

Belief Control Barrier Functions for Risk-aware Control

Matti Vahs, Christian Pek and Jana Tumova

Abstract—Ensuring safety in real-world robotic systems is often challenging due to unmodeled disturbances and noisy sensors. To account for such stochastic uncertainties, many robotic systems leverage probabilistic state estimators such as Kalman filters to obtain a robot’s *belief*, i.e. a probability distribution over possible states. We propose belief control barrier functions (BCBFs) to enable risk-aware control, leveraging all information provided by state estimators. This allows robots to stay in predefined safety regions with desired confidence under these stochastic uncertainties. BCBFs are general and can be applied to a variety of robots that use extended Kalman filters as state estimator. We demonstrate BCBFs on a quadrotor that is exposed to external disturbances and varying sensing conditions. Our results show improved safety compared to traditional state-based approaches while allowing control frequencies of up to 1kHz.

Index Terms—Robot Safety, Sensor-based Control

I. INTRODUCTION

AUTONOMOUS robotic systems are exposed to various sources of uncertainty, such as noise in the robot’s sensor readings or external disturbances, e.g., an unknown wind force acting on a quadrotor, as shown in Fig. 1. Considering uncertainty is thus crucial for safe operation of such robots. Safe and robust control synthesis under external disturbances has been thoroughly explored for a variety of robots with full state information [1]. Observer-based Control Barrier Functions (CBFs) consider measurement uncertainties through bounded estimation errors [2]–[4]. However, these methods do not consider the stochastic uncertainties that are commonly captured in probabilistic state estimation techniques deployed in practice, such as Kalman Filters (KFs). Such stochastic uncertainties have been considered for instance in the state-of-the-art chance-constrained nonlinear Model Predictive Control (MPC) [5], which bounds the probability of undersired events. However, to make the method computationally tractable, theoretical safety guarantees are sacrificed. In this paper, we consider both uncertainties in the robot’s motion and its observations. We aim to offer both safety guarantees under uncertainty and a practical solution.

Manuscript received: June, 12, 2023; Revised September, 12, 2023; Accepted November, 15, 2023.

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

This work was partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation. This research has been carried out as part of the Vinnova Competence Center for Trustworthy Edge Computing Systems and Applications at KTH Royal Institute of Technology.

Matti Vahs and Jana Tumova are with the Division of Robotics, Perception and Learning, KTH Royal Institute of Technology, Stockholm, Sweden and also affiliated with Digital Futures. Christian Pek is with the Department of Cognitive Robotics (CoR) at Delft University of Technology (TUD). Mail addresses: {vahs, tumova}@kth.se, C.Pek@tudelft.nl

Digital Object Identifier (DOI): see top of this page.




Fig. 1. An illustration of a quadrotor that is exposed to noisy measurements and unknown wind forces. The uncertainty in the state estimate is shown as a purple ellipsoid and the cuboid-shaped safe set is colored in red.

In practice, state estimation pipelines are used to obtain a robot’s belief that accounts for the uncertainty arising from stochastic motion and observations. Up to now, the Kalman filter (KF) is still one of the gold standards in robotic state estimation because of its simplicity and robustness [6]. The extended KF (EKF), as its extension to nonlinear systems, is widely used for state estimation of robots such as quadrotors, legged robots or autonomous underwater vehicles [7]. One of the key properties of KFs is that they not only provide an estimate about the robot’s state but also quantify the uncertainty of that estimate through the covariance matrix. For instance, Fig. 1 illustrates a drone’s localization uncertainty as a Gaussian level set (purple ellipsoid), i.e. a set containing the state with certain probability. In this work, we leverage both the mean state estimate and the covariance matrix.

Due to Gaussian state uncertainty, hard safety constraints on the system states are generally infeasible. This gives rise to a risk-aware perspective on safety: We consider specifications that bound the probability of violating safety constraints on the state. Fig. 1 shows a scenario in which a safety specification is encoded as “the probability of leaving the safe region (red cuboid) should be less than 1%”.

However, designing controllers that ensure satisfaction of such risk-aware safety specifications is challenging as it requires us to reason about the robot’s belief instead of its state. Belief spaces in these applications suffer from the curse of dimensionality and are hybrid in nature as robots evolve in continuous time while sensors only provide measurements at discrete timesteps. To enable risk-aware control, we introduce Belief CBFs (BCBFs) that 1) serve as safety filters in the presence of real-world stochastic uncertainties, 2) provide theoretical safety guarantees under the hybrid nature of belief spaces, and 3) overcome the curse of dimensionality and allow real-time control for general robotic systems that employ EKFs as a state estimator. We evaluate our approach in experiments

with a quadrotor that is exposed to external wind disturbances and varying sensing conditions, see Fig. 1.

A. Related Work

We review two relevant safety-critical control approaches for robotic systems under uncertainties – CBFs and MPC.

Related work on observer-based CBFs assume a bounded uncertainty around the current estimate that is used to account for potential measurement errors [2], [3]. However, these methods do not consider stochastic uncertainties that are commonly used in probabilistic state estimators. Measurement-robust CBFs use bounded errors in measurements and assume a given mapping from measurements to state estimates [4]. In contrast, our BCBFs explicitly provide this mapping to handle stochastic measurements.

Control of stochastic systems using Kalman Filters as state estimator has been addressed in various settings [5], [8]. Chance-constrained nonlinear MPC (CCNMPC) [5] enforce obstacle avoidance with a desired confidence level in which the uncertainty in the state estimate originates from an unscented KF. Unfortunately, solving NMPC problems in belief spaces is computationally expensive due to the curse of dimensionality. The problem was made tractable by neglecting the dynamics of the covariance through linearization around the last trajectory solution [5]. However, this linearization is only valid when the subsequent trajectory does not significantly change. The belief space planning approach SACBP operates in continuous-time through sequential action control (SAC) [9]. The authors model the belief dynamics as a hybrid dynamical system which we build upon. However, SACBP cannot ensure safety as constraints can only be included in the objective function.

In [10], invariance properties of deterministic CBFs are extended to systems that are described by stochastic differential equations (SDEs). The proposed stochastic CBFs handle uncertainty in the system's dynamics and also in the measurements. They guarantee safety under stochastic uncertainties by bounding the estimation errors in an EKF. This CBF is used for systems under sensor faults and attacks in [11] as well as risk-bounded control in highway scenarios in [12]. Probabilistic safety barrier certificates (PrSBC) have been proposed in [13] for multi-robot collision avoidance under uncertainty. However, the guarantees only hold for bounded uniform additive noise on the system dynamics and the observation model. Furthermore, the guarantees in [10]–[13] only hold for continuous-time observations, which is generally not consistent with real-world robotic systems, e.g., when global positioning data is not available.

II. PRELIMINARIES

A. Control Barrier Functions

Consider a dynamical system in control affine form

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{u} \quad (1)$$

with state $\mathbf{x} \in \mathcal{X} \subseteq \mathbb{R}^n$ and control input $\mathbf{u} \in \mathcal{U} \subseteq \mathbb{R}^m$. A safe set \mathcal{C} is constructed as the superlevel set of a continuously differentiable function $h: \mathcal{X} \rightarrow \mathbb{R}$ such that

$$\begin{aligned} \mathcal{C} &= \{\mathbf{x} \in \mathcal{X} \mid h(\mathbf{x}) \geq 0\}, \\ \partial\mathcal{C} &= \{\mathbf{x} \in \mathcal{X} \mid h(\mathbf{x}) = 0\}. \end{aligned} \quad (2)$$

Definition 1. A safe set \mathcal{C} is forward invariant with respect to the system (1) if for every initial condition $\mathbf{x}(t_0) \in \mathcal{C}$ it holds that $\mathbf{x}(t) \in \mathcal{C}, \forall t \geq t_0$.

