_t, u_t) = & \left\langle \frac{f_p(x_t) + g_p(x_t)u_t + d_p(x_t) - f_p^h(x_t) - d_p^h(x_t)}{\|f_p(x_t) + g_p(x_t)u_t + d_p(x_t) - f_p^h(x_t) - d_p^h(x_t)\|}, \right. \\
& \left. f_v(x_t) + g_v(x_t)u_t + d_v(x_t) - f_v^h(x_t) - d_v^h(x_t) \right\rangle + \\
& \sqrt{a_{max}(\|f_p(x_t) + g_p(x_t)u_t + d_p(x_t) - f_p^h(x_t) - d_p^h(x_t)\| - D_s)} + \\
& (\eta - 1)\sqrt{a_{max}(\|\Delta p_t\| - D_s)} + (\eta - 1)\frac{\Delta p_t^T \Delta v_t}{\|\Delta p_t\|}. \tag{14}
\end{aligned}$$

If we can **(a)** determine bounds on the dynamic uncertainties, d , in (1) and (2), and **(b)** compute control actions that satisfy $CBC(x, u) \geq 0$ in an online fashion, then we can obtain robust safety guarantees utilizing the multi-agent CBF. Ideally, we could incorporate (14) into an efficiently solvable program as follows,

$$\begin{aligned}
u = \operatorname{argmin}_{u_t \in \mathcal{U}} & \|u - u_{des}\|_2 \\
\text{s.t.} & \max_{d(x_t)} CBC(x_t, u, d_t) \geq 0 \\
& \text{where } d(x_t) \in \mathcal{D} \\
& \|u\|_2 \leq u_{max}. \tag{15}
\end{aligned}$$

However, the CBC constraint (14) in (15) is clearly not linear (nor convex). Therefore, the resulting program is non-convex and cannot be solved at high frequency for adequate safety assurances. However, recall that our system has relative degree 2, which allows us to derive the following bound,

$$CBC(x_t, u_t, d_t) \geq k_c(x_t) - H_1(x_t)d_t - u_t^T H_2(x_t)d_t - H_3(x_t)u_t, \tag{16}$$

where the definitions of the terms (k_c, H_1, H_2, H_3) are given in (27) in the Appendix. We also move the derivation of the bound (16) to the Appendix due to space limitations.

The following lemma allows us to utilize CBC bound (16) to obtain safety guarantees *under polytopic uncertainties*.

Lemma 2: *Suppose the uncertainty in our dynamics d is bounded in the polytope $\{d \in \mathbb{R}^n \mid Gd \leq g\}$. Then the action, u , obtained from solving the following optimization problem (17) robustly satisfies the CBC condition (14) (i.e. renders the set \mathcal{C} forward invariant).*

$$\begin{aligned}
\min_{u \in \mathcal{U}, \xi \in \mathbb{R}_+^{4n}} & \|u - u_{des}\|_2 \\
\text{s.t.} & H_3(x_t)u + \xi g \leq k_c(x_t) \\
& H_1(x_t) + u^T H_2(x_t) = \xi G \\
& \xi \geq 0 \\
& \|u\|_2 \leq u_{max} \quad (\text{actuation limits}), \tag{17}
\end{aligned}$$

Proof: The robust optimization problem (17) can be equivalently represented by the following optimization problem (i.e. (17) is the dual to (18) with no duality gap [39]

where ξ is the dual variable):

$$\begin{aligned}
\min_{u \in \mathcal{U}} & \|u - u_{des}\|_2 \\
\text{s.t.} & \forall d \in \{d \in \mathbb{R}^n \mid Gd \leq g\} \\
& H_1(x_t)d + u^T H_2(x_t)d + H_3(x_t)u \leq k_c(x_t) \\
& \|u\|_2 \leq u_{max} \quad (\text{actuation limits}), \tag{18}
\end{aligned}$$

where u is the decision vector, d is the uncertainty variable, and $\{d \in \mathbb{R}^n \mid Gd \leq g\}$ is the uncertainty bound. If the inequality in (18) is satisfied such that $H_1d + u^T H_2d + H_3u \leq k_c$, then it follows directly from (16) that the CBC condition is satisfied for all d in our polytopic uncertainty set. Therefore, the set \mathcal{C} is rendered forward invariant. ■

This lemma shows that if we bound $d(x_t)$ in a polytope, we can transform our robust multi-agent CBF (15) into a quadratic program (17), which gives us a computationally efficient way to provide robust guarantees of safety under robot and environment uncertainties. Hence, in the following section, we examine the problem of learning accurate polytopic bounds on the uncertainty d in an online fashion.

