Probabilistic Control Barrier Functions for Systems with State Estimation Uncertainty using Sub-Gaussian Concentration

Kazuya Echigo[1], David E. J. van Wijk[1], Pol Mestres[1], *Member, IEEE*, Ersin Daş[2], *Member, IEEE*, Joel W. Burdick[1], *Member, IEEE*, and Aaron D. Ames[1], *Fellow, IEEE*

***Abstract*— Safety-critical control systems, such as spacecraft performing proximity operations, must provide formal safety guarantees despite stochastic uncertainties from state estimation and unmodeled dynamics. Although Control Barrier Functions (CBFs) have been extended to stochastic systems, existing approaches typically face a trade-off between the tightness of probabilistic guarantees and computational tractability. This paper presents a particle-based probabilistic CBF framework that overcomes this limitation by exploiting the sub-Gaussian structure of the barrier function increment under sub-Gaussian uncertainties. We establish that sub-Gaussian uncertainties propagating through Lipschitz-continuous control-affine dynamics preserve sub-Gaussianity of the barrier function increment, with explicit tail bounds. Leveraging this structure, we derive finite-sample bounds on the approximation error between particle-based Conditional Value at Risk (CVaR) estimates and ground-truth probabilistic constraints; applying this yields a tractable optimization problem formulation with finite-sample safety certificates. We show through numerical experiments how the proposed approach provides tight yet provably valid probabilistic safety guarantees.**

***Index Terms*— Optimization, safety-critical control, stochastic systems.**

I. Introduction

Designing safe controllers using Control Barrier Functions (CBFs) has been an extensive area of study within the control community. CBF-based controllers can be synthesized with provably safe guarantees and low computational overhead, with effectiveness validated in multiple applications [1]. However, conventional CBFs have rarely been applied to systems subject to both model and state estimation uncertainty—such as deep space explorations [2], [3]. This limitation arises because they are fundamentally designed for deterministic systems or systems with bounded disturbances. When the system is subject to stochastic disturbances with unbounded support—as commonly arises when using state estimators—such deterministic guarantees fail: the system will almost surely violate the safety constraints in finite time [4].

Motivated by this challenge, considerable research has focused on stochastic extensions of CBFs, which provide probabilistic safety guarantees [5], [6]. However, a key remaining challenge is obtaining a tractable representation of state estimation uncertainty as it propagates through nonlinear affine dynamics—the typical dynamical model in CBF applications [7]. When state uncertainty propagates through such dynamics, the resulting distribution generally lacks a closed-form expression. This forces existing methods to either adopt overly conservative approximations or rely on computationally expensive representations.

With that in mind, we focus our attention on extending probabilistic CBFs to deal with both state estimation and model uncertainty, and formulate a tractable optimization-based algorithm to compute a probabilistically safe control input. Specifically, we exploit the sub-Gaussian structure of the barrier function increment under sub-Gaussian uncertainties to derive tight, explicit concentration bounds without requiring conservative assumptions. This yields a tractable optimization problem (a Quadratic Program (QP) when the safety and input constraints are linear) with rigorous probabilistic guarantees suitable for real-time implementation.

a) Related Works: While stochastic extensions of other control techniques [8], [9] have been actively developed, we focus on CBF synthesis as a real-time safety filter, where per-step computational tractability is central. There has been substantial effort to adapt CBF techniques for systems with uncertainty. Initial work focused on robust methods for bounded, deterministic disturbances [10], [11], [12], [13]. Building on these foundations while relaxing those assumptions, other works have explored the problem under stochastic and possibly unbounded disturbances using modeling techniques such as martingale-based analysis [14] or scenario sampling [6], [15]. While these approaches account for uncertainty propagation through system dynamics, they either do not account for state estimation uncertainty or require restrictive assumptions.

Some works have examined this problem by specializing to specific classes of estimators. In continuous time, [16], [17] use martingale arguments to generate probabilistic safety certificates, while in discrete time, [18] apply similar martingale-based approaches. Alternatively, [19], [20] specialize to the extended Kalman filter in discrete time. However, those ap-

The work was supported by Air Force Office of Scientific Research Grant No. 113535-19668.

[1] Mechanical and Civil Engineering, California Institute of Technology, Pasadena, CA 91125, USA; {kazuyae, vanwijk, mestres, jburdick, ames}@caltech.edu.

[2] Department of Mechanical, Materials, and Aerospace Engineering, Illinois Institute of Technology, Chicago, IL 60616, USA; edas2@illinois.edu.

proaches only leverage the first moment of the uncertainty distribution and can be conservative when more distributional information is available. Sample-based approaches offer another direction, but have notable limitations for a CBF framework: sample average approximations (SAAs) [21] introduce per-sample binary variables, yielding $\mathcal{NP}$-hard mixed-integer programs; scenario approaches [22] are outlier-sensitive with a sample size growing in the control dimension and under stricter safety demands; and smooth approximations [8] yield non-convex problems requiring sequential linearization—all unsuitable for real-time per-step control. A more promising sample-based direction is to reformulate the stochastic CBF conditions via CVaR, which provides a convex tight upper bound [23] with continuity and outlier robustness. Recent work [24] adopts this reformulation and bounds the sample-based approximation error using the distribution-free Dvoretzky–Kiefer–Wolfowitz (DKW) inequality, accommodating state estimation uncertainty in nonlinear control-affine systems. However, this approach requires the random variable to have bounded support: the bound must be known *a priori*, requiring artificial truncation for unbounded distributions, introducing an approximation error that does not decay with sample size.

