

# Risk-aware Control for Robots with Non-Gaussian Belief Spaces

Matti Vahs and Jana Tumova

**Abstract**— This paper addresses the problem of safety-critical control of autonomous robots, considering the ubiquitous uncertainties arising from unmodeled dynamics and noisy sensors. To take into account these uncertainties, probabilistic state estimators are often deployed to obtain a belief over possible states. Namely, Particle Filters (PFs) can handle arbitrary non-Gaussian distributions in the robot’s state. In this work, we define the belief state and belief dynamics for continuous-discrete PFs and construct safe sets in the underlying belief space. We design a controller that provably keeps the robot’s belief state within this safe set. As a result, we ensure that the risk of the unknown robot’s state violating a safety specification, such as avoiding a dangerous area, is bounded. We provide an open-source implementation as a ROS2 package and evaluate the solution in simulations and hardware experiments involving high-dimensional belief spaces.

## I. INTRODUCTION

Autonomous robots are often exposed to various uncertainties in real-world applications such as unmodeled dynamics, noisy sensor readings or partially known environments. Ultimately, these uncertainties prevent us from accurately knowing the state of a robot. This requires us to consider probabilistic state estimators to obtain a robot’s *belief*, i.e. a probability distribution over possible states [1].

Oftentimes, a Kalman Filter (KF) and its variations are deployed in practice. These methods assume that the underlying probability distribution is Gaussian. Although they are reliable in many scenarios, KFs might fail in the case of highly nonlinear systems or multi-modal distributions. Particle filters (PFs) [2], on the other hand, can represent an arbitrary non-Gaussian belief as a discrete set of samples. Consider for example the robot depicted in Fig. 1 that is using a 2D LiDAR for localization in a given map. Due to nonlinear dynamics and observation models, the belief is non-Gaussian as illustrated by the red samples of the PF.

Guaranteeing safety, such as avoiding dangerous areas in the map that are difficult to detect by local sensing (untraversable terrain, glass walls, ...) is challenging. This is because it requires us to consider the uncertainty in the robot’s state due to the fact that we do not know exactly where the robot is located. Existing works for safety-critical control under stochastic uncertainties typically as-

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. Our experiments were carried out in the WASP Research Arena (WARA) - Robotics, hosted by ABB Corporate Research Center in Västerås, Sweden.

The authors are with the Division of Robotics, Perception and Learning, KTH Royal Institute of Technology, Stockholm, Sweden and also affiliated with Digital Futures. Mail addresses: {vahs, tumova}@kth.se




Fig. 1. An illustration of our experiments with the ABB Mobile YuMi Research Platform operating in an uncertain environment. The robot is equipped with a LiDAR sensor used for localization. Due to uncertainties in motions and observations, we only have access to a belief which is provided by a PF shown as red dots. The robot is supposed to avoid a dangerous area that cannot be detected by local sensing such as e.g. an untraversable area.

sume Gaussian beliefs [3]–[6], bounded estimation errors [7]–[9], or only consider external uncertainty but also assume deterministic dynamics [10], [11]. Applying these methods to the described scenario might not yield desired safety.

Providing rigorous safety guarantees for general nonlinear systems under stochastic uncertainties is challenging for several reasons: 1) The true underlying distribution over states is unknown as the standard Bayes Filter is intractable for continuous nonlinear systems, 2) the curse of dimensionality in belief spaces complicates the design of computationally efficient solutions and 3) hard safety constraints on the state cannot be ensured due to unbounded probability distributions. The latter point requires us to consider safety specifications in a risk-aware sense, meaning that small violations are allowed where the admissible level of risk is ultimately defined by the user. Measures of risk enable us to reason about specification violation in the presence of uncertainty [12]–[14]. Especially measures of risk such as *Conditional-Value-at-Risk* (CVaR) have gained a lot of attention in the robotics community as these account for less likely tail events that could result in unsafe robot behavior.

To address the three above mentioned challenges and offer safety guarantees in real-world robotic scenarios, we propose a solution enhancing Control Barrier Functions (CBFs) to non-Gaussian belief spaces. Namely, we

- 1) enable guaranteed risk-awareness by using CVaR for arbitrary non-Gaussian beliefs described by PFs,
- 2) provide computationally efficient control synthesis under high dimensionalities of belief spaces,
- 3) offer an open-source ROS2 package<sup>1</sup> that can be easily integrated with existing navigation stacks.

<sup>1</sup>available at [https://github.com/mattivahs/pfcbf\\_ros2](https://github.com/mattivahs/pfcbf_ros2)

## II. RELATED WORK

There is a large body of literature on safety-critical control under bounded uncertainties. These works are based on tools from set-based reachability analysis [15], [16], robust model predictive control (MPC) [17]–[20] or robust CBFs [7], [8], [21], [22]. While these methods have proven effective in many scenarios, it is important to consider that most estimators deployed on robots consider disturbances of stochastic nature which are typically unbounded.

Stochastic disturbances on the system dynamics have been addressed in [23]–[30] which are based on MPC and CBF approaches. Stochastic CBFs (SCBFs) ensure safety with probability one for systems that are described by stochastic differential equations (SDEs) [24]. In [30], risk measures such as CVaR are leveraged for risk-aware trajectory optimization under stochastic dynamics. Their approach is based on sample-average approximations of risk constraints which is real-time applicable if the number of samples is low. Although [23]–[30] consider the stochastic evolution of the system dynamics, they assume precise knowledge of the state, thus neglecting the uncertainties in sensor observations.

Approaches based on a belief instead of a state have been proposed in [3]–[6], [31], [32]. Chance-constrained nonlinear MPC (CCNMP) [5] enforces obstacle avoidance with a desired confidence in which the uncertainty in the state estimate originates from an unscented KF. In [31], the authors focused on providing safety guarantees through risk-aware control problem for systems which use an extended KF as state estimator. However, all of these methods assume that the belief follows a Gaussian distribution which might result in unsafe behaviors in scenarios in which a Gaussian is a poor approximation of the true distribution.