V. LEARNING UNCERTAINTY BOUNDS

In this section, our goal is to learn accurate confidence supports for the uncertainties d and $d^{(i)}$ (for all agents i) in an online manner, which will allow us to guarantee safety with high probability. To this end, we utilize Matrix-Variate Gaussian Processes which provide multivariate Gaussian distributions over the uncertainties, d and $d^{(i)}$.

Using (1), we have $d(x_t) = x_{t+1} - f(x_t) - g(x_t)u_t$. A similar relation holds based on (2) for other agents. Thus, given a sequence of measurements (x_t, u_t, x_{t+1}) over a horizon T , we compute the uncertain variable, d_{t-T}, \dots, d_{t-1} over that horizon. Then, we infer a distribution over the query point, d_t (i.e. next time point), as described in Equation (8).

Learning Kernel Parameters: Direct application of the MVG (8) to our multi-agent setup will be problematic without first training the MVG model hyperparameters. This is easy to see by noting that the covariance, $\Sigma(X) \otimes \Omega$, does not depend on the observed values, Y . Furthermore, the coupling between uncertainties, captured by Ω , is completely independent of our online measurements. Instead, much of the uncertainty prediction is baked into the kernel parameters, $\kappa(l, \sigma)$, and matrix Ω . Thus, to obtain accurate estimates of d , we must learn MVG model parameters offline from data. In other words, some agents might behave predictably and others might behave more erratically, and hyperparameter optimization is necessary to capture these uncertainty profiles in our Bayesian inference.

Based on the probability density function (6), we obtain the negative log-likelihood of a given set of training data X ,

$$\begin{aligned}
L(\mathbf{X}, \mathbf{Y}; K, \Omega) = & -\ln p(\mathbf{X}, \mathbf{Y}; K, \Omega) = \\
& \frac{nm}{2} \ln(2\pi) + \frac{m}{2} \ln |K| + \frac{n}{2} \ln |\Omega| + \frac{1}{2} \operatorname{tr}[(K)^{-1} Y \Omega^{-1} Y^T] \tag{19}
\end{aligned}$$

which we optimize (over Σ, Ω) using Stochastic Gradient Descent (see hyperparameter optimization in Fig. 1) [40]. We run the optimization several times with different initializations to avoid local optima. The gradient update expressions are shown in Equation (20) below. Note that we use projected gradient updates for Ω , in order to enforce the condition that Ω must be positive definite.

$$\begin{aligned} \frac{dL}{dl} &= \frac{d}{2} \text{tr} \left(K^{-1} \frac{dK}{dl} \right) + \frac{1}{2} \text{tr} \left(-K^{-1} \frac{dK}{dl} K^{-1} Y \Omega^{-1} Y^T \right) \\ \frac{dL}{d\sigma} &= \frac{d}{2} \text{tr} \left(K^{-1} \frac{dK}{d\sigma} \right) + \frac{1}{2} \text{tr} \left(-K^{-1} \frac{dK}{d\sigma} K^{-1} Y \Omega^{-1} Y^T \right) \\ \frac{dL}{d\Omega} &= \frac{n}{2} \Omega^{-1} - \frac{1}{2} \Omega^{-1} Y^T K^{-1} Y \Omega^{-1} \end{aligned} \quad (20)$$

Converting Gaussian Process Uncertainty to Polytopic Bound: After learning the kernel parameters, we can obtain the mean, $\mu_d = \hat{M}$, and variance, $\Sigma_d = \hat{\Sigma} \otimes \hat{\Omega}$, from data observed online based on the multivariate Gaussian Process (8). Then, the uncertainties should follow the distribution,

$$(d - \mu_d)^T \Sigma_d^{-1} (d - \mu_d) \sim \chi_N^2, \quad (21)$$

where χ_N^2 represents the chi-squared distribution with N degrees of freedom (equal to dimension of d). This allows us to obtain the confidence support,

$$(d - \mu_d)^T \Sigma_d^{-1} (d - \mu_d) \leq k_\delta \quad \text{with probability } 1 - \delta. \quad (22)$$

However, this set defines an *ellipsoid* over d rather than a polytope, which we require for the robust optimization. To obtain a polytope, we compute the minimum bounding box surrounding this ellipsoid.

