

Safe Perception-Based Control under Stochastic Sensor Uncertainty using Conformal Prediction

Shuo Yang, George J. Pappas, Rahul Mangharam, and Lars Lindemann

Abstract—We consider perception-based control using state estimates that are obtained from high-dimensional sensor measurements via learning-enabled perception maps. However, these perception maps are not perfect and result in state estimation errors that can lead to unsafe system behavior. Stochastic sensor noise can make matters worse and result in estimation errors that follow unknown distributions. We propose a perception-based control framework that i) quantifies estimation uncertainty of perception maps, and ii) integrates these uncertainty representations into the control design. To do so, we use conformal prediction to compute valid state estimation regions, which are sets that contain the unknown state with high probability. We then devise a sampled-data controller for continuous-time systems based on the notion of measurement robust control barrier functions. Our controller uses idea from self-triggered control and enables us to avoid using stochastic calculus. Our framework is agnostic to the choice of the perception map, independent of the noise distribution, and to the best of our knowledge the first to provide probabilistic safety guarantees in such a setting. We demonstrate the effectiveness of our proposed perception-based controller for a LiDAR-enabled F1/10th car.

I. INTRODUCTION

Perception-based control has received much attention lately [1]–[4]. System states are usually not directly observable and can only be estimated from complex and noisy sensors, e.g., cameras or LiDAR. Learning-enabled perception maps can be utilized to estimate the system’s state from such high-dimensional measurements. However, these estimates are usually imperfect and may lead to estimation errors, which are detrimental to the system safety.

The above observation calls for perception-based control with safety guarantees as it is crucial for many autonomous and robotic systems like self-driving cars. Recent work has been devoted to addressing these safety concerns while applying perception-based control using perception maps, see, e.g., [3], [5]–[8]. These work, however, either assume simple or no sensor noise models, consider specific perception maps, or lack end-to-end safety guarantees. In realistic settings, stochastic sensor noise may be unknown and follow skewed and complex distributions that do not resemble a Gaussian distribution, as is often assumed. Additionally, perception maps can be complex, e.g., deep neural networks, making it difficult to quantify estimation uncertainty.

S. Yang, G. J. Pappas, and R. Mangharam are with the Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104, USA. Email: {yangs1, pappasg, rahulm}@seas.upenn.edu

L. Lindemann is with the Department of Computer Science, University of Southern California, Los Angeles, CA 90089, USA. Email: llindema@usc.edu

In this paper, we study perception-based control under stochastic sensor noise that follows arbitrary and unknown distributions. To provide rigorous safety guarantees, we have to account for estimation uncertainty caused by i) imperfect learning-enabled perception maps, and ii) noisy sensor measurements. As shown in Figure 1, to perform safety-critical control, we first leverage conformal prediction [9], a statistical tool for uncertainty quantification, to obtain state estimation regions that are valid with high probability. We then integrate these uncertain state estimation regions into the control design inspired by the notion of measurement robust control barrier functions from [5]. Specifically, we design a sampled-data controller using idea from self-triggered control to ensure safety for continuous-time systems while avoiding the use of stochastic calculus.

To summarize, we make the following contributions:

- We use conformal prediction to quantify state estimation uncertainty of complex learning-enabled perception maps under arbitrary sensor and noise models;
- We use these uncertainty quantifications to design a sampled-data controller for continuous-time systems. We provide probabilistic safety guarantees which, to our knowledge, is the first work to do so in such a setting;
- We demonstrate the effectiveness of our framework in the LiDAR-enabled F1/10th vehicle simulations.

II. RELATED WORK

Perception-based control: Control from high-dimensional sensor measurements such as cameras or LiDAR has lately gained attention. While empirical success has been reported, e.g., [1], [2], [10], [11], there is a need for designing safe-by-construction perception-enabled systems. Resilience of perception-enabled systems to sensor attacks has been studied in [12], [13], while control algorithms that provably generalize to novel environments are proposed in [14], [15]. In another direction, the authors in [16] plan trajectories that actively reduce estimation uncertainty.

Control barrier functions under estimation uncertainty: Control barrier functions (CBFs) have been widely used for autonomous systems since system safety can be guaranteed [17]–[21]. For example, an effective data-driven approach for synthesizing safety controllers for unknown dynamic systems using CBFs is proposed in [22]. Perception maps are first presented in combination with measurement-robust control barrier functions in [5], [8] when true system states are not available but only imprecise measurements. In these works, the perception error is quantified for the specific choice of the Nadarya-Watson regressor. Our approach is




Fig. 1: Overview of the system and robust safe controller. The stochastic sensor noise and imperfect perception module result in state estimation error. Conformal prediction is used to obtain the estimation error upper bound, which is then integrated into the sampled-data safe controller.

agnostic to the perception map and, importantly, allows to consider arbitrary stochastic sensor noise which poses challenges for continuous-time control. Measurement robust control barrier functions are learned in different variations in [6], [23], [24]. Perception maps are further used to design sampled-data controllers [25]–[27] without explicit uncertainty quantification of the sensor and perception maps.

The works in [28], [29] consider state observers, e.g., extended Kalman filters, for barrier function-based control of stochastic systems. On the technical level, our approach is different as we avoid dealing with Itô calculus using sampled-data control. Similarly, bounded state observers were considered in [30], [31]. However, state observer-based approaches are generally difficult to use in perception-systems as models of high-dimensional sensors are difficult to obtain. The authors in [32] address this challenge by combining perception maps and state observers. However, the authors assume a bound on the sensor noise and do not explicitly consider the effect of stochastic noise distributions.

