

Gaussian Process Implicit Surfaces as Control Barrier Functions for Safe Robot Navigation

Mouhyemen Khan¹, Tatsuya Ibuki², Abhijit Chatterjee¹

Abstract—Level set methods underpin modern safety techniques such as control barrier functions (CBFs), while also serving as implicit surface representations for geometric shapes via distance fields. Inspired by these two paradigms, we propose a unified framework where the implicit surface itself acts as a CBF. We leverage Gaussian process (GP) implicit surface (GPIS) to represent the safety boundaries, using *safety samples* which are derived from sensor measurements to condition the GP. The GP posterior mean defines the implicit safety surface (safety belief), while the posterior variance provides a robust safety margin. Although GPs have favorable properties such as uncertainty estimation and analytical tractability, they scale cubically with data. To alleviate this issue, we develop a sparse solution called *sparse Gaussian CBFs*. To the best of our knowledge, GPIS have not been explicitly used to synthesize CBFs. We validate the approach on collision avoidance tasks in two settings: a simulated 7-DOF manipulator operating around the Stanford bunny, and a quadrotor navigating in 3D around a physical chair. In both cases, Gaussian CBFs (with and without sparsity) enable safe interaction and collision-free execution of trajectories that would otherwise intersect the objects. The experiment video link is: <https://youtu.be/lXjaU3u1OVE>.

I. INTRODUCTION

Safety-critical robotic systems are rapidly expanding across domains such as medicine, agriculture, warehouses, disaster response, and defense, where dependability is crucial since failures can endanger lives and cause major economic loss. Control barrier functions (CBFs), a level-set approach to enforcing safety [1], have been demonstrated on legged robots [2], [3], aerial vehicles [4], [5], manipulators [6]. CBFs guarantee forward invariance of a prescribed *safe set* using three elements: a candidate scalar function, a nominal system model, and a nominal control input. The candidate's 0-superlevel set defines the CBF, from which quadratic program (QP) constraints rectify the nominal input into a safe control. However, most CBFs are hand-crafted from domain knowledge and heuristics, which is often infeasible in many real-world applications. *This motivates our goal of achieving safe control for dynamical systems, particularly robots, by constructing CBFs through a data-driven approach.*

Data-driven approaches for constructing CBFs are actively pursued. Expert demonstrations of state and control trajectories have been used to generate CBFs [7], but relying on such demonstrations is impractical in unseen environments,

and these methods lack hardware validation. Neural certificates have also been proposed [8], [9], providing formal correctness guarantees to learning-based controllers, yet they remain limited to offline training and simulation. Episodic learning has been applied to update controllers under safety constraints [10], but the focus was on modeling system uncertainty rather than constructing CBFs. Support vector machines have been used to characterize CBFs from sensor data [11], though they require carefully tuned weights and were demonstrated only in 2D simulation. CBFs have also been constructed directly from point cloud observations using geometric robot-envelope models [12]. Gaussian processes (GPs) have been used to synthesize CBFs directly from data [13]. Other studies involved using Bayesian meta-learning for fast adaptation across environments in simulation [14], and a LiDAR-based GP-CBF for 2D navigation of unicycle robots [15]. In contrast to [12], [13], [14], [15], we characterize safe implicit surfaces in 3D using GPs, introduce *sparse Gaussian CBFs* to reduce complexity while retaining safety guarantees, and provide hardware results in 3D.

To construct CBFs in 3D from data, we draw on signed distance functions (SDFs) as implicit surface representations. An SDF determines whether a point $\mathbf{x} \in \Omega$ lies inside or outside the set, and has been shown to be a natural representation for navigation and planning [16]. SDFs have also been used to construct CBFs with memory [17], though only with 2D sensing and in simulation. In contrast, we employ Gaussian processes (GPs) to form GP implicit surfaces (GPIS) [18], which have been studied for surface representation, navigation, shape estimation, and grasping [19], [20], [21], [22]. To the best of our knowledge, GPIS has not been explicitly used to synthesize CBFs. Motivated by SDFs for surface representation and CBFs for safe control, we propose a unified framework that directly infers implicit surfaces while enforcing safety control.

In summary, our **major contributions** are the following:

- We introduce a unified approach for implicit surface representation and safe control, where a GPIS serves as the candidate CBF. To our knowledge, GPIS have not previously been used for synthesizing safe control.
- We develop *sparse Gaussian CBFs*, which retain the uncertainty estimation and analytical tractability of Gaussian CBFs while reducing computational complexity.
- We validate the approach in both simulation and hardware: a 7-DOF robotic manipulator and a quadrotor platform. We believe this is the first work to demonstrate collision avoidance based safe control through implicit surface based CBFs on hardware.

¹School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332, USA.

²Department of Electronics and Bioinformatics, School of Science and Technology, Meiji University, Kanagawa 214-8571, Japan.

Corresponding author: M. Khan, mouhyemen.khan@gmail.com.

This work is supported in part by NSF grant 2128419.

II. BACKGROUND PRELIMINARIES

We model safety for a dynamical system using GPs as smooth surfaces, combining CBFs and GPIS to design probabilistic implicit surfaces. Below, we briefly review CBFs, GPs, and implicit surfaces (see [23], [24], [25] for details).

A. Control Barrier Function

Consider a general control affine dynamical system,

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

where $\mathbf{x}(t) \in \mathbb{R}^n$ is the state and $\mathbf{u}(t) \in \mathbb{R}^m$ is the control input, $f: \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $g: \mathbb{R}^n \rightarrow \mathbb{R}^{n \times m}$ are assumed to be locally Lipschitz continuous. Let safety for (1) be encoded as the superlevel set \mathcal{S} of a smooth function $h: \mathbb{R}^n \rightarrow \mathbb{R}$ as,

$$\mathcal{S} = \{\mathbf{x} \in \mathbb{R}^n \mid h(\mathbf{x}) \geq 0\}. \quad (2)$$

Definition 1 (Control Barrier Function [23]). *The function $h(\mathbf{x}): \mathbb{R}^n \rightarrow \mathbb{R}$ is defined as a control barrier function (CBF), if there exists an extended class- κ function α ($\alpha(0) = 0$ and strictly increasing) such that for any $\mathbf{x} \in \mathcal{S}$,*

$$\sup_{\mathbf{u} \in \mathbb{R}^m} L_f h(\mathbf{x}) + L_g h(\mathbf{x})\mathbf{u} + \alpha(h(\mathbf{x})) \geq 0, \quad (3)$$

where $L_f h(\mathbf{x}) = \frac{\partial h}{\partial \mathbf{x}} f(\mathbf{x})$ and $L_g h(\mathbf{x}) = \frac{\partial h}{\partial \mathbf{x}} g(\mathbf{x})$ are the Lie derivatives of $h(\mathbf{x})$ along $f(\mathbf{x})$ and $g(\mathbf{x})$ respectively.