Lemma 3: *Suppose our robot/environment uncertainty can be described by our MVG model (described by the distribution (21)). With probability $1 - \delta$, the following polytopic bound on the uncertainty d holds:*

$$-\sqrt{k_\delta \lambda_i} + v_i^T \mu_d \leq v_i^T d \leq \sqrt{k_\delta \lambda_i} + v_i^T \mu_d, \quad (23)$$

where v_i and λ_i represent the eigenvectors and eigenvalues of Σ_d , respectively.

Proof: Since Σ_d is a positive symmetric covariance matrix, the eigendecomposition of Σ_d exists, and we have $\Sigma_d = \Psi^T \Lambda \Psi$, where Λ is a diagonal matrix containing positive eigenvalues of Σ_d and Ψ is the orthogonal eigenvector matrix. Thus, (22) can be rewritten as follows:

$$\left[\Psi^T (d - \mu_d) \right]^T \Lambda^{-1} \left[\Psi^T (d - \mu_d) \right] \leq k_\delta \quad \text{w.p. } 1 - \delta. \quad (24)$$

We can bound the left hand side:

$$\begin{aligned} &\left[\Psi^T (d - \mu_d) \right]^T \Lambda^{-1} \left[\Psi^T (d - \mu_d) \right] \\ &= \sum_i \left[v_i^T (d - \mu_d) \right]^T \lambda_i^{-1} \left[v_i^T (d - \mu_d) \right] \\ &\geq \left[v_i^T (d - \mu_d) \right]^T \lambda_i^{-1} \left[v_i^T (d - \mu_d) \right] \quad \forall i = 1, \dots, N \end{aligned} \quad (25)$$

where v_i represent the eigenvectors of Σ_d contained in Ψ , and λ_i are the eigenvalues of Σ_d contained in Λ . We can then conclude that with probability $1 - \delta$, the following relations hold giving us our polytopic bound,

$$\begin{aligned} &\left[v_i^T (d - \mu_d) \right]^T \lambda_i^{-1} \left[v_i^T (d - \mu_d) \right] \leq k_\delta \quad \forall i = 1, \dots, N \\ &-\sqrt{k_\delta \lambda_i} + v_i^T \mu_d \leq v_i^T d \leq \sqrt{k_\delta \lambda_i} + v_i^T \mu_d \quad \text{for } i = 1, \dots, N \end{aligned} \quad (26)$$

High-Confidence Safety Guarantee: Combining the uncertainty bound on d with our result from Sec. IV leads us to the main result, summarized in the following Theorem.

Theorem 4: *Using the polytopic bounds (23), the control action obtained from the quadratic program (17) guarantees robust safety (i.e. collision avoidance between agents) with probability at least $1 - \delta$.*

Proof: We can represent (23) in the form $\{Gd \leq g\}$; therefore, with probability $1 - \delta$, the uncertainty d is contained in the set $\{d \in \mathbb{R}^n | Gd \leq g\}$ (by Lemma 3). From Lemma 2 and Equation (16), if we solve the quadratic program (17), we are guaranteed that $CBC(x, u, d) \geq k_c - H_1 d - u^T H_2 d - H_3 u \geq 0$ for all $d \in \{d \in \mathbb{R}^n | Gd \leq g\}$. Therefore, the CBF condition is satisfied with probability $1 - \delta$, so safety is guaranteed with probability at least $1 - \delta$ (by Definition 1 and the forward invariance property of CBFs [26]).

VI. RESULTS

A. Navigation in Unstructured Environment

We test our algorithm in a simulated multi-agent environment in which our robot, with nonlinear dynamics satisfying Assumption 1, navigates from a start to goal position while avoiding collisions, in the presence of a random number of other agents (3-12 agents). Each of the other agents has a randomized (unknown) goal that they try to reach. Approximately half of them blindly travel from their start to goal position without accounting for others, while the other half exhibit some collision avoidance behavior through their own control barrier functions (with random CBF parameters). An example simulation instance is shown in Fig. 2.