Only few works consider non-Gaussian beliefs for safety-critical control. In the seminal work [33], the authors propose a chance constrained planning framework for obstacle avoidance that leverages PFs for belief propagation and approximates chance constraints through sample averages. In [34], samples from a PF are used in planning to disambiguate a multiple potential hypotheses but safety constraints are not considered. Online POMDP solvers proposed in [35] combine PFs and monte carlo tree search (MCTS) for planning under uncertainty which allows to maximize a reward function but does not include rigorous safety constraints. Additionally, to make [33]–[35] computationally tractable, only a subset of particles from the actual PF is used. Consequently, there can be a mismatch between the sample-based distribution and the true distribution which can lead to unsafe behavior at deployment.

In this paper, we consider the problem of risk-aware control where the only thing we have access to is a PF belief. In contrast to other existing works, we consider the entire PF belief which could consist of thousands of particles and still offer a computationally effective solution. Furthermore, we account for the mismatch between the sample-based distribution and the true underlying continuous distribution by using an underapproximation of CVaR.

## III. PRELIMINARIES

Consider a robot that is modeled by a stochastic differential equation (SDE) of the form

$$d\mathbf{x} = (\mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{u})dt + \boldsymbol{\sigma}(\mathbf{x})d\mathbf{W} \quad (1)$$

where  $\mathbf{x} \in \mathcal{X} \subseteq \mathbb{R}^{n_x}$  is the state,  $\mathbf{u} \in \mathcal{U} \subseteq \mathbb{R}^m$  is the control input and  $\mathbf{W} \in \mathbb{R}^q$  is  $q$ -dimensional Brownian motion. We assume that the diffusion term  $\boldsymbol{\sigma}(\mathbf{x})$  is a globally Lipschitz non-degenerate diagonal matrix and the drift term  $(\mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{u})$  is a locally Lipschitz function that can have jumps. These assumptions ensure that Eq. (1) has a unique global strong solution [36].

### A. Continuous-discrete Particle Filters

In practice the state of a robot cannot be measured precisely as measurements are corrupted by noise. We assume that the robot's observations can be modeled by

$$\mathbf{z}_{t_k} = \boldsymbol{\ell}(\mathbf{x}_{t_k}, \mathbf{v}_{t_k}) \quad (2)$$

with measurement noise  $\mathbf{v}_{t_k} \sim p(\mathbf{v})$ . Measurements only occur at discrete timesteps  $t_1, \dots, t_k$  due to limited update rates of sensors. Thus, the exact state at time  $t$  is unknown, but a Bayesian posterior  $p(\mathbf{x}_t | \mathbf{z}_{t_1:t_k})$ , for  $t_k \leq t < t_{k+1}$  describes the probability distribution over states conditioned on past measurements. Unfortunately, this Bayesian posterior does not admit closed form solutions except for specialized cases such linear Gaussian systems. PF techniques stem from the idea of approximating the Bayesian posterior by means of a finite set of  $N$  weighted samples, i.e. particles,  $\{(\mathbf{x}_t^{(i)}, w_t^{(i)})\}_{i=1}^N$ , where  $w_t^{(i)} \in \mathbb{R}_{\geq 0}$ . The PF is a non-parametric approach that approximates the true posterior as

$$p(\mathbf{x}_t) \approx \sum_{i=1}^N w_t^{(i)} \delta(\mathbf{x}_t - \mathbf{x}_t^{(i)})$$

where  $\delta(\cdot)$  denotes the Dirac delta function. It is known that as  $N \rightarrow \infty$ , the approximated posterior converges to the true posterior [37].

Since we are considering a continuous time SDE as motion model and a discrete-time observation model, we distinguish between two different evolutions of the posterior: 1) the posterior evolves in continuous time in periods  $t \in [t_{k-1}, t_k)$  between measurements and 2) a discrete-time Bayesian update step at times  $t = t_k$ . In the former period, the evolution of the belief is obtained by propagating all particles through the SDE in Eq. (1). In the update step, each particle  $\mathbf{x}_{t_k}^{(i)}$  is weighted and resampled according to the observation likelihood  $p(\mathbf{z}_{t_k} | \mathbf{x}_{t_k}^{(i)})$ , see, e.g., [38] for more details on continuous-discrete PFs.

### B. Risk Measures

Given a scalar random variable  $y \in \mathcal{Y}$  with  $p(y)$  denoting its probability density function (pdf), a risk measure is a function that assigns a real value to the random variable, i.e.,  $\mathcal{R} : \mathcal{Y} \mapsto \mathbb{R}$ . To evaluate tail events of a random variable, we consider the following risk measures.

**Definition 1.** The Value-at-Risk (VaR) of a random variable  $y \in \mathbb{R}$  with pdf  $p(y)$  at level  $\alpha \in (0, 1]$  is given by

$$\text{VaR}_\alpha(y) = \inf_{\tau \in \mathbb{R}} \{\tau \mid \Pr[y \leq \tau] \geq \alpha\}.$$

The Conditional-Value-at-Risk (CVaR) of a random variable  $y \in \mathbb{R}$  with pdf  $p(y)$  at level  $\alpha$  is the expected value of the  $\alpha$  quantile, i.e.

$$\text{CVaR}_\alpha(y) = \mathbb{E}\{y \mid y \leq \text{VaR}_\alpha(y)\}.$$

VaR provides a measure that is qualitatively equivalent to chance constraints, i.e.  $\text{VaR}_\delta(y) \geq 0 \Leftrightarrow \Pr[y \leq 0] \leq \delta$ . CVaR, on the other hand, provides a risk measure that considers the mean of the tail of a distribution, which is more risk-averse than VaR. In this work, we choose to use CVaR as the risk measure suitable for safety-critical systems.

### C. Stochastic Control Barrier Functions

Consider the SDE defined in Eq. (1) and a safe set

$$\begin{aligned} \mathcal{C}_x &= \{\mathbf{x} \in \mathcal{X} \mid h_x(\mathbf{x}) \geq 0\}, \\ \partial\mathcal{C}_x &= \{\mathbf{x} \in \mathcal{X} \mid h_x(\mathbf{x}) = 0\}. \end{aligned} \quad (3)$$

in which we want the state to remain.