A prominent approach to render a safe set forward invariant is to use CBFs.

Definition 2. Given a set \mathcal{C} , defined by Eq. (2), h serves as a zeroing CBF for the system (1) if $\forall \mathbf{x}$ satisfying $h(\mathbf{x}) \geq 0, \exists \mathbf{u} \in \mathcal{U}$ such that

$$\frac{\partial h}{\partial \mathbf{x}} (\mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{u}) \geq -h(\mathbf{x}). \quad (3)$$

In case a valid CBF exists, it follows that a controller satisfying Eq. (3) renders \mathcal{C} forward invariant [14]. However, this only holds for systems with relative degree $r_b = 1$ because $(\partial h / \partial \mathbf{x})\mathbf{g} = \mathbf{0}$ otherwise, thus \mathbf{u} would not appear in Eq. (3) [15]. For systems with $r_b > 1$, exponential CBFs (ECBFs) can be used to ensure forward invariance of \mathcal{C} . We leverage second order CBFs as a special case of ECBFs.

Definition 3. Given a set \mathcal{C} in Eq. (2) and a system (1) with relative degree $r_b = 2$, h is called an exponential CBF, if there exists a gain vector $\zeta \in \mathbb{R}^{r_b}$ and $\mathbf{u} \in \mathcal{U}$ such that

$$\ddot{h}(\mathbf{x}) + \zeta^T [h(\mathbf{x}) \quad \dot{h}(\mathbf{x})]^T \geq 0. \quad (4)$$

The gain vector can be obtained using classical tools from control theory such as, e.g., pole placement. A proof of forward invariance under ECBFs can be found in [15].

B. Gaussian Belief States

Gaussian filters are a family of state estimators that describe a Bayesian approach in which the belief is constrained to follow a multivariate Gaussian (MVG). Its probability density function (pdf) is described by

$$p(\mathbf{x}(t)) = \mathcal{N}(\boldsymbol{\mu}(t), \boldsymbol{\Sigma}(t)), \quad (5)$$

where $\boldsymbol{\mu} \in \mathbb{R}^n$ is the mean vector and $\boldsymbol{\Sigma} = \boldsymbol{\Sigma}^T \in \mathbb{R}^{n \times n}$ is the positive semidefinite covariance matrix. The pdf in (5) is uniquely described by the belief state $\mathbf{b} = [\boldsymbol{\mu}, \text{vec}(\boldsymbol{\Sigma})]^T$ where, due to symmetry, only the upper triangular matrix is stored which is encoded in the $\text{vec}(\cdot)$ operator [16]. Thus, the dimensionality of the belief state n_b increases quadratically with the state dimension n , i.e. $n_b = (n^2 + 3n)/2$. If not mentioned explicitly otherwise, we refer to the belief state \mathbf{b} as the belief.

C. Chance Constraints and Risk Measures

Chance constraints handle safety constraints under uncertainty by bounding the probability of undesired events. We use safety specifications in the form of half-spaces $\boldsymbol{\alpha}^T \mathbf{x} \geq \beta$. Consider a Gaussian distributed random variable \mathbf{x} with belief state \mathbf{b} . We calculate the probability of satisfying a half-space constraint as [5]

$$\Pr[\boldsymbol{\alpha}^T \mathbf{x} \geq \beta] = \frac{1}{2} \left(1 - \text{erf} \left(\frac{\boldsymbol{\alpha}^T \boldsymbol{\mu} - \beta}{\sqrt{2\boldsymbol{\alpha}^T \boldsymbol{\Sigma} \boldsymbol{\alpha}}} \right) \right), \quad (6)$$

where $\text{erf}(\cdot)$ is the standard error function.

To quantify the outcome of violating a chance constraint $\Pr[\boldsymbol{\alpha}^T \mathbf{x} \geq \beta] \geq 1 - \delta$, we use the Value-at-Risk (VaR).

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.

Definition 4. The VaR of a random variable $x \in \mathbb{R}$ with pdf $p(x)$ at level $\delta \in (0, 1]$ is the $(1 - \delta)$ quantile, i.e.

$$\text{VaR}_\delta(x) = \inf_{\tau \in \mathbb{R}} \{ \tau \mid \Pr[x \geq \tau] \geq 1 - \delta \} \quad (7)$$

VaR allows us to formulate constraints that are qualitatively equivalent to chance constraints, i.e.

$$\Pr[\alpha^T x \geq \beta] \geq 1 - \delta \Leftrightarrow \text{VaR}_\delta(\alpha^T x - \beta) \geq 0. \quad (8)$$

III. PROBLEM SETTING

We consider the robot's stochastic motion and observations in the form

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{u} + \mathbf{w}, \quad \mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}) \quad (9)$$

$$\mathbf{z}_k = \boldsymbol{\ell}(\mathbf{x}_k) + \mathbf{v}_k, \quad \mathbf{v}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{R}_k) \quad (10)$$

where $\mathcal{Z} \subseteq \mathbb{R}^\ell$ is the observation space and \mathbf{w}, \mathbf{v} are the motion and observation noise, respectively. We model the robot's motion as a continuous-time differential equation. The observations are always provided in discrete time due to the sensor's sampling time. Especially for exteroceptive sensors like GPS, measurements occur much less frequently than the robot's control rate, encouraging us to consider discrete-time formulations. A reference controller, e.g. a controller that drives the robot to a goal state, is given as \mathbf{u}_{ref} .

Problem 1. Given the stochastic dynamics in (9), a reference controller \mathbf{u}_{ref} , a safe set $\mathcal{C}_x = \{\mathbf{x} \in \mathcal{X} \mid \alpha^T \mathbf{x} \geq \beta\}$ defined in state space and a confidence level $\delta \in (0, 1]$, synthesize a control law \mathbf{u} that maps $\mathbf{x} \times \mathbf{u}_{\text{ref}} \mapsto \mathbf{u}$ such that at any time $\Pr[\mathbf{x}(t) \in \mathcal{C}_x] \geq 1 - \delta$.

Ideally, to solve Problem 1, we would use the Bayes filter to capture the exact time evolution of the belief given an initial belief $p(\mathbf{x}(t_0))$, controls \mathbf{u} and observations \mathbf{z} . However, exact belief calculations only exist in specialized cases which is why approximations need to be considered [6]. Thus, we use an EKF as a tractable implementation of the Bayes filter in which beliefs are Gaussian. While the Gaussian belief of an EKF is exact for linear systems, it is only an approximation of the true belief in the nonlinear case. This assumption, however, is common in many practical scenarios [7], especially if the true probability distribution is unimodal [6]. In future work, we aim to consider the mismatch between the modeled belief and the true belief.

To approach Problem 1, we reason about Gaussian belief states \mathbf{b} instead of states \mathbf{x} so that we can solve a relaxed problem under stochastic uncertainties. We propagate the belief through the nonlinear model in Eq. (9)-(10) by exploiting the fundamental EKF step that systems are linearized around the current mean.