In contrast to these existing approaches, we propose a particle-based CBF framework that exploits the sub-Gaussian structure of the barrier function increment. Our approach requires no distribution truncation or bounded support assumptions, yielding an approximation error that decays to zero as the number of samples increases at a rate of $O((n \ln n)^{-1/2})$.

b) Statement of Contributions: We develop a particle-based probabilistic CBF framework that simultaneously handles both state estimation and dynamics uncertainty with rigorous finite-sample safety guarantees. Our contributions are threefold:

1) We establish that the barrier function increment is sub-Gaussian when sub-Gaussian uncertainties propagate through Lipschitz-continuous nonlinear control-affine dynamics, with an explicit parameter bound.
2) We derive explicit, finite-sample bounds on the particle-based CVaR approximation error without requiring the uncertainty to have bounded support.
3) We reformulate the probabilistic CBF condition as a tractable optimization problem, resolving the theory-computation tradeoff in existing approaches.

The remainder of this paper is organized as follows. Section II introduces the problem formulation and reviews probabilistic CBFs and risk measures. Section III presents our approach: establishing sub-Gaussianity of the barrier function increment, deriving finite-sample CVaR bounds, and reformulating the probabilistic CBF condition as a tractable optimization problem. Section IV demonstrates the efficacy of our approach through numerical experiments.

II. Problem Setup

We begin with a discrete-time stochastic system[1]

$$x_{t+1} = f(x_t) + g(x_t)u_t + d_t := F(x_t, u_t, d_t), \qquad (1)$$

where $x_t \in \mathbb{R}^{n_x}$ is the unknown true state, $u_t \in \mathbb{R}^{n_u}$ is the control input, and $d_t \in \mathbb{R}^{n_x}$ is an additive, state-independent process disturbance signal. The functions $f : \mathbb{R}^{n_x} \to \mathbb{R}^{n_x}$ and $g : \mathbb{R}^{n_x} \to \mathbb{R}^{n_x \times n_u}$ are Lipschitz continuous with constants L_f and L_g, respectively.

Assumption 1. The true state $x_t \sim \mathcal{N}(\mu_{x_t}, \Sigma_{x_t})$ and disturbance $d_t \sim \mathcal{N}(\mu_{d_t}, \Sigma_{d_t})$ are independent Gaussian random vectors with known moments. We assume the disturbance distribution is state-independent. The state estimate μ_{x_t} and uncertainty covariance Σ_{x_t} may come from an estimator such as an Extended Kalman Filter (EKF).

Remark 1. While we state the theory under the Gaussian assumption for clarity, all our results extend to sub-Gaussian distributions simply by replacing $\sqrt{\lambda_{\max}(\Sigma)}$ in Lemma 1 and Theorem 1 with the sub-Gaussian parameter. This extension provides robustness against filtering modeling errors, such as the Gaussian posterior assumption made by the EKF.

Let $h : \mathbb{R}^{n_x} \to \mathbb{R}$ be a barrier function defining the safe set as the zero sublevel set (opposite to standard CBFs [1]):

$$\mathcal{C} = \{x \in \mathbb{R}^{n_x} : h(x) \le 0\}.$$

Assumption 2. The barrier function h is Lipschitz continuous with constant L_h and is convex.

Our objective is to design an output feedback controller $k : \mathbb{R}^{n_x} \to \mathbb{R}^{n_u}$ that keeps the system trajectories within the safe set $\mathcal{C}$ at all times. Since the true state x_t is unobservable, the controller must operate using the state estimate and account for its uncertainty. In deterministic systems, the discrete-time CBF condition $h(x_{t+1}) \le \gamma h(x_t)$ with $\gamma \in [0, 1]$ ensures forward invariance of $\mathcal{C}$ [1], [25]. However, under stochastic disturbances with unbounded support, the system exits the safe set in finite time with probability one [4], making deterministic safety guarantees impossible. This necessitates a probabilistic framework for safety. Following [6], we say k is ϵ-safe over horizon H if, starting from any $x_0 \in \mathcal{C}$, the probability of remaining in $\mathcal{C}$ for all $t \in [0, H]$ is at least $1 - \epsilon$. To construct such controllers, we leverage the notion of probabilistic CBF developed in [6]. The key idea is to enforce a single-step probabilistic constraint: for each $x_t \in \mathcal{C}$, there exists a $u_t \in \mathcal{U}$ satisfying

$$\mathbb{P}\big(h(F(x_t, u_t, d_t)) \le \gamma h(x_t)\big) \ge 1 - \alpha, \qquad (2)$$

where $\alpha \in (0, 1)$ is the single-step risk parameter and $\gamma \in [0, 1]$ is a decay rate. It has been shown [6] that if (2) is satisfied at each time step with $\alpha \le 1 - (1 - \epsilon)^{1/H}$, the controller achieves ϵ-safety over horizon H.

With this theoretical foundation for safe control, we implement an optimization problem with a probabilistic CBF, which is solved at each time step. Given a desired control $u_{\text{des}} \in \mathbb{R}^{n_u}$ from a nominal policy, we seek to solve

$$\min_{u_t \in \mathcal{U}} \quad \|u_t - u_{\text{des}}\|^2 \quad \text{s.t.} \quad (2) \qquad (3)$$

where $\mathcal{U} \subseteq \mathbb{R}^{n_u}$ is a compact admissible control set with $u_{\max} := \sup_{u \in \mathcal{U}} \|u\|$. In practice, this problem presents two challenges: the chance constraint is generally non-convex, and

[1] We use $\|\cdot\|$ for the Euclidean norm. For a symmetric matrix Σ, we denote its largest and smallest eigenvalues by $\lambda_{\max}(\Sigma)$ and $\lambda_{\min}(\Sigma)$. For any $a \in \mathbb{R}$, the positive part is $(a)_+ := \max(0, a)$. Also, let $[n] := \{1, 2, \ldots, n\}$ be the index set.

evaluating the probability requires the distribution of

$$\Delta h(x_t, u_t, d_t) := h(F(x_t, u_t, d_t)) - \gamma h(x_t), \qquad (4)$$

which is intractable for nonlinear f, g, and h. We henceforth suppress the arguments of Δh when there is no ambiguity.