**Definition 2.** A safe set  $\mathcal{C}_x$  is forward invariant with respect to the system (1) if for every initial condition  $\mathbf{x}_{t_0} \in \mathcal{C}_x$  it holds that  $\mathbf{x}_t \in \mathcal{C}_x, \forall t \geq t_0$ .

Stochastic CBFs have been proposed in [24] which render the safe set  $\mathcal{C}_x$  forward invariant assuming that the function  $h_x(\mathbf{x})$  is twice continuously differentiable everywhere. These invariance results can be extended to consider safe sets that are described by non-smooth functions  $h_x(\mathbf{x})$ .

**Definition 3.** Given a safe set  $\mathcal{C}_x$  defined by (3), the function  $B_x(\mathbf{x})$  serves as a non-smooth stochastic CBF (NSCBF) for the system (1), if  $\forall \mathbf{x} \in \mathcal{C}_x$

- 1) There exist class-K functions  $\alpha_1$  and  $\alpha_2$  such that

$$\frac{1}{\alpha_1(h_x(\mathbf{x}))} \leq \frac{1}{B_x(\mathbf{x})} \leq \frac{1}{\alpha_2(h_x(\mathbf{x}))}. \quad (4)$$

- 2) There exists a class-K function  $\alpha_3$  and a control input  $\mathbf{u} \in \mathcal{U}$  such that

$$\begin{aligned} \frac{\partial B_x}{\partial \mathbf{x}} (\mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x}) \mathbf{u}) + \frac{1}{2} \text{tr} \left[ \boldsymbol{\sigma}^T \frac{\partial^2 B_x}{\partial \mathbf{x}^2} \boldsymbol{\sigma} \right] \\ \leq \alpha_3(h_x(\mathbf{x})). \end{aligned}$$

**Theorem 1.** If a locally Lipschitz control input  $\mathbf{u}(t)$  satisfies Def. 3 for a given safe set, then  $\Pr[\mathbf{x}(t) \in \mathcal{C}_x, \forall t \geq t_0] = 1$ , provided that  $\mathbf{x}(t_0) \in \mathcal{C}_x$ .

*Proof.* The proof can be found in [39].  $\square$

## IV. PROBLEM SETTING

We consider robotic systems described by continuous time SDEs in Eq. (1) and discrete stochastic observation models in Eq. (2). Ideally, the state of the robot should remain in a safe set  $\mathcal{C}_x$  at all times. However, the state is unknown and we only have access to the belief, i.e. an unbounded probability distribution over states. Due to the unboundedness of the




Fig. 2. **Left:** A drone with uncertain position  $x$  and its pdf  $p(x)$  is moving in one dimension. A safety specification is defined as not colliding with the wall, i.e.  $h_x = 2 - x$ . The true pdf is a Gaussian shown in red and the PF is shown in blue. **Right:** The distributions over  $h_x$  are shown on the right with their corresponding risk measures visualized.

belief, there is always the possibility that the state might be outside the safe set, hence we need to formulate safety in terms of risk-awareness.

**Example 1.** Consider a drone operating in one dimension as shown in Fig. 2 in the bottom left, where  $x$  denotes its position. A safe set is given as  $\mathcal{C}_x = \{x \in \mathbb{R} \mid x \leq 2\}$  which is the set of collision free states. The true posterior follows a Gaussian distribution, i.e.  $x \sim \mathcal{N}(\mu, \sigma)$  where  $\mu$  and  $\sigma$  denote the mean and standard deviation, respectively. The true posterior, is generally unknown and we only have access to the belief, which is modeled by a PF (in our example shown by the blue circles). Consequently, we have a distribution over possible values of  $h_x(x) = 2 - x$ , as shown in Fig. 2 in the right. The true distribution is unbounded (shown by the red curve), whereas the resulting distribution based on the PF belief is given by a histogram distribution. We want to ensure that  $\text{CVaR}(h_x(x)) \geq 0$  which is our risk-aware objective.

Although the motivating example is linear and Gaussian, we consider general nonlinear motions and observations, nonlinear safe sets as well as non-Gaussian beliefs.

**Problem 1.** Given stochastic motion and observations (1)-(2), a reference control input  $\mathbf{u}_{\text{ref}}$  and a safe set  $\mathcal{C}_x$  defined over states, synthesize control inputs that ensure that at any time  $\text{CVaR}_\alpha(h_x(\mathbf{x}(t))) \geq 0$ . If exact satisfaction cannot be guaranteed, provide probabilistic satisfaction bounds.

A challenge is that we only know the empirical  $\widehat{\text{CVaR}}$  based on the PF, which can mismatch the true CVaR as not enough particles are available, see Fig. 2. Thus, decisions based on  $\widehat{\text{CVaR}}$  might result in unsafe behavior in the real world. In the following, we address this challenge.

## V. RISK-AWARE CONTROL

In the following, we present our solution to Problem 1 which follows three steps: 1) We define the belief state and the corresponding belief dynamics for continuous-discrete PFs, 2) We construct a safe set in belief space which, if forward invariant, ensures provable risk-aware safety for the unknown state  $\mathbf{x}$ , and, 3) We formulate a controller that renders the proposed safe set forward invariant, i.e. it ensures that the belief state remains in the safe set.

### A. Particle Filter Belief Dynamics

As the PF used for state estimation consists of a weighted particle set  $\{(\mathbf{x}_t^{(i)}, w_t^{(i)})\}_{i=1}^N$ , we define a belief state

$$\mathbf{b}_t = \begin{bmatrix} \mathbf{x}_t^{(1)} & \dots & \mathbf{x}_t^{(N)} \end{bmatrix}^T \in \mathbb{R}^{N \cdot n_x},$$

which uniquely describes the current belief at time  $t$ . The continuous time dynamics of the belief state  $\mathbf{b}$  during the period where no observations are available is obtained by propagating each particle through the nonlinear SDE in Eq. (1), resulting in