Assumption 1. To propagate the mean and covariance of a random variable \mathbf{x} through a nonlinear function $\boldsymbol{\eta}$, we use a first-order Taylor series expansion

$$\begin{aligned} \mathbb{E}\{\boldsymbol{\eta}(\mathbf{x})\} &\approx \boldsymbol{\eta}(\mathbb{E}\{\mathbf{x}\}) \\ \text{Var}\{\boldsymbol{\eta}(\mathbf{x})\} &\approx \left(\frac{\partial \boldsymbol{\eta}}{\partial \mathbf{x}}(\mathbb{E}\{\mathbf{x}\}) \right) \text{Var}\{\mathbf{x}\} \left(\frac{\partial \boldsymbol{\eta}}{\partial \mathbf{x}}(\mathbb{E}\{\mathbf{x}\}) \right)^T. \end{aligned}$$

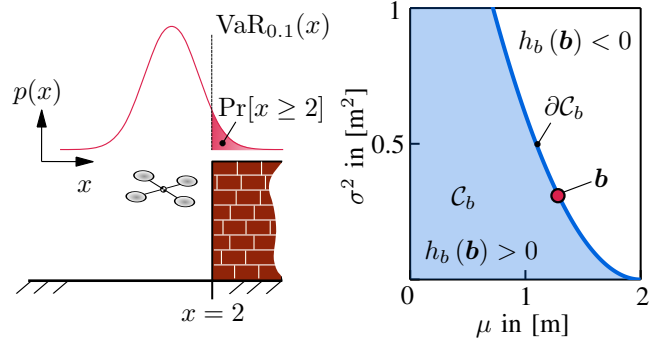


Fig. 2. A drone with uncertain position x and its pdf $p(x)$ is moving in one dimension. A safety specification is defined as a bounded probability of collision with the wall. The resulting safe set over belief states is shown in blue and the current belief of the robot is depicted in red.

We translate the safety specification over states in Problem 1 to a safe set over belief states

$$\mathcal{C}_b = \{\mathbf{b} \in \mathbb{R}^{n_b} \mid h_b(\mathbf{b}) \geq 0\}, \quad (11)$$

where h_b defines a risk-aware half-space

$$\begin{aligned} h_b(\mathbf{b}) &:= \text{VaR}_\delta(\alpha^T \mathbf{x} - \beta) \\ &= \alpha^T \boldsymbol{\mu} - \beta - \text{erf}^{-1}(1 - 2\delta) \sqrt{2\alpha^T \boldsymbol{\Sigma} \alpha}. \end{aligned} \quad (12)$$

Problem 2. Given the model (9)-(10) and a risk-aware safe set \mathcal{C}_b defined over beliefs, find a control input \mathbf{u} that renders \mathcal{C}_b forward invariant.

By solving Problem 2, we ensure that the belief satisfies a VaR formulation which is qualitatively equivalent (see Eq. (8)) to satisfying the chance constraint in Problem 1.

Example. Consider a drone operating in one dimension with position $x \sim \mathcal{N}(\mu, \sigma^2)$, as shown in Fig. 2 and the Gaussian belief state $\mathbf{b} = [\mu, \sigma^2]^T \in \mathbb{R}^2$. A safety specification over states is to stay within the collision-free space with 90% probability, given by $\delta = 0.1$ and $\mathcal{C}_x = \{x \in \mathbb{R} \mid x \leq 2\}$. The corresponding safe set \mathcal{C}_b in belief space is defined through

$$h(\mathbf{b}) = \text{VaR}_\delta(2 - x) = 2 - \mu + \text{erf}^{-1}(1 - 2\delta) \sqrt{2\sigma^2} \geq 0,$$

which is illustrated in blue in Fig. 2. In the depicted point in time, the belief state, shown as red circle, is right at the boundary of the safe set $\partial \mathcal{C}_b$ as the probability of colliding with the wall is exactly 10%. Solving Problem 2 keeps the belief state (red circle) in \mathcal{C}_b which in turn satisfies the original safety specification over states that 90% of the drone's probability mass should be left of the wall.

IV. RISK-AWARE CONTROL

Our solution to Problem 2 is divided into two main steps. First, we derive a stochastic hybrid system that accounts for uncertainties in the robot's motion and observations. Given this hybrid system describing the evolution of the robot's belief, we propose BCBFs to ensure forward invariance of our safe set defined over beliefs. BCBFs serve as a computationally efficient risk-aware safety filter that can be applied to various robotic systems under stochastic uncertainties.

A. Hybrid Belief Dynamics

In the belief dynamics, it is important to distinguish between two evolutions: 1) The belief advances in continuous time as the state dynamics (9) also evolve in continuous time. 2) At discrete timesteps t_k when a new sensor observation is available, the belief changes instantaneously since the state distribution is conditioned on the measurement. This naturally leads to a hybrid dynamical system which is derived from a continuous-discrete EKF [9], [17].

1) *Continuous-time Belief Dynamics:* The continuous-time evolution of the belief directly follows from the derivation of the continuous EKF and is given by a set of ordinary differential equations (ODEs)

$$\dot{\boldsymbol{\mu}} = \mathbf{f}(\boldsymbol{\mu}) + \mathbf{g}(\boldsymbol{\mu}) \mathbf{u}, \quad (13)$$

$$\dot{\boldsymbol{\Sigma}} = \mathbf{A}\boldsymbol{\Sigma} + \boldsymbol{\Sigma}\mathbf{A}^T + \mathbf{Q}, \quad (14)$$

where \mathbf{A} is the Jacobian of the noise-free motion model in Eq. (9) evaluated at the current mean $\boldsymbol{\mu}$. For the design of BCBFs, the dynamics are required to be in control-affine form which has been proven to be true in [9].

Since the belief vector \mathbf{b} is comprised of the mean and the covariance matrix, we can express the continuous-time belief dynamics as a vector-valued ODE

$$\dot{\mathbf{b}} = \begin{bmatrix} \dot{\boldsymbol{\mu}} \\ \text{vec}(\dot{\boldsymbol{\Sigma}}) \end{bmatrix} = \underbrace{\begin{bmatrix} \mathbf{f}(\boldsymbol{\mu}) \\ \mathbf{f}_{\boldsymbol{\Sigma}}(\mathbf{b}) \end{bmatrix}}_{\mathbf{f}_b(\mathbf{b})} + \underbrace{\begin{bmatrix} \mathbf{g}(\boldsymbol{\mu}) \\ \mathbf{g}_{\boldsymbol{\Sigma}}(\mathbf{b}) \end{bmatrix}}_{\mathbf{g}_b(\mathbf{b})} \mathbf{u}, \quad (15)$$

where $\mathbf{f}_{\boldsymbol{\Sigma}} = \text{vec}(\mathbf{F}_{\boldsymbol{\Sigma}})$ and $\mathbf{g}_{\boldsymbol{\Sigma}} = \text{vec}(\mathbf{G}_{\boldsymbol{\Sigma}})$ are the vector fields of Eq. (14) in control-affine form. Note that $\mathbf{G}_{\boldsymbol{\Sigma}} \in \mathbb{R}^{n \times n \times m}$ and, thus, the $\text{vec}(\cdot)$ operator maps it to a matrix $\mathbf{g}_{\boldsymbol{\Sigma}} \in \mathbb{R}^{n(n+1)/2 \times m}$.

Consequently, if the control signal \mathbf{u} is known, the belief at any point in time is obtained by forward integrating the belief dynamics in Eq. (15). However, when a new sensor reading is available, the belief changes instantaneously which is covered in the discrete-time Kalman update.