Uncertainty quantification of perception maps is vital. In similar spirit to our paper, [7], [33] use (self-)supervised learning for uncertainty quantification of vision-based systems. While success is empirically demonstrated, no formal guarantees are provided as we pursue in this paper.

Conformal prediction for control: Conformal prediction is a statistical method that provides probabilistic guarantees on prediction errors of machine learning models. It has been applied in computer vision [34], [35], protein design [36], and system verification [37], [38]. Recently, there are works that use conformal prediction for safe planning in dynamic environments, e.g., [39], [40]. However, conformal prediction is only used for quantifying the prediction and not perception uncertainty, as we do in this work. To our knowledge, our work is the first to integrate uncertainty quantification from conformal prediction into perception-based control.

III. PRELIMINARIES AND PROBLEM FORMULATION

We denote by \mathbb{R} , \mathbb{N} , and \mathbb{R}^n the set of real numbers, natural numbers, and real vectors, respectively. Let $\beta: \mathbb{R} \rightarrow \mathbb{R}$ denote an extended class \mathcal{K}_∞ function, i.e., a strictly increasing

function with $\beta(0) = 0$. For a vector $v \in \mathbb{R}^n$, let $\|v\|$ denote its Euclidean norm.

A. System Model

We consider nonlinear control-affine systems of the form

$$\dot{x}(t) = f(x(t)) + g(x(t))u(t) =: F(x(t), u(t)) \quad (1)$$

where $x(t) \in \mathbb{R}^n$ and $u(t) \in \mathcal{U}$ are the state and the control input at time t , respectively, with $\mathcal{U} \subseteq \mathbb{R}^m$ denoting the set of permissible control inputs. The functions $f: \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $g: \mathbb{R}^n \rightarrow \mathbb{R}^{n \times m}$ describe the internal and input dynamics, respectively, and are assumed to be locally Lipschitz continuous. We assume that the dynamics in (1) are bounded, i.e., that there exists an upper bound \bar{F} such that $\|F(x, u)\| \leq \bar{F}$ for every $(x, u) \in \mathbb{R}^n \times \mathcal{U}$. For an initial condition $x(0) \in \mathbb{R}^n$ and a piecewise continuous control law $u: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^m$, we denote the unique solution to the system in (1) as $x: \mathcal{I} \rightarrow \mathbb{R}^n$ where $\mathcal{I} \subseteq \mathbb{R}_{\geq 0}$ is the maximum time interval on which the solution x is defined.

In this paper, we assume that we do not have knowledge of $x(t)$ during testing time, but that we observe potentially high-dimensional measurements $y(t) \in \mathbb{R}^l$ via an unknown locally Lipschitz continuous sensor map $p: \mathbb{R}^n \times \mathbb{R}^d \rightarrow \mathbb{R}^l$ as

$$y(t) = p(x(t), \delta(x(t), t)), \quad (2)$$

where $\delta(x(t), t)$ is a disturbance modeled as a state-dependent random variable that is drawn from an unknown distribution \mathcal{D}_x over \mathbb{R}^d , i.e., $\delta(x, t) \sim \mathcal{D}_x$.¹ A special case that equation (2) covers is those imperfect and noisy sensors that can be modeled as $y(t) = x(t) + \delta(t)$, e.g., as considered in [4], [41]. The function $p(x, \delta(x, t))$ can also encode a simulated image plus noise emulating a real camera. In general, the function p can model high-dimensional sensors such as camera images or LiDAR point clouds. A common assumption in recent work that we adopt implicitly in this paper is that there exists a hypothetical inverse sensor map $q: \mathbb{R}^l \rightarrow \mathbb{R}^n$ that can recover the state x as $q(p(x, 0)) = x$ when

¹To increase readability, we omit time indices when there is no risk of ambiguity, i.e., in this case we mean $\delta(x(t), t) \sim \mathcal{D}_x(t)$.

there is no disturbance [5], [42]. This inverse sensor map q is, however, rarely known and hard to model. One can instead learn perception map $\hat{q} : \mathbb{R}^l \rightarrow \mathbb{R}^n$ that approximately recovers the state x such that $\|q(y, 0) - \hat{q}(y)\|$ is small and bounded, which can then be used for control design [5], [32], [42]. Note that learning an approximation of p is much harder than learning the approximation \hat{q} of q when $l \gg n$.

Remark 1. *The assumption on the existence of an inverse map q is commonly made, as in [5], [32], [42], and realistic when the state x consists of positions and orientations that can, for instance, be recovered from a single camera image. If the state x additionally consists of other quantities such as velocities, one can instead assume that q partially recovers the state as $q(p(x, 0)) = Cx$ for a selector matrix C while using a contracting Kalman filter to estimate the remaining states when the system is detectable [32]. For the sake of simplicity, we leave this consideration for future work.*

Based on this motivation, we assume that we have obtained such a perception map $\hat{q} : \mathbb{R}^l \rightarrow \mathbb{R}^n$ that estimates our state $x(t)$ at time t from measurements $y(t)$, and is denoted as

$$\hat{x}(t) := \hat{q}(y(t)).$$

Note that \hat{q} could be any state estimator, such as a convolutional neural network. In our case study, we used a multi-layer perceptron (MLP) as the estimator.

