Journal of Cyber-Physical Security and Robotics
Volume 02, Issue 4, 46-54, 2026

Article

Control barrier functions coupled with an ensemble transform Kalman filter for calibrated chance-constrained safety filtering

Toheeb Babatunde Ibrahim[1,*], **Olasubomi Akanbi**[2], **Joseph Oladejo**[1], **Warith Akinade**[1]

[1]*Department of Mathematics and Statistics, Bowling Green State University, Bowling Green, OH 43403, USA*
[2]*Department of Systems Engineering, Bowling Green State University, Bowling Green, OH 43403, USA*

[*]Corresponding author. *E-mail addresses:* toheeb.ibrahim@aims-cameroon.org (T.B. Ibrahim), olasubomi.akanbi1@gmail.com (O. Akanbi), oladejojoseph0@gmail.com (J. Oladejo), warith@aims.edu.gh (W. Akinade).

Received: 06 September 2026 Accepted: 01 October 2026 Publication: 12 November 2026

Abstract

Control barrier functions (CBFs) certify safety for the true state, yet controllers act on estimates. We propose the Ensemble-Transform Control Barrier Function (ET-CBF), a safety filter that evaluates the discrete-time barrier condition on every member of an ensemble transform Kalman filter (ETKF) analysis and enforces a one-sided chance constraint on the resulting ensemble distribution. For control-affine systems the filter is a second-order cone program whose size is set by the ensemble size rather than the state dimension, and it needs no Jacobians of the dynamics, the measurement model or the barrier. We prove that ET-CBF coincides with the Kalman-filter chance-constrained CBF for affine barriers under linear-Gaussian models, show that it retains the curvature (Jensen) term that linearised belief CBFs omit, with a sign rule that predicts when linearisation is optimistic, and give a probabilistic safety certificate under ensemble consistency. We also identify a failure mode specific to ensemble beliefs, miscalibration at small ensemble sizes, and correct it with a Student-t predictive gain. In 5,600 closed-loop Monte Carlo runs of a unicycle with nonlinear range-bearing sensing, certainty-equivalent CBFs on EKF or ETKF means entered the obstacle in 69–79% of runs, while ET-CBF limited this to 0–4.5% with one-step failure rates of 3.0–3.6% against a 5% design level, on par with a linearised EKF chance constraint. With four members, the corrected gain reduced one-step failure from 14.7% to 5.1%. In a static test on a keep-in barrier, linearisation let the realised risk reach 10.5% against a 5% target, while a 200-member ensemble held it to 7.0%.

Keywords: Control barrier functions; Ensemble transform Kalman filter; Chance constraints; Safety-critical control; State estimation; Second-order cone programming

DOI: 10.64235/bnd7mdjkx0

1. Introduction

A control barrier function (CBF) safety filter can be no safer than the state estimate it receives. Classical CBF quadratic programs [1–3] certify forward invariance of a safe set for the true state, but deployed controllers act on an estimate. When the estimate is used as if it were the truth (certainty equivalence), the barrier condition can hold for the estimate while failing for the true state; in our experiments this occurred in more than half of the closed-loop runs (Section 6).

Two families of remedies exist. Robust, bounded-error methods tighten the barrier by a worst-case estimation margin [4, 5]. Stochastic and risk-aware methods instead reason about the estimator's belief: they replace the barrier condition with a chance or risk constraint whose mean and spread come from an extended Kalman filter (EKF) or a Kalman filter [6–10], or from a weighted particle set [11, 12]. Almost all of them evaluate the barrier and its Lie derivatives through a Jacobian linearisation around the mean, or

© *Journal of Cyber-Physical Security and Robotics*

require resampled particles whose weights degenerate in higher dimension.

The ensemble transform Kalman filter (ETKF) [13, 14] sits between these options. It carries an ensemble of N equally weighted members whose sample mean and covariance reproduce the Kalman analysis exactly within the ensemble subspace, uses no Jacobians, and performs its update with an $N \times N$ transform whose cost does not depend on the state dimension. It was designed in numerical weather prediction precisely to quantify how observations change error variance [13, 15], yet, to our knowledge, it has not been used to drive a CBF safety filter.

This paper examines what changes when the belief inside a chance-constrained CBF is an ETKF ensemble rather than an EKF mean and covariance. The result is a filter that we call the Ensemble-Transform Control Barrier Function (ET-CBF). Our contributions are as follows.

1. **ET-CBF formulation.** We evaluate the discrete-time barrier condition on every analysis member and enforce a one-sided chance constraint on its ensemble distribution. For control-affine systems this is a single second-order cone (SOC) constraint whose data are N barrier values and N input gains, giving a convex SOCP solved in about 1 ms (Proposition 1).
2. **Exactness and consistency results.** For affine barriers under linear-Gaussian models, ET-CBF reduces exactly to the Kalman-filter chance-constrained CBF (Proposition 2). For quadratic barriers, the ensemble mean of the barrier captures the Jensen (curvature) term that EKF-linearised constraints drop (Proposition 3).
3. **Probabilistic safety certificate.** Under ensemble consistency, a Cantelli or Gaussian quantile gain bounds the one-step probability of violating the true-state barrier condition by ε and the finite-horizon probability of leaving the safe set by $K\varepsilon$ (Theorem 1).
4. **Finite-ensemble correction.** We show that the nominal Gaussian quantile is miscalibrated for small ensembles and give a Student-t predictive correction that depends only on N. In closed loop it lowers the one-step failure rate at $N = 4$ from 14.7% to 5.1% against a 5% target (Section 6.3).
5. **A head-to-head study.** We compare ET-CBF with certainty-equivalent EKF and ETKF filters, a linearised EKF chance-constrained CBF and a true-state reference over 5,600 Monte Carlo closed-loop runs, plus a static calibration test against 200,000-sample ground truth.