$$\begin{aligned} d\mathbf{b}_t &= \left( \begin{bmatrix} \mathbf{f}(\mathbf{x}_t^{(1)}) \\ \vdots \\ \mathbf{f}(\mathbf{x}_t^{(N)}) \end{bmatrix} + \begin{bmatrix} \mathbf{g}(\mathbf{x}_t^{(1)}) \\ \vdots \\ \mathbf{g}(\mathbf{x}_t^{(N)}) \end{bmatrix} \mathbf{u} \right) dt + \begin{bmatrix} \boldsymbol{\sigma} d\mathbf{W}_t^{(1)} \\ \vdots \\ \boldsymbol{\sigma} d\mathbf{W}_t^{(N)} \end{bmatrix} \\ &:= (\mathbf{f}_b(\mathbf{b}) + \mathbf{g}_b(\mathbf{b}) \mathbf{u}) dt + \boldsymbol{\Sigma} d\tilde{\mathbf{W}} \end{aligned} \quad (5)$$

where  $\boldsymbol{\Sigma} = \text{BD}(\{\boldsymbol{\sigma}\}_{i=1}^N)$  is a block diagonal matrix and  $\tilde{\mathbf{W}}$  is a Brownian motion of dimension  $N \cdot q$ . The belief dynamics are of very high dimension due to the curse of dimensionality in belief spaces which complicates the use of common control techniques such as MPC directly in belief space. However, it should be noted that the individual state particles are entirely decoupled which means that the SDE in Eq. (5) can be solved efficiently using parallelization. Further, as  $\boldsymbol{\Sigma}$  is globally Lipschitz and non-degenerate it is ensured that a solution to Eq. (5) exists.

At discrete times  $t_k = t_k^+$  where observations are available, the PF belief is conditioned on the measurement which leads to a discrete update of the belief state, i.e.  $\mathbf{b}(t_k^+) = \boldsymbol{\Delta}(\mathbf{b}(t_k^-), \mathbf{z}_k)$  where  $t_k^-$  is infinitesimal smaller than  $t_k^+$ . This discrete map  $\boldsymbol{\Delta}(\cdot)$ , however, cannot be obtained in closed form as the conditioning includes a stochastic resampling step. In this step,  $N$  particles are resampled from the prior belief  $\mathbf{b}(t_k^-)$  according to their likelihood weights. For a detailed description of this step, the reader is referred to [1].

### B. Construction of Safe Sets

Now that we have a model of how the belief state  $\mathbf{b}$  changes over time, we construct a safe set  $\mathcal{C}_b$  that we want  $\mathbf{b}$  to remain in. This safe set should be defined in a way such that  $\mathbf{b}(t) \in \mathcal{C}_b$  implies risk-awareness for the unknown system state, i.e.  $\text{CVaR}_\alpha(h_x(\mathbf{x}(t))) \geq 0$ .

The main difficulty lies in the fact that the true distribution  $p(h_x(\mathbf{x}))$  is unknown, thus the exact CVaR cannot be calculated. Yet, we have access to the PF belief  $\mathbf{b}$  which represents a discrete sample set from the unknown true posterior  $p(\mathbf{x})$ . Thus, we can propagate all the particles through the function  $h_x(\mathbf{x})$ , i.e.  $h_x^{(i)} = h_x(\mathbf{x}^{(i)})$ , to obtain a set of samples from the true unknown distribution  $p(h_x(\mathbf{x}))$ . We leverage a theorem that allows us to draw conclusions about the CVaR of an unknown continuous distribution based on a set of samples that are drawn from that distribution. In contrast to the original version, we reformulate it focusing on the lower-tail distribution.

**Theorem 2** ([40], Thm. 3). *Let  $y^{(1)}, \dots, y^{(N)}$  be i.i.d. samples from a random variable  $y$  and  $\Pr[y \geq b] = 1$  for some finite  $b$ , then for any  $\delta \in (0, 0.5]$ ,*

$$\Pr \left[ \overline{\text{CVaR}}_\alpha \left( y^{(1)}, \dots, y^{(N)} \right) \leq \text{CVaR}_\alpha(y) \right] \geq 1 - \delta$$

where the lower bound is obtained as

$$\begin{aligned} \overline{\text{CVaR}}_\alpha \left( y^{(1)}, \dots, y^{(N)} \right) &= \xi_{N+1} + \frac{1}{\alpha} \sum_{i=1}^N \left( (\xi_i - \xi_{i+1}) \cdot \right. \\ &\quad \left. \left[ \frac{i}{N} - \sqrt{\frac{\ln(1/\delta)}{2N}} - (1 - \alpha) \right]^+ \right) \end{aligned} \quad (6)$$

with  $\xi_1, \dots, \xi_N$  as order statistics (i.e.  $y^{(1)}, \dots, y^{(N)}$  in descending order),  $\xi_{N+1} = b$  and  $(\cdot)^+ = \max\{0, \cdot\}$ .

This theorem states that  $\overline{\text{CVaR}}$  is a lower bound of the true CVaR with desired probability  $1 - \delta$ . We leverage this result to construct a safe set in belief space, reading

$$\mathcal{C}_b = \{ \mathbf{b} \in \mathcal{B} \mid h_b(\mathbf{b}) \geq 0 \}, \text{ where}$$

$$h_b(\mathbf{b}) = \overline{\text{CVaR}}_\alpha \left( h_x(\mathbf{x}^{(1)}), \dots, h_x(\mathbf{x}^{(N)}) \right). \quad (7)$$

In turn, if  $\mathbf{b}(t) \in \mathcal{C}_b \implies \Pr[\text{CVaR}(h_x(\mathbf{x}(t))) \geq 0] \geq 1 - \delta$ , we solve Problem 1. The only problem that remains is to design a controller that ensures that the belief state stays within  $\mathcal{C}_b$  at all times.

### C. Control Synthesis

Ultimately, we want to guarantee forward invariance of the safe set  $\mathcal{C}_b$ . Popular approaches to render sets forward invariant are CBFs and especially stochastic CBFs that have been designed for stochastic systems as in Eq. (5). However, the standard formulation of stochastic CBFs cannot be applied here since our safe set defined by Eq. (6) is a continuous but non-smooth function. Due to the sorting operation included in Eq. (6), the gradient of  $\overline{\text{CVaR}}$  with respect to the belief state  $\mathbf{b}$  can be discontinuous.