B. Implicit Surfaces

An implicit surface represents the shape of a volumetric object in n -dimensional Euclidean space via a function that determines whether a point belongs to the object. Formally, it is defined as the 0-level set (isosurface) of a real-valued implicit function $f_{\text{is}}: \mathbb{R}^n \rightarrow \mathbb{R}$ for $\mathbf{x} \in \mathbb{R}^n$:

$$f_{\text{is}}(\mathbf{x}) \begin{cases} > 0, & \mathbf{x} \text{ outside the surface} \\ = 0, & \mathbf{x} \text{ on the surface} \\ < 0, & \mathbf{x} \text{ inside the surface} \end{cases}. \quad (4)$$

The 0-level set $f_{\text{is}}^{-1}(0)$ of an implicit function f_{is} defines a surface $\mathcal{S} \subset \mathbb{R}^n$, which in our case represents the safety boundary. Conversely, for any surface \mathcal{S} in \mathbb{R}^n , there exists a function $f_{\text{is}}: \mathbb{R}^n \rightarrow \mathbb{R}$ such that $f_{\text{is}}^{-1}(0) = \mathcal{S}$ (Prop. 2 in [26]). Hence, given samples on \mathcal{S} , one can construct f_{is} . We approximate the implicit surface of a volumetric object in $d = 3$ ($d \leq n$) using a GP regressor, referred to in the literature as a Gaussian Process Implicit Surface (GPIS).

C. Gaussian Process Regression

GP regression is a well established nonparametric approach which relies on kernels for solving non-linear regression tasks. Kernels provide a notion of similarity between pairs of input points, $\mathbf{x}_i, \mathbf{x}_j \in \mathbb{R}^n$. A popular kernel is the squared exponential (SE) kernel given by,

$$k(\mathbf{x}_i, \mathbf{x}_j) = \sigma_f^2 \exp\left(-\frac{(\mathbf{x}_i - \mathbf{x}_j)^\top \mathbf{L}^{-2}(\mathbf{x}_i - \mathbf{x}_j)}{2}\right) + \delta_{ij} \sigma_y^2, \quad (5)$$

where $\delta_{ij} = 1$ if $i = j$ and 0 otherwise, $\mathbf{l} \in \mathbb{R}^n$ is the characteristic length scale, with $\mathbf{L} = \text{diag}(\mathbf{l}) \in \mathbb{R}^{n \times n}$. The signal scale and observation noise are σ_f^2 and σ_y^2 respectively.




Fig. 1: The safety surface of a volumetric object of interest is modeled as a sparse Gaussian CBF using safety samples.

Together, these parameters constitute the SE kernel's hyperparameters, $\Theta = \{\mathbf{L}, \sigma_f^2, \sigma_y^2\}$. These hyperparameters for a dataset can be optimized by maximizing the log marginal likelihood using quasi-Newton methods [24].

We are interested in constructing an implicit surface, which will later be our safety function of interest, for which we will carefully design scalar targets y in Section IV-A. Given a set of N data points, with input vectors $\mathbf{x} \in \mathbb{R}^n$, and scalar targets $y \in \mathbb{R}$, we compose the dataset $\mathcal{D}_N = \{\mathbf{X}_N, \mathbf{y}_N\}$, where $\mathbf{X}_N = \{\mathbf{x}_i\}_{i=1}^N$ and $\mathbf{y}_N = \{y_i\}_{i=1}^N$. GPs can compute the posterior mean and variance for an arbitrary deterministic query point $\mathbf{x}_* \in \mathbb{R}^n$, by conditioning on previous measurements. The posterior mean $\mu \in \mathbb{R}$ and variance $\sigma^2 \in \mathbb{R}$ are [24],

$$\mu(\mathbf{x}_*) = \mathbf{k}(\mathbf{x}_*)^\top \bar{\mathbf{K}}^{-1} \mathbf{y}_N, \quad (6)$$

$$\sigma^2(\mathbf{x}_*) = k(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}(\mathbf{x}_*)^\top \bar{\mathbf{K}}^{-1} \mathbf{k}(\mathbf{x}_*), \quad (7)$$

where $\mathbf{k}(\mathbf{x}_*) = [k(\mathbf{x}_1, \mathbf{x}_*), \dots, k(\mathbf{x}_N, \mathbf{x}_*)]^\top \in \mathbb{R}^N$ is the covariance vector between \mathbf{X}_N and \mathbf{x}_* , $\bar{\mathbf{K}} \in \mathbb{R}^{N \times N}$, with entries $[\bar{\mathbf{K}}]_{(i,j)} = k(\mathbf{x}_i, \mathbf{x}_j)$, $i, j \in \{1, \dots, N\}$, is the covariance matrix between pairs of input points in \mathbf{X}_N , and $k(\mathbf{x}_*, \mathbf{x}_*) \in \mathbb{R}$ is the prior covariance.

III. PROBLEM STATEMENT

Given a dynamical system (1) and a 3D object, our goal is to recover the object's surface using GPs and perform collision avoidance through safe control. The surface is learned from on- and off-surface points derived from sensor data, which we term safety samples (Fig. 1), and these samples define the Gaussian CBF.

Assumption 1. *We assume access to sensor data providing surface points and normals of the object.*

This assumption is reasonable, since many sensing modalities (e.g., laser, haptic, or vision-based) provide point clouds and surface normals [21].

Remark 1. *Safety surface is modeled in $d = 3$ (Euclidean space) while the system state $\mathbf{x} \in \mathbb{R}^n$ with $d \leq n$, the $(n-d)$ free dimensions can be set to zero during GP training.*

$$\underbrace{f_{\text{is}}(\mathbf{x})}_{\text{Implicit Surface}} = \underbrace{h_{\text{sgp}}(\mathbf{x})}_{\text{Sparse Gaussian CBF}} := \underbrace{h_{\text{b}}(\mathbf{x})}_{\text{Safety Belief}} + \underbrace{h_{\text{u}}(\mathbf{x})}_{\text{Safety Margin}}. \quad (8)$$

The sparse Gaussian CBF provides an implicit surface approximation of a volumetric object. Because this surface is data-driven, it is necessary to include a safety margin in the safety function estimate, especially when exploiting sparsity, where uncertainty must be addressed. The safety belief

denotes the best estimate of system safety (or the implicit surface), while the safety margin captures uncertainty in the estimate, increasing in regions with limited data.

Problem 1. *Given system (1) and online measurements of the system state \mathbf{x}_* , synthesize $h_{\text{sgp}}(\mathbf{x})$ to model an implicit surface with a safety belief and associated safety uncertainty, conditioned on past safety samples and observations for a volumetric object in the dataset, $\mathcal{D}_N = \{\mathbf{X}_N, \mathbf{y}_N\}$, where $\mathbf{X}_N = \{\mathbf{x}_i\}_{i=1}^N$ and $\mathbf{y}_N = \{y_i\}_{i=1}^N$, such that system (1) is safe.*

Given a nominal controller, the system may not be safe under its influence. Hence, we want to modify the nominal input using the synthesized $h_{\text{sgp}}(\mathbf{x})$ to achieve safety.

Assumption 2. *We assume a nominal control input \mathbf{u}_{nom} exists that drives system (1) state \mathbf{x} to a desired state \mathbf{x}_{des} .*

Problem 2. *Given system (1), the synthesized $h_{\text{sgp}}(\mathbf{x})$ with safe set \mathcal{S} , and nominal control input $\mathbf{u}_{\text{nom}} \in \mathbb{R}^m$, rectify the control input $\mathbf{u}_{\text{rect}} \in \mathbb{R}^m$ ensuring system (1) is safe.*

However, designing \mathbf{u}_{rect} with a nonparametric CBF is difficult, since Lie derivatives are ill-posed without assumptions. We address this by using a kernel representation to encode both the safety belief and its margin.