Section 2 positions the work, Section 3 recalls discrete-time CBFs and the ETKF, Section 4 develops ET-CBF and its theory, Section 5 sets out similarities and differences between the ETKF and CBF routes, Section 6 reports the numerical study, and Section 7 discusses limits and open problems.

2. Related work and positioning

Existing CBF-under-estimation methods either bound the error deterministically or describe it with a Kalman-type or particle belief; none builds the barrier constraint from a deterministic square-root ensemble.

Bounded-error CBFs. Measurement-robust CBFs add a margin proportional to a known bound on the estimation or perception error [4, 5]. They give deterministic guarantees, but the margin is worst-case and must be known.

Kalman-filter stochastic CBFs. Clark introduced stochastic CBFs for complete and incomplete information, using an EKF estimate and bounding the safety probability through the filter's asymptotic error [6, 7]. Yaghoubi et al. derived linear sufficient conditions on the input from EKF confidence intervals [8]. Kishida combined a Kalman filter with a worst-case CVaR to obtain risk-aware discrete-time CBFs [10], and Zhao et al. give an output-feedback, Jensen-based reformulation with finite-horizon chance guarantees [16].

Belief-space CBFs. Vahs, Pek and Tumova defined belief CBFs on the EKF mean–covariance state, enforcing probabilistic constraints without fixed error bounds [9]. Later work moved to non-Gaussian particle-filter beliefs with CVaR [11] and to sample-based beliefs over obstacle states with concentration bounds on tail risk [12].

Discrete-time stochastic CBFs. Cosner and co-authors analysed discrete-time CBFs under unbounded noise with martingale bounds, highlighted the Jensen gap in expectation conditions [17], tightened the bounds with Freedman's inequality [18], and introduced probabilistic CBFs with concentration, scenario and conformal certificates [19].

SOCP chance constraints. Gaussian-process CBFs turn uncertainty that is affine in the input into a second-order cone constraint [20, 21]; distributionally robust,

Table 1: Positioning of ET-CBF among CBF methods under state uncertainty.

Approach	Belief used	Jacobians needed	Barrier statistics	Constraint size
Certainty-equivalent CBF [1]	Point estimate	Barrier gradient	h at the mean	1 linear row
Measurement-robust CBF [4, 5]	Error bound	Yes	Worst-case margin	1 linear row
EKF stochastic / risk-aware CBF [7–10]	Mean + covariance	Dynamics, measurement, barrier	First-order (linearised)	Linear row or SOC of size $n+1$
Particle belief CBF [11, 12]	Weighted particles	Barrier gradient	Sample CVaR	Grows with particle count
ET-CBF (this paper)	**ETKF ensemble**	**None**	**Sample mean and variance; Jensen term kept**	**SOC of size N + 2 (reducible to m + 2)**

sample-based CBF and CLF constraints also lead to SOCPs [22]. ET-CBF inherits this cone structure but obtains the uncertainty from an estimator ensemble rather than a learned model.

Ensemble Kalman methods. The ETKF [13] and its local form, the LETKF [14], are square-root filters [23] in which analysis perturbations are the forecast perturbations multiplied by a symmetric transform; the symmetric root keeps the ensemble mean unbiased [24]. We found no CBF safety filter in the literature, as of October 2026, that is driven by an ETKF or another ensemble Kalman filter. Table 1 summarises the positioning.

3. Preliminaries

3.1. System, safe set and discrete-time CBFs

We consider sampled control-affine dynamics with additive Gaussian process noise and nonlinear measurements,

$$\begin{aligned} x_{k+1} &= F(x_k) + G(x_k)u_k + w_k, \quad w_k \sim \mathcal{N}(0, Q), \\ y_k &= H(x_k) + v_k, \qquad\quad v_k \sim \mathcal{N}(0, R), \end{aligned} \tag{1}$$

with $u_k \in \mathcal{U} = [u_{\min}, u_{\max}]$. The safe set is $\mathcal{C} = \{x : h(x) \geq 0\}$ for a differentiable barrier h. A discrete-time CBF (DT-CBF) [3, 17] with rate $\gamma \in (0,1]$ requires

$$h(x_{k+1}) \geq (1-\gamma)\, h(x_k). \tag{2}$$

If (2) holds at every step and $h(x_0) \geq 0$, then $x_k \in \mathcal{C}$ for all k. The condition is generally nonlinear in u; we use a tangent form that is an exact sufficient condition when h is convex.

Lemma 1 (Affine-in-input sufficient condition). *Define*

$$\psi(x, u, w) = \gamma h(x) + \nabla h(x)^{\top}(F(x) - x + G(x)u + w). \tag{3}$$

If h is convex, then $\psi(x_k, u_k, w_k) \geq 0$ implies $h(x_{k+1}) \geq (1-\gamma)h(x_k)$.