We simulate several instances of the other agents moving and interacting, and use this data for hyperparameter optimization of an MVG model as described in Sec. V. We then equip our robot with the robust CBF described in Sec. IV, using the optimized MVG for uncertainty prediction.

By running 1000 simulated tests in randomized environments, we show that the robust CBF avoids collision in 98.5% of cases (when we set $\delta = 0.05$ [cf. (22)]), performing much better than the nominal multi-agent CBF, which avoids collisions in 85.0% of cases. The simulation results are summarized in Table I.

Robustness must always come at the cost of performance (e.g. we can reach the goal faster if we do not care about collisions). To investigate the conservativeness of our

	Robust Multi-Agent CBF	Nominal Multi-Agent CBF
Collision Rate	1.5 %	15.0%
Distance to Collision	7.4 ± 2.3	7.3 ± 2.1

TABLE I: Performance statistics for the robust vs nominal multi-agent CBF across 1000 randomized trials. For fair comparison, the robust and nominal CBFs were tested in the same randomized 1000 trials. **Collision Rate:** Percentage of trials that ended in collision. **Distance to Collision:** For trials without collision, the robot’s margin from collision. The closer the robust CBF is to the nominal CBF, the less conservativeness is introduced by the uncertainty prediction.

approach, we looked at the uncertainty predictions of the MVG; Fig. 3 shows the uncertainty ellipse (over the 4-dimensional disturbance, d) projected onto the two velocity dimensions, as well as the true disturbances, d_v . We found that ($\approx 97\%$, 99%) of disturbances, d , were within the $(2, 3)\sigma$ confidence ellipsoid, respectively, in line with the expectations of the MVG model. Furthermore, the results in Table I show that the robust CBF only introduces slight conservativeness, as the margin from collision (in instances where the CBF was active) was very similar when utilizing the robust CBF vs. the nominal CBF. This suggests that the MVG model does well at modeling the uncertainties.

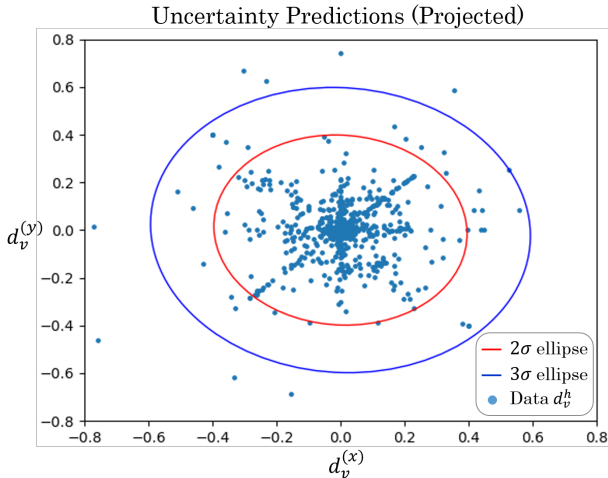


Fig. 3: The normalized 2σ (red) and 3σ (blue) uncertainty ellipsoids over the other agents’ dynamics, d^h , projected onto the velocity dimensions, d_v^h . The true disturbances over 1000 time steps (across different trials) are plotted as the blue dots. We found that the percentage of disturbances within the 2σ and 3σ uncertainty ellipsoids were consistent with expectations based on the MVG model. Note that the uncertainty ellipse is state-dependent, so we normalize the ellipsoid for each point, (x^h, d_v^h) , for fair comparison.

The code for implementing the robust multi-agent CBF in our simulated environment can be found at https://github.com/rcheng805/robust_cbf. A video of the simulations can be found at <https://youtu.be/hXg5kZ086Lw>.

VII. CONCLUSION

Robot navigation in unstructured environments with humans must be safe, but such environments are fraught with uncertainty due to the unpredictability of agents. In this work, we have introduced a robust multi-agent control barrier formulation, which guarantees safety with high probability in the presence of multiple uncontrolled, uncertain agents. We learn uncertainties online for the agents in the environment using Matrix-Variate Gaussian Processes, and design our CBF to be robust to the learned uncertainties.