IV. PROPOSED METHODOLOGY

Here, we present our approach to approximate the GPIS f_{is} with a sparse Gaussian CBF $h_{\text{sgp}}(\mathbf{x})$. GPs are preferable to models such as neural networks, polynomial chaos, or radial basis functions because their Bayesian nonparametric formulation provides flexible priors and closed-form probabilistic posterior estimates. Moreover, GPs naturally yield uncertainty through posterior variance, whereas neural networks require careful design to capture uncertainty [27].

A. Data Processing

Given N_s sensor measurements in an n -dimensional Euclidean space, $\mathcal{P} \in \mathbb{R}^{N_s \times n} \times \mathbb{R}^{N_s \times n}$ denotes the set of point cloud data and corresponding surface normals. The point cloud gives the on-surface points, while off-surface points, both external and internal, are generated using the surface normals; collectively called *safety samples*,

$$\begin{aligned} \Omega_{\text{ext}} &:= \{ \mathbf{x}_+ \mid f_{\text{is}}(\mathbf{x}_+) > 0 \} \\ \Omega_0 &:= \{ \mathbf{x}_0 \mid f_{\text{is}}(\mathbf{x}_0) = 0 \} \\ \Omega_{\text{int}} &:= \{ \mathbf{x}_- \mid f_{\text{is}}(\mathbf{x}_-) < 0 \} \end{aligned} .$$

Since only the point cloud coordinates and surface normals in \mathcal{P} are available, we construct the sets Ω_0 , Ω_{ext} , and Ω_{int} .

- Ω_0 is formed by taking N_0 points from the point cloud \mathcal{P} ($N_0 \leq N_s$), each labeled with 0, i.e., $\mathbf{y}_{N_0} = \{0\}_{i=1}^{N_0}$.
- Ω_{ext} is formed by taking N_+ synthetic points ($N_+ < N_0$) along the surface normals with a fixed offset.
- Ω_{int} is formed by creating N_- points ($N_- \leq N_+$) in the opposite normal direction.
- Targets are +1 for Ω_{ext} and -1 for Ω_{int} , creating the labels, $\mathbf{y}_{N_+} = \{+1\}_{i=1}^{N_+}$ and $\mathbf{y}_{N_-} = \{-1\}_{i=1}^{N_-}$.
- The GP training set is $\mathbf{X}_N = [\mathbf{x}_0^\top, \mathbf{x}_+^\top, \mathbf{x}_-^\top]^\top \in \mathbb{R}^{N \times n}$ with labels $\mathbf{y}_N = [\mathbf{y}_{N_0}, \mathbf{y}_{N_+}, \mathbf{y}_{N_-}]$, where $N = N_0 + N_+ + N_-$.

B. Sparse Gaussian Control Barrier Function

We use a GP prior on the desired safety function, $h_{\text{sgp}} \sim \mathcal{GP}(0, k(\mathbf{x}, \mathbf{x}'))$ where the inputs and targets to the GP are \mathbf{X}_N and \mathbf{y}_N as discussed above. By using the GP prior, we fully specify the candidate safety function, unlike [28] where only the posterior variance was modeled by GP. Additionally, we also present a sparse approximation to our approach.

Assumption 3. *The safe set is nonempty with at least one data point, the initial state $\mathbf{x}(0)$ and associated safety value $h_{\text{sgp}}(\mathbf{x}(0)) \in \mathbb{R}_{\geq 0}$, to synthesize h_{sgp} .*

We assume that the system begins in an initial compact safe set. Safety for h_{sgp} is encoded as,

$$\mathcal{S} = \{ \mathbf{x} \in \mathbb{R}^n \mid h_{\text{sgp}}(\mathbf{x}) \geq 0 \}, \quad (9)$$

$$\partial\mathcal{S} = \{ \mathbf{x} \in \mathbb{R}^n \mid h_{\text{sgp}}(\mathbf{x}) = 0 \}. \quad (10)$$

Despite GPs being very powerful regressors, as the dataset grows larger, they become computationally intractable. GP prediction complexity has a cost of $\mathcal{O}(N^3)$. Even if one stores the covariance matrix to save costs, the complexity per test case is $\mathcal{O}(N)$ for predictive mean and $\mathcal{O}(N^2)$ for predictive variance. Hence, many sparse approximations of GPs have been developed to reduce the complexity cost while retaining accuracy [29], [30], [31]. We focus on the variant called *Sparse Pseudo-Input Gaussian Process* (SPGP) [30]. Here, we simply present its mean and variance:

$$\mu(\mathbf{x}_*) = \mathbf{k}_M^\top(\mathbf{x}_*) \mathbf{Q}_M^{-1} \mathbf{K}_{MN} (\mathbf{\Lambda}_N + \sigma_y^2 \mathbf{I}_N)^{-1} \mathbf{y}_N \quad (11)$$

$$\sigma^2(\mathbf{x}_*) = k(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}_M^\top(\mathbf{x}_*) \mathbf{P}_M \mathbf{k}_M(\mathbf{x}_*), \quad (12)$$

where $\mathbf{k}_M(\mathbf{x}_*) = [k(\bar{\mathbf{x}}_1, \mathbf{x}_*), \dots, k(\bar{\mathbf{x}}_M, \mathbf{x}_*)]^\top \in \mathbb{R}^M$ is the covariance between $\bar{\mathbf{X}}_M$ and \mathbf{x}_* , $\mathbf{Q}_M = \mathbf{K}_{NM}^\top (\mathbf{\Lambda}_N + \sigma_y^2 \mathbf{I}_N)^{-1} \mathbf{K}_{NM} + \mathbf{K}_M \in \mathbb{R}^{M \times M}$, $\mathbf{\Lambda}_N = \text{diag}[\mathbf{K}_N - \mathbf{K}_{NM} \mathbf{K}_M^{-1} \mathbf{K}_{MN}] \in \mathbb{R}^{N \times N}$ is a diagonal matrix, $\mathbf{K}_M \in \mathbb{R}^{M \times M}$ is the covariance between pairs of pseudo-inputs $\bar{\mathbf{X}}_M$, $\mathbf{P}_M = \mathbf{K}_M^{-1} - \mathbf{Q}_M^{-1}$, and $[\mathbf{K}_{NM}]_{(i,j)} = k(\mathbf{x}_i, \bar{\mathbf{x}}_j)$, $i \in \{1, \dots, N\}$, $j \in \{1, \dots, M\}$ is the covariance matrix between \mathbf{X}_N and $\bar{\mathbf{X}}_M$. Computation cost for \mathbf{Q}_M is dominated by the inversion operation which is $\mathcal{O}(M^2 N)$ [30]. By precomputing the inverse, the cost per test case is $\mathcal{O}(M)$ and $\mathcal{O}(M^2)$ for predictive mean and variance respectively. We can jointly optimize for the kernel hyperparameters, Θ , and pseudo-inputs, $\bar{\mathbf{X}}_M$, by maximizing the log marginal likelihood using quasi-Newton gradient methods:

$$\log p(y \mid \mathbf{X}, \bar{\mathbf{X}}) = \frac{-N \log(2\pi) + \log |\bar{\mathbf{K}}_N| + \mathbf{y}_N \bar{\mathbf{K}}_N^{-1} \mathbf{y}_N}{2} \quad (13)$$