III. Proposed Approach

We address the intractability of (3) through two key steps: reformulating the constraint using CVaR and estimating it through particle propagation. Recall that CVaR provides a sufficient condition for chance constraints: [2]

$$\mathrm{CVaR}_\alpha(\Delta h) \le 0 \implies \mathbb{P}(\Delta h \le 0) \ge 1 - \alpha. \qquad (5)$$

This reformulation replaces the discontinuous chance constraint with a continuous sufficient condition, amenable to tractable empirical approximations. Evaluating $\mathrm{CVaR}_\alpha(\Delta h)$ requires knowledge of the distribution of Δh. Under Assumption 1, we represent the joint uncertainty in state and disturbance using n i.i.d. samples:

$$\{(x_t^i, d_t^i)\}_{i=1}^n, \text{ where } x_t^i \sim \mathcal{N}(\mu_{x_t}, \Sigma_{x_t}),\ d_t^i \sim \mathcal{N}(\mu_{d_t}, \Sigma_{d_t}).$$

Accordingly, for a given control input u_t, each particle propagates through the dynamics as

$$x_{t+1}^i = f(x_t^i) + g(x_t^i)u_t + d_t^i, \quad i \in [n] \qquad (6)$$

yielding barrier function increments

$$\Delta h^i(u_t) = h(x_{t+1}^i) - \gamma h(x_t^i), \quad i \in [n]. \qquad (7)$$

These particle-based realizations, $\{\Delta h^i(u_t)\}_{i=1}^n$, yield a tractable constraint upper-bounding the true CVaR.

A. Sub-Gaussianity of the Barrier Function Increment

To tightly bound the gap between the particle-based CVaR approximation and the ground truth, we exploit the distributional properties of Δh. In this section, we establish that Δh is sub-Gaussian with an explicit parameter that depends on Lipschitz constants and covariance matrices [3]. We will use the following standard Gaussian concentration result [26].

Lemma 1 (Gaussian concentration)**.** *Let $\xi \sim \mathcal{N}(\mu, \Sigma)$, and let $\phi : \mathbb{R}^n \to \mathbb{R}^m$ be L-Lipschitz. Then $\phi(\xi)$ is sub-Gaussian with parameter $\sigma \le L\sqrt{\lambda_{\max}(\Sigma)}$.*

We now prove Δh follows a sub-Gaussian distribution and its parameter is bounded.

Theorem 1 (Sub-Gaussianity of Δh, uniform in u)**.** *Δh in (4) is sub-Gaussian with u-independent parameter*

$$\sigma_{\Delta h} \le C\sqrt{L_x^2 \lambda_{\max}(\Sigma_x) + L_d^2 \lambda_{\max}(\Sigma_d)}, \quad \forall u \in \mathcal{U},$$

where $L_x = L_h(L_f + L_g u_{\max} + |\gamma|)$ and $L_d = L_h$.

Proof. Since both terms in Δh depend on x, it is not a sum of independent sub-Gaussians. Treating Δh as a function of (x, d), Lipschitz continuity of h, f, g give

$$|\Delta h(x_1, d_1) - \Delta h(x_2, d_2)| \le L_x \|x_1 - x_2\| + L_d \|d_1 - d_2\|.$$

Since x and d are statistically independent, the joint covariance is block-diagonal, and Lemma 1 applied after standardization yields the bound. □

[2] $\mathrm{CVaR}_\alpha(\xi) = \mathbb{E}[\xi | \xi \ge q_\alpha]$ represents the expected value of a random variable ξ in its worst $1-\alpha$-percent of outcomes, where q_α is the α-quantile.

[3] A random vector X is *sub-Gaussian with parameter* σ if, for all $\lambda \in \mathbb{R}$ and unit vectors $\|v\| = 1$, $\mathbb{E}[\exp(\lambda\langle X - \mathbb{E}[X], v\rangle)] \le e^{\lambda^2\sigma^2/2}$.

B. Finite-Sample CVaR Approximation

Having established the sub-Gaussian property of Δh, we now derive finite-sample bounds that relate the particle-based CVaR to the true CVaR. This enables replacing the intractable CVaR constraint in (3) with a computable sample-based condition. We begin with a technical lemma bounding tail integrals for sub-Gaussian random variables.