This nonsmoothness of the safe set, however, does not affect the forward invariance property, see Thm. 1. Thus, to synthesize control inputs, we solve the quadratic program

$$\begin{aligned} \mathbf{u}^* &= \arg \min_{\mathbf{u} \in \mathcal{U}} (\mathbf{u} - \mathbf{u}_{\text{ref}})^T \mathbf{Q} (\mathbf{u} - \mathbf{u}_{\text{ref}}) \\ \text{s.t.} \quad & \frac{\partial B_b}{\partial \mathbf{b}} (\mathbf{f}_b(\mathbf{b}) + \mathbf{g}_b(\mathbf{b}) \mathbf{u}) \\ & + \frac{1}{2} \text{tr} \left[ \boldsymbol{\Sigma}^T \frac{\partial^2 B_b}{\partial \mathbf{b}^2} \boldsymbol{\Sigma} \right] \leq \gamma h_b(\mathbf{b}). \end{aligned} \quad (8)$$

where  $B_b(\mathbf{b}) = 1/h_b(\mathbf{b})$ ,  $\gamma \geq 0$ ,  $\mathbf{Q}$  is a weighting matrix and  $\mathbf{u}_{\text{ref}}$  is a reference control input. In order to provide guarantees for the forward invariance of  $\mathcal{C}_b$  under the continuous-discrete evolution of the PF, we need to make the following assumption.

**Assumption 1.** *The belief state  $\mathbf{b}$  does not leave the safe set under a discrete transition  $\mathbf{b}(t_k^+) = \boldsymbol{\Delta}(\mathbf{b}(t_k^-), \mathbf{z}_k)$ , i.e.  $\mathbf{b}(t_k^-) \in \mathcal{C}_b \implies \mathbf{b}(t_k^+) \in \mathcal{C}_b$ .*

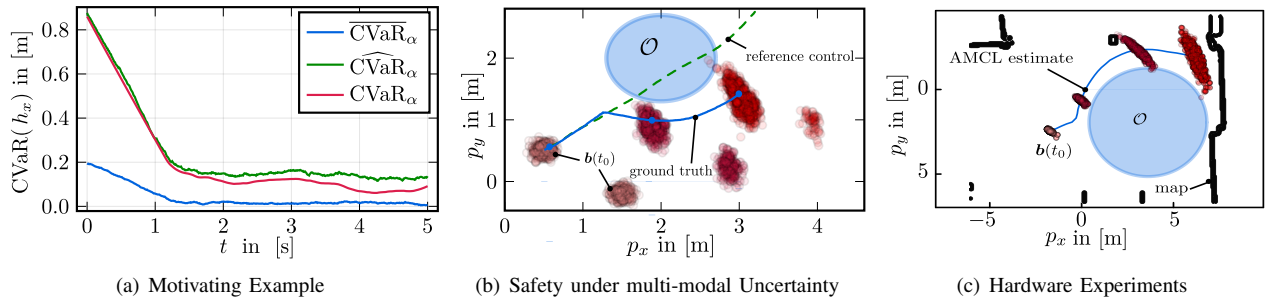


Fig. 3. (a) Values of the true CVaR, the empirical  $\widehat{\text{CVaR}}$  and our lower bound  $\overline{\text{CVaR}}$  over time for a single simulation of Example 1. (b) Illustration of the PF belief over time for a 2D avoidance task where the initial belief follows a mixture of Gaussians. The ground truth trajectory is shown in blue and the trajectory for the reference controller is shown in green. (c) Illustration of the PF belief and the AMCL estimate for our hardware experiment.

TABLE I  
COMPARISON OF CVaR MEASURES WITH VARYING PARTICLES.

$N$	$\widehat{\text{CVaR}}_\alpha$		$\overline{\text{CVaR}}_\alpha$		$t_c$ in [s]
	$e(\mu \pm \sigma)$	$\Pr[e \leq 0]$	$e(\mu \pm \sigma)$	$\Pr[e \leq 0]$	
100	$-0.025 \pm 0.008$	100	$0.18 \pm 0.022$	0.	$0.0005 \pm 1.3 \cdot 10^{-5}$
1000	$-0.001 \pm 0.004$	69.5	$0.025 \pm 0.005$	0	$0.0006 \pm 2.2 \cdot 10^{-5}$
5000	$-0.0004 \pm 0.002$	53.7	$0.015 \pm 0.018$	0.	$0.001 \pm 0.003$

Since measurements are stochastic and cannot be controlled, this assumption states that we do not receive extreme outlier measurements which make the PF diverge.<sup>2</sup>

**Theorem 3.** *Under Assumption 1, if the controller in Eq. (8) is feasible at all times with the nonsmooth stochastic CBF  $h_b(\mathbf{b})$  defined in Eq. (7), then  $\mathcal{C}_b$  is forward invariant.*

*Proof Sketch:* Forward invariance follows directly from Thm. 1 in which we replace the state  $\mathbf{x}$  by the belief state  $\mathbf{b}$ , the safe set  $h_x$  by  $h_b$  and the stochastic state dynamics in Eq. (1) by the stochastic belief dynamics in Eq. (5).

Thus, the proposed controller solves Problem 1 with desired confidence  $1 - \delta$  that can be chosen by the user.

## VI. EXPERIMENTS

We evaluate our approach in simulations as well as hardware experiments in which uncertainties naturally occur. Specifically, we show

- 1) correctness of the risk-aware guarantees in an example where the true posterior is known, in Sec. VI-A,
- 2) improved adherence to safety requirements over baselines under multi-modal distributions, in Sec. VI-B,
- 3) integration with existing robot navigation stacks on hardware, in Sec. VI-C.