We propose the following sparse Gaussian CBF $h_{\text{sgp}}(\mathbf{x})$ which models both the safety belief and associated margin using the sparse predictive mean (11) and variance (12),

$$\begin{aligned} h_{\text{sgp}}(\mathbf{x}) &:= \mu(\mathbf{x}) + \sigma^2(\mathbf{x}) \\ &= \underbrace{\mathbf{k}_M^\top(\mathbf{x}) \bar{\mathbf{Q}}_{MN} \mathbf{y}_N}_{\text{safety belief}} + \underbrace{k(\mathbf{x}, \mathbf{x}) - \mathbf{k}_M^\top(\mathbf{x}) \mathbf{P}_M \mathbf{k}_M(\mathbf{x})}_{\text{safety margin}}, \end{aligned} \quad (14)$$

where $\bar{\mathbf{Q}}_{MN} = \mathbf{Q}_M^{-1} \mathbf{K}_{MN} (\boldsymbol{\Lambda}_N + \sigma_y^2 \mathbf{I}_N)^{-1} \in \mathbb{R}^{M \times N}$. As can be seen from the formulation above, we fully realize the safety function using GPs, as well as exploit sparsity, to model any arbitrary volumetric object in an n -dimensional Euclidean space while addressing safety. The admissible control space for the sparse Gaussian CBF ensures that the system (1) remains forward invariant in the safe set characterized by h_{sgp} (Theorem 1 in [23]).

Remark 2. Sparse Gaussian CBFs offer several advantages for modeling implicit surfaces with safety considerations. First, sensor data directly informs the safety function, enabling data-driven surfaces rather than hand-crafted candidates. Second, GPs yield posterior variance, providing principled uncertainty estimates for predictions. Third, sparsity ensures real-time feasibility when handling large datasets.

C. Lie Derivatives of Sparse Gaussian CBF

To achieve safe control using sparse Gaussian CBFs, we require the necessary Lie derivatives for ensuring forward invariance in the safe set. A key advantage of using GPs is the use of kernel representations. This makes it easy to compute the partial derivatives in closed-form analytically for the safety belief and margin using (11), (12). The partial derivative of (14) with respect to \mathbf{x} at a query point \mathbf{x}_* is,

$$\begin{aligned} \left. \frac{\partial h_{\text{sgp}}(\mathbf{x})}{\partial \mathbf{x}} \right|_{\mathbf{x}_*} &= \left. \frac{\partial \mu(\mathbf{x})}{\partial \mathbf{x}} \right|_{\mathbf{x}_*} + \left. \frac{\partial \sigma^2(\mathbf{x})}{\partial \mathbf{x}} \right|_{\mathbf{x}_*} \\ &= \left(\mathbf{y}_N^\top \bar{\mathbf{Q}}_{NM} - 2\mathbf{k}_M^\top(\mathbf{x}_*) \mathbf{P}_M \right) \left. \frac{\partial \mathbf{k}_M(\mathbf{x})}{\partial \mathbf{x}} \right|_{\mathbf{x}_*}, \end{aligned} \quad (15)$$

where $\bar{\mathbf{Q}}_{MN}$ and \mathbf{P}_M are defined in (14). The derivative of the SE kernel (5) in (15) is given by,

$$\left. \frac{\partial \mathbf{k}_M(i)(\mathbf{x})}{\partial \mathbf{x}} \right|_{\mathbf{x}_*} = (\mathbf{x}_{(i)} - \mathbf{x}_*)^\top k(\mathbf{x}_{(i)}, \mathbf{x}_*) \mathbf{L}^{-2}, \quad (16)$$

where $\mathbf{k}_{(i)}$ is the i^{th} element of $\mathbf{k}_M(\mathbf{x})$, and (16) is the i^{th} row of $\frac{\partial \mathbf{k}_M(\mathbf{x})}{\partial \mathbf{x}} \in \mathbb{R}^{M \times n}$. Now, we can compute the Lie derivatives of $h_{\text{sgp}}(\mathbf{x})$ by taking its time derivative as follows,

$$\begin{aligned} \dot{h}_{\text{sgp}}(\mathbf{x}) &= \frac{\partial h_{\text{sgp}}(\mathbf{x})}{\partial \mathbf{x}} f(\mathbf{x}) + \frac{\partial h_{\text{sgp}}(\mathbf{x})}{\partial \mathbf{x}} g(\mathbf{x}) \mathbf{u} \\ &= L_f h_{\text{sgp}}(\mathbf{x}) + L_g h_{\text{sgp}}(\mathbf{x}) \mathbf{u}, \end{aligned} \quad (17)$$

where (15) is used to get $L_f h_{\text{sgp}}(\mathbf{x})$ and $L_g h_{\text{sgp}}(\mathbf{x})$.

D. Safe Control using Sparse Gaussian CBF

Given \mathbf{u}_{nom} and \mathbf{x}_{des} outside the safe set \mathcal{S} , the system would exit \mathcal{S} and violate safety. To ensure forward invariance, \mathbf{u}_{nom} is rectified by solving an online QP with constraints from the Lie derivatives in (17) [23]:

Sparse Gaussian CBF QP: *Input modification*

$$\begin{aligned} \mathbf{u}_{\text{rect}} &= \arg \min_{\mathbf{u} \in \mathbb{R}^m} \frac{1}{2} \|\mathbf{u} - \mathbf{u}_{\text{nom}}\|^2 \quad \text{subject to} \quad (18) \\ L_f h_{\text{sgp}}(\mathbf{x}) + L_g h_{\text{sgp}}(\mathbf{x}) \mathbf{u} + \alpha(h_{\text{sgp}}(\mathbf{x})) &\geq 0, \end{aligned}$$

where \mathbf{u}_{rect} is the rectified control input. This guarantees forward invariance of the system within the safe set \mathcal{S} . The procedure for generating the safe control input using the synthesized CBF is summarized in Algorithm 1.

Algorithm 1 Sparse Gaussian CBF Synthesis & Safe Control

Input:

SYSTEM (1) & NOMINAL INPUT \mathbf{u}_{nom}
 GP PRIOR $h_{\text{sgp}}(\mathbf{x}) \sim \mathcal{GP}(0, k(\mathbf{x}, \mathbf{x}'))$
 DATASET \mathcal{D}_N & number of pseudo-points M

- 1: **procedure** SAFECONTROL
- 2: INITIALIZE $\mathcal{D}_M = \{\bar{\mathbf{X}}_M, \bar{\mathbf{y}}_M\}$ randomly from \mathcal{D}_N
- 3: OPTIMIZE $h_{\text{sgp}}(\mathbf{X}_N, \mathbf{y}_N; \bar{\mathbf{X}}_M, \bar{\mathbf{y}}_M)$ using (13)
- 4: SYNTHESIZE $h_{\text{sgp}}(\mathbf{X}_N, \mathbf{y}_N)$ using (14)
- 5: COMPUTE $\frac{\partial h_{\text{sgp}}(\mathbf{x})}{\partial \mathbf{x}}$ using (15) & (16)
- 6: SETUP QP constraint using (1) & (17)
- 7: RECTIFY \mathbf{u}_{nom} using (18)

return \mathbf{u}_{rect}

A subset of \mathcal{D}_N is first sampled to initialize the pseudo-dataset \mathcal{D}_M . Hyperparameters and pseudo-inputs are optimized in step 3, followed by the synthesis of the sparse Gaussian CBF in step 4, which defines the safety surface for control. The corresponding derivatives are computed in steps 5–6 to construct the QP constraint, and the nominal control input is rectified in step 7.