2) *Discrete-time Kalman Update:* Every time a new observation z_k is available, we obtain the posterior distribution $p(\mathbf{x}(t_k^+) | \mathbf{z}_{1:k-1}, z_k) = \mathcal{N}(\boldsymbol{\mu}^+, \boldsymbol{\Sigma}^+)$ by conditioning the prior distribution parameterized by $\mathbf{b}^- = \mathbf{b}(t_k^-)$ on z_k . Note that t_k^- is infinitesimal smaller than $t_k = t_k^+$. The conditioning leads to a discrete belief transition which is governed by the discrete-time Kalman update [6]

$$\mathbf{b}^+ = \boldsymbol{\Delta}(\mathbf{b}^-) = \begin{bmatrix} \boldsymbol{\mu}^- + \mathbf{K}(z_k - \ell(\boldsymbol{\mu}^-)) \\ \text{vec}((\mathbf{I} - \mathbf{K}\mathbf{H})\boldsymbol{\Sigma}^-) \end{bmatrix}, \text{ where} \quad (16)$$

$$\mathbf{K} = \boldsymbol{\Sigma}^- \mathbf{H}^T (\mathbf{H}\boldsymbol{\Sigma}^- \mathbf{H}^T + \mathbf{R})^{-1} \text{ and } \mathbf{H} = \left. \frac{\partial \ell}{\partial \mathbf{x}} \right|_{\boldsymbol{\mu}^-}$$

denote the Kalman gain and the Jacobian of the observation model, respectively. The measurement obtained at a discrete timestep is not known in advance which makes it generally difficult to do planning or control in belief space. However, the unknown measurement can be treated as a random variable making the measurement update (16) stochastic [16].

Proposition 1. *Under Assumption 1, the innovation term $\boldsymbol{\theta} = \mathbf{K}(z_k - \ell(\boldsymbol{\mu}^-))$ in Eq. (16) is a random variable with distribution $\boldsymbol{\theta} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Lambda})$ where*

$$\boldsymbol{\Lambda} = \mathbf{K}(\mathbf{H}\boldsymbol{\Sigma}\mathbf{H}^T + \mathbf{R})\mathbf{K}^T. \quad (17)$$

The proof is straight-forward by applying Assumption 1 to the innovation term and is omitted for brevity.

Finally, combining the ODE in Eq. (15) and the discrete-time Kalman update in Eq. (16) leads to the hybrid system

$$\mathcal{S} = \begin{cases} \dot{\mathbf{b}} = \mathbf{f}_b(\mathbf{b}) + \mathbf{g}_b(\mathbf{b}) \mathbf{u}, & \forall t \in [t_{k-1}, t_k) \\ \mathbf{b}^+ = \boldsymbol{\Delta}(\mathbf{b}^-) & t = t_k, \end{cases} \quad (18)$$

describing the evolution of the robot's belief over time. Fig. 3 illustrates an example trajectory of the belief. Our formulation has the advantage that we can allow arbitrary sampling times $\delta t = t_k - t_{k-1}$. This is particularly useful for the analysis of scenarios where e.g., a sensor does not provide any information for a certain duration as in the case of GPS data inside a building.

B. Belief Control Barrier Functions

The dynamical system \mathcal{S} of the belief allows us to introduce BCBFs for safe control under stochastic uncertainties. BCBFs are defined similarly as CBFs (2), but they are defined over beliefs instead of states.

Definition 5. (BCBF) *Given a safe set \mathcal{C}_b defined by Eq. (11), $h_b(\mathbf{b})$ serves as a Belief Control Barrier Function (BCBF) for the stochastic dynamical system (9)-(10) if $\forall \mathbf{b}$ satisfying $h_b(\mathbf{b}) \geq 0, \exists \mathbf{u} \in \mathcal{U}$ such that*

$$\frac{\partial h_b}{\partial \mathbf{b}}(\mathbf{f}_b(\mathbf{b}) + \mathbf{g}_b(\mathbf{b}) \mathbf{u}) \geq -h_b(\mathbf{b}). \quad (19)$$

Using this definition of a BCBF, the following theorem provides a condition for the forward invariance of the safe set \mathcal{C}_b under the continuous-time evolution of the belief.

Theorem 1. *If a locally Lipschitz control input $\mathbf{u}(t)$ satisfies Eq. (19) $\forall t \in [t_k, t_{k+1})$ for a given safe set with valid BCBF $h_b(\mathbf{b})$, then $\Pr[h_b(\mathbf{b}(t)) \geq 0] = 1$, provided that $\mathbf{b}(t_k) \in \mathcal{C}_b$.*

The proof is straight-forward since we only reason about the continuous-time evolution in the interval $[t_k, t_{k+1})$ in which the belief dynamics in Eq. (15) are deterministic. Thus, interested readers are referred to the proof for traditional state-based CBFs [14].

Next, we analyze the effects of measurements on the forward invariance of \mathcal{C}_b . Since measurements are treated as random variables, we do not have control of the discrete belief transition. As a result, outlier measurements may cause the belief to leave the safe set \mathcal{C}_b even though the considered state is safe, see Fig. 3. We show how to ensure forward invariance even in the presence of outlier measurements. To that end, we exploit the asymptotic stability of zeroing CBFs.

Remark 1. *Our BCBFs are zeroing CBFs that not only render \mathcal{C}_b forward invariant but also asymptotically stable as the continuous-time belief dynamics are deterministic [1]. Thus, if an initial belief is outside the safe set, it will be driven back to \mathcal{C}_b over time.*

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.

In the following theorem, we derive an upper bound on the probability of leaving \mathcal{C}_b under a discrete transition.

Theorem 2. *If the control input $\mathbf{u}(t)$ satisfies Eq. (19), the probability of leaving the safe set under a discrete transition, i.e. $\Pr[h_b(\mathbf{b}^+) < 0]$, is bounded by*

$$\Pr[h_b(\mathbf{b}^+) < 0] \leq \frac{1}{2} \left(1 - \operatorname{erf} \left(\frac{\xi(\mathbf{b}^-)}{\sqrt{2\boldsymbol{\alpha}^T \boldsymbol{\Lambda} \boldsymbol{\alpha}}} \right) \right), \quad (20)$$

with $\boldsymbol{\Lambda}$ defined as in Proposition 1 and

$$\xi(\mathbf{b}^-) = \operatorname{erf}^{-1}(1 - 2\delta) \left(\sqrt{2\boldsymbol{\alpha}^T \boldsymbol{\Sigma}^- \boldsymbol{\alpha}} \right) \quad (21)$$

$$- \sqrt{2\boldsymbol{\alpha}^T (\mathbf{I} - \mathbf{K}\mathbf{H}) \boldsymbol{\Sigma}^- \boldsymbol{\alpha}}. \quad (22)$$

Proof: Since \mathbf{u} satisfies Eq. (19) we know that $h_b(\mathbf{b}^-) = h_b(\mathbf{b}(t_k^-)) \geq 0$. Consequently, the value of the BCBF is

$$h_b(\mathbf{b}^+) = \boldsymbol{\alpha}^T \boldsymbol{\mu}^+ - \beta - \operatorname{erf}^{-1}(1 - 2\delta) \sqrt{2\boldsymbol{\alpha}^T \boldsymbol{\Sigma}^+ \boldsymbol{\alpha}}$$

in which we replace the posterior mean $\boldsymbol{\mu}^+$ by its stochastic update equation from Proposition 1. Further, we add a covariance term that adds up to zero,

$$h_b(\mathbf{b}^+) = \boldsymbol{\alpha}^T (\boldsymbol{\mu}^- + \boldsymbol{\theta}) - \beta - \operatorname{erf}^{-1}(1 - 2\delta) \sqrt{2\boldsymbol{\alpha}^T \boldsymbol{\Sigma}^+ \boldsymbol{\alpha}} + \operatorname{erf}^{-1}(1 - 2\delta) \left(\sqrt{2\boldsymbol{\alpha}^T \boldsymbol{\Sigma}^- \boldsymbol{\alpha}} - \sqrt{2\boldsymbol{\alpha}^T \boldsymbol{\Sigma}^+ \boldsymbol{\alpha}} \right).$$

which we rearrange to obtain $h_b(\mathbf{b}^-)$. This yields a Gaussian distribution over possible BCBF values

$$h_b(\mathbf{b}^+) = h_b(\mathbf{b}^-) + \boldsymbol{\alpha}^T \boldsymbol{\theta} + \xi(\mathbf{b}^-). \quad (23)$$