B. Safe Perception-Based Control Problem

We are interested in designing control inputs u from measurements y that guarantee safety with respect to a continuously differentiable constraint function $h : \mathbb{R}^n \rightarrow \mathbb{R}$, i.e., so that $h(x(t)) \geq 0$ for all $t > 0$ if initially $h(x(0)) \geq 0$. Safety here can be framed as the controlled forward invariance of the system (1) with respect to the safe set $\mathcal{C} := \{x \in \mathbb{R}^n | h(x) \geq 0\}$ which is the superlevel set of the function h . The difficulty in this paper is that we are not able to measure the state $x(t)$ directly during runtime, and that we have only sensor measurements $y(t)$ from the unknown and noisy sensor map p available.

Problem 1. *Consider the system in (1) with initial state $x(0) \in \mathbb{R}^n$ and sensor model in (2). Let $h : \mathbb{R}^n \rightarrow \mathbb{R}$ be a continuously differentiable constraint function, $\mathcal{T} \subset \mathbb{R}_{\geq 0}$ be a time interval, and α be a failure probability. Design a control input u from sensor measurements y such that $\text{Prob}(x(t) \in \mathcal{C}, \forall t \in \mathcal{T}) \geq 1 - \alpha$.*

C. Uncertainty Quantification via Conformal Prediction

In our solution to Problem 1, we use conformal prediction which is a statistical tool introduced in [9], [43] to obtain valid uncertainty regions for complex prediction models without making assumptions on the underlying distribution or the prediction model [44], [45]. Let $Z, Z^{(1)}, \dots, Z^{(k)}$ be $k+1$ independent and identically distributed real-valued random variables, known as the nonconformity scores. Our goal is to obtain an uncertainty region for Z defined via a function $\bar{Z} : \mathbb{R}^k \rightarrow \mathbb{R}$ so that Z is bounded by $\bar{Z}(Z^{(1)}, \dots, Z^{(k)})$ with high probability. Formally, given a failure probability

$\alpha \in (0, 1)$, we want to construct an uncertainty region \bar{Z} such that $\text{Prob}(Z \leq \bar{Z}) \geq 1 - \alpha$ where we omitted the dependence of \bar{Z} on $Z^{(1)}, \dots, Z^{(k)}$ for convenience.

By a surprisingly simple quantile argument, see [46, Lemma 1], the uncertainty region \bar{Z} is obtained as the $(1-\alpha)$ th quantile of the empirical distribution over the values of $Z^{(1)}, \dots, Z^{(k)}$ and ∞ . We recall this result next.

Lemma 1 (Lemma 1 in [46]). *Let $Z, Z^{(1)}, \dots, Z^{(k)}$ be $k+1$ independent and identically distributed real-valued random variables. Without loss of generality, let $Z^{(1)}, \dots, Z^{(k)}$ be sorted in non-decreasing order and define $Z^{(k+1)} := \infty$. For $\alpha \in (0, 1)$, it holds that $\text{Prob}(Z \leq \bar{Z}) \geq 1 - \alpha$ where*

$$\bar{Z} := Z^{(r)} \text{ with } r := \lceil (k+1)(1-\alpha) \rceil$$

and where $\lceil \cdot \rceil$ is the ceiling function.

Some clarifying comments are in order. First, we remark that $\text{Prob}(Z \leq \bar{Z})$ is a marginal probability over the randomness in $Z, Z^{(1)}, \dots, Z^{(k)}$ and not a conditional probability. Second, note that $\lceil (k+1)(1-\alpha) \rceil > k$ implies that $\bar{Z} = \infty$.

IV. SAFE PERCEPTION-BASED CONTROL WITH CONFORMAL PREDICTION

Addressing Problem 1 is challenging for two reasons. First, the perception map \hat{q} may not be exact, e.g., even in the disturbance-free case, it may not hold that $\hat{q}(p(x, 0)) = x$. Second, even if we have accurate state estimates in the disturbance-free case, i.e., when $\hat{q}(p(x, 0))$ is close to x , this does not imply that we have the same estimation accuracy with disturbances, i.e., $\hat{q}(p(x, \delta))$ may not necessarily be close to x . Our setting is thus distinctively different from existing approaches and requires uncertainty quantification of the noisy error between $\hat{x}(t)$ and $x(t)$.

A. Conformal Prediction for Perception Maps

Let us now denote the stochastic state estimation error as

$$e(x, t) := \|\hat{x} - x\| = \|\hat{q}(\underbrace{p(x, \delta(x, t))}_{=y}) - x\|.$$