Remark 3. Owing to its nonparametric nature, the sparse Gaussian CBF can be treated as a black-box model for control synthesis, allowing Algorithm 1 to operate without requiring an explicit reformulation of the Lie derivatives. Unlike traditional CBFs, where altering the function changes the derivative form, Gaussian CBFs preserve the same structure, with their characterization depending only on the data and the dynamical system.

V. TEST CASE I : MANIPULATOR SIMULATION

We apply Gaussian CBFs, with and without sparsity, to control a 7-DOF robot manipulator in simulation. Such manipulators are widely used in warehouses, manufacturing, and medicine, where safe manipulation is essential.

A. Stanford Bunny Implicit Surface Modeling

We model the Stanford bunny surface as the safe set boundary using both full and sparse GPs. The bunny, a




Fig. 2: The Stanford bunny defines the safe set boundary via Gaussian CBFs: (left) point cloud with surface normals, (middle) Gaussian CBF, (right) sparse Gaussian CBF.

standard 3D graphics test model, provides a complex non-convex geometry for safety evaluation. From 34,817 data points, we down-sampled to 3,500 and scaled the model to a $0.45\text{m} \times 0.45\text{m} \times 0.55\text{m}$ cuboid.

1) *Offline Training*: The safety function is synthesized as the implicit surface of the bunny, with the Gaussian CBF formulation given in Section V-C. Training was performed with the `gpm1` toolbox [32] on an Intel i7-9800X CPU (16 GB RAM, 4.4 GHz). For the Gaussian CBF, 2,178 surface points and normals were sampled, giving an average training time of 14.15s. The sparse Gaussian CBF used one-fifth of these as pseudo-points, reducing training time to 9.46s. Online training is described in Section V-D. Figure 2 shows the 0-isosurfaces: the Gaussian CBF captures fine facial contours, while the sparse variant preserves shape with less detail, showing the accuracy–speed trade-off.

2) *Evaluation Metric*: We quantify the modeling performance using the Chamfer distance, a symmetric metric that measures the distance between two point clouds \mathcal{P}_1 and \mathcal{P}_2 sampled from the surfaces.

$$d_{\text{ch}}(\mathcal{P}_1, \mathcal{P}_2) = \sum_{x \in \mathcal{P}_1} \arg \min_{y \in \mathcal{P}_2} \|x - y\| + \sum_{y \in \mathcal{P}_2} \arg \min_{x \in \mathcal{P}_1} \|y - x\| \quad (19)$$

We compute the Chamfer distance between the bunny point cloud $\mathcal{P}_{\text{bunny}}$ and the 0-isosurfaces of h_{gp} and h_{sgp} , obtaining $d_{\text{ch}}(\mathcal{P}_{\text{bunny}}, \mathcal{P}_{\text{gp}}) = 0.011$ and $d_{\text{ch}}(\mathcal{P}_{\text{bunny}}, \mathcal{P}_{\text{sgp}}) = 0.042$. These values indicate good surface approximation, with the GP capturing finer detail than the sparse model.

B. Robot Manipulator Kinematics

We control a 7-DOF manipulator’s end-effector around complex volumetric objects (e.g., the bunny) using kinematic control with direct joint-velocity inputs. The dynamics are

$$\dot{\mathbf{q}} = \mathbf{u}, \quad \mathbf{u}_{\text{nom}} = \dot{\mathbf{q}}_{\text{ref}}, \quad (20)$$

where $\mathbf{q} \in \mathbb{R}^7$ are the joint positions and $\mathbf{u} \in \mathbb{R}^7$ the control inputs. Reference velocities are obtained by interpolating the end-effector’s homogeneous transformation matrix between initial and desired poses. We simulate the Kinova Gen3 manipulator in MATLAB’s Robotic Systems Toolbox, creating trajectories that pass through the bunny. The synthesized CBFs and reference trajectories are shown in Figure 3.

C. Safe Kinematic Control Synthesis for Manipulator

We rectify the nominal joint-velocity control to enforce safety while the manipulator tracks its reference trajectory around the bunny without collision. Safety is ensured using Gaussian CBFs (with and without sparsity) trained to represent the bunny’s implicit surface. Each candidate CBF is

$$h_{(\cdot)}(\mathbf{x}) = \mu(\mathbf{x}) + 4\sigma^2(\mathbf{x}), \quad (21)$$

$$\dot{h}_{(\cdot)}(\mathbf{x}) = \frac{\partial h_{(\cdot)}(\mathbf{x})}{\partial \mathbf{x}} \frac{\partial \mathbf{x}}{\partial \mathbf{q}} \dot{\mathbf{q}} = \frac{\partial h_{(\cdot)}(\mathbf{x})}{\partial \mathbf{x}} \mathbf{J}(\mathbf{q}) \dot{\mathbf{q}}, \quad (22)$$

where $h_{(\cdot)}(\mathbf{x})$ ’s 0-isosurface is the bunny, $\mathbf{J}(\mathbf{q}) : \mathbb{R}^7 \rightarrow \mathbb{R}^{3 \times 7}$ is the manipulator Jacobian, and $\frac{\partial h_{(\cdot)}(\mathbf{x})}{\partial \mathbf{x}}$ is computed from

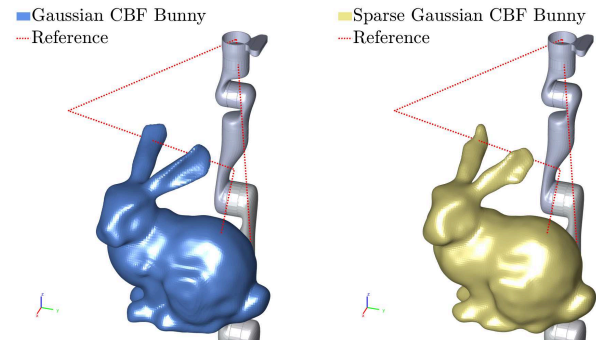


Fig. 3: The reference trajectory passes through the bunny’s back and ear regions for both synthesized CBFs without (left) and with (right) sparsity.

(15) (sparse GP) or (6)–(7) (regular GP). The QP safety constraint is,

$$\underbrace{\frac{\partial h_{(\cdot)}(\mathbf{x})}{\partial \mathbf{x}} \mathbf{J}(\mathbf{q})}_{L_g h_{(\cdot)}(\mathbf{x})} \underbrace{\dot{\mathbf{q}}}_{\mathbf{u}} \geq -k_0 h_{(\cdot)}(\mathbf{x}). \quad (23)$$

With the Lie derivatives, and the decision variable, $\mathbf{u} = \dot{\mathbf{q}}$, appearing linearly in the inequality (23), we can rectify the nominal control, $\mathbf{u}_{\text{nom}} = \dot{\mathbf{q}}_{\text{ref}}$, using the QP in (18).

D. Simulation Scenario : Proximal Sensing of the Bunny

In practice, robots must sense obstacles and avoid collisions online. However, training and generating the implicit surfaces offline is too slow; instead, we learn Gaussian CBF implicit surfaces directly from local data in real time.

Assumption 4. *The proximal sensor can sense samples within its field-of-view (FOV) and scanning range, and provide the point cloud locations and surface normals.*

1) *Proximal Sensing*: We equip the manipulator’s end-effector with a sensor, modeled as a 3D LiDAR using a spherical cone configuration (110° FOV, 0.8m range). A sample is detected if it is within the FOV and scanning range.