### A. Motivating Example Continued

1) *Setup:* In this simulation we refer back to our motivating example shown in Fig. 2. The drone's dynamics are described by a single integrator  $d\mathbf{x} = \mathbf{u} dt + 0.1 d\mathbf{W}$  with a

<sup>2</sup>One example of a discrete transition leading to the belief leaving the safe set is when the worst state particle would get resampled  $N$  times, meaning that the PF belief would collapse to a single particle. Although the probability of this event happening is non-zero, it is a highly unlikely event.

reference controller as  $u_{\text{ref}} = 1$ , i.e. a controller that would steer the drone into the wall. In this simulation we do not consider observations, so the uncertainty will increase over time. As the initial belief is Gaussian and the dynamics are linear, the true posterior is in this case given by a Kalman Filter (KF). We use the KF to calculate the mismatch of the CVaR between the true posterior and the approximated PF belief. We evaluate the mismatch as  $e = \text{CVaR} - \phi$  where  $\phi \in \{\widehat{\text{CVaR}}, \overline{\text{CVaR}}\}$ . Note that  $e < 0$  means that CVaR is overapproximated which should be avoided at all times. For this simulation we picked  $\alpha = 0.2$  (thus we evaluate the mean of the worst 20 % quantile) and  $\delta = 0.05$  which means that our  $\overline{\text{CVaR}}$  measure should be underapproximating in 95 % of the cases.

2) *Discussion:* Table I shows the error  $e$  for varying numbers of particles  $N$ . It can be seen that the empirical  $\widehat{\text{CVaR}}$  generally overapproximates the true CVaR which stresses that empirical measures should not be used for safety-critical systems. Our  $\overline{\text{CVaR}}$  measure on the other hand is a lower bound in all cases which means that the bound in Thm. 2 might not be very tight. However, the error  $e$  decreases with the number of particles as the PF belief approaches the true belief. Fig. 3(a) shows the evolution of the different CVaR measures over time for  $N = 100$ . It can be seen that our controller keeps  $h_b(\mathbf{b}) \geq 0$ , thus, rendering the set  $\mathcal{C}_b$  forward invariant.

Additionally, Table I shows the computation times for solving the QP in Eq. (8). Since we are only solving a QP over control inputs, the controller can still be run at a rate of almost 1kHz for a belief dimension of 5000 which shows that it is a real-time applicable approach.

### B. Safety under multi-modal Uncertainty

1) *Setup:* We compare our proposed method to various baselines in an example including nonlinear dynamics and observations as well as a non-Gaussian initial belief. We consider a unicycle robot that uses noisy radio signals for localization which can be described by

$$\begin{bmatrix} dp_x \\ dp_y \\ d\varphi \end{bmatrix} = \begin{bmatrix} c\varphi & 0 \\ c\varphi & 0 \\ 0 & s\varphi \end{bmatrix} \mathbf{u} + \boldsymbol{\sigma} d\mathbf{W}, \quad \begin{aligned} z_k &= \|\mathbf{p} - \boldsymbol{\ell}\|_2 + v_k, \\ v_k &\sim \mathcal{N}(0, r). \end{aligned} \quad (9)$$

where  $s\varphi = \sin(\varphi)$ ,  $c\varphi = \cos(\varphi)$ ,  $\mathbf{x} = [p_x, p_y, \varphi]^T$  is the state with  $\mathbf{p} = [p_x, p_y]^T$  being the 2D position and

TABLE II  
COMPARISON TO BASELINES.

	$\mu$ -SCBF	ML-SCBF	BE-SCBF	Ours ( $\alpha = 0.2$ )	Ours ( $\alpha = 0.05$ )
# coll.	46	42	0	2	0
$h_x(\mathbf{x})$	$0.85 \pm$	$0.8 \pm$	$2.45 \pm$	$0.95 \pm$	$1.71 \pm$
$(\mu \pm \sigma)$	0.59	0.56	0.49	0.46	0.43

$\mathbf{u} = [v, \omega]^T$  is the control input. The antenna from which we obtain noisy distance measurements  $z$  is located at  $\ell = [4, 4]^T$  m. The motion noise is  $\sigma = \text{diag}([0.3, 0.3, 0.1])$  and observation noise  $r = 0.3$  with a sensor update rate of 1 Hz. The initial belief is given as a mixture of Gaussians  $p(\mathbf{x}(t_0)) = 0.5 \cdot \mathcal{N}(\mathbf{p}_1, \Sigma_1) + 0.5 \cdot \mathcal{N}(\mathbf{p}_2, \Sigma_2)$  which is visualized in Fig. 3(b). A proportional controller that drives the robot to a goal based on the mean state of the PF belief serves as reference controller  $\mathbf{u}_{\text{ref}}$ , see Fig. 3(b).

A safe set which enables smooth navigation for unicycle dynamics with circular stay-out regions (from [41]) is defined as  $h_x(\mathbf{x}) = \|\hat{\mathbf{p}} - \mathcal{O}\|_2 - (r_o + d)$  where

$$\hat{\mathbf{p}} = \begin{bmatrix} p_x + d \cdot \cos(\varphi) \\ p_y + d \cdot \sin(\varphi) \end{bmatrix} \quad (10)$$

with  $d > 0$ . This function is chosen because it allows to control both velocities,  $v$  and  $\omega$ , simultaneously. For risk-aware control, we embed  $h_x(\mathbf{x})$  in our safe set defined over beliefs in Eq. (7) and solve the QP in Eq. (8).

2) *Baselines*: We compare our approach to three different baselines that are based on stochastic CBFs (SCBFs): 1.) A standard SCBF called  $\mu$ -SCBF that uses the mean state, i.e. the weighted sum of all particles. 2.) A standard SCBF called ML-SCBF based on the most likely particle, i.e. the particle with the highest observation likelihood. 3.) An SCBF that considers a known time-varying estimation error. Here we use Chebyshev inequality [42] to obtain a ball  $\mathcal{B}_\eta = \{\mathbf{x} \in \mathcal{X} \mid \|\mathbf{x} - \boldsymbol{\mu}\|_2^2 \leq \text{tr}[\Sigma]/\eta\}$  which is centered around the mean  $\boldsymbol{\mu}$  s.t.  $\Pr[\mathbf{p} \in \mathcal{B}_\eta] \geq 1 - \eta$  with  $\eta = 0.05$ .