Finally, the probability that an outlier measurement causes the belief to leave \mathcal{C}_b is calculated using (6) as

$$\begin{aligned} \Pr[h_b(\mathbf{b}^+) < 0] &= \Pr[-\boldsymbol{\alpha}^T \boldsymbol{\theta} \geq h_b(\mathbf{b}^-) + \xi(\mathbf{b}^-)] \\ &= \frac{1}{2} \left(1 - \operatorname{erf} \left(\frac{\xi(\mathbf{b}^-) + h_b(\mathbf{b}^-)}{\sqrt{2\boldsymbol{\alpha}^T \boldsymbol{\Lambda} \boldsymbol{\alpha}}} \right) \right) \\ &\leq \frac{1}{2} \left(1 - \operatorname{erf} \left(\frac{\xi(\mathbf{b}^-)}{\sqrt{2\boldsymbol{\alpha}^T \boldsymbol{\Lambda} \boldsymbol{\alpha}}} \right) \right) \quad \blacksquare \end{aligned}$$

Theorem 2 provides us with a natural bound on the probability of leaving the safe set \mathcal{C}_b under discrete sensor observations when the belief state \mathbf{b}^- is on the boundary $\partial\mathcal{C}_b$. The derived natural bound only depends on our observation model $\ell(\mathbf{x})$, the noise covariance \mathbf{R} and the prior covariance $\boldsymbol{\Sigma}^-$. Thus, the specific bound varies for different robotic systems and sensor types.

Example. (Cont.) Consider the one-dimensional drone example introduced before. The drone is equipped with a sensor that provides noisy state observations $z_k = x(t_k) + v$ where $v \sim \mathcal{N}(0, r = 0.1)$. The prior variance is $\sigma^- = 0.3$ and a confidence level is set to $1 - \delta = 0.99$, then

$$\begin{aligned} \Pr[h_b(\mathbf{b}^+) < 0] &\leq \\ &\frac{1}{2} \left(1 - \operatorname{erf} \left(\operatorname{erf}^{-1}(1 - 2\delta) \left(\frac{\sqrt{\sigma^- + r} - \sqrt{r}}{\sqrt{\sigma^-}} \right) \right) \right) \approx 0.09 \end{aligned}$$

meaning that the probability of leaving \mathcal{C}_b is at most 9%. Note that this is a tight bound if the prior belief state is right at the boundary of the safe set and strictly smaller otherwise.




Fig. 3. An example belief trajectory subject to an outlier measurement and the corresponding safe set shown as half-space. The initial belief \mathbf{b}^- is within the safe set \mathcal{C}_b whereas the posterior belief \mathbf{b}^+ leaves the safe set under the discrete transition. An asymptotically stable belief trajectory is shown in red while the true evolution of the state is shown in blue. The plot on the right shows value of $h_b(\mathbf{b})$, depicted in green.

Figure 3 shows a scenario in which the posterior belief \mathbf{b}^+ leaves the safe set due to an outlier measurement. Although the state \mathbf{x} still satisfies the original half-space constraint at time t_k , the belief switches instantaneously from $h_b(\mathbf{b}^-) \geq 0$ to $h_b(\mathbf{b}^+) < 0$. This highlights the importance of keeping the belief state inside \mathcal{C}_b : Initially, the state \mathbf{x} is considered to be safe, i.e. $\mathbf{x} \in \mathcal{C}_x$, but as time evolves, the only statement we can make is that the belief trajectory $\mathbf{b}(t)$ is asymptotically stable, see Remark 1. However, during this period \mathbf{x} could leave \mathcal{C}_x while the belief is driven back to \mathcal{C}_b as shown in Fig. 3. We thus need to ensure that the belief state does not leave \mathcal{C}_b in the first place, in order to guarantee that \mathbf{x} remains in \mathcal{C}_x with desired probability.

To ensure that the belief state does not leave \mathcal{C}_b , we modify \mathcal{C}_b such that the probability of leaving \mathcal{C}_b is bounded with a desired confidence. We define an augmented safe set $\tilde{\mathcal{C}}_b \subseteq \mathcal{C}_b$

$$\tilde{\mathcal{C}}_b = \left\{ \mathbf{b} \in \mathbb{R}^{n_b} \mid \tilde{h}_b(\mathbf{b}) \geq 0 \right\}, \quad (24)$$

with $\tilde{h}_b(\mathbf{b}) = h_b(\mathbf{b}) - \gamma$ for some $\gamma \geq 0$. This safe set is essentially a shrunk version of the original safe set \mathcal{C}_b . In the following theorem, we choose γ such that the belief stays within \mathcal{C}_b with desired probability.

Theorem 3. *The belief state \mathbf{b}^+ remains in the safe set \mathcal{C}_b with probability $\Pr[\mathbf{b}^+ \in \mathcal{C}_b] \geq 1 - \varepsilon$ under the discrete reset map $\mathbf{b}^+ = \boldsymbol{\Delta}(\mathbf{b}^-)$ if the control input $\mathbf{u}(t)$ satisfies Eq. (19) for the function $h_b(\mathbf{b})$ (thus \tilde{h}_b is a valid BCBF) and $\tilde{h}_b(\mathbf{b}^-) \geq 0$ for*

$$\gamma \geq \sqrt{2\boldsymbol{\alpha}^T \boldsymbol{\Lambda} \boldsymbol{\alpha}} \left(\operatorname{erf}^{-1}(1 - 2\varepsilon) - \xi(\mathbf{b}^-) \right).$$

Proof: Similarly as for Thm. 2, we calculate the probability of staying in the safe set \mathcal{C}_b using Eq. (23),

$$\begin{aligned} \Pr[h(\mathbf{b}^+) \geq 0] &= \frac{1}{2} \left(1 + \operatorname{erf} \left(\frac{\xi(\mathbf{b}^-) + h(\mathbf{b}^-)}{\sqrt{2\boldsymbol{\alpha}^T \boldsymbol{\Lambda} \boldsymbol{\alpha}}} \right) \right) \\ &\geq \frac{1}{2} \left(1 + \operatorname{erf} \left(\frac{\xi(\mathbf{b}^-) + \gamma}{\sqrt{2\boldsymbol{\alpha}^T \boldsymbol{\Lambda} \boldsymbol{\alpha}}} \right) \right) = 1 - \varepsilon. \quad \blacksquare \end{aligned}$$

Note that Theorem 3 provides probabilistic safety guarantees for the case that $\mathbf{b}^- \in \tilde{\mathcal{C}}_b$. If the posterior belief state ends up in the set $\mathcal{C}_b \setminus \tilde{\mathcal{C}}_b$, it will be driven back to the augmented safe set if \tilde{h}_b is a valid BCBF. If an outlier measurement occurs during this asymptotically stable convergence to $\tilde{\mathcal{C}}_b$, the natural bound in Theorem 2 holds. Once $\mathbf{b} \in \tilde{\mathcal{C}}_b$, the desired confidence in Theorem 3 holds again.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.