2) *Online Training & Rectification*: The rectified control law is applied for both the Gaussian CBFs (Fig. 3). The manipulator follows the reference trajectory when safe, but relaxes tracking near the bunny to avoid collisions. Gaussian CBFs are trained online as samples arrive from local sensing: a local dataset is created if ≥ 100 points are detected, where half the points are used as pseudo-inputs for sparse GP. The full GP formed 48 datasets with an average training time of 124ms, while the sparse GP formed 53 datasets with 52ms. Figure 4 also shows the time plots of the CBFs, which are always non-negative, indicating no safety violations. The average per-query inference time at the end-effector is approximately 5ms for both models.

3) *Discussion*: Figure 5 shows four simulation instances for both the synthesized CBFs. As expected, Gaussian CBF yields finer isosurfaces than the sparse variant. For example, Gaussian CBF captures the bunny’s neck and ear more precisely than the sparse Gaussian CBF. Yet the sparse model still delineates the safety boundary and prevents collisions.

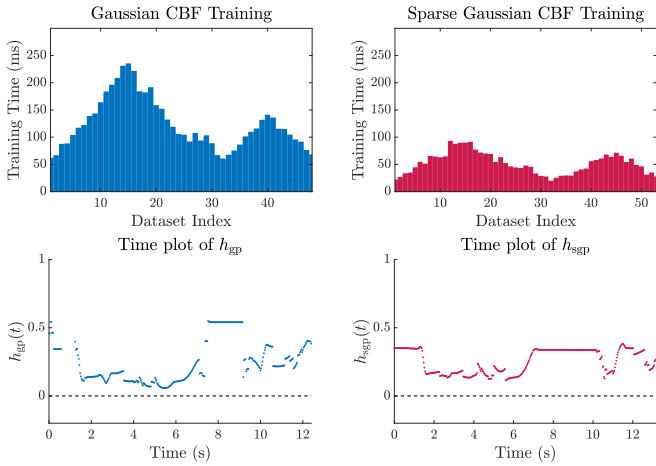


Fig. 4: The training times per local dataset are shown for both the CBFs (top). The time plots (bottom) are non-negative, indicating no safety violations occurred.

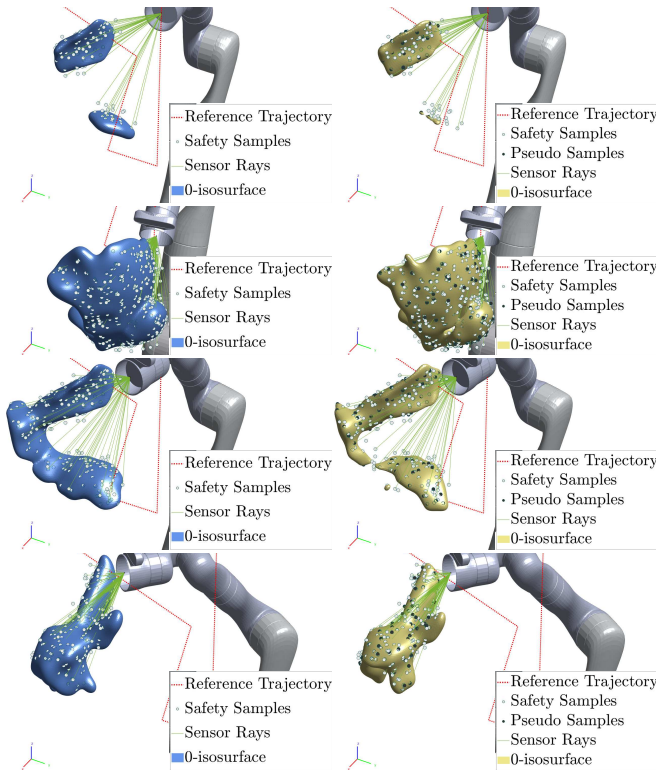


Fig. 5: (Left) Gaussian CBFs trained online from local point cloud samples (white discs). (Right) Sparse Gaussian CBFs using pseudo-inputs (black discs) with local data; both show 0-isosurfaces modeling the bunny surface.

Notice that the data-driven generation of the CBFs is not limited to convex or connected surfaces. Disconnected iso-surfaces can also be modeled using GPs.

VI. TEST CASE II : 3D QUADROTOR HARDWARE

We next validate our method on a hardware quadrotor, a challenging platform due to its nonlinear, inherently unstable dynamics. We construct implicit surfaces as Gaussian CBFs from sensed data and apply them for safe control, ensuring collision avoidance around a chair.

A. Experimental Hardware Overview

We use the Crazyflie 2.1, a 27g open-source quadrotor with a 15g payload limit. Due to its limited payload limit, it is unsuitable for any onboard depth sensing; thus, we generate Gaussian CBF surfaces offline. State estimation runs onboard aided with the help of an external infrared positioning system, which requires line-of-sight. For collision-avoidance experiments we use an IKEA ADDE chair, chosen for its large backrest gap with multiple holes, enabling state estimation even when flying beneath or behind it (otherwise the lighthouse signal is occluded). GP training (with and without sparsity) and control rectification are executed on the same ground station as discussed in V-A.

B. Chair Implicit Surface Modeling

We first model the static chair for hardware experiments. As the Crazyflie cannot carry a depth or 3D LiDAR sensor, we use virtual point cloud data. The chair’s OBJ file is processed in MeshLab to extract 10,000 point and surface normal samples. We train a GP with the Matérn kernel (a generalization of the SE kernel) using $\nu = 3/2$:

$$k_{\nu=3/2}(\mathbf{x}_i, \mathbf{x}_j) = \sigma_f^2 (1 + t) \exp(-t) + \delta_{ij} \sigma_\omega^2, \quad (24)$$

where $t = \frac{\sqrt{3} \|\mathbf{x}_i - \mathbf{x}_j\|}{l} \in \mathbb{R}_{\geq 0}$, $\mathbf{x}_i, \mathbf{x}_j \in \mathbb{R}^n$, $i, j \in \{1, \dots, N\}$, N is the number of samples, and $\sigma_f \in \mathbb{R}$, $l \in \mathbb{R}$, σ_ω are signal variance, length scale, and signal noise hyperparameters respectively. Unlike the SE kernel, which assumes infinite smoothness and may be unrealistic for physical phenomena, the Matérn class allows explicit control of smoothness and better captures real-world structures [33].

We downsample the chair’s point cloud to 2,201 points and use one-third as pseudo-inputs for the sparse GP. Training times were 4.8s for GP and 2.72s for the sparse GP, with 15 iterations of optimization defined in (13).

We use the same CBF candidate as (21). Figure 6 shows the point cloud with surface normals and the GP-based implicit surfaces used as Gaussian CBFs. Gaps smaller than 0.05m are ignored since the quadrotor cannot pass through them. The resulting Chamfer distances using (19) between $\mathcal{P}_{\text{chair}}$ and $\mathcal{P}_{\text{gp}}/\mathcal{P}_{\text{sgp}}$ are 0.078m and 0.0952m, respectively.