Proof. Convexity gives $h(x_{k+1}) \geq h(x_k) + \nabla h(x_k)^{\top}(x_{k+1} - x_k)$, and substituting the dynamics yields the result. □

For non-convex h, $\psi \geq 0$ is the usual first-order (Euler) approximation of (2), with an $O(\| \Delta x \|^2)$ residual. Distance barriers $h = \| p - o \|^2 - r^2$, written in the coordinates of the constrained point p, are convex. For each fixed x, ψ is affine in both u and w, which is what makes the chance constraint of Section 4 a cone.

3.2. The ensemble transform Kalman filter

We use the ETKF [13] in the symmetric square-root form of Hunt et al. [14]. Let the forecast members be x_i^f, $i = 1, \ldots, N$, with mean $\bar{x}^f$ and anomaly matrix $X^f = \left[x_1^f - \bar{x}^f, \ldots, x_N^f - \bar{x}^f\right] \in \mathbb{R}^{n \times N}$. Map each member through H, take the mean $\bar{y}^f$ and the observation anomalies $Y^f \in \mathbb{R}^{p \times N}$. With multiplicative inflation $\rho \geq 1$,

$$\tilde{P} = \left[\tfrac{N-1}{\rho} I_N + (Y^f)^{\top} R^{-1} Y^f\right]^{-1}, \quad \bar{w} = \tilde{P}\,(Y^f)^{\top} R^{-1}(y - \bar{y}^f), \tag{4}$$

$$W = [(N-1)\tilde{P}]^{1/2} \text{ (symmetric)}, \quad x_i^a = \bar{x}^f + X^f(\bar{w} + W e_i). \tag{5}$$

Three properties of this update are relevant here. (i) All linear algebra is $N \times N$; the state enters only through X^f, so the update costs $O(N^3 + nN^2 + pN^2)$ and needs no Jacobian of F or H. (ii) For linear H, the analysis mean $\bar{x}^f + X^f \bar{w}$ and sample covariance $X^f \tilde{P} (X^f)^{\top}$ equal the Kalman analysis restricted to the ensemble subspace [13, 14, 23]. (iii) The symmetric root makes the analysis anomalies sum to zero, so the ensemble mean is the analysis mean [14, 24]. The forecast step propagates each member through the full nonlinear model with sampled process noise.

4. Ensemble-transform control barrier functions

ET-CBF evaluates the barrier condition on each ETKF analysis member and accepts an input only if the ensemble estimate of the probability that the condition fails is at most ε.

4.1. Member-wise barrier data

At step k, after the analysis, compute for every member $x_i = x_i^a$

$$\begin{aligned} a_i &= \gamma h(x_i) + \nabla h(x_i)^\top (F(x_i) - x_i), \\ b_i &= G(x_i)^\top \nabla h(x_i) \in \mathbb{R}^m, \qquad q_i = \nabla h(x_i)^\top Q\, \nabla h(x_i). \end{aligned} \tag{6}$$

Then $\psi(x_i, u, w) = a_i + b_i^\top u + \nabla h(x_i)^\top w$. Let $\bar{a}, \bar{b}, \bar{q}$ be ensemble means and $a' \in \mathbb{R}^N$, $B' \in \mathbb{R}^{N\times m}$ the anomalies. By the law of total variance (state uncertainty plus process noise), the ensemble estimates of the mean and variance of ψ under input u are

$$\hat{\mu}(u) = \bar{a} + \bar{b}^\top u, \qquad \hat{\sigma}^2(u) = \tfrac{1}{N-1} \parallel a' + B'u \parallel_2^2 + \bar{q}. \tag{7}$$

4.2. The ET-CBF program

Given a nominal input u_{nom}, a weight $M \succ 0$, a risk level ε and a gain $\kappa(\varepsilon, N)$, ET-CBF solves

$$\begin{aligned} \min_{u\in\mathcal{U},\, s\geq 0} \quad & (u - u_{nom})^\top M (u - u_{nom}) + \lambda s \\ \text{s.t.} \quad & \parallel \begin{bmatrix} (a' + B'u)/\sqrt{N-1} \\ \sqrt{\bar{q}} \end{bmatrix} \parallel_2 \leq \tfrac{\bar{a} + \bar{b}^\top u + s}{\kappa}. \end{aligned} \tag{8}$$

The slack s, with a large penalty λ, keeps the program feasible; any step with $s > 0$ is logged as uncertified.

Proposition 1 (Convexity and size). *Problem (8) is a second-order cone program. Its cone has dimension $N + 2$, independent of the state dimension n, and its data require N evaluations of h, ∇h, F and G and no Jacobian of F, G or H. A thin QR factorisation of $[a'\ B']$ reduces the cone to dimension $m + 2$ at $O(Nm^2)$ cost.*

Proof. $\hat{\mu}$ is affine in u and $\hat{\sigma}$ is the Euclidean norm of an affine map of u, so the constraint is a standard SOC; the remaining statements follow by counting the problem data. □

Algorithm (one control step).

1. Forecast each member through the nonlinear model with sampled process noise.
2. If a measurement arrives, apply the ETKF update.
3. Compute a_i, b_i, q_i on the analysis members.
4. Solve (8); apply u_k and record s_k.