Lemma 2 (Sub-Gaussian Tail Integral Bound)**.** *Let X be a sub-Gaussian random variable with mean μ and parameter $\bar{\sigma}^2$. Let $\epsilon \in (0, 0.5)$ and define $z^*(\epsilon) := \mu + \bar{\sigma}\sqrt{2\ln(1/\epsilon)}$. Then, for any $z \ge z^*(\epsilon)$,*

$$\int_z^\infty \mathbb{P}(X > y)\, dy \ \le\ \frac{\bar{\sigma}\,\epsilon}{\sqrt{2\ln(1/\epsilon)}}.$$

Proof. Let $g(y) := \exp(-(y-\mu)^2/(2\bar{\sigma}^2))$. By the sub-Gaussian assumption and $z \ge \mu$, we have $\mathbb{P}(X > y) \le g(y)$ for all $y \ge z$. Thus:

$$\int_z^\infty \mathbb{P}(X > y)\, dy \le \int_z^\infty g(y)\, dy \le \int_{z^*(\epsilon)}^\infty g(y)\, dy,$$

where the last inequality uses monotonicity (g decreasing on $[\mu, \infty)$ and $z \ge z^*(\epsilon) \ge \mu$). With substitution $t = (y-\mu)/\bar{\sigma}$ and $t^* = \sqrt{2\ln(1/\epsilon)}$, $\int_{z^*}^\infty g\, dy = \bar{\sigma}\int_{t^*}^\infty e^{-t^2/2}\, dt$. Mills' inequality $\int_{t^*}^\infty e^{-t^2/2}\, dt \le e^{-(t^*)^2/2}/t^*$ with $e^{-(t^*)^2/2} = \epsilon$ yields the claim. □

Having established this tail bound, we now state our main result, a finite-sample upper bound on CVaR in terms of order statistics. While the theorem uses order statistics, Section III-C bypasses their explicit computation.

Theorem 2 (Data-Dependent Sub-Gaussian CVaR Bound)**.** *Let $X_1, \ldots, X_n$ be i.i.d. sub-Gaussian random variables with mean μ and sub-Gaussian parameter $\sigma^2 > 0$. Let $Z_1 \le \cdots \le Z_n$ denote the order statistics of $X_1, \ldots, X_n$. Assume $Z_n \ge \mu + \bar{\sigma}\sqrt{2\ln(1/\varepsilon_n(\delta))}$ (see Remark 2). Then for any $\alpha \in (0, 1)$ and $\delta \in (0, 0.5]$ with $\varepsilon_n(\delta) < 0.5$ (which holds when $n > 2\ln(2/\delta)$),*

$$\mathbb{P}\Bigg[\mathrm{CVaR}_\alpha(X) \le Z_n + C_{\text{tail}}(n, \delta) - \frac{1}{\alpha}\sum_{i=1}^{n-1}(Z_{i+1} - Z_i)\left(\frac{i}{n} - \varepsilon_n(\delta) - (1-\alpha)\right)_+\Bigg] \ge 1 - \delta,$$

where $\varepsilon_n(\delta) := \sqrt{\frac{\ln(2/\delta)}{2n}}$ and $C_{\text{tail}}(n, \delta) := \frac{\bar{\sigma}\,\varepsilon_n(\delta)}{\alpha\sqrt{2\ln(1/\varepsilon_n(\delta))}}$

Proof. Let F and F_ω denote the true and empirical Cumulative Distribution Functions (CDFs), respectively. By the DKW inequality with Massart's tight constants [27],

$$\mathbb{P}\left(\sup_{x \in \mathbb{R}}(F_\omega(x) - F(x)) \le \varepsilon_n(\delta)\right) \ge 1 - \delta.$$

Define the event $\mathcal{E}_{\text{DKW}} := \{\sup_x (F_\omega(x) - F(x)) \le \varepsilon_n(\delta)\}$. On this event, since $F_\omega(Z_n) = 1$, we have $F(Z_n) \ge 1 - \varepsilon_n(\delta)$, and hence $\mathbb{P}(X > Z_n) \le \varepsilon_n(\delta)$. Applying Lemma 2 with $z = Z_n$ then yields

$$I_{\text{tail}} := \int_{Z_n}^{\infty} (1 - F(y))\, dy \le \frac{\bar{\sigma}\, \varepsilon_n(\delta)}{\sqrt{2 \ln(1/\varepsilon_n(\delta))}} =: \bar{I}_{\text{tail}}.$$

Now define the lower bound CDF $G_\omega^-(x) := \max\{0, F_\omega(x) - \varepsilon_n(\delta)\}$, which satisfies $G_\omega^-(x) \le F(x)$ for all x on $\mathcal{E}_{\text{DKW}}$. Using the integral representation of CVaR:

$$\text{CVaR}_\alpha(X) \le \inf_{c \in \mathbb{R}} \left\{ c + \frac{1}{\alpha} \left(\int_c^{Z_n} (1 - G_\omega^-(y)) dy + I_{\text{tail}} \right) \right\}.$$

Applying Lemma 1 in [28], this yields

$$\text{CVaR}_\alpha(X) \le Z_n - \frac{1}{\alpha} \int_{-\infty}^{Z_n} (G_\omega^-(y) - (1-\alpha))_+ dy + \frac{I_{\text{tail}}}{\alpha}.$$

Since $G_\omega^-(x) = \max\{0, i/n - \varepsilon_n(\delta)\}$ for $x \in (Z_i, Z_{i+1}]$, this integral equals $\sum_{i=1}^{n-1} (Z_{i+1} - Z_i)(i/n - \varepsilon_n(\delta) - (1-\alpha))_+$. Setting $C_{\text{tail}} := \bar{I}_{\text{tail}}/\alpha$ completes the proof. $\square$

Remark 2 (Sample Augmentation)**.** The hypothesis is enforced by appending $\Delta h^n(u_t) := \bar{\mu}(u_t) + \bar{\sigma}\sqrt{2\ln(1/\varepsilon_n(\delta))}$ to $n-1$ i.i.d. realizations of Δh, where $\bar{\mu}(u_t) \ge \mathbb{E}[\Delta h]$ is obtained via Lipschitz analysis (exact and affine in u_t for linear h) or by sample-mean estimation.

Remark 3. $\bar{I}_{tail}$ decays as $O((n \ln n)^{-1/2})$.