3) *Discussion*: Figure 3(b) shows an exemplary trajectory of the PF for the avoidance task using our proposed method. It can be clearly seen that our method corrects the reference controller such that the majority of particles avoid the area  $\mathcal{O}$ . In this simulation, we used  $\text{CVaR}_{0.2}$ , i.e. we look at the mean of the worst 20 % quantile. We used  $N = 1000$  particles and repeated the simulation 100 times. Table II summarizes the results of our proposed method as well as the baselines. The two baselines  $\mu$ -SCBF and ML-SCBF fail to ensure safety in almost half of the cases which is due to the fact that they do not consider the entire distribution. On the other hand, BE-SCBF achieves zeros safety violations but is overly conservative as the distance to the boundary of the safe set ( $h_x$ ) is more than twice as high as for our method. This is because reliable estimation errors are difficult to obtain, thus Chebyshev inequality is rather conservative if the distribution is not Gaussian. Our method ensures risk-aware safety where the level of riskiness can be defined by the user. Note that, since  $\text{CVaR}_\alpha(h_x) \leq \text{VaR}_\alpha(h_x)$  it is also ensured that  $\text{CVaR}_\alpha(h_x) \geq 0 \implies \Pr[h_x \geq 0] \geq 1 - \alpha$  which is satisfied in our simulations.

### C. Hardware Experiments

We implemented our approach as a ROS2 node which subscribes to a PF belief as well as a reference control input and publishes a minimally deviating control input that ensures risk-aware safety. The package is open-source for anyone who wishes to make it a part of their navigation.

1) *Setup*: In our hardware experiments, we used the ABB Mobile YuMi Research Platform in the WASP Research Arena for Robotics (WARA Robotics) at ABB, depicted in Fig. 1. The robot is running the ROS2 navigation stack (NAV2 [43]) that enables functionalities such as localization and planning. For state estimation, two LiDARs are used to localize the robot in a given map using Adaptive Monte Carlo Localization (AMCL) which is a PF approach with varying number of particles. The robots dynamics are modeled as

$$\begin{bmatrix} dp_x \\ dp_y \\ d\varphi \end{bmatrix} = \begin{bmatrix} c\varphi & -s\varphi & 0 \\ s\varphi & c\varphi & 0 \\ 0 & 0 & 1 \end{bmatrix} \mathbf{u} dt + \sigma d\mathbf{W}$$

where  $\mathbf{u} = [v_x, v_y, \omega]^T$  are the control inputs consisting of a reference linear and angular velocity in the robot's local frame. The observation model is a likelihood field model which is the standard observation model for LiDAR localization in the NAV2 stack. Our node is subscribing to the AMCL PF belief as well as the velocity command provided by the local planner. In the AMCL node, the particle size varies between 3000 – 4000 particles and our risk-aware controller publishes a control signal at a rate of 30 Hz.

We specify our safe set as a circular area, shown in Fig. 3(c), that we want to avoid, i.e.  $h_x(\mathbf{x}) = \|\mathbf{p} - \mathcal{O}\|_2 - r_o$ . As we can directly control the velocity in  $p_x$  and  $p_y$  direction, the modification used in the previous section is not necessary.

2) *Discussion*: We investigate a scenario in which a sensor failure happens, i.e. lidar measurements are not available and the robot purely relies on odometry information. Figure 3(c) shows the PF belief as well as the AMCL estimate over time. It can be seen that the majority of particles stays within the safe area as enforced by our controller. Further, it is interesting to see that the distance of the AMCL estimate to  $\mathcal{O}$  gradually increases since the uncertainty increases<sup>3</sup>. In this experiment, we cannot evaluate the actual distance to the safe set as ground truth data is not available.

## VII. CONCLUSIONS AND FUTURE WORK

Our work enables risk-aware safety for robots with non-linear stochastic dynamics and observations by leveraging particle filters and control barrier functions. We define safe sets using Conditional-Value-at-Risk underapproximations that account for the mismatch between the approximated belief and the true unknown posterior which has not been considered before. We offer a computationally efficient control technique which we have analyzed in simulations as well as real world experiments, showcasing improved adherence to safety specifications over baselines. In future work, we aim to use our results for safe navigation in dynamic environments with probabilistic trajectory forecasting models.