C. Risk-aware control synthesis

To synthesize risk-aware control inputs under stochastic uncertainties arising from state estimation, we formulate the quadratic program (QP)

$$\begin{aligned} \mathbf{u}^* &= \arg \min_{\mathbf{u} \in \mathcal{U}} (\mathbf{u} - \mathbf{u}_{\text{ref}})^T (\mathbf{u} - \mathbf{u}_{\text{ref}}) \\ \text{s.t.} \quad & \frac{\partial \tilde{h}}{\partial \mathbf{b}} (\mathbf{f}_b(\mathbf{b}) + \mathbf{g}_b(\mathbf{b}) \mathbf{u}) \geq -\tilde{h}(\mathbf{b}), \end{aligned} \quad (25)$$

where \mathbf{u}_{ref} is a reference control input. This synthesis problem can be solved efficiently using off-the-shelf QP solvers.

Consider for example a quadrotor operating in 3D: It has 6 degrees of freedom (3D pose and orientation) and 4 control inputs (thrusters) which leads to a 12 dimensional state. Consequently, the belief state is of dimension $\mathbf{b} \in \mathbb{R}^{90}$ which showcases the curse of dimensionality and complicates the use of common optimization-based controllers such as MPC. We overcome the dimensionality problem of belief spaces by only optimizing over control inputs \mathbf{u} which enables real-time applicability. The resulting control inputs satisfy the entire belief dynamics.

V. EXPERIMENTS

We evaluate our BCBFs in challenging scenarios in which uncertainties cannot be neglected¹. Specifically, we show

- 1) improved adherence to safety specifications through risk-awareness and computational efficiency over stochastic CBFs and CCNMPC in simulations in Sec. V-A,
- 2) safety under changing sensing conditions, such as sensor rates and measurement variances, in Sec. V-C,
- 3) robustness to external disturbances as well as the generality to different sensor systems in Sec. V-D.

A. Safety analysis of BCBFs

a) *Setup:* We compare our BCBFs to two baselines in an obstacle avoidance scenario. We consider a unicycle robot [18] with state $\mathbf{x} = [p_x, p_y, v, \varphi]^T$ and dynamics and observations

$$\begin{aligned} \dot{p}_x &= v \cos(\varphi) + w_x, & \dot{v} &= a + w_v \\ \dot{p}_y &= v \sin(\varphi) + w_y, & \dot{\varphi} &= \omega + w_\varphi \\ \mathbf{z}_k &= \mathbf{x}_k + \mathbf{v}_k, \mathbf{v}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{R}). \end{aligned} \quad (26)$$

The nominal dynamics are corrupted by Gaussian noise $[w_x, w_y, w_v, w_\varphi]^T \sim \mathcal{N}(\mathbf{0}, \mathbf{Q})$ with motion noise $\mathbf{Q} = \text{diag}([0.1^2, 0.1^2, 0.005^2, 0.005^2])$ and observation noise $\mathbf{R} = \text{diag}([0.2^2, 0.2^2, 0.1^2, 0.1^2])$ with a sensor update rate of 10 Hz. The control input is given as $\mathbf{u} = [a, \omega]^T$. We obtain the belief dynamics in Eq. (18) by applying the method described in Sec. IV-A where $\mathbf{b} \in \mathbb{R}^{14}$. The objective is to steer the robot towards a goal at $[8, 0]$ while avoiding collisions with a circular obstacle $\mathcal{O} = \{\mathbf{p} \in \mathbb{R}^2 \mid \|\mathbf{p} - \mathbf{c}\|_2 < r\}$ where $\mathbf{c} = [5, 0]^T$ and $r = 1$ as shown in Fig. 4. The safety specification for collision avoidance is given as

$$\Pr[\mathbf{x} \notin \mathcal{O}] = \Pr[\|\mathbf{x} - \mathbf{c}\|_2 - r \geq 0] \geq 1 - \delta \quad (27)$$

¹Video available at https://www.youtube.com/watch?v=va_wXkasOPQ

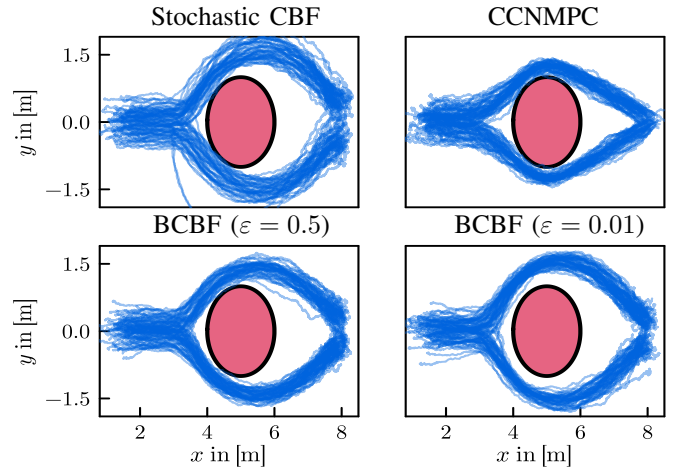


Fig. 4. Comparison of different control strategies for a two-dimensional avoidance scenario. The initial state follows a Gaussian distribution $\mathbf{x}_0 \sim \mathcal{N}([0, 1.5, 0, 0], \text{diag}([0.3^2, 0.2^2, 0.1^2, 0.1^2]))$ and the goal is at $[8, 0]$. Simulated trajectories are shown in blue and the obstacle is depicted in red.

with $\delta = 0.01$. Due to the non-linearity of the obstacle description, we use the common approach to linearize Eq. (27) around the mean state which is a strict overapproximation of the true collision condition as shown in [5]. The resulting risk-aware safe set is given by

$$\mathcal{C}_b = \{\mathbf{b} \in \mathbb{R}^{14} \mid \text{VaR}_\delta(\boldsymbol{\alpha}^T(\mathbf{x} - \mathbf{c}) - r) \geq 0\} \quad (28)$$

where $\text{VaR}_\delta(\cdot)$ with $\boldsymbol{\alpha} = (\boldsymbol{\mu} - \mathbf{c}) / \|\boldsymbol{\mu} - \mathbf{c}\|_2$ serves as a BCBF candidate. Due to a relative degree $r_b = 2$ in our BCBF candidate, we use second order formulations as in Eq.(4).

We solve the resulting SDE (26) using the DifferentialEquations.jl package [19] in the Julia programming language. The reference control \mathbf{u}_{ref} is given by a Linear-quadratic Regulator (LQR) that uses the mean estimate $\boldsymbol{\mu}$ to steer the system towards the goal without any knowledge of the obstacle. The parameters chosen for the LQR are $\mathbf{Q}_L = \text{diag}[10, 5, 5, 5]$ and $\mathbf{R}_L = \text{diag}[5, 10]$. We compare two versions of BCBFs, namely one with only the natural bound ($\varepsilon = 0.5$) and one bounding the probability of leaving \mathcal{C}_b by $\varepsilon = 0.01$.

b) *Baselines:* We compare our BCBFs to two different baselines: 1.) Stochastic CBFs with incomplete state information which require cont.-time observations as well as a bounded estimation error. To enable continuous estimation, we use maximum likelihood observations between actual measurements which is not true in practice. The bounded estimation error was obtained by running a monte carlo study and take the maximum error. 2.) chance-constrained NMPC (CCNMPC) [5] that formulates an MPC problem which includes Eq. (28) as a constraint. Since MPC operates in discrete-time while we simulate a continuous system, we use a zero order hold to apply the control input. We run the MPC at 30 Hz and use a planning horizon of $N = 40$. The cost function is the same quadratic cost as in the reference LQR controller. 3.) In hardware experiments, we compare to a traditional state CBF that only considers the mean dynamics.

c) *Safety and computational efficiency:* Figure 4 shows 100 simulated trajectories for different initial conditions sampled from a Gaussian while the quantitative results are summa-

TABLE I
COMPARISON OF BCBFs TO BASELINES.