4.3. Structural results

Proposition 2 (Linear-Gaussian equivalence). *Let $F(x) = Ax$, G constant, H linear, the noises Gaussian and $h(x) = c^\top x + d$. If the analysis ensemble has full rank ($N \geq n + 1$) and $\rho = 1$, ET-CBF is identical to the Kalman-filter chance-constrained CBF with mean $\mu(u) = \ell^\top \hat{x} + \gamma d + c^\top G u$ and variance $\ell^\top P \ell + c^\top Q c$, where $\ell = (\gamma I + A - I)^\top c$ and $(\hat{x}, P)$ is the Kalman analysis.*

Proof. $b_i = G^\top c$ is constant, so $B' = 0$; $a_i = \ell^\top x_i + \gamma d$ is affine in x_i, so $\bar{a} = \ell^\top \bar{x}^a$ and $\parallel a' \parallel^2/(N-1) = \ell^\top \hat{P} \ell$; by property (ii), $\bar{x}^a = \hat{x}$ and $\hat{P} = P$. □

ET-CBF therefore matches the Kalman-filter risk-aware CBFs of [9, 10] in the setting where those are exact.

Proposition 3 (The Jensen term). *Let $h(x) = (x - o)^\top \Lambda (x - o) - r^2$ with symmetric Λ, and $F(x) = x$. Then*

$$\bar{a} = \gamma \left[h(\bar{x}^a) + \tfrac{N-1}{N} tr\big(\Lambda \hat{P}\big) \right], \tag{9}$$

whereas a constraint linearised at the mean uses $\gamma h(\bar{x}^a)$ and omits $\gamma tr\big(\Lambda \hat{P}\big)$. For a keep-out barrier ($\Lambda \succcurlyeq 0$) the omitted term is non-negative, so linearisation is conservative; for a keep-in barrier ($\Lambda \preccurlyeq 0$) it is non-positive, so linearisation is optimistic and can certify unsafe inputs. Linearisation also drops the variance contribution $2\gamma^2 tr((\Lambda P)^2)$, which the ensemble retains.

Proof. Expand $h(\bar{x} + \delta_i)$, average over members, and use $\sum_i \delta_i = 0$. □

This is the estimation-side counterpart of the Jensen gap that Cosner et al. identified for process noise [17].

4.4. Probabilistic safety

Assumption 1 (Ensemble consistency). Conditioned on the measurement history Y_k, the true mean and variance of $\psi(x_k, u_k, w_k)$ equal $\hat{\mu}(u_k)$ and $\hat{\sigma}^2(u_k)$.

Theorem 1. *Let h be convex, $h(x_0) \geq 0$ and Assumption 1 hold. Choose $\kappa = \sqrt{(1-\varepsilon)/\varepsilon}$ (Cantelli), or $\kappa = \Phi^{-1}(1-\varepsilon)$ if ψ is also conditionally Gaussian. If (8) is solved with $s_k = 0$ for $k = 0, \dots, K-1$, then*

$$\begin{aligned} \mathbb{P}[h(x_{k+1}) < (1-\gamma) h(x_k) \mid Y_k] \leq \varepsilon, \\ \mathbb{P}[\exists\, k \leq K\colon\ x_k \notin \mathcal{C}] \leq K\varepsilon. \end{aligned} \tag{10}$$

Proof. u_k is a function of Y_k. With $s_k = 0$ the constraint gives $\hat{\mu} \geq \kappa \hat{\sigma}$, so Cantelli's inequality [25] (or the Gaussian quantile) bounds $\mathbb{P}[\psi < 0 \mid Y_k]$ by ε. By Lemma 1 the failure event $E_k = \{h(x_{k+1}) < (1-\gamma) h(x_k)\}$ is contained in $\{\psi < 0\}$. Taking expectations, $\mathbb{P}[E_k] \leq \varepsilon$. Outside $\bigcup_{k<K} E_k$, induction from $h(x_0) \geq 0$ keeps every $h(x_k) \geq 0$; the union bound finishes the proof. □

The $K\varepsilon$ bound is deliberately simple and loose; the martingale and Freedman-type bounds of [17–19] can be substituted once the one-step condition is in place.

4.5. Finite-ensemble correction

Assumption 1 fails most visibly when N is small, because $\hat{\mu}$ and $\hat{\sigma}$ are themselves estimated from N members. Treating the members as a sample from the predictive distribution of ψ, a new draw satisfies a Student-t predictive law [26], which gives

$$\kappa_N(\varepsilon) = t_{N-1,\,1-\varepsilon} \sqrt{1 + 1/N}. \tag{11}$$

Table 2: Per-step cost of the two estimator–safety-filter pipelines.

Component	EKF + linearised chance CBF	ETKF + ET-CBF
Filter forecast	$O(n^3)$ covariance propagation; Jacobian of F	N model runs; no Jacobian
Filter update	$O(p^3 + n^2 p)$; Jacobian of H	$O(N^3 + (n+p)N^2)$; no Jacobian
Barrier statistics	Gradient of ψ in x (second derivatives of h, Jacobians of F, G)	N evaluations of h, ∇h, F, G
Safety program	SOC built from $P \in \mathbb{R}^{n\times n}$	SOC built from N members
Memory	$O(n^2)$	$O(nN)$

For $\varepsilon = 0.05$ this gives $\kappa = 2.63$ at $N = 4$, 2.18 at $N = 6$, 1.92 at $N = 10$, 1.77 at $N = 20$ and 1.71 at $N = 40$, against 1.645 for the Gaussian quantile. ETKF members are not independent draws (the transform is deterministic), so this correction is a heuristic; Section 6.3 tests it in closed loop. A distribution-free alternative is the Chebyshev inequality with estimated mean and variance [27].

4.6. Computational cost per step

For $n \gg N$ the ensemble route avoids every $n \times n$ object; for small n both are cheap and the choice should rest on calibration (Section 6). Table 2 compares the two pipelines.