<sup>3</sup>video available at <https://www.youtube.com/watch?v=YUdRtdFbYc>

## REFERENCES

- [1] S. Thrun, W. Burgard, and D. Fox, *Probabilistic robotics*. Cambridge, Mass.: MIT Press, 2005.
- [2] “Monte carlo localization for mobile robots,” in *International Conference on Robotics and Automation (ICRA)*, vol. 2. IEEE, 1999, pp. 1322–1328.
- [3] L. Blackmore, H. Li, and B. Williams, “A probabilistic approach to optimal robust path planning with obstacles,” in *American Control Conference*. IEEE, 2006, pp. 7–pp.
- [4] 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.
- [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] M. Ono and B. C. Williams, “Iterative risk allocation: A new approach to robust model predictive control with a joint chance constraint,” in *Conference on Decision and Control*. IEEE, 2008, pp. 3427–3432.
- [7] 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.
- [8] Y. Wang and X. Xu, “Observer-based control barrier functions for safety critical systems,” in *American Control Conference (ACC)*. IEEE, 2022, pp. 709–714.
- [9] 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.
- [10] O. de Groot, B. Brito, L. Ferranti, D. Gavrilu, and J. Alonso-Mora, “Scenario-based trajectory optimization in uncertain dynamic environments,” *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 5389–5396, 2021.
- [11] O. de Groot, L. Ferranti, D. Gavrilu, and J. Alonso-Mora, “Scenario-based motion planning with bounded probability of collision,” *arXiv preprint arXiv:2307.01070*, 2023.
- [12] T. Nyberg, C. Pek, L. Dal Col, C. Norén, and J. Tumova, “Risk-aware motion planning for autonomous vehicles with safety specifications,” in *2021 IEEE Intelligent Vehicles Symposium (IV)*, 2021, pp. 1016–1023.
- [13] F. S. Barbosa, B. Lacerda, P. Duckworth, J. Tumova, and N. Hawes, “Risk-aware motion planning in partially known environments,” in *60th IEEE Conference on Decision and Control (CDC)*, 2021, pp. 5220–5226.
- [14] A. Majumdar and M. Pavone, “How should a robot assess risk? towards an axiomatic theory of risk in robotics,” in *Robotics Research: The 18th International Symposium ISRR*. Springer, 2020, pp. 75–84.
- [15] M. Althoff and J. M. Dolan, “Online verification of automated road vehicles using reachability analysis,” *IEEE Transactions on Robotics*, vol. 30, no. 4, pp. 903–918, 2014.
- [16] C. Pek, S. Manzinger, M. Koschi, and M. Althoff, “Using online verification to prevent autonomous vehicles from causing accidents,” *Nature Machine Intelligence*, vol. 2, no. 9, pp. 518–528, 2020.
- [17] P. J. Campo and M. Morari, “Robust model predictive control,” in *1987 IEEE American control conference*. IEEE, 1987, pp. 1021–1026.
- [18] W. Langson, I. Chrysochoos, S. Raković, and D. Q. Mayne, “Robust model predictive control using tubes,” *Automatica*, vol. 40, no. 1, pp. 125–133, 2004.
- [19] P. Roque, W. S. Cortez, L. Lindemann, and D. V. Dimarogonas, “Corridor mpc: Towards optimal and safe trajectory tracking,” in *American Control Conference (ACC)*. IEEE, 2022, pp. 2025–2032.
- [20] M. N. Zeilinger, C. N. Jones, D. M. Raimondo, and M. Morari, “Real-time mpc-stability through robust mpc design,” in *Conference on Decision and Control (CDC)*. IEEE, 2009, pp. 3980–3986.
- [21] X. Xu, P. Tabuada, J. W. Grizzle, and A. D. Ames, “Robustness of control barrier functions for safety critical control,” *IFAC-PapersOnLine*, vol. 48, no. 27, pp. 54–61, 2015.
- [22] S. Kolathaya and A. D. Ames, “Input-to-state safety with control barrier functions,” *IEEE control systems letters*, vol. 3, no. 1, pp. 108–113, 2018.
- [23] A. T. Schwarm and M. Nikolaou, “Chance-constrained model predictive control,” *AICHE Journal*, vol. 45, no. 8, pp. 1743–1752, 1999.
- [24] A. Clark, “Control barrier functions for complete and incomplete information stochastic systems,” in *American Control Conference (ACC)*, 2019, pp. 2928–2935.
- [25] R. K. Cosner, P. Culbertson, A. J. Taylor, and A. D. Ames, “Robust safety under stochastic uncertainty with discrete-time control barrier functions,” *arXiv preprint arXiv:2302.07469*, 2023.
- [26] M. Black, G. Fainekos, B. Hoxha, D. Prokhorov, and D. Panagou, “Safety under uncertainty: Tight bounds with risk-aware control barrier functions,” *arXiv preprint arXiv:2304.01040*, 2023.
- [27] P. Li, M. Wendt, and G. Wozny, “A probabilistically constrained model predictive controller,” *Automatica*, vol. 38, no. 7, pp. 1171–1176, 2002.
- [28] F. Oldewurtel, C. N. Jones, A. Parisio, and M. Morari, “Stochastic model predictive control for building climate control,” *IEEE Transactions on Control Systems Technology*, vol. 22, no. 3, pp. 1198–1205, 2013.
- [29] A. Singletary, M. Ahmadi, and A. D. Ames, “Safe control for nonlinear systems with stochastic uncertainty via risk control barrier functions,” *IEEE Control Systems Letters*, vol. 7, pp. 349–354, 2022.
- [30] T. Lew, R. Bonalli, and M. Pavone, “Risk-averse trajectory optimization via sample average approximation,” *arXiv preprint arXiv:2307.03167*, 2023.
- [31] M. Vahs, C. Pek, and J. Tumova, “Belief control barrier functions for risk-aware control,” *Robotics and Automation Letters*, 2023.
- [32] C. J. Ostafew, A. P. Schoellig, and T. D. Barfoot, “Robust constrained learning-based nmpc enabling reliable mobile robot path tracking,” *The International Journal of Robotics Research*, vol. 35, no. 13, pp. 1547–1563, 2016.
- [33] L. Blackmore, M. Ono, A. Bektassov, and B. C. Williams, “A probabilistic particle-control approximation of chance-constrained stochastic predictive control,” *IEEE transactions on Robotics*, vol. 26, no. 3, pp. 502–517, 2010.
- [34] R. Platt, L. Kaelbling, T. Lozano-Perez, and R. Tedrake, “Efficient planning in non-gaussian belief spaces and its application to robot grasping,” in *Robotics Research*. Springer, 2017, pp. 253–269.
- [35] Z. Sunberg and M. Kochenderfer, “Online algorithms for pomdps with continuous state, action, and observation spaces,” in *Proceedings of the International Conference on Automated Planning and Scheduling*, vol. 28, 2018, pp. 259–263.
- [36] G. Leobacher and M. Szölgényi, “A strong order 1/2 method for multidimensional sdes with discontinuous drift,” *The Annals of Applied Probability*, vol. 27, no. 4, pp. 2383–2418, 2017.
- [37] D. Crisan and A. Doucet, “A survey of convergence results on particle filtering methods for practitioners,” *IEEE Transactions on Signal Processing*, vol. 50, no. 3, pp. 736–746, 2002.
- [38] M. K. Nielsen, T. K. Ritschel, I. Christensen, J. Dragheim, J. K. Huusom, K. V. Gernaey, and J. B. Jørgensen, “State estimation for continuous-discrete-time nonlinear stochastic systems,” *arXiv preprint arXiv:2212.02139*, 2022.
- [39] M. Vahs and J. Tumova, “Non-smooth control barrier functions for stochastic dynamical systems,” *Accepted at European Control Conference (ECC)*, 2024.
- [40] P. Thomas and E. Learned-Miller, “Concentration inequalities for conditional value at risk,” in *International Conference on Machine Learning*. PMLR, 2019, pp. 6225–6233.
- [41] J. Huang, Z. Liu, J. Zeng, X. Chi, and H. Su, “Obstacle avoidance for unicycle-modelled mobile robots with time-varying control barrier functions,” 2023.
- [42] X. Chen, “A new generalization of chebyshev inequality for random vectors,” *arXiv preprint arXiv:0707.0805*, 2007.
- [43] S. Macenski, F. Martín, R. White, and J. G. Clavero, “The marathon 2: A navigation system,” in *International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 2718–2725.