C. Matérn Kernel and Partial Derivatives

Here, we provide the equations of Matérn kernel’s Jacobian and Hessian for parameter $\nu = 3/2$. The first and second




Fig. 6: The chair is modeled as the safe set boundary with Gaussian CBFs: (left) point cloud and surface normals for training, (middle) GP and (right) sparse GP implicit surface.

order partial derivatives of t in (24) with respect to \mathbf{x} at a query point \mathbf{x}_* are,

$$\left. \frac{\partial t(\mathbf{x})}{\partial \mathbf{x}} \right|_{\mathbf{x}_*} = \frac{\sqrt{3}}{l} \frac{(\mathbf{x}_* - \mathbf{x})^\top}{\|\mathbf{x}_* - \mathbf{x}\|}, \quad (25)$$

$$\left. \frac{\partial^2 t(\mathbf{x})}{\partial \mathbf{x}^2} \right|_{\mathbf{x}_*} = \frac{\sqrt{3}}{l} \frac{\mathbf{I}}{\|\mathbf{x}_* - \mathbf{x}\|} - \frac{\sqrt{3}}{l} \frac{(\mathbf{x}_* - \mathbf{x})(\mathbf{x}_* - \mathbf{x})^\top}{\|\mathbf{x}_* - \mathbf{x}\|^3}, \quad (26)$$

where $\left. \frac{\partial t(\mathbf{x})}{\partial \mathbf{x}} \right|_{\mathbf{x}_*} \in \mathbb{R}^{1 \times n}$, $\left. \frac{\partial^2 t(\mathbf{x})}{\partial \mathbf{x}^2} \right|_{\mathbf{x}_*} \in \mathbb{R}^{n \times n}$, and $\mathbf{I} \in \mathbb{R}^{n \times n}$ is the identity matrix. The kernel's Jacobian and Hessian in (24) with respect to \mathbf{x} at a query point \mathbf{x}_* are,

$$\left. \frac{\partial \mathbf{k}(\mathbf{x})}{\partial \mathbf{x}} \right|_{\mathbf{x}_*} = -\sigma_f^2 t \exp(-t) \left. \frac{\partial t(\mathbf{x})}{\partial \mathbf{x}} \right|_{\mathbf{x}_*}, \quad (27)$$

$$\begin{aligned} \left. \frac{\partial^2 \mathbf{k}(\mathbf{x})}{\partial \mathbf{x}^2} \right|_{\mathbf{x}_*} &= -\sigma_f^2 t \exp(-t) \left. \frac{\partial^2 t(\mathbf{x})}{\partial \mathbf{x}^2} \right|_{\mathbf{x}_*} \\ &\quad + \sigma_f^2 (t-1) \exp(-t) \left. \frac{\partial t(\mathbf{x})}{\partial \mathbf{x}} \right|_{\mathbf{x}_*} \left. \frac{\partial t(\mathbf{x})}{\partial \mathbf{x}} \right|_{\mathbf{x}_*}^\top, \end{aligned} \quad (28)$$

where (25) and (26) are substituted.

D. Safe Control Synthesis for Quadrotor

We use thrust-attitude setpoints to control the Crazyflie [34]. We use the same second-order integrator and rectification process as [35]. The relative degree for the second-order integrator is $\rho = 2$. Hence, the associated Lie derivatives for the Gaussian CBF (with and without sparsity) in (21) are,

$$\begin{aligned} L_f h_{(\cdot)}(\mathbf{x}) &= (\nabla \mu(\mathbf{x}) - \nabla \sigma^2(\mathbf{x}))^\top f(\mathbf{x}), \\ L_f^2 h_{(\cdot)}(\mathbf{x}) &= \alpha^\top f(\mathbf{x}) + \beta^\top f(\mathbf{x}), \\ L_g L_f h_{(\cdot)}(\mathbf{x}) &= \alpha^\top g(\mathbf{x}) + \beta^\top g(\mathbf{x}), \end{aligned} \quad (29)$$

where $\nabla \mu(\mathbf{x}) = \frac{\partial \mu(\mathbf{x})}{\partial \mathbf{x}}$ and $\nabla \sigma^2(\mathbf{x}) = \frac{\partial \sigma^2(\mathbf{x})}{\partial \mathbf{x}}$ are the gradients of GP mean and variance (with and without sparsity), $\nabla f(\mathbf{x}) = \frac{\partial f(\mathbf{x})}{\partial \mathbf{x}}$ is the Jacobian of $f(\mathbf{x})$, $\alpha = f(\mathbf{x})^\top (\mathbf{H}_\mu(\mathbf{x}) - \mathbf{H}_{\sigma^2}(\mathbf{x}))$, and $\beta = (\nabla \mu(\mathbf{x}) - \nabla \sigma^2(\mathbf{x}))^\top \cdot \nabla f(\mathbf{x})$. The corresponding Hessians of both Gaussian and sparse Gaussian CBFs, $\mathbf{H}_\mu(\mathbf{x})$ and $\mathbf{H}_{\sigma^2}(\mathbf{x})$, are:

$$\begin{aligned} \mathbf{H}_{\mu_{\text{gp/sgp}}}(\mathbf{x}) &= \sum_i a_{\text{gp/sgp}}^i \frac{\partial^2 \mathbf{k}_{(i)}(\mathbf{x})}{\partial \mathbf{x}^2}, \\ \mathbf{H}_{\sigma^2_{\text{gp/sgp}}}(\mathbf{x}) &= -2\mathbf{A}_{\text{gp/sgp}} - 2 \sum_i b_{\text{gp/sgp}}^i(\mathbf{x}) \frac{\partial^2 \mathbf{k}_{(i)}(\mathbf{x})}{\partial \mathbf{x}^2}, \end{aligned}$$

where a_{gp}^i and a_{sgp}^i are the i^{th} entries of $\mathbf{y}_N^\top \bar{\mathbf{K}}^{-1} \in \mathbb{R}^{1 \times N}$ and $\mathbf{y}_N^\top (\mathbf{Q}_M^{-1} \mathbf{K}_{MN} (\boldsymbol{\Lambda}_N + \sigma_y^2 \mathbf{I}_N)^{-1})^\top \in \mathbb{R}^{1 \times M}$, respectively. $b_{\text{gp}}^i(\mathbf{x})$ and $b_{\text{sgp}}^i(\mathbf{x})$ are the i^{th} entries of $\mathbf{k}(\mathbf{x})^\top \bar{\mathbf{K}}^{-1} \in \mathbb{R}^{1 \times N}$ and $\mathbf{k}_M(\mathbf{x})^\top (\mathbf{K}_M^{-1} - \mathbf{Q}_M^{-1}) \in \mathbb{R}^{1 \times M}$, respectively. $\mathbf{A}_{\text{gp}} = \nabla \mathbf{k}(\mathbf{x}) \bar{\mathbf{K}}^{-1} \nabla \mathbf{k}(\mathbf{x})^\top \in \mathbb{R}^{N \times N}$, $\mathbf{A}_{\text{sgp}} = \nabla \mathbf{k}_M(\mathbf{x}) (\mathbf{K}_M^{-1} - \mathbf{Q}_M^{-1}) \nabla \mathbf{k}_M(\mathbf{x})^\top \in \mathbb{R}^{M \times M}$, $\nabla \mathbf{k}(\mathbf{x}) = \frac{\partial \mathbf{k}(\mathbf{x})}{\partial \mathbf{x}} \in \mathbb{R}^{n \times N}$, $\nabla \mathbf{k}_M(\mathbf{x}) = \frac{\partial \mathbf{k}_M(\mathbf{x})}{\partial \mathbf{x}} \in \mathbb{R}^{n \times M}$, and $\frac{\partial^2 \mathbf{k}_{(i)}(\mathbf{x})}{\partial \mathbf{x}^2}$ is the partial derivative of (16) with respect to \mathbf{x} . Since the relative degree is 2, we adopt higher-order CBFs, in particular the exponential CBF [1], [36]. Using the Lie derivatives in (29), we can design a QP of the form (18) to yield rectified control inputs that guarantee forward invariance of the safe set.