5. Similarities and differences between the CBF and ETKF routes

CBFs and the ETKF solve structurally similar problems: both are one-step, minimum-change corrections built from a quadratic cost, and pairing them puts the estimator's uncertainty directly into the safety constraint.

Shared structure.

- *Both are minimum-norm projections.* The CBF-QP finds the input closest to u_{nom} that satisfies a linear constraint; the ETKF analysis finds the ensemble-space weight closest to the forecast that fits the observation, weighted by $\tilde{P}$.
- *Both are recursive and myopic.* Each acts at one sampling instant, which keeps them cheap enough to run at the control rate.
- *Both rely on second-order statistics.* The Kalman update uses means and covariances; the chance-constrained CBF needs the mean and variance of the barrier condition. ET-CBF supplies the latter directly from the former.
- *Both have a robustness parameter.* Inflation ρ in the ETKF and the gain κ in the CBF both widen an uncertainty set, and they interact: inflation enlarges $\hat{\sigma}$ and hence the safety margin.

What the combination adds over an EKF-based chance CBF. Linearised belief CBFs [9, 10] and ET-CBF agree exactly in the linear-Gaussian, affine-barrier case (Proposition 2), so the differences lie in the nonlinear regime and in scale (Table 3): (1) *Curvature*: ET-CBF keeps the Jensen term of the barrier mean and the second-order part of its variance (Proposition 3). (2) *Input-dependent uncertainty*: because b_i varies across members, the variance of the barrier condition depends on u; ET-CBF captures this through B' without differentiating G. (3) *No derivatives*: black-box simulators can supply the model. (4) *Scale*: every object is N-dimensional. (5) *A new failure mode*: small ensembles under-represent spread, which requires the correction κ_N; this issue does not arise for the EKF.

What it shares with particle-filter belief CBFs. Both are sample-based and Jacobian-free [11, 12]. ET-CBF uses equal weights produced by a deterministic transform rather than resampling, so it avoids weight degeneracy, and it uses a Gaussian-moment chance constraint rather than an empirical CVaR. In exchange, it cannot represent strongly multimodal beliefs.

6. Numerical study

Both chance-constrained filters cut the share of runs that hit the obstacle from about 70–80% under certainty equivalence to 0–4.5%, with one-step failure rates of 2.8–3.6% against a 5% design level. In this three-state benchmark, ET-CBF and the linearised EKF constraint are statistically close in closed loop; ET-CBF's calibration advantage appears where barrier curvature is large (Section 6.4), and its small-ensemble miscalibration is corrected by κ_N (Section 6.3).

6.1. Setup

The plant is a unicycle written in look-ahead-point coordinates $x = (q_1, q_2, \theta)$, with $\dot{q} = R(\theta)diag(1, \ell)\,u$ and $\dot{\theta} = \omega$, Euler-discretised at $\Delta t = 0.1$ s with process noise. The robot drives from $(0,0)$ to $(10,0)$ m past a disc of radius 1.0 m centred at $(5.0,0.6)$ m, so the straight path cuts 0.4 m into the obstacle. The barrier $h = \| q - o \|^2 - r^2$ is convex in x, so Lemma 1 is exact.

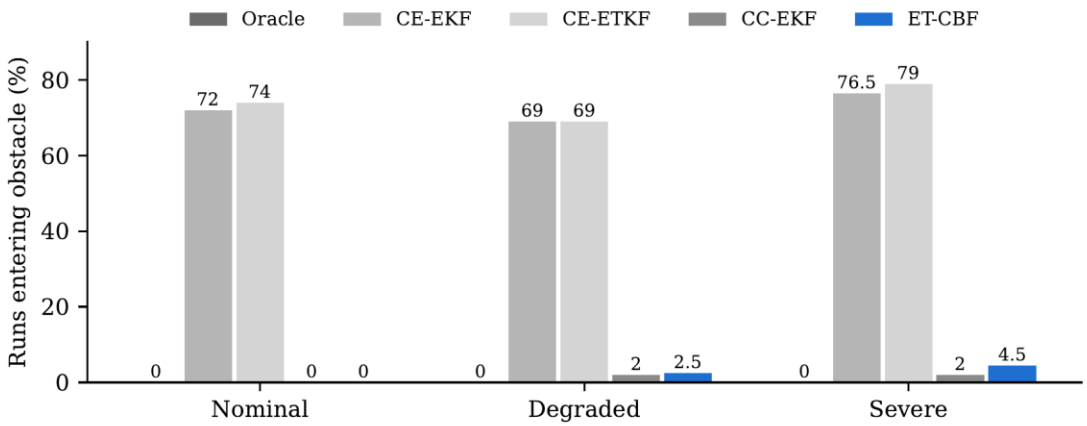


Figure 1: Share of 200 closed-loop runs in which the true state entered the obstacle, by method and sensing scenario ($\varepsilon = 0.05$, Gaussian gain).

Table 3: Key differences between the CBF and ETKF components.

Aspect	Control barrier function	ETKF
Object it acts on	Input u	Belief over state x
Quantity it controls	Invariance of the safe set $\mathcal{C}$	Mean and covariance of the estimate
Constraint type	Inequality on the barrier rate	Observation fit, soft via R
Uncertainty it sees	None in the classical form	Explicit, through the ensemble
Nonlinearity handling	Lie derivatives at a point	Member-wise model runs, no Jacobians
Typical failure	Unsafe under estimation error	Under-dispersion with small N
Cost driver	Number of inputs m	Ensemble size N

Table 4: Sensing scenarios (1σ values).