	Stochastic CBF	CCNMPC	BCBF ($\varepsilon = 0.5$)	BCBF ($\varepsilon = 0.01$)
# collisions	3	3	1	0
% $h(\mathbf{b}) < 0$	6.52	0.78	0.42	0
Avg. t_{comp} in [s] ($\mu \pm \sigma$)	0.0008 \pm 0.004	0.35 \pm 0.052	0.0008 \pm 0.003	0.0017 \pm 0.0027
Avg. $\ \mathbf{u}\ _2$ ($\mu \pm \sigma$)	1.63 \pm 2.54	0.67 \pm 0.4	1.62 \pm 1.65	1.52 \pm 1.6
t_g in [s] ($\mu \pm \sigma$)	4.01 \pm 0.5	3.35 \pm 0.19	4.1 \pm 0.41	4.16 \pm 0.44

rized in Table I. BCBFs have the lowest number of collisions (0 for $\varepsilon = 0.01$) and outperform both baselines. Interestingly, SCBFs cannot ensure safety almost surely as the assumption of cont-time observations does not apply. Similarly, CCNMPC results in multiple collisions although it tries to satisfy the same safety specification. This is due to the fact that the optimization problem sometimes becomes infeasible in which case we need to relax the safety constraint. In [5], the authors reported that the percentage of infeasible solutions in their experiments was 2.8%. Infeasibility is not an issue for BCBFs, since they allow to formulate the control synthesis as QP over controls instead of a nonlinear program over states and controls, overcoming the curse of dimensionality. QP formulations are computationally efficient and can be run with up to 1kHz, see Table I.

d) Effect of augmented safe set: We analyze the effect of the augmented safe set defined in Theorem 3. If no additional safety margin is used ($\varepsilon = 0.5$), the belief leaves \mathcal{C}_b in 0.42% of all simulated trajectories. When bounding the probability of leaving the safe set by $\Pr[h(\mathbf{b}^+) < 0] \leq 0.01$, the belief remains in \mathcal{C}_b for all trajectories while keeping larger distances to the obstacle.

e) Control efficiency: Lastly, we evaluate the efficiency in terms of the average norm of the control input and the time t_g to reach the goal. In this case CCNMPC is superior to all other approaches since it optimizes controls over a planning horizon whereas CBF approaches are purely reactive. This motivates to combine state-based MPC with BCBFs to achieve both, control efficiency and rigorous safety properties. In future work, we aim to explore this.

B. Setup of hardware experiment

We use the Bitcraze Crazyflie 2.1 quadrotor inside an OptiTrack motion capture (mocap) system, as shown in Fig. 1. The state of the drone is given as its 3D position \mathbf{p} and velocity $\dot{\mathbf{p}}$. Its motion and observations are modeled as

$$\dot{\mathbf{x}} = \begin{bmatrix} \dot{\mathbf{p}} \\ \dot{\mathbf{p}} \end{bmatrix} = \begin{bmatrix} \mathbf{0} & \mathbf{I} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{p} \\ \dot{\mathbf{p}} \end{bmatrix} + \begin{bmatrix} \mathbf{0} \\ \mathbf{I} \end{bmatrix} \mathbf{u} + \mathbf{w}, \quad (29)$$

$$\mathbf{z}_k = \mathbf{p}_k + \mathbf{v}_k \quad (30)$$

where $\mathbf{u} \in \mathbb{R}^3$ are the drone's desired accelerations and $\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{Q})$ with a variance of 0.05^2 along the diagonal of \mathbf{Q} . In our setting, we can only observe the drone's position \mathbf{p} . We add Gaussian noise $\mathbf{v} \sim \mathcal{N}(\mathbf{0}, \mathbf{R})$ to the mocap data and reduce the sampling rate of the measurements. These

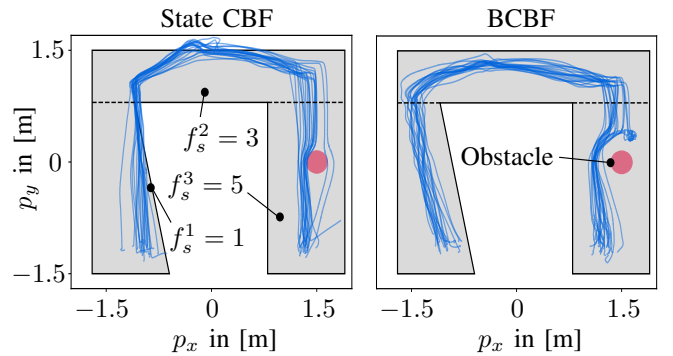


Fig. 5. Ground truth trajectories of 20 different runs of Exp. I for CBFs and BCBFs. The sampling rate f_s changes throughout the corridor segments. The measurement variances are given as $\mathbf{R}_1 = \text{diag}[0.08^2, 0.08^2]$, $\mathbf{R}_2 = \text{diag}[0.2^2, 0.2^2]$ and $\mathbf{R}_3 = \text{diag}[0.15^2, 0.15^2]$, respectively.

parameters \mathbf{R} and f_s are varied throughout the experiments. We use the uncorrupted mocap data as ground truth to evaluate the performance of BCBFs.

We use the EKF in Sec. IV-A to obtain the belief dynamics with $\mathbf{b} \in \mathbb{R}^{27}$, where the ODE is discretized using Euler's scheme. We convert the desired acceleration \mathbf{u} synthesized by our BCBFs into a setpoint consisting of desired roll and pitch angles as well as thrust for the real quadrotor using [20]. Setpoints are sent to the Crazyflie at 100 Hz and tracked by the onboard PID controller which serves as \mathbf{u}_{ref} . We compare BCBFs only to traditional CBFs as we were not able to run CCNMPCs in real-time with our setup.

C. Experiment I - Changing sensing conditions

In our first experiment, we showcase the robustness of BCBFs to changing sensing conditions. For that purpose, we navigate the drone along a U-shaped corridor as shown in Fig. 5. The corridor is given as the union of three polytopes. In each of the three polytopes, we change the measurement noise \mathbf{R} and sampling rates f_s according to Fig. 5. In the third polytope, we additionally place a circular obstacle to increase the difficulty of the scenario. The safety specification is given as an intersection of risk-aware half-spaces with a collision probability of $\delta = 0.05$. The half-spaces are given by the polytope representation of the corridor.

Figure 5 shows the ground truth trajectories of 20 different runs for a state-based CBF and a BCBF. It can be observed that the state-based approach cannot satisfy the safety specification whereas all 20 trajectories are collision free for the BCBF. Interestingly, there are no trajectories for the BCBF passing the obstacle on the right side since it is too risky to navigate through the narrow passage. As a consequence, the drone stops in front of the obstacle and slowly moves around the left side which took about 5-10 seconds. Although this behavior ensures safety, it highlights that a robot can get stuck in a local minimum when using a CBF approach. This could be overcome by combining BCBFs with a motion planner such as MPC.