Remark 4 (A Posteriori Verification)**.** Theorem 2 assumes i.i.d. samples, which may be violated when the same particles are used for both optimization and evaluation. In our experiments, we evaluate with new samples. Although no verification failures were observed in our case, standard remedies include resampling or tightening the risk level [29].

Alternatively, when the barrier function increment admits a separable structure, the i.i.d. assumption holds directly without requiring sample splitting as detailed in Remark 4.

Corollary 1 (Separable Structure)**.** *Suppose $\Delta h = a(u_t) + b(x_t, d_t)$, and let $\{Z_i\}_{i=1}^n$ be the order statistics of $\{b(x_t^i, d_t^i)\}_{i=1}^n$ from i.i.d. samples. By the translation invariance of* CVaR*, applying Theorem 2 to b shows that any u_t satisfying $a(u_t) + C \le 0$ guarantees* $\text{CVaR}_\alpha(\Delta h) \le 0$ *with probability at least $1 - \delta$, where C denotes the CVaR upper bound in Theorem 2.*

We note that for the control-affine dynamics (1), this separable structure holds when h is linear and g is constant. It arises naturally in linear systems or fully-actuated mechanical systems with position-based safety constraints.

Remark 5 (Safety Over a Finite Time Horizon)**.** Suppose Theorem 2 holds with confidence $1 - \delta$ at each $t \in [H]$ using independent particles. Let $q_0 := \mathbb{P}(h(x_0) \le 0)$ denote the probability that the true initial state lies in $\mathcal{C}$ under Assumption 1. Then, setting $\alpha \le \epsilon/H$, an analogue of Proposition 6 of [6] holds under state uncertainty:

$$\mathbb{P}_{\text{ptcl}} \left(\mathbb{P}\left(\bigcap_{t=0}^{H} \{h(x_t) \le 0\} \right) \ge q_0 - \epsilon \right) \ge 1 - H\delta, \tag{8}$$

where $\mathbb{P}_{\text{ptcl}}$ denotes the probability over particle randomness.

C. CVaR Reformulation as a Linear Program

Theorem 2 provides a high-probability upper bound on $\text{CVaR}_\alpha(\Delta h)$ in terms of order statistics and correction terms. However, this form requires explicit sorting, making it unsuitable for direct optimization. We resolve this by reformulating the bound as a linear program. Recall the following classical result expressing sample CVaR as an optimization problem over unordered samples [23].

Theorem 3 (Sample CVaR Formula [23])**.** *Let $\{w_1, \dots, w_n\}$ be n real numbers with order statistics $w_{(1)} \le \dots \le w_{(n)}$, and let $\alpha \in (0,1)$. Then*

$$\begin{aligned} &\inf_{\theta \in \mathbb{R}} \left\{ \theta + \frac{1}{n\alpha} \sum_{i=1}^n (w_i - \theta)_+ \right\} \\ &= w_{(n)} - \frac{1}{\alpha} \sum_{j=1}^{n-1} (w_{(j+1)} - w_{(j)}) \left(\frac{j}{n} - (1-\alpha) \right)_+. \end{aligned}$$

Crucially, the left-hand side depends only on the unordered samples $\{w_i\}$, while the right-hand side involves order statistics—enabling computation without explicit sorting. To accommodate the correction term $\varepsilon_n(\delta)$ from Theorem 2, we extend this classical formula.

Corollary 2 (Shifted Threshold CVaR Formula)**.** *Under the setting of Theorem 3, let $B \in (0, \alpha)$. Then*

$$\begin{aligned} &\frac{B}{\alpha} w_{(n)} + \frac{\alpha - B}{\alpha} \inf_{\theta \in \mathbb{R}} \left\{ \theta + \frac{1}{n(\alpha - B)} \sum_{i=1}^n (w_i - \theta)_+ \right\} \\ &= w_{(n)} - \frac{1}{\alpha} \sum_{j=1}^{n-1} (w_{(j+1)} - w_{(j)}) \left(\frac{j}{n} - B - (1-\alpha) \right)_+. \end{aligned}$$

Proof. Factor $w_{(n)}$ from the right-hand side and apply Theorem 3 with risk level $\alpha - B$. $\square$

Setting $B = \varepsilon_n(\delta)$ yields the tractable constraint for our probabilistic CBF formulation. The infimum in Corollary 2 admits a linear-constraint reformulation: the constraint

$$\frac{B}{\alpha} w_{(n)} + \frac{\alpha - B}{\alpha} \inf_{\theta} \left\{ \theta + \frac{1}{n(\alpha - B)} \sum_{i=1}^n (w_i - \theta)_+ \right\} \le 0, \tag{9}$$

holds iff there exist $\theta \in \mathbb{R}$ and $\{s_i \ge 0\}_{i=1}^n$ satisfying

$$\begin{aligned} \frac{B}{\alpha} w_i + \frac{\alpha - B}{\alpha} \left(\theta + \frac{1}{n(\alpha - B)} \sum_{j=1}^n s_j \right) \le 0,\ i \in [n], \\ s_i \ge w_i - \theta, \quad s_i \ge 0,\ i \in [n]. \end{aligned} \tag{10}$$

Here, s_i encode the positive parts $(w_i - \theta)_+$.