Future work will look at learning and designing safe controllers for a larger class of uncertainties, uncaptured by our Matrix Variate GP. Particularly, many agents in the real world exhibit multi-modal uncertainties, and it will be important for us to design safe, robust controllers for such uncertainties.

REFERENCES

- [1] F. Berkenkamp, M. Turchetta, A. Schoellig, and A. Krause, “Safe model-based reinforcement learning with stability guarantees,” in *Advances in Neural Information Processing Systems*, 2017.
- [2] J. F. Fisac, A. K. Akametalu, M. N. Zeilinger, S. Kaynama, J. Gillula, and C. J. Tomlin, “A general safety framework for learning-based control in uncertain robotic systems,” *IEEE Transactions on Automatic Control*, 2018.
- [3] T. Koller, F. Berkenkamp, M. Turchetta, and A. Krause, “Learning-based model predictive control for safe exploration,” in *2018 IEEE Conference on Decision and Control (CDC)*. IEEE, 2018.
- [4] K. P. Wabersich and M. N. Zeilinger, “Safe exploration of nonlinear dynamical systems: A predictive safety filter for reinforcement learning,” *arXiv preprint arXiv:1812.05506*, 2018.
- [5] L. Janson, T. Hu, and M. Pavone, “Safe motion planning in unknown environments: Optimality benchmarks and tractable policies,” in *Robotics: Science and Systems*, Pittsburgh, USA, June 2018.
- [6] U. Borrmann, L. Wang, A. D. Ames, and M. Egerstedt, “Control Barrier Certificates for Safe Swarm Behavior,” *IFAC Conference on Analysis and Design of Hybrid Systems*, 2015.
- [7] L. Wang, A. Ames, and M. Egerstedt, “Safety barrier certificates for heterogeneous multi-robot systems,” in *Proceedings of the American Control Conference*, 2016.
- [8] A. Singletary, P. Nilsson, T. Gurriet, and A. D. Ames, “Online Active Safety for Robotic Manipulators,” 2020.
- [9] F. Bartoli, G. Lisanti, L. Ballan, and A. Del Bimbo, “Context-Aware Trajectory Prediction,” in *International Conference on Pattern Recognition*, 2018.
- [10] P. Fiorini and Z. Shiller, “Motion planning in dynamic environments using velocity obstacles,” *International Journal of Robotics Research*, 1998.
- [11] J. D. Van Berg, M. Lin, and D. Manocha, “Reciprocal velocity obstacles for real-time multi-agent navigation,” in *Proceedings - IEEE International Conference on Robotics and Automation*, 2008.
- [12] D. Wilkie, J. Van Den Berg, and D. Manocha, “Generalized velocity obstacles,” in *International Conference on Intelligent Robots and Systems*, 2009.
- [13] D. Bareiss and J. Van Den Berg, “Generalized reciprocal collision avoidance,” *International Journal of Robotics Research*, 2015.
- [14] D. Zhou, Z. Wang, S. Bandyopadhyay, and M. Schwager, “Fast, On-line Collision Avoidance for Dynamic Vehicles Using Buffered Voronoi Cells,” *IEEE Robotics and Automation Letters*, 2017.
- [15] M. Wang, Z. Wang, S. Paudel, and M. Schwager, “Safe Distributed Lane Change Maneuvers for Multiple Autonomous Vehicles Using Buffered Input Cells,” in *International Conference on Robotics and Automation*, 2018.
- [16] J. F. Fisac, M. Chen, C. J. Tomlin, and S. Shankar Sastry, “Reach-avoid problems with time-varying dynamics, targets and constraints,” in *International Conference on Hybrid Systems: Computation and Control, HSCC*, 2015.