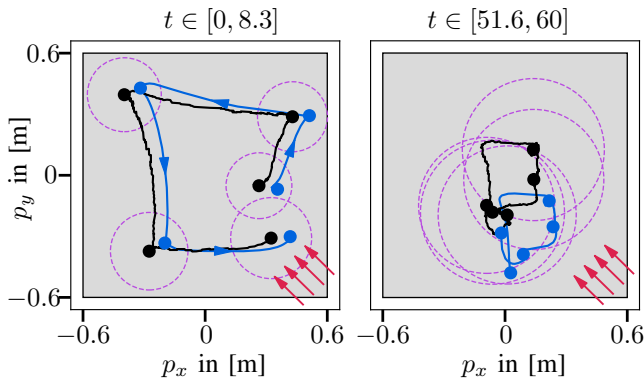


Fig. 6. Illustration of Exp. II. The drone should stay within the safety region colored in grey while a wind disturbance (shown as red arrows) is acting on the drone. The ground truth trajectory is shown in black and the mean belief trajectory is shown in blue. The 95% confidence ellipses at selected points in time are depicted in purple.

D. Experiment II - External disturbances

In this experiment illustrated in Figure 1, we additionally study the effect of disturbances on the nominal dynamics. To that end, we define a safe set as a cuboid that we want the drone to stay in with 95% probability. The drone is exposed to both, sensing uncertainties as well as an external wind disturbance from a fan. Instead of position measurements, we measure the drone's velocity through an added optical flow sensor. These velocity measurements are given by

$$z_k = \dot{p}_k + v_k, \quad v \sim \mathcal{N}(\mathbf{0}, \text{diag}(0.2^2)).$$

Since we only measure the drone's velocity and neglect the global position measurements from mocap, there is an inevitable drift in the position estimate obtained from integrating the velocity. A human operator is generating reference acceleration commands u_{ref} using a gamepad and actively tries to steer the drone outside the safe set.

Figure 6 shows the mean belief trajectory as well as ground truth for the described setting. The ellipses indicate the position uncertainty at selected points in time. The drone moves towards the corners of the safety region in which the CBF prevents the belief from leaving the safe set. Over time, the uncertainty in the position estimate grows due to the lack of a global positioning system and, thus, the drone gets increasingly cautious. The mean belief moves further towards the center of the safety region as this increases the likelihood of satisfying the safety specification. At all times, the ground-truth state stays within the safety region.

VI. CONCLUSIONS AND FUTURE WORK

Our work enables risk-aware control synthesis for stochastic dynamical systems with incomplete state information by combining continuous-discrete EKF with CBFs defined over Gaussian belief states. Instead of defining safety specifications as hard constraints on the state, we consider a risk-aware approach in which we bound the probability of violation. CBFs are applicable to any robotic system in which the state estimate is provided by an EKF. Our simulation and hardware experiments show that CBFs ensure that robots adhere to

safety specifications in the presence of both, real-world motion and observation uncertainties.

In future, we aim to explore the combination of motion planners such as state-based MPC that only consider the mean estimate with our proposed CBFs. In that way, we can ensure safety while enhancing control efficiency. We are also interested in extending our results to arbitrary belief distributions that can be represented using particles filters.

REFERENCES

- [1] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: Theory and applications," in *2019 18th European control conference (ECC)*. IEEE, 2019, pp. 3420–3431.
- [2] Y. Zhang, S. Walters, and X. Xu, "Control Barrier Function Meets Interval Analysis: Safety-Critical Control with Measurement and Actuation Uncertainties," in *American Control Conference (ACC)*. IEEE, 2022.
- [3] Y. Wang and X. Xu, "Observer-based Control Barrier Functions for Safety Critical Systems," 2021.
- [4] S. Dean, A. Taylor, R. Cosner, B. Recht, and A. Ames, "Guaranteeing safety of learned perception modules via measurement-robust control barrier functions," in *Conference on Robot Learning*. PMLR, 2021, pp. 654–670.
- [5] H. Zhu and J. Alonso-Mora, "Chance-Constrained Collision Avoidance for MAVs in Dynamic Environments," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 776–783, 2019.
- [6] S. Thrun, W. Burgard, and D. Fox, *Probabilistic robotics*. Cambridge, Mass.: MIT Press, 2005.
- [7] C. Urrea and R. Agramonte, "Kalman filter: historical overview and review of its use in robotics 60 years after its creation," *Journal of Sensors*, vol. 2021, pp. 1–21, 2021.
- [8] M. Castillo-Lopez, P. Ludvig, S. A. Sajadi-Alamdari, J. L. Sanchez-Lopez, M. A. Olivares-Mendez, and H. Voos, "A real-time approach for chance-constrained motion planning with dynamic obstacles," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 3620–3625, 2020.
- [9] H. Nishimura and M. Schwager, "SACBP: Belief space planning for continuous-time dynamical systems via stochastic sequential action control," *The International Journal of Robotics Research*, vol. 40, no. 10-11, pp. 1167–1195, 2021.
- [10] A. Clark, "Control Barrier Functions for Complete and Incomplete Information Stochastic Systems," in *American Control Conference (ACC)*, 2019, pp. 2928–2935.
- [11] A. Clark, Z. Li, and H. Zhang, "Control Barrier Functions for Safe CPS Under Sensor Faults and Attacks," in *59th IEEE Conference on Decision and Control (CDC)*, 2020, pp. 796–803.
- [12] S. Yaghoubi, G. Fainekos, T. Yamaguchi, D. Prokhorov, and B. Hoxha, "Risk-Bounded Control with Kalman Filtering and Stochastic Barrier Functions," 2021.
- [13] W. Luo, W. Sun, and A. Kapoor, "Multi-Robot Collision Avoidance under Uncertainty with Probabilistic Safety Barrier Certificates," in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., vol. 33. Curran Associates, Inc., 2020, pp. 372–383.
- [14] 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*, vol. 62, no. 8, pp. 3861–3876, 2016.
- [15] Q. Nguyen and K. Sreenath, "Exponential Control Barrier Functions for enforcing high relative-degree safety-critical constraints," in *American Control Conference (ACC)*, 2016, pp. 322–328.
- [16] J. Van Den Berg, S. Patil, and R. Alterovitz, "Motion planning under uncertainty using iterative local optimization in belief space," *The International Journal of Robotics Research*, vol. 31, no. 11, pp. 1263–1278, 2012.
- [17] F. Lewis, L. Xie, and D. Popa, "Optimal and robust estimation," *International Journal of Robust and Nonlinear Control*, vol. 19, no. 10, pp. 1197–1199, 2009.
- [18] W. S. Cortez, X. Tan, and D. V. Dimarogonas, "A robust, multiple control barrier function framework for input constrained systems," *IEEE Control Systems Letters*, vol. 6, pp. 1742–1747, 2021.
- [19] C. Rackauckas and Q. Nie, "DifferentialEquations.jl—a performant and feature-rich ecosystem for solving differential equations in julia," *Journal of Open Research Software*, vol. 5, no. 1, 2017.
- [20] B. Xu and K. Sreenath, "Safe Teleoperation of Dynamic UAVs Through Control Barrier Functions," in *IEEE International Conference on Robotics and Automation*, 2018, pp. 7848–7855.