D. Tractable Probabilistic CBF Optimization

We now apply the reformulation from Section III-C to construct a computationally efficient version of (3). Theorem 2 guarantees with probability at least $1-\delta$ that

$$\mathrm{CVaR}_\alpha(\Delta h) \le \Delta h_{(n)} + C_{\text{tail}}(n,\delta) - \frac{1}{\alpha}\sum_{j=1}^{n-1}(\Delta h_{(j+1)} - \Delta h_{(j)})\left(\frac{j}{n} - \varepsilon_n(\delta) - (1-\alpha)\right)_+,$$

where $\{\Delta h^i\}_{i=1}^n$ are particle realizations and $\Delta h_{(1)} \le \cdots \le \Delta h_{(n)}$ denote their order statistics. Applying Corollary 2 with $w_i = \Delta h^i$, $B = \varepsilon_n(\delta)$, and the reformulation (10), the sufficient condition $\mathrm{CVaR}_\alpha(\Delta h) \le 0$ holds iff there exist $\theta \in \mathbb{R}$ and $\{s_i \ge 0\}_{i=1}^n$ such that

$$C_{\text{tail}}(n,\delta) + \frac{\varepsilon_n(\delta)}{\alpha}\Delta h^i(u_t) + \frac{\alpha - \varepsilon_n(\delta)}{\alpha}\Bigg(\theta + \frac{1}{n(\alpha - \varepsilon_n(\delta))}\sum_{j=1}^n s_j\Bigg) \le 0, \quad i \in [n], \tag{11}$$

$$s_i \ge \Delta h^i(u_t) - \theta, \quad s_i \ge 0, \quad i \in [n]. \tag{12}$$

This yields the following tractable approximation to (3).

Problem 1. Given particles $\{(x_t^i, d_t^i)\}_{i=1}^n$ sampled according to Assumption 1 and a desired control $u_{\text{des}} \in \mathbb{R}^{n_u}$ from a nominal controller, define $\Delta h^i(u_t)$ for $i \in [n-1]$ via (7) and $\Delta h^n(u_t)$ via Remark 2, then solve

$$\min_{u_t \in \mathcal{U}, \theta, \{s_i\}} \quad \|u_t - u_{\text{des}}\|^2 \quad \text{s.t.} \quad (11),\ (12). \tag{13}$$

We remark that Problem 1 is convex with quadratic objective when $\Delta h^i(u_t)$ and $\mathcal{U}$ are convex, reducing to a QP when Δh^i is affine in u_t and $\mathcal{U}$ is polytopic.

Remark 6 (Feasibility Verification)**.** Several methods can verify feasibility; e.g., when h is affine and $\mathcal{U}$ is a polytope, the feasibility of Problem 1 can be verified *a priori* by solving a dual Linear Programming (LP) (via Farkas' Lemma) with only n_u+3 coupled constraints, independent of the particle count n. In practice, α is set from the mission safety budget ($\alpha \approx \epsilon/H$); $\delta < \alpha$ and n are jointly tuned to satisfy $n > 2\ln(2/\delta)$ while balancing feasibility and solver time.

Remark 7. By applying Corollary 2 to $\{b(x_t^i, d_t^i)\}_{i=1}^n$, the constraint in Corollary 1 can be reformulated as a linear program in auxiliary variables $(\theta, \{s_i\})$. When $a(u_t)$ is affine in u_t, the resulting problem remains a QP with the same computational complexity as Problem 1.

IV. Numerical Experiments

We demonstrate the efficacy of our proposed approach through Monte Carlo simulations on a non-holonomic mobile robot navigating under uncertainty[4]. The robot must reach a goal while respecting a geofence constraint [5], subject to both localization uncertainty and stochastic dynamics. We compare our approach against five baselines: **(i)** the nominal deterministic CBF ignoring uncertainty (*Det. CBF*); **(ii)** a naively robustified deterministic CBF that adds a constant margin $m \ge 0$ to the CBF condition without principled tuning (*Det. CBF (Robust)*); **(iii)** a scenario approach [22] directly approximating the chance constraint with $N = 500$ samples, satisfying an (α,δ)-guarantee analogous to Theorem 2 (*Scenario*); **(iv)** a SAA [21] with Big-M continuous relaxation of the chance-constraint indicator $\mathbf{1}\{\Delta h^i > 0\}$ (*SAA*); and **(v)** a probabilistic CBF using the conservative DKW inequality (*DKW*). The robot operates in a 2D workspace with state $x = [r_x, r_y, \theta]^\top \in \mathbb{R}^3$, where (r_x, r_y) is the position and θ is the heading angle. Its control input is $u = [v, \omega]^\top$, where v and ω are linear and angular velocities with $|v| \le v_{\text{m}}$, $|\omega| \le \omega_{\text{m}}$, where v_{m}, ω_{m} are derived from an Ackermann robot with a wheelbase of w_l and a maximum steering angle of η. To obtain a formulation suitable for CBF design, we follow [30] and introduce a shifted reference point. Define

$$R(\theta) = \begin{bmatrix}\cos\theta & -\sin\theta \\ \sin\theta & \cos\theta\end{bmatrix}, \quad p = s + \ell R(\theta) e_1, \tag{14}$$

where $s = [r_x, r_y]^\top$ is the robot's center position, $e_1 = [1, 0]^\top$, and $\ell > 0$ is a design parameter. This shifted point p lies orthogonal to the wheel axis and satisfies $\dot{p} = R(\theta)L^{-1}u$, where $L = \mathrm{diag}(1, 1/\ell)$. Thus, both linear and angular velocity appear in the dynamics of p, yielding relative degree one for position-based barrier functions.