- [17] M. Chen and C. J. Tomlin, "Hamilton–Jacobi Reachability: Some Recent Theoretical Advances and Applications in Unmanned Airspace Management," *Annual Review of Control, Robotics, and Autonomous Systems*, 2018.
- [18] S. Bansal, M. Chen, S. Herbert, and C. J. Tomlin, "Hamilton-jacobi reachability: A brief overview and recent advances," in *Conference on Decision and Control, CDC 2017*.
- [19] Y. F. Chen, M. Everett, M. Liu, and J. P. How, "Socially aware motion planning with deep reinforcement learning," in *IEEE International Conference on Intelligent Robots and Systems*, 2017.
- [20] M. Everett, Y. F. Chen, and J. P. How, "Motion Planning among Dynamic, Decision-Making Agents with Deep Reinforcement Learning," in *International Conference on Intelligent Robots and Systems*, 2018.
- [21] W. Zhang, O. Bastani, and V. Kumar, "MAMPS: Safe Multi-Agent Reinforcement Learning via Model Predictive Shielding," *arXiv e-prints*, p. arXiv:1910.12639, Oct 2019.
- [22] 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," *Proceedings of the AAAI Conference on Artificial Intelligence*, 2019.
- [23] L. Wang, E. A. Theodorou, and M. Egerstedt, "Safe learning of quadrotor dynamics using barrier certificates," in *International Conference on Robotics and Automation (ICRA)*. IEEE, 2018.
- [24] M. J. Khojasteh, V. Dhiman, M. Franceschetti, and N. Atanavov, "Probabilistic safety constraints for learned high relative degree system dynamics," *arXiv preprint arXiv:1912.10116*, 2019.
- [25] D. D. Fan, J. Nguyen, R. Thakker, N. Alatur, A.-a. Agha-mohammadi, and E. A. Theodorou, "Bayesian learning-based adaptive control for safety critical systems," *arXiv preprint arXiv:1910.02325*, 2019.
- [26] A. D. Ames, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs with application to adaptive cruise control," in *Proceedings of the IEEE Conference on Decision and Control*, 2014.
- [27] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, "Control Barrier Function Based Quadratic Programs for Safety Critical Systems," *IEEE Transactions on Automatic Control*, 2017.
- [28] Q. Nguyen and K. Sreenath, "Exponential control barrier functions for enforcing high relative-degree safety-critical constraints," in *2016 American Control Conference (ACC)*. IEEE, 2016, pp. 322–328.
- [29] A. Clark, "Control barrier functions for complete and incomplete information stochastic systems," in *2019 American Control Conference (ACC)*. IEEE, 2019, pp. 2928–2935.
- [30] G. Yang, C. Belta, and R. Tron, "Self-triggered control for safety critical systems using control barrier functions," in *2019 American Control Conference (ACC)*. IEEE, 2019, pp. 4454–4459.
- [31] M. Srinivasan, N.-s. P. Hyun, and S. Coogan, "Weighted polar finite time control barrier functions with applications to multi-robot systems," in *IEEE Conference on Decision and Control*, 2019.
- [32] T. Gurriet, P. Nilsson, A. Singletary, and A. D. Ames, "Realizable set invariance conditions for cyber-physical systems," in *Proceedings of the American Control Conference*, 2019.
- [33] M. Srinivasan, A. Dabholkar, S. Coogan, and P. Vela, "Synthesis of Control Barrier Functions using a Supervised Machine Learning Approach," *arXiv*, 2020.
- [34] A. Agrawal and K. Sreenath, "Discrete control barrier functions for safety-critical control of discrete systems with application to bipedal robot navigation," in *Robotics: Science and Systems*, 2017.
- [35] A. K. Gupta and D. K. Nagar, *Matrix variate distributions*. Chapman and Hall/CRC, 2018.
- [36] C. Louizos and M. Welling, "Structured and efficient variational deep learning with matrix gaussian posteriors," in *International Conference on Machine Learning*, 2016, pp. 1708–1716.
- [37] S. Sun, C. Chen, and L. Carin, "Learning Structured Weight Uncertainty in Bayesian Neural Networks," in *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2017, pp. 1283–1292.
- [38] C. K. Williams and C. E. Rasmussen, *Gaussian processes for machine learning*. MIT press Cambridge, MA, 2006, vol. 2, no. 3.
- [39] Y. Chen and J. Anderson, "System level synthesis with state and input constraints," in *Conference on Decision and Control (CDC)*, 2019.
- [40] S. Crepey and M. F. Dixon, "Gaussian Process Regression for Derivative Portfolio Modeling and Application to CVA Computations," *arXiv*, 2019.