E. Scenario: Safe Autonomous Navigation

In this scenario, the Crazyflie was commanded to follow unsafe reference trajectories that passed through the chair. The goal was to track the reference while ensuring safety by preventing collisions. We synthesized the CBFs across three separate runs, each using a different unsafe reference trajectory. The references were generated by specifying a final desired position, with the thrust-attitude controller computing the desired setpoint at every sampling step. Gaussian CBF rectification was applied at 50 Hz.

We evaluated three unsafe reference trajectories for the Crazyflie: (i) flying diagonally through the chair legs, (ii) flying straight through the front-left and rear-left legs, and (iii) flying directly through the chair's headrest. These scenarios enforced reference paths that would naturally result in collisions, providing diverse and challenging conditions for assessing the effectiveness of Gaussian and sparse Gaussian CBFs in ensuring safe navigation.

Figure 7 shows the flights, CBF temporal plots, and average per-iteration compute times. In all three runs, the quadrotor relaxed trajectory tracking when near the chair to maintain safety, avoiding collisions. The h_{gp} and h_{sgp} traces are nonnegative, and the average inference time was 11ms for Gaussian CBF and about 6.5ms for the sparse variant.

VII. CONCLUSION

In summary, we used Gaussian CBFs to construct safe implicit surfaces, where the surface of a volumetric object is defined as the boundary of the safe set. We presented sparse Gaussian CBFs to take advantage of reduced computational complexity from $\mathcal{O}(N^3)$ to $\mathcal{O}(M^2N)$ (N is number of training points, M is number of pseudo-inputs). We presented two robotic test cases, first, in simulation for a 7-DOF manipulator, and second, for a hardware quadrotor. For the first study, we estimated the Stanford bunny's safety boundary as Gaussian CBFs (with and without sparsity), where the safe implicit surface was the bunny's surface. We demonstrated safe proximal sensing without collision on the manipulator in simulation. For the quadrotor, we demonstrated safe 3D navigation on the Crazyflie hardware around a physical IKEA ADDE chair, using its OBJ model to extract point cloud and surface normal data. Even when reference trajectories were unsafe and would have intersected the chair, the Gaussian CBF ensured collision-free autonomous flight in all trials.

REFERENCES

- [1] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: Theory and applications," in *IEEE European control conference*, pp. 3420–3431, 2019.
- [2] R. Grandia, A. J. Taylor, A. D. Ames, and M. Hutter, "Multi-layered safety for legged robots via control barrier functions and model predictive control," in *IEEE International Conference on Robotics and Automation*, pp. 8352–8358, 2021.
- [3] S. Teng, Y. Gong, J. W. Grizzle, and M. Ghaffari, "Toward safety-aware informative motion planning for legged robots," *arXiv preprint arXiv:2103.14252*, 2021.
- [4] M. Khan, M. Zafar, and A. Chatterjee, "Barrier functions in cascaded controller: Safe quadrotor control," in *2020 American Control Conference (ACC)*, pp. 1737–1742, 2020.

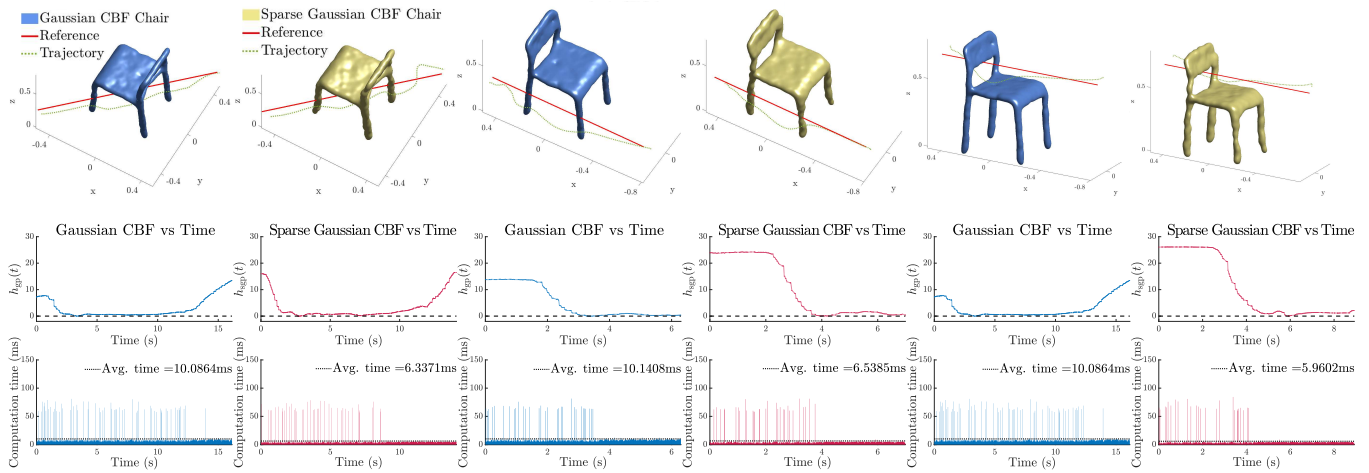


Fig. 7: Gaussian (odd columns) and sparse Gaussian (even columns) CBFs in three unsafe quadrotor scenarios: (i) diagonal through chair legs, (ii) straight through front–rear legs, (iii) through the headrest. Each column shows the trajectories, CBF time plots, and compute times; in all cases safety was maintained, with average inference of 11 ms (GP) and 6.5 ms (sparse).

- [5] A. Singletary, A. Swann, Y. Chen, and A. D. Ames, “Onboard safety guarantees for racing drones: High-speed geofencing with control barrier functions,” *IEEE Robotics and Automation Letters*, 2022.
- [6] M. Rauscher, M. Kimmel, and S. Hirche, “Constrained robot control using control barrier functions,” in *IEEE International Conference on Intelligent Robots and Systems*, pp. 279–285, 2016.
- [7] A. Robey, L. Lindemann, S. Tu, and N. Matni, “Learning robust hybrid control barrier functions for uncertain systems,” *IFAC-PapersOnLine*, vol. 54, no. 5, pp. 1–6, 2021.
- [8] N. Gaby, F. Zhang, and X. Ye, “Lyapunov-net: A deep neural network architecture for lyapunov function approximation,” *arXiv preprint arXiv:2109.13359*, 2021.
- [9] H. Tsukamoto and S.-J. Chung, “Neural contraction metrics for robust estimation and control: A convex optimization approach,” *IEEE Control Systems Letters*, vol. 5, no. 1, pp. 211–216, 2020.
- [10] N. Csomay-Shanklin, R. K. Cosner, M. Dai, A. J. Taylor, and A. D. Ames, “Episodic learning for safe bipedal locomotion with control barrier functions and projection-to-state safety,” in *Learning for Dynamics and Control*, pp. 1041–1053, PMLR, 2021.
- [11] M. Srinivasan, A. Dabholkar, S. Coogan, and P. A. Vela, “Synthesis of control barrier functions using a supervised machine learning approach,” in *IEEE International Conference on Intelligent Robots and Systems*, pp. 7139–7145, 2020.
- [12] B. Dai, R. Khorrambakht, P. Krishnamurthy, and F. Khorrami, “Sailing through point clouds: Safe navigation using point cloud