The barrier function $h(x)$ is a linear geofence: $h(x) = r_y$. The robot has Gaussian localization noise and process disturbances with covariances Σ_x, Σ_d, respectively. A nominal PID controller drives the robot toward the goal at $(0, -0.05)$, while our CBF-based safety filter modifies the control input when necessary. In our simulation, the state and its time-varying uncertainty covariance are dynamically estimated via an EKF with process noise $Q = \Sigma_d$, measurement noise R, and initial covariance P_0, using partial measurements of r_y and θ. We perform a Monte Carlo analysis with 10000 trials on an Apple M1 Pro. [6] Each trial begins at $x_0 \sim \mathcal{N}(\mu_0, P_0)$ and runs for up to 15 s with a time step of Δt. A trial is deemed successful if the robot reaches within 0.02 m of the goal. Table I summarizes all trials, and Figure 1 shows representative trajectories. The unsafe baselines (*Det. CBF*, *Scenario*, *SAA*) violate the boundary in nearly every trial, while the overly conservative baselines (*Det. CBF (Rob.)*, *DKW*) avoid violations only by stalling before the goal. Only the proposed method achieves both safety (within the $\alpha = 10\%$ bound) and 100% goal-reaching with the smallest final distance. This also empirically confirms robustness to the EKF's approximation error.

Method	Violation	Reach	Mean Dist. to Goal [m]
Det. CBF	100%	100%	0.05
Det. CBF (Rob.)	0%	0%	0.07
Scenario	60%	100%	0.05
SAA	100%	100%	0.06
DKW	0%	0%	0.11
Proposed	**1%**	**100%**	**0.02**

TABLE I
Summary of Monte Carlo simulations.

Focusing on approaches approximating CVaR, Table II reports the trade-off between conservatism and computational cost as the particle count n varies. Note that the CVaR estimation error was measured as the difference between the particle-

[4] Implementation available at: https://github.com/kazuyae/Stochastic_CBF

[5] Nonlinear barriers can be incorporated via local linearization.

[6] The simulations use the constants $\Sigma_d = Q = 10^{-6}\,\mathrm{diag}(1,1,25)$, $P_0 = 10^{-4}\,\mathrm{diag}(1,1,9)$, $R = 10^{-6}\,\mathrm{diag}(4,49)$, $\mu_0 = [0, -0.5, \pi/2]^\top$, $v_{\text{m}} = 0.3$ m/s, $\omega_{\text{m}} \approx 0.10$ rad/s, $w_l = 2.5$ m, $\ell = 0.2$ m, $n = 500$, $\eta = 40°$, $\alpha = 0.1$, $\delta = 0.1$, $\gamma = 0.2$, $\Delta t = 0.5$ s, $m = 0.1$.

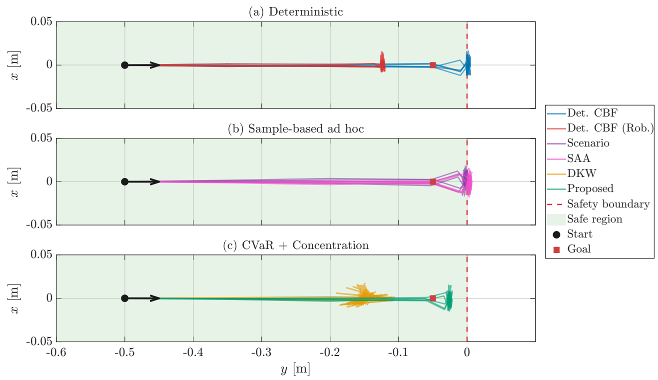

Fig. 1. Representative trajectories for the six tested methods. The plots display the robot's path from the start position (black circle) to the goal position (red box), with the safety boundary at $\boldsymbol{y} = \mathbf{0}$ (red dotted line).

n	10	50	200	500	1000
Gap, Prop. [m]	0.0337	0.0276	0.0241	0.0233	0.0231
Gap, DKW [m]	0.112	0.166	0.124	0.116	0.115
Time, Prop. [ms]	4.30	5.92	12.1	34.5	86.6
Time, DKW [ms]	1.23	5.90	8.37	19.4	46.4

TABLE II

PROPOSED VS. DKW: CVAR ESTIMATION ERROR AND SOLVER TIME

based CVaR estimate when running probabilistic CBF-based methods and a Monte Carlo ground truth computed with 10^6 samples. Unlike *DKW*, whose error does not decrease monotonically with n and consistently overestimates the true risk, the proposed method monotonically tightens the CVaR approximation as n grows. This improvement is achieved while keeping the per-step solver time below $100\,\mathrm{ms}$ for all tested n.

V. Conclusion

We presented a particle-based probabilistic CBF framework that provides probabilistic safety guarantees for systems subject to both state estimation uncertainty and process disturbances. Our approach exploits the sub-Gaussian structure of the barrier function increment to derive explicit concentration bounds, enabling the reformulation of the intractable probabilistic CBF condition as a tractable optimization program. Monte Carlo simulations on a non-holonomic mobile robot navigation problem demonstrated that the proposed approach outperforms both the deterministic baseline and the existing truncation-based probabilistic CBF method by providing tighter yet provably valid risk approximations. A current limitation is the double probability structure—the CVaR bound holds with probability $1-\delta$ while the safety constraint holds with probability $1-\alpha$—unifying them remains an important direction for future work.

References

[1] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs for safety critical systems," *IEEE Trans. Autom. Control*, vol. 62, no. 8, pp. 3861–3876, 2017.

[2] I. A. D. Nesnas *et al.*, "Autonomous exploration of small bodies toward greater autonomy for deep space missions," *Frontiers in Robotics and AI*, vol. 8, no. 11, pp. 1872–1873, 2021.

[3] K. Echigo, A. Cauligi, S. Bandyopadhyay, D. Scharf, G. Lantoine, B. Acikmese, and I. Nesnas, "Autonomy in the real-world: Autonomous trajectory planning for asteroid reconnaissance via stochastic optimization," in *AIAA SciTech Forum*, 2025.