Scenario	Landmarks	Range / bearing noise	Heading noise/step	Initial heading
Nominal	3	0.35 m / 0.06 rad	0.02 rad	20
Degraded	2, far from path	1.0 m / 0.12 rad	0.04 rad	30
Severe	2, near path	3.0 m / 0.15 rad	0.06 rad	45

A proportional go-to-goal law computed on the estimate gives u_{nom}; all filters minimise $\| R(\theta)S(u - u_{nom}) \|^2$, the deviation in point velocity. Measurements are ranges and bearings to fixed landmarks. Remaining parameters: $\ell = 0.2$ m, $\gamma = 0.25$, $v \in [-0.5, 1.5]$ m/s, $\omega \in [-5, 5]$ rad/s, $\varepsilon = 0.05$, $N = 20$, $\rho = 1.05$, horizon 160 steps, 200 runs per configuration with common random numbers across methods. Table 4 lists the three sensing scenarios.

Five filters are compared, all with the Gaussian gain $\kappa = 1.645$ unless stated: **Oracle** (true state, chance constraint on process noise only; a non-implementable reference); **CE-EKF** and **CE-ETKF** (deterministic DT-CBF on the EKF mean or ETKF ensemble mean); **CC-EKF** (linearised Gaussian chance constraint from the EKF mean and covariance, representing [8–10]); and **ET-CBF** (proposed, 20 members). We report the share of runs in which the true state ever left $\mathcal{C}$, and the *one-step failure rate*: the fraction of steps on which the filter was active and the true-state condition (2) nevertheless failed. Theorem 1 predicts the latter is at most ε under Assumption 1.

6.2. Closed-loop safety and performance

Three findings stand out (Fig. 1, Table 5). First, certainty equivalence fails regardless of estimator: the ETKF mean is no safer than the EKF mean, so the benefit comes from using the spread, not from a better point estimate. Second, both chance-constrained filters hold one-step failure below the 5% design level in every scenario, matching the true-state oracle. Third, ET-CBF and CC-EKF are within run-to-run noise of each other: with 200 runs, the severe-scenario gap (4.5% vs 2.0% of runs) is about 1.4 combined standard errors (± 1.5 and ± 1.0 points each), so it is not significant. The price of safety is about 0.5–1.0 s of extra

Table 5: Closed-loop results, 200 runs per row. Clearance is the closest true distance to the obstacle edge averaged over runs (negative = inside); arrival time is averaged over runs reaching within 0.3 m of the goal in 16 s.

Scenario	Method	One-step failure (%)	Min. clearance (m)	Arrival time (s)	Reached goal (%)	Uncertified steps/run	Solve (ms)
Nominal	Oracle	3.8	+0.176	11.08	100	0.00	0.67
	CE-EKF	46.4	−0.044	10.76	100	0.00	0.66
	CE-ETKF	45.6	−0.043	10.77	100	0.00	0.63
	CC-EKF	3.2	+0.193	11.24	100	0.01	0.83
	ET-CBF	3.6	+0.190	11.23	100	0.04	1.30
Degraded	Oracle	4.2	+0.175	11.12	100	0.00	0.68
	CE-EKF	45.0	−0.065	11.15	100	0.00	0.74
	CE-ETKF	44.4	−0.077	11.14	100	0.00	0.60
	CC-EKF	2.8	+0.237	11.80	98	0.00	0.76
	ET-CBF	3.0	+0.236	11.90	100	0.36	1.13
Severe	Oracle	4.0	+0.227	11.34	100	0.00	0.65
	CE-EKF	46.2	−0.126	11.75	82	0.00	0.59
	CE-ETKF	44.4	−0.142	11.89	82	0.00	0.60
	CC-EKF	3.0	+0.259	12.47	80	0.07	0.75
	ET-CBF	3.2	+0.244	12.76	78	0.19	1.16

travel time over certainty equivalence.

With the distribution-free Cantelli gain ($\kappa = 4.36$) in the degraded scenario, no run of the oracle, CC-EKF or ET-CBF violated the safe set and no one-step failure occurred, but arrival time rose to 12.0, 13.9 and 14.2 s, and only 96% (CC-EKF) and 88% (ET-CBF) of runs reached the goal in 16 s. Cantelli is the choice when a guarantee must hold without a Gaussian assumption; the Gaussian gain is the practical default.

6.3. Ensemble size and the finite-ensemble correction

With the plain Gaussian gain, small ensembles are miscalibrated (Fig. 2): at $N = 4$, 14.7% of active steps failed and 30% of runs touched the obstacle. The Student-t gain κ_N brought one-step failure to 5.1% at $N = 4$ and 3.3% at $N = 6$, and run violations to 14.0% and 6.5%. From $N = 20$ the two gains differ by less than one point, so κ_N matters chiefly when ensemble size is constrained, as it is in high-dimensional applications.

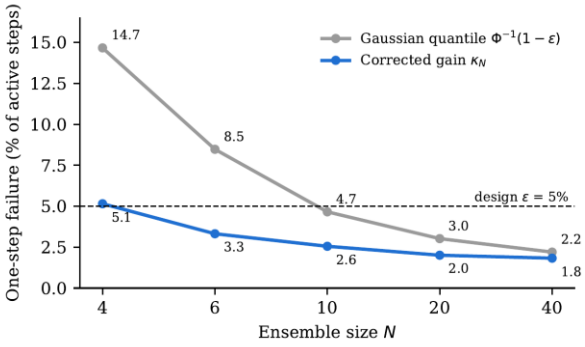


Figure 2: One-step failure rate of ET-CBF versus ensemble size (degraded scenario, 200 runs per point).