For a fixed state $x \in \mathbb{R}^n$, our first goal is to construct a prediction region \bar{E}_x so that

$$\text{Prob}(e(x, t) \leq \bar{E}_x) \geq 1 - \alpha \quad (3)$$

holds uniformly over $t \in \mathbb{R}_{\geq 0}$. Note that the distribution \mathcal{D}_x of δ is independent of time t so that we will get uniformity automatically. While we do not know the sensor map p , we assume here that we have an oracle that gives us $N \geq \lceil (N+1)(1-\alpha) \rceil$ state-measurement data pairs $(x, y^{(i)})$ called calibration dataset, where $i \in \{1, \dots, N\}$ and $y^{(i)} = p(x, \delta^{(i)})$ with $\delta^{(i)} \sim \mathcal{D}_x$. This is a common assumption, see, e.g., [5], [32], and such an oracle can, for instance, be a simulator that we can query data from. By defining the nonconformity score $Z^{(i)} := \|\hat{q}(y^{(i)}) - x\|$, and assuming that $Z^{(i)}$ are sorted in non-decreasing order, we can now obtain the guarantees in equation (3) by applying Lemma 1. In other words, we obtain $\bar{E}_x := Z^{(r)}$ with r from Lemma 1 so that $\text{Prob}(\hat{x} \in \{\zeta \in \mathbb{R}^n | \|x - \zeta\| \leq \bar{E}_x\}) \geq 1 - \alpha$

holds. Note that this gives us information about the estimate \hat{x} , but not about the state x which was, in fact, fixed a-priori. To revert this argument and obtain a prediction region for x from \hat{x} , we have to ensure that equation (3) holds for a set of states instead of only a single state x , which will be presented next. To do so, we use a covering argument next.

Consider now a compact subset of the workspace $\mathcal{X} \subseteq \mathbb{R}^n$ that should include the safe set \mathcal{C} . Let $\epsilon > 0$ be a gridding parameter and construct an ϵ -net $\bar{\mathcal{X}}$ of \mathcal{X} , i.e., construct a finite set $\bar{\mathcal{X}}$ so that for each $x \in \mathcal{X}$ there exists an $x_j \in \bar{\mathcal{X}}$ such that $\|x - x_j\| \leq \epsilon$. For this purpose, simple gridding strategies can be used as long as the set \mathcal{X} has a convenient representation. Alternatively, randomized algorithms can be used that sample from \mathcal{X} [47]. We can now again apply a conformal prediction argument for each grid point $x_j \in \bar{\mathcal{X}}$ and then show the following proposition².

Proposition 1. *Consider the Lipschitz continuous sensor map p in (2) and a perception map \hat{q} with respective Lipschitz constants \mathcal{L}_p and $\mathcal{L}_{\hat{q}}$.³ Assume that we constructed an ϵ -net $\bar{\mathcal{X}}$ of \mathcal{X} . For each $x_j \in \bar{\mathcal{X}}$, let $(x_j, y_j^{(i)})$ be $N \geq \lceil (N+1)(1-\alpha) \rceil$ data pairs where $y_j^{(i)} := p(x_j, \delta^{(i)})$ with $\delta^{(i)} \sim \mathcal{D}_{x_j}$. Define $Z_j^{(i)} := \|\hat{q}(y_j^{(i)}) - x_j\|$, and assume that $Z_j^{(i)}$ are sorted in non-decreasing order, and let $\bar{E}_{x_j} := Z_j^{(r)}$ with r from Lemma 1. Then, for any $x \in \mathcal{X}$, it holds that*

$$\text{Prob}\left(e(x, t) \leq \sup_j \bar{E}_{x_j} + (\mathcal{L}_p \mathcal{L}_{\hat{q}} + 1)\epsilon\right) \geq 1 - \alpha, \quad (4)$$

The above result says that the state estimation error $e(x, t)$ can essentially be bounded, with probability $1 - \alpha$, by the worst case of conformal prediction region \bar{E}_{x_j} within the grid $\bar{\mathcal{X}}$ and by the gridding parameter ϵ . Under the assumption that our system operates in the workspace \mathcal{X} and based on inequality (4), we can hence conclude that

$$\text{Prob}\left(x \in \{\zeta \in \mathbb{R}^n \mid \|\zeta - \hat{x}\| \leq \sup_j \bar{E}_{x_j} + (\mathcal{L}_p \mathcal{L}_{\hat{q}} + 1)\epsilon\}\right) \geq 1 - \alpha.$$

Remark 2. *We note that the Lipschitz constants of the sensor and perception maps are used in the upper bound in (4) (as commonly done in the literature [5], [25], [32]), which may lead to a conservative bound. One practical way to mitigate this conservatism is to decrease the gridding parameter ϵ , i.e., to increase the sampling density in the workspace \mathcal{X} .*

B. Sampled-Data Controller using Conformal Estimation Regions

After bounding the state estimation error in Proposition 1, we now design a uncertainty-aware controller based on equation (4). However, a technical challenge in doing so is that the measurements are stochastic. By designing a sampled-data controller, we can avoid difficulties dealing with stochastic calculus. To do so, we first present a slightly modified version

²Due to space constraints, all proofs have been omitted, but they can be found in our online technical report [48].

³We assume that the Lipschitz constant of the sensor map p is uniform over the parameter δ , i.e., that δ does not affect the value of \mathcal{L}_p .

of measurement robust control barrier function (MR-CBF) introduced in [5].