[4] R. K. Cosner, P. Culbertson, A. J. Taylor, and A. D. Ames, "Robust safety under stochastic uncertainty with discrete-time control barrier functions," in *Proc. Robot. Sci. Syst. (RSS)*, 2023.

[5] A. Clark, "Control barrier functions for complete and incomplete information stochastic systems," in *Proc. Amer. Control Conf. (ACC)*, 2019, pp. 2928–2935.

[6] P. Mestres, B. Werner, R. K. Cosner, and A. D. Ames. (2025) Probabilistic control barrier functions: Safety in probability for discrete-time stochastic systems. Available at https://arxiv.org/abs/2510.01501.

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

[8] Y. Wang, X. Shen, and H. Qian, "Stochastic model predictive control with probabilistic control barrier functions and smooth sample-based approximation," in *Proc. IEEE Conf. Decis. Control (CDC)*, 2024.

[9] X. Shen, Y. Wang, K. Hashimoto, Y. Wu, and S. Gros, "Probabilistic reachable sets of stochastic nonlinear systems with contextual uncertainties," *Automatica*, vol. 176, p. 112237, 2025.

[10] S. Kolathaya and A. D. Ames, "Input-to-state safety with control barrier functions," *IEEE Control Syst. Lett.*, vol. 3, no. 1, pp. 108–113, 2019.

[11] A. Anil, A. J. Taylor, C. R. He, G. Orosz, and A. D. Ames, "Safe controller synthesis with tunable input-to-state safe control barrier functions," *IEEE Control Syst. Lett.*, vol. 6, pp. 908–913, 2022.

[12] M. Jankovic, "Robust control barrier functions for constrained stabilization of nonlinear systems," *Automatica*, vol. 96, pp. 359–367, 2018.

[13] W. Wang and X. Xu, "Disturbance observer-based robust control barrier functions," in *Proc. Amer. Control Conf. (ACC)*, 2023, pp. 3681–3687.

[14] C. Santoyo, M. Dutreix, and S. Coogan, "Verification and control for finite-time safety of stochastic systems via barrier functions," in *Proc. IEEE Conf. Control Technol. Appl.*, 2019, pp. 712–717.

[15] A. A. do Nascimento, A. Papachristodoulou, and K. Margellos, "Probabilistically safe controllers based on control barrier functions and scenario model predictive control," in *Proc. IEEE Conf. Decis. Control (CDC)*, 2024, pp. 1814–1819.

[16] A. Clark, "Control barrier functions for stochastic systems," *Automatica*, vol. 130, p. 109688, 2021.

[17] S. Yaghoubi, G. Fainekos, T. Yamaguchi, D. Prokhorov, and B. Hoxha, "Risk-bounded control with Kalman filtering and stochastic barrier functions," in *Proc. IEEE Conf. Decis. Control (CDC)*, 2021, pp. 5213–5219.

[18] J. Steinhardt and R. Tedrake, "Finite-time regional verification of stochastic non-linear systems," *The International Journal of Robotics Research*, vol. 31, no. 7, pp. 901–923, 2012.

[19] M. Vahs, C. Pek, and J. Tumova, "Belief control barrier functions for risk-aware control," *IEEE Robotics and Automation Letters*, vol. 8, no. 12, pp. 8565–8572, 2023.

[20] M. Kishida, "Risk-aware control of discrete-time stochastic systems: Integrating Kalman filter and worst-case CVaR in control barrier functions," in *Proc. IEEE Conf. Decis. Control (CDC)*, 2024, pp. 2019–2024.

[21] B. K. Pagnoncelli, S. Ahmed, and A. Shapiro, "Sample average approximation method for chance constrained programming: theory and applications," *Journal of Optimization Theory and Applications*, vol. 142, no. 2, pp. 399–416, 2009.

[22] G. C. Calafiore and M. C. Campi, "The scenario approach to robust control design," *IEEE Trans. Autom. Control*, vol. 51, no. 5, pp. 742–753, 2006.

[23] R. T. Rockafellar and S. Uryasev, "Optimization of conditional value-at-risk," *J. Risk*, vol. 2, no. 3, pp. 21–41, 2000.

[24] M. Vahs and J. Tumova, "Risk-aware control for robots with non-Gaussian belief spaces," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, 2024, pp. 11 661–11 667.

[25] A. Agrawal and K. Sreenath, "Discrete control barrier functions for safety-critical control of discrete systems with application to bipedal robot navigation," in *Proc. Robot. Sci. Syst. (RSS)*, 2017, pp. 1–10.

[26] R. Vershynin, *High-Dimensional Probability: An Introduction with Applications in Data Science*. Cambridge University Press, 2018.

[27] P. Massart, "The tight constant in the Dvoretzky-Kiefer-Wolfowitz inequality," *The Annals of Probability*, vol. 18, no. 3, pp. 1269 – 1283, 1990.

[28] P. Thomas and E. Learned-Miller, "Concentration inequalities for conditional value at risk," in *Proc. Int. Conf. Mach. Learn.*, 2019, pp. 6225–6233.

[29] Y. Zhao, X. Yu, M. Sesia, J. V. Deshmukh, and L. Lindemann. (2025) Conformal predictive programming for chance constrained optimization. Available at https://arxiv.org/abs/2402.07407.

[30] P. Glotfelter, I. Buckley, and M. Egerstedt, "Hybrid nonsmooth barrier functions with applications to provably safe and composable collision avoidance for robotic systems," *IEEE Robot. Autom. Lett.*, vol. 4, no. 2, pp. 1303–1310, 2019.