APPENDIX I

PARAMETERS OF THE CONTROL BARRIER CONDITION

Here we define the terms used in the lower bound of the Control Barrier Condition (CBC) in (16):

$$\begin{aligned}
 H_1(1 \times P) &= \left[-\frac{f_v(x) - f_v^h(x)}{\|f_p(x) - f_p^h(x)\| - \zeta_p(x) - \zeta_p^h(x)}, \right. \\
 &\quad \left. -\frac{f_p(x) - f_p^h(x)}{\|f_p(x) - f_p^h(x)\| - \zeta_p(x) - \zeta_p^h(x)}, \right. \\
 &\quad \left. \frac{f_v(x) - f_v^h(x)}{\|f_p(x) - f_p^h(x)\| - \zeta_p(x) - \zeta_p^h(x)}, \right. \\
 &\quad \left. \frac{f_p(x) - f_p^h(x)}{\|f_p(x) - f_p^h(x)\| - \zeta_p(x) - \zeta_p^h(x)} \right] \\
 H_2(M \times P) &= \left[-\frac{g_v(x)}{\|f_p(x) - f_p^h(x)\| - \zeta_p(x) - \zeta_p^h(x)}, 0, \right. \\
 &\quad \left. \frac{g_v(x)}{\|f_p(x) - f_p^h(x)\| - \zeta_p(x) - \zeta_p^h(x)}, 0 \right] \\
 H_3(1 \times M) &= \left[-\frac{(f_p(x) - f_p^h(x))^T g_v(x)}{\|f_p(x) - f_p^h(x)\| + \zeta_p(x) + \zeta_p^h(x)} \right] \\
 k_c &= \min \left(\frac{(f_p(x) - f_p^h(x))^T (f_v(x) - f_v^h(x))}{\|f_p(x) - f_p^h(x)\| \pm (\zeta_p(x) + \zeta_p^h(x))} \right) + \\
 &\quad \sqrt{a_{\max}(\|f_p(x) - f_p^h(x)\| - \zeta_p(x) - \zeta_p^h(x) - D_s)} + \\
 &\quad (\eta - 1) \sqrt{a_{\max}(\|\Delta p_t\| - D_s)} + (\eta - 1) \frac{\Delta p_t^T \Delta v_t}{\|\Delta p_t\|} - \\
 &\quad \frac{\zeta_p(x) \zeta_v(x) + \zeta_p(x) \zeta_v^h(x) + \zeta_p^h(x) \zeta_v^h(x) + \zeta_v(x) \zeta_p^h(x)}{\|f_p(x) - f_p^h(x)\| - \zeta_p(x) - \zeta_p^h(x)} \tag{27}
 \end{aligned}$$

APPENDIX II
PROOF OF CBC LOWER BOUND

In this section, we prove that the lower bound defined in (16) holds. We begin by expanding out the full CBC condition in (14), using our assumption of a relative degree 2 system, we reach the following expression:

$$\begin{aligned}
CBC(x, u, d) = & \frac{(f_p(x) - f_p^h(x))^T (f_v(x) - f_v^h(x))}{\|f_p(x) + d_p(x) - f_p^h(x) - d_p^h(x)\|} + \sqrt{a_{max}(\|f_p(x) + d_p(x) - f_p^h(x) - d_p^h(x)\| - D_s)} + \\
& (\gamma - 1) \sqrt{a_{max}(\|\Delta p\| - D_s)} + (\gamma - 1) \frac{\Delta p^T \Delta v}{\|\Delta p\|} + \left[\frac{(f_p(x) - f_p^h(x))^T g_v(x)}{\|f_p(x) + d_p(x) - f_p^h(x) - d_p^h(x)\|}, \right. \\
& \frac{f_v(x) - f_v^h(x)}{\|f_p(x) + d_p(x) - f_p^h(x) - d_p^h(x)\|}, \quad \frac{f_p(x) - f_p^h(x)}{\|f_p(x) + d_p(x) - f_p^h(x) - d_p^h(x)\|}, \quad - \frac{f_v(x) - f_v^h(x)}{\|f_p(x) + d_p(x) - f_p^h(x) - d_p^h(x)\|}, \\
& \left. - \frac{f_p(x) - f_p^h(x)}{\|f_p(x) + d_p(x) - f_p^h(x) - d_p^h(x)\|} \right] \times \left[u_R, d_p, d_v, d_p^h, d_v^h \right]^T + \\
& u_R^T \left[\frac{g_v(x)}{\|f_p(x) + d_p(x) - f_p^h(x) - d_p^h(x)\|}, 0, \frac{-g_v(x)}{\|f_p(x) + d_p(x) - f_p^h(x) - d_p^h(x)\|}, 0 \right] \times \\
& \left[d_p, d_v, d_p^h, d_v^h \right]^T + \left[0, \frac{1}{\|f_p(x) + d_p(x) - f_p^h(x) - d_p^h(x)\|}, \frac{-1}{\|f_p(x) + d_p(x) - f_p^h(x) - d_p^h(x)\|}, \right. \\
& \left. \frac{1}{\|f_p(x) + d_p(x) - f_p^h(x) - d_p^h(x)\|}, \frac{-1}{\|f_p(x) + d_p(x) - f_p^h(x) - d_p^h(x)\|} \right] \times \left[u_R^T u_R, d_p^T d_v, d_p^T d_v^h, d_v^T d_p^h, d_v^T d_p^h \right]^T
\end{aligned}$$