Definition 1. *Let $\mathcal{C} \subseteq \mathbb{R}^n$ be the zero-superlevel set of a continuously differentiable function $h : \mathbb{R}^n \rightarrow \mathbb{R}$. The function h is a measurement robust control barrier function (MR-CBF) for the system in (1) with parameter function pair $(a, b) : \mathbb{R}^l \rightarrow \mathbb{R}_{\geq 0}^2$ if there exists an extended class K_∞ function β such that*

$$\sup_{u \in \mathcal{U}} [L_f h(\hat{x}) + L_g h(\hat{x})u - (a(y) + b(y)\|u\|)] \geq -\beta(h(\hat{x})) \quad (5)$$

for all $(y, \hat{x}) \in V(\mathcal{C})$, where $V(\mathcal{C}) := \{(y, \hat{x}) \in \mathbb{R}^l \times \mathbb{R}^n \mid \exists (x, \delta) \in \mathcal{C} \times \mathcal{D}_x \text{ s.t. } \hat{x} = \hat{q}(p(x, \delta))\}$, and $L_f h(\hat{x})$ and $L_g h(\hat{x})$ denote the Lie derivatives.

Compared to regular CBFs [17], a MR-CBF introduces a non-positive robustness term $-(a(y) + b(y)\|u\|)$ which makes the constraint in (5) more strict. Now, given a MR-CBF $h(x)$, the set of MR-CBF consistent control inputs is

$$K_{CBF}(y) := \{u \in \mathcal{U} \mid L_f h(\hat{x}) + L_g h(\hat{x})u - (a(y) + b(y)\|u\|) + \beta(h(\hat{x})) \geq 0\}. \quad (6)$$

Note that we can not simply follow [5, Theorem 2] to obtain a safe control law as $u(t) \in K_{CBF}(y(t))$ since $y(t)$ and consequent $u(t)$ are stochastic. We hence propose a sampled-data control law that keeps the trajectory $x(t)$ within the set \mathcal{C} with high probability. The sampled-data control law \hat{u} is piecewise continuous and defined as

$$\hat{u}(t) := u(t_i), \quad \forall t \in [t_i, t_{i+1}), \quad (7)$$

where $u(t_i)$ at triggering time t_i is computed by solving the following quadratic optimization problem

$$u(t_i) = \underset{u \in K_{CBF}(y)}{\text{argmin}} \|u - u_{nom}(t_i)\|^2, \quad (8)$$

where $u_{nom}(t_i)$ is any nominal control law that may not necessarily be safe. Then, we select the triggering instances t_i as follows:

$$t_0 := 0, \\ t_{i+1} := (\Delta - \sup_j \bar{E}_{x_j} - (\mathcal{L}_p \mathcal{L}_{\hat{q}} + 1)\epsilon) / \bar{F} + t_i, \quad (9)$$

where Δ is a user-defined parameter that will define the parameter pair (a, b) of the MR-CBF and that has to be $\Delta > \sup_j \bar{E}_{x_j} + (\mathcal{L}_p \mathcal{L}_{\hat{q}} + 1)\epsilon$. Naturally, larger Δ lead to less frequent control updates, but will require more robustness and reduce the set of permissible control inputs in $K_{CBF}(y)$. Based on the computation of triggering times in (9), the following lemma holds.

Lemma 2. *Consider the sampled-data control law $\hat{u}(t)$ in (7) with the triggering rules (9), it holds that*

$$\text{Prob}\left(\|x(t) - \hat{x}(t_i)\| \leq \Delta, \forall t \in [t_i, t_{i+1})\right) \geq 1 - \alpha. \quad (10)$$

Intuitively, the above lemma says that $\|x(t) - \hat{x}(t_i)\| \leq \Delta$ holds with high probability in between triggering times if

the sampled-data control law $\hat{u}(t)$ in (7) with the triggering rules (9) is executed. Then, we can obtain the following probabilistic safety guarantees.

Theorem 1. Consider a MR-CBF h with parameter pair $(a(y), b(y)) = ((\mathcal{L}_{L_f h} + \mathcal{L}_{\beta \circ h})\Delta, \mathcal{L}_{L_g h}\Delta)$ where $\mathcal{L}_{L_f h}$, $\mathcal{L}_{\beta \circ h}$, and $\mathcal{L}_{L_g h}$ are the Lipschitz constants of the functions $L_f h$, $\beta \circ h$ and $L_g h$, respectively. Then, for any nominal control law u_{nom} , the sampled-data law $\hat{u}(t)$ in (7) with the triggering rule in (9) will render the set \mathcal{C} forward invariant with a probability of at least $1 - \alpha$. In other words, we have that

$$\text{Prob}(x(t) \in \mathcal{C}, \forall t \in [t_i, t_{i+1})) \geq 1 - \alpha. \quad (11)$$

The above theorem solves Problem 1 for the time interval $\mathcal{T} = [t_i, t_{i+1})$. If we want to consider a larger time interval $\mathcal{T} = [0, T)$ under the sampled-data control law, we have the following guarantees.

Proposition 2. Under the same condition as in Theorem 1, for a time interval $\mathcal{T} = [0, T)$, we have that:

$$\text{Prob}(x(t) \in \mathcal{C}, \forall t \in [0, T)) \geq (1 - \alpha)^m, \quad (12)$$

where $m \in \mathbb{N}_{>0}$ such that $t_{m-1} \leq T < t_m$.

Note that if we want to achieve any probability guarantee $p \in (0, 1)$, we can just let $(1 - \alpha)^m = p$ and obtain $\alpha = 1 - p^{1/m}$.