6.4. Static calibration against Monte Carlo ground truth

To isolate the effect of curvature, we fixed a Gaussian belief, found for each method the forward speed at which its chance constraint is exactly tight, and measured the true probability that $\psi < 0$ at that speed with 200,000 samples. Ensemble results average 40 independent ensembles (no κ_N correction). The target is 5% (Table 6).

The pattern follows Proposition 3. For the convex keep-out barrier, linearisation is conservative and becomes more so as heading uncertainty grows; the ensemble sits closer to the target. For the concave keep-in barrier ($h = r^2 - \| q - c \|^2$, $r = 2$ m, robot 0.5 m from the wall and heading outward), linearisation is

Table 6: Realised risk $\mathbb{P}[\psi < 0]$ when each method's constraint is tight (target 5%).

Barrier and uncertainty	Linearised (EKF-style)	ET, $N = 10$	ET, $N = 20$	ET, $N = 200$
Keep-out, heading 1σ 5–45 (12 cases)	2.4–3.2%	3.5–5.0%	3.7–4.6%	3.6–4.0%
Keep-in, position 1σ 0.1 m	5.1%	6.6%	5.8%	5.5%
Keep-in, position 1σ 0.2 m	6.8%	7.1%	5.9%	6.0%
Keep-in, position 1σ 0.3 m	8.5%	8.2%	6.6%	6.5%
Keep-in, position 1σ 0.4 m	10.5%	11.5%	10.1%	7.0%

optimistic and its realised risk doubles to 10.5% as position uncertainty grows, while a large ensemble stays within 2 points of the target. Small uncorrected ensembles inherit sampling error, which is the effect Section 6.3 corrects.

7. Discussion, limitations and conclusion

The results support a specific claim: ET-CBF delivers calibrated chance-constrained safety from an ETKF, matches the EKF chance constraint where linearisation is adequate, and is better calibrated where barrier curvature matters, provided the ensemble-size correction is used.

What is novel. To our knowledge this is the first CBF safety filter built on an ensemble transform Kalman filter. The new elements are the member-wise SOC formulation whose size is set by N rather than n (Proposition 1), the exact equivalence to Kalman-filter chance CBFs in the linear-Gaussian case (Proposition 2), the sign rule for the Jensen term that tells when linearised belief CBFs are optimistic (Proposition 3), and the finite-ensemble gain κ_N, which addresses a failure mode specific to ensemble beliefs.

What it is not. ET-CBF is not a cure for non-Gaussian beliefs; it uses two moments and inherits the ETKF's Gaussian update. Where beliefs are multimodal, particle-based belief CBFs [11, 12] are the better tool. In low-dimensional, well-observed problems, the EKF chance constraint performs about as well and is slightly cheaper.

Limitations.

- Theorem 1 rests on Assumption 1. Inflation and κ_N make it hold approximately in our tests, but it is not proven for nonlinear models.
- The $K\varepsilon$ horizon bound is loose; martingale bounds [17–19] would give a tighter horizon bound.
- Lemma 1 is exact only for convex barriers; keep-in barriers need the first-order approximation or a curvature margin.
- The study uses one robotic benchmark with $n = 3$, so it shows calibration effects rather than the scale advantage; a high-dimensional benchmark is the natural next experiment.
- The certificate lapses on steps where the slack is active; their frequency is reported but not bounded.

Open directions. (i) Exploit the ETKF's original purpose, predicting how a planned observation will shrink variance [13, 15], to anticipate post-measurement barrier spread inside a horizon. (ii) Localise the constraint, as the LETKF localises the update [14], for spatially distributed safety constraints. (iii) Learn ρ and κ jointly from closed-loop calibration data.

Conclusion. Driving a CBF with an ETKF ensemble replaces Jacobian-based uncertainty propagation with member-wise evaluation, turns the safety filter into a small SOCP whose size depends on the ensemble rather than the state, keeps second-order barrier effects, and, with a simple ensemble-size correction, yields one-step violation rates close to the designed risk level.

Reproducibility. All results come from a single Python implementation (NumPy, SciPy, Clarabel [28]) of the plant, EKF, ETKF (symmetric square root, $\rho = 1.05$), the five safety filters and the Monte Carlo drivers, with common random numbers across methods.

CRediT authorship contribution statement

Toheeb Babatunde Ibrahim: Conceptualization, Methodology, Software, Writing – original draft. **Olasubomi Akanbi:** Formal analysis, Validation, Writing – review & editing. **Joseph Oladejo:** Visualization. **Warith Akinade:** Software.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported

in this paper.

Data availability

The simulation code and all results are available from the corresponding author on reasonable request.

Declaration of generative AI and AI-assisted technologies

During the preparation of this work the authors used Claude (Anthropic) to assist with literature search and to check the text for grammatical errors and mistakes. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

References

[1] A. D. Ames, X. Xu, J. W. Grizzle, P. Tabuada, Control barrier function based quadratic programs for safety critical systems, IEEE Transactions on Automatic Control 62 (8) (2017) 3861–3876.