By bounding the positional uncertainty terms $\|d_p(x)\| \leq \zeta_p(x)$ and $\|d_p^h(x)\| \leq \zeta_p^h(x)$, we obtain a lower bound on $CBC(x, u, d)$:

$$\begin{aligned}
CBC(x) \geq & \min \left(\frac{(f_p(x) - f_p^h(x))^T (f_v(x) - f_v^h(x))}{\|f_p(x) - f_p^h(x)\| \pm (\zeta_p(x) + \zeta_p^h(x))} \right) + \sqrt{a_{max}(\|f_p(x) - f_p^h(x)\| - \zeta_p(x) - \zeta_p^h(x) - D_s)} + \\
& (\gamma - 1) \sqrt{a_{max}(\|\Delta p\| - D_s)} + (\gamma - 1) \frac{\Delta p^T \Delta v}{\|\Delta p\|} + \left[\frac{(f_p(x) - f_p^h(x))^T g_v(x)}{\|f_p(x) - f_p^h(x)\| + \zeta_p(x) + \zeta_p^h(x)}, \frac{f_v(x) - f_v^h(x)}{\|f_p(x) - f_p^h(x)\| - \zeta_p(x) - \zeta_p^h(x)}, \right. \\
& \left. \frac{f_p(x) - f_p^h(x)}{\|f_p(x) - f_p^h(x)\| - \zeta_p(x) - \zeta_p^h(x)}, - \frac{f_v(x) - f_v^h(x)}{\|f_p(x) - f_p^h(x)\| - \zeta_p(x) - \zeta_p^h(x)}, - \frac{f_p(x) - f_p^h(x)}{\|f_p(x) - f_p^h(x)\| - \zeta_p(x) - \zeta_p^h(x)} \right] * \\
& \left[u_R, d_p, d_v, d_p^h, d_v^h \right]^T + u_R^T \left[\frac{g_v(x)}{\|f_p(x) - f_p^h(x)\| - \zeta_p(x) - \zeta_p^h(x)}, 0, \frac{-g_v(x)}{\|f_p(x) - f_p^h(x)\| - \zeta_p(x) - \zeta_p^h(x)}, 0 \right] * \\
& \left[d_p, d_v, d_p^h, d_v^h \right]^T - \frac{\zeta_p(x) \zeta_v(x)}{\|f_p(x) - f_p^h(x)\| - \zeta_p(x) - \zeta_p^h(x)} - \frac{\zeta_p(x) \zeta_v^h(x)}{\|f_p(x) - f_p^h(x)\| - \zeta_p(x) - \zeta_p^h(x)} - \\
& \frac{\zeta_v^h(x) \zeta_p^h(x)}{\|f_p(x) - f_p^h(x)\| - \zeta_p(x) - \zeta_p^h(x)} - \frac{\zeta_v(x) \zeta_p^h(x)}{\|f_p(x) - f_p^h(x)\| - \zeta_p(x) - \zeta_p^h(x)}
\end{aligned}$$

Grouping the terms, this can be written in simplified form using the parameters defined in Appendix I:

$$CBC(x, u, d) \geq k_c(x) - H_1(x) \mathbf{d} - \mathbf{u}^T H_2(x) \mathbf{d} - H_3(x) \mathbf{u}$$