V. SIMULATION RESULTS

To demonstrate our proposed safe perception-based control law, we consider navigating an F1/10th autonomous vehicle in a structured environment [49], which is shown in Figure 2. The vehicle system has the state $x = [p_x, p_y, \theta]$, where $[p_x, p_y]$ denotes its position and θ denotes its orientation. We have $[\dot{p}_x, \dot{p}_y] = [u_x, u_y]$, where u_x and u_y are control inputs denoting velocities, and $\theta = \arctan(u_y/u_x)$. The control input constraint is $(u_x, u_y) \in [-1, 1] \times [-1, 1]$. Thus, the assumption that there exists an upper bound \bar{F} for dynamics holds for this system.

Observation model: The vehicle is equipped with a 2D LiDAR scanner from which it obtains LiDAR measurements as its observations. Specifically, the measurement include 64 LiDAR rays uniformly ranging from $-\frac{3\pi}{4}$ to $\frac{3\pi}{4}$ relative to the vehicle's heading direction. To model the uncertainty of measurements, unknown noise conforming to exponential distribution is added to each ray:

$$y_n^k = y^k + \delta, \quad \delta \sim \text{exp}(\lambda),$$

where y^k is the ground truth for ray k , y_n^k is the corrupted observed ray k , and λ is the parameter of exponential distribution, where the noise δ is drawn from. In our experiments, we let $\lambda := 2/3$.

Perception map: We trained a feedforward neural network to estimate the state of the vehicle. The input is the 64-dimensional LiDAR measurement and output is the vehicle's state. The training dataset D_{train} contains 4×10^5 data points, and the calibration dataset D_{cal} for conformal prediction




Fig. 2: The F1/10 vehicle is equipped with a 2D LiDAR sensor that outputs an array of 64 laser scans. The vehicle starts at a random position on the starting line.

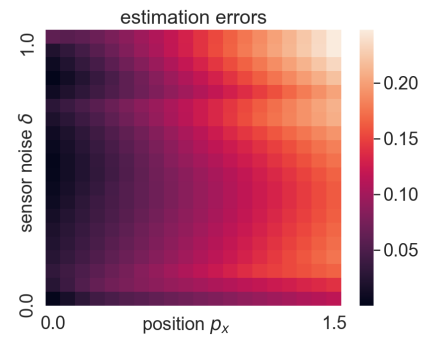


Fig. 3: The empirical model errors $e(x)$ w.r.t. p_x and δ measured on a validation set. p_y and θ are fixed.

contains 1.25×10^4 data points. For illustration, under a fixed heading θ and longitudinal position p_y , the errors $e(x)$ of the learned perception map with respect to sensor noise δ and horizontal position p_x is shown in Figure 3.

Barrier functions: To prevent collision with the walls, when the vehicle is traversing the long hallway, the CBF is chosen as $h(x) = \min\{h_1(x), h_2(x)\}$, where $h_1(x) = p_x$ and $h_2(x) = 1.5 - p_x$. Then we have the safe set $\mathcal{C} = \{x \in \mathbb{R}^3 | h(x) \geq 0\}$. CBFs can be similarly defined when the vehicle is operating in the corner. To demonstrate the effectiveness of our method, we compare the following two cases in simulations:

- 1) *Measurement robust CBF:* as shown in Theorem 1, we choose the parameters pair $(a(y), b(y)) = ((\mathcal{L}_{L_f h} + \mathcal{L}_{\beta \circ h})\Delta, \mathcal{L}_{L_g h}\Delta)$ to ensure robust safety.
- 2) *Vanilla CBF:* we choose the parameters pair $(a(y), b(y)) = (0, 0)$, which essentially reduces to the vanilla non-robust CBF [17]. However, the perceived state is from perceptual estimation rather than real state, so this CBF cannot provide any safety guarantee.

Note that we obtain necessary Lipschitz constants using sampling-based estimation method in simulations.

Uncertainty and results: The vehicle is expected to track along the hallway and make a successful turn in the corner as




Fig. 4: Nonconformity scores R_x histogram during runtime. We select the coverage rate as 75%.




Fig. 5: Traces for the sampled-data measurement robust CBF and vanilla CBF (5 traces are presented). All traces are tested with horizon $T = 30s$. We run 100 traces totally, and the safety rates are 93% and 16%, respectively.

shown in Figure 2. The nominal controller is a PID controller. We set the coverage error $\alpha = 0.25$, so we desire $P(\|x_t - \hat{x}_t\| \leq \epsilon') \geq 1 - \alpha = 75\%$. Based on calibration of conformal prediction and Proposition 1, we calculate that $\epsilon' = 0.34$, and we choose $\Delta = 0.35 > \epsilon'$. The nonconformity score histogram is presented in Figure 4, in which the 75% quantile value is $R_x^{0.75} = 0.32 < \epsilon'$, so our Proposition 1 holds in practice. As presented in Figure 5, the safety rate of sampled-data measurement robust CBF is 93%, which is significantly higher than vanilla non-robust CBF case (16%).

VI. CONCLUSION

In this paper, we consider the safe perception-based control problem under stochastic sensor noise. We use conformal prediction to quantify the state estimation uncertainty, and then integrate this uncertainty into the design of sampled-data safe controller. We obtain probabilistic safety guarantees for continuous-time systems. Note that, in this work, the perception map only depends on current observation, which might limit its accuracy in some cases. We plan to incorporate history observations into perception maps in the future. Also, we are interested in providing a more sample-efficient scheme while constructing calibration dataset.