[2] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, P. Tabuada, Control barrier functions: Theory and applications, in: Proceedings of the 18th European Control Conference (ECC), 2019, pp. 3420–3431.

[3] A. Agrawal, K. Sreenath, Discrete control barrier functions for safety-critical control of discrete systems with application to bipedal robot navigation, in: Robotics: Science and Systems XIII, 2017.

[4] S. Dean, A. J. Taylor, R. K. Cosner, B. Recht, A. D. Ames, Guaranteeing safety of learned perception modules via measurement-robust control barrier functions, in: Proceedings of the 2020 Conference on Robot Learning (CoRL), PMLR 155, 2021, pp. 654–670.

[5] R. K. Cosner, A. W. Singletary, A. J. Taylor, T. G. Molnar, K. L. Bouman, A. D. Ames, Measurement-robust control barrier functions: Certainty in safety with uncertainty in state, in: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2021, pp. 6286–6291.

[6] A. Clark, Control barrier functions for complete and incomplete information stochastic systems, in: Proceedings of the American Control Conference (ACC), 2019, pp. 2928–2935.

[7] A. Clark, Control barrier functions for stochastic systems, Automatica 130 (2021) 109688.

[8] S. Yaghoubi, G. Fainekos, T. Yamaguchi, D. Prokhorov, B. Hoxha, Risk-bounded control with Kalman filtering and stochastic barrier functions, in: Proceedings of the 60th IEEE Conference on Decision and Control (CDC), 2021, pp. 5213–5219.

[9] M. Vahs, C. Pek, J. Tumova, Belief control barrier functions for risk-aware control, IEEE Robotics and Automation Letters 8 (12) (2023) 8565–8572.

[10] M. Kishida, Risk-aware control of discrete-time stochastic systems: Integrating Kalman filter and worst-case CVaR in control barrier functions, arXiv:2312.15638 (2023).

[11] M. Vahs, J. Tumova, Risk-aware control for robots with non-Gaussian belief spaces, in: Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), 2024, pp. 11661–11667.

[12] S. Han, M. Vahs, J. Tumova, Risk-aware robot control in dynamic environments using belief control barrier functions, arXiv:2504.04097 (2025).

[13] C. H. Bishop, B. J. Etherton, S. J. Majumdar, Adaptive sampling with the ensemble transform Kalman filter. Part I: Theoretical aspects, Monthly Weather Review 129 (3) (2001) 420–436.

[14] B. R. Hunt, E. J. Kostelich, I. Szunyogh, Efficient data assimilation for spatiotemporal chaos: A local ensemble transform Kalman filter, Physica D: Nonlinear Phenomena 230 (1–2) (2007) 112–126.

[15] S. J. Majumdar, C. H. Bishop, B. J. Etherton, I. Szunyogh, Z. Toth, Can an ensemble transform Kalman filter predict the reduction in forecast-error variance produced by targeted observations?, Quarterly Journal of the Royal Meteorological Society 127 (2001) 2803–2820.

[16] J. Zhao, Z. Cai, X. Yin, Output-feedback safe control of discrete-time stochastic systems with chance constraints, arXiv:2604.12956 (2026).

[17] R. K. Cosner, P. Culbertson, A. J. Taylor, A. D. Ames, Robust safety under stochastic uncertainty with discrete-time control barrier functions, in: Robotics: Science and Systems XIX, 2023.

[18] R. K. Cosner, P. Culbertson, A. D. Ames, Bounding stochastic safety: Leveraging Freedman's inequality with discrete-time control barrier functions, IEEE Control Systems Letters 8 (2024) 1937–1942.

[19] P. Mestres, B. Werner, R. K. Cosner, A. D. Ames, Probabilistic control barrier functions: Safety in probability for discrete-time stochastic systems, arXiv:2510.01501 (2025).

[20] F. Castañeda, J. J. Choi, B. Zhang, C. J. Tomlin, K. Sreenath, Pointwise feasibility of Gaussian process-based safety-critical control under model uncertainty, in: Proceedings of the 60th IEEE Conference on Decision and Control (CDC), 2021, pp. 6762–6769.

[21] M. Aali, J. Liu, Learning high-order control barrier functions for safety-critical control with Gaussian processes, arXiv:2403.09573 (2024).

[22] P. Mestres, K. Long, N. Atanasov, J. Cortés, Feasibility analysis and regularity characterization of distributionally robust safe stabilizing controllers, IEEE Control Systems Letters 8, doi:10.1109/LCSYS.2023.3348688 (2024).

[23] M. K. Tippett, J. L. Anderson, C. H. Bishop, T. M. Hamill, J. S. Whitaker, Ensemble square root filters, Monthly Weather Review 131 (7) (2003) 1485–1490.

[24] P. Sakov, P. R. Oke, Implications of the form of the ensemble transformation in the ensemble square root filters, Monthly Weather Review 136 (3) (2008) 1042–1053.

[25] G. Grimmett, D. Stirzaker, Probability and Random Processes, 4th Edition, Oxford University Press, Oxford, 2020.

[26] S. Geisser, Predictive Inference: An Introduction, Chapman & Hall, New York, 1993.

[27] J. G. Saw, M. C. K. Yang, T. C. Mo, Chebyshev inequality with estimated mean and variance, The American Statistician 38 (2) (1984) 130–132.

[28] P. J. Goulart, Y. Chen, Clarabel: An interior-point solver for conic programs with quadratic objectives, arXiv:2405.12762 (2024).