REFERENCES

[1] S. Tang, V. Wüest, and V. Kumar, “Aggressive flight with suspended payloads using vision-based control,” *IEEE Robotics and Automation Letters*, vol. 3, no. 2, pp. 1152–1159, 2018.

[2] Y. Lin, F. Gao, T. Qin, W. Gao, T. Liu, W. Wu, Z. Yang, and S. Shen, “Autonomous aerial navigation using monocular visual-inertial fusion,” *Journal of Field Robotics*, vol. 35, no. 1, pp. 23–51, 2018.

[3] H. Zhou and V. Tzoumas, “Safe perception-based control with minimal worst-case dynamic regret,” *arXiv preprint arXiv:2208.08929*, 2022.

[4] Y. Kantaros, S. Kalluraya, Q. Jin, and G. J. Pappas, “Perception-based temporal logic planning in uncertain semantic maps,” *IEEE Transactions on Robotics*, 2022.

[5] S. Dean, A. J. Taylor, R. K. Cosner, B. Recht, and A. D. Ames, “Guaranteeing safety of learned perception modules via measurement-robust control barrier functions,” *arXiv preprint arXiv:2010.16001*, 2020.

[6] D. Sun, N. Musavi, G. Dullerud, S. Shakkottai, and S. Mitra, “Learning certifiably robust controllers using fragile perception,” *arXiv preprint arXiv:2209.11328*, 2022.

[7] R. K. Cosner, I. D. J. Rodriguez, T. G. Molnar, W. Ubellacker, Y. Yue, A. D. Ames, and K. L. Bouman, “Self-supervised online learning for safety-critical control using stereo vision,” in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 11 487–11 493.

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

[9] V. Vovk, A. Gammerman, and G. Shafer, *Algorithmic learning in a random world*. Springer Science & Business Media, 2005.

[10] M. Abu-Khalaf, S. Karaman, and D. Rus, “Feedback from pixels: Output regulation via learning-based scene view synthesis,” in *Learning for Dynamics and Control*. PMLR, 2021, pp. 828–841.

[11] X. Sun, M. Zhou, Z. Zhuang, S. Yang, J. Betz, and R. Mangharam, “A benchmark comparison of imitation learning-based control policies for autonomous racing,” in *2023 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2023, pp. 1–5.

[12] A. Khazraei, H. Pfister, and M. Pajic, “Attacks on perception-based control systems: Modeling and fundamental limits,” *arXiv preprint arXiv:2206.07150*, 2022.

[13] —, “Resiliency of perception-based controllers against attacks,” in *Learning for Dynamics and Control Conference*. PMLR, 2022, pp. 713–725.

[14] S. Veer and A. Majumdar, “Probably approximately correct vision-based planning using motion primitives,” in *Conference on Robot Learning*. PMLR, 2021, pp. 1001–1014.

[15] A. Majumdar, A. Farid, and A. Sonar, “Pac-bayes control: learning policies that provably generalize to novel environments,” *The International Journal of Robotics Research*, vol. 40, no. 2-3, pp. 574–593, 2021.

[16] M. Ostertag, N. Atanasov, and T. Rosing, “Trajectory planning and optimization for minimizing uncertainty in persistent monitoring applications,” *Journal of Intelligent & Robotic Systems*, vol. 106, no. 1, p. 2, 2022.

[17] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, “Control barrier function based quadratic programs for safety critical systems,” *IEEE Transactions on Automatic Control*, vol. 62, no. 8, pp. 3861–3876, 2016.

[18] S. Yang, S. Chen, V. M. Preciado, and R. Mangharam, “Differentiable safe controller design through control barrier functions,” *IEEE Control Systems Letters*, vol. 7, pp. 1207–1212, 2022.

[19] W. Xiao, T.-H. Wang, R. Hasani, M. Chahine, A. Amini, X. Li, and D. Rus, “BarrierNet: Differentiable control barrier functions for learning of safe robot control,” *IEEE Transactions on Robotics*, 2023.

[20] L. Lindemann and D. V. Dimarogonas, “Control barrier functions for signal temporal logic tasks,” *IEEE control systems letters*, vol. 3, no. 1, pp. 96–101, 2018.

[21] J. Wang, S. Yang, Z. An, S. Han, Z. Zhang, R. Mangharam, M. Ma, and F. Miao, “Multi-agent reinforcement learning guided by signal temporal logic specifications,” *arXiv preprint arXiv:2306.06808*, 2023.

[22] Y. Chen, C. Shang, X. Huang, and X. Yin, “Data-driven safe controller synthesis for deterministic systems: A posteriori method with validation tests,” in *2023 IEEE 62nd Conference on Decision and Control (CDC)*, 2023.

[23] L. Lindemann, A. Robey, L. Jiang, S. Tu, and N. Matni, “Learning robust output control barrier functions from safe expert demonstrations,” *arXiv preprint arXiv:2111.09971*, 2021.

- [24] C. Dawson, B. Lowenkamp, D. Goff, and C. Fan, "Learning safe, generalizable perception-based hybrid control with certificates," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 1904–1911, 2022.
- [25] L. Cothren, G. Bianchin, and E. Dall'Anese, "Online optimization of dynamical systems with deep learning perception," *IEEE Open Journal of Control Systems*, vol. 1, pp. 306–321, 2022.
- [26] L. Cothren, G. Bianchin, S. Dean, and E. Dall'Anese, "Perception-based sampled-data optimization of dynamical systems," *arXiv preprint arXiv:2211.10020*, 2022.
- [27] D. R. Agrawal and D. Panagou, "Safe and robust observer-controller synthesis using control barrier functions," *IEEE Control Systems Letters*, vol. 7, pp. 127–132, 2022.
- [28] A. Clark, "Control barrier functions for complete and incomplete information stochastic systems," in *2019 American Control Conference (ACC)*. IEEE, 2019, pp. 2928–2935.
- [29] —, "Control barrier functions for stochastic systems," *Automatica*, vol. 130, p. 109688, 2021.
- [30] Y. Wang and X. Xu, "Observer-based control barrier functions for safety critical systems," in *2022 American Control Conference (ACC)*. IEEE, 2022, pp. 709–714.
- [31] Y. Zhang, S. Walters, and X. Xu, "Control barrier function meets interval analysis: Safety-critical control with measurement and actuation uncertainties," in *2022 American Control Conference (ACC)*. IEEE, 2022, pp. 3814–3819.
- [32] G. Chou, N. Ozay, and D. Berenson, "Safe output feedback motion planning from images via learned perception modules and contraction theory," in *Algorithmic Foundations of Robotics XV: Proceedings of the Fifteenth Workshop on the Algorithmic Foundations of Robotics*. Springer, 2022, pp. 349–367.
- [33] R. Römer, A. Lederer, S. Tesfazgi, and S. Hirche, "Uncertainty-aware visual perception for safe motion planning," *arXiv preprint arXiv:2209.06936*, 2022.
- [34] A. Angelopoulos, S. Bates, J. Malik, and M. I. Jordan, "Uncertainty sets for image classifiers using conformal prediction," *arXiv preprint arXiv:2009.14193*, 2020.
- [35] A. N. Angelopoulos, A. P. Kohli, S. Bates, M. Jordan, J. Malik, T. Al-shaabi, S. Upadhyayula, and Y. Romano, "Image-to-image regression with distribution-free uncertainty quantification and applications in imaging," in *International Conference on Machine Learning*. PMLR, 2022, pp. 717–730.
- [36] C. Fannjiang, S. Bates, A. Angelopoulos, J. Listgarten, and M. I. Jordan, "Conformal prediction for the design problem," *arXiv preprint arXiv:2202.03613*, 2022.
- [37] L. Bortolussi, F. Cairoli, N. Paoletti, and S. D. Stoller, "Conformal predictions for hybrid system state classification," in *From Reactive Systems to Cyber-Physical Systems*. Springer, 2019, pp. 225–241.
- [38] F. Cairoli, L. Bortolussi, and N. Paoletti, "Neural predictive monitoring under partial observability," in *Runtime Verification: 21st International Conference, RV 2021, Virtual Event, October 11–14, 2021, Proceedings 21*. Springer, 2021, pp. 121–141.
- [39] L. Lindemann, M. Cleaveland, G. Shim, and G. J. Pappas, "Safe planning in dynamic environments using conformal prediction," *arXiv preprint arXiv:2210.10254*, 2022.
- [40] A. Dixit, L. Lindemann, S. Wei, M. Cleaveland, G. J. Pappas, and J. W. Burdick, "Adaptive conformal prediction for motion planning among dynamic agents," *arXiv preprint arXiv:2212.00278*, 2022.
- [41] V. M. H. Bennetts, A. J. Lilienthal, A. A. Khaliq, V. P. Sese, and M. Trincavelli, "Towards real-world gas distribution mapping and leak localization using a mobile robot with 3d and remote gas sensing capabilities," in *2013 IEEE International Conference on Robotics and Automation*. IEEE, 2013, pp. 2335–2340.
- [42] S. Dean, N. Matni, B. Recht, and V. Ye, "Robust guarantees for perception-based control," in *Learning for Dynamics and Control*. PMLR, 2020, pp. 350–360.
- [43] G. Shafer and V. Vovk, "A tutorial on conformal prediction," *Journal of Machine Learning Research*, vol. 9, no. 3, 2008.
- [44] A. N. Angelopoulos and S. Bates, "A gentle introduction to conformal prediction and distribution-free uncertainty quantification," *arXiv preprint arXiv:2107.07511*, 2021.
- [45] J. Lei, M. G'Sell, A. Rinaldo, R. J. Tibshirani, and L. Wasserman, "Distribution-free predictive inference for regression," *Journal of the American Statistical Association*, vol. 113, no. 523, pp. 1094–1111, 2018.
- [46] R. J. Tibshirani, R. Foygel Barber, E. Candès, and A. Ramdas, "Conformal prediction under covariate shift," *Advances in neural information processing systems*, vol. 32, 2019.
- [47] R. Vershynin, *High-dimensional probability: An introduction with applications in data science*. Cambridge university press, 2018, vol. 47.
- [48] S. Yang, G. J. Pappas, R. Mangharam, and L. Lindemann, "Safe perception-based control under stochastic sensor uncertainty using conformal prediction," *arXiv preprint arXiv:2304.00194*, 2023.
- [49] R. Ivanov, T. J. Carpenter, J. Weimer, R. Alur, G. J. Pappas, and I. Lee, "Case study: verifying the safety of an autonomous racing car with a neural network controller," in *Proceedings of the 23rd International Conference on Hybrid Systems: Computation and Control*, 2020, pp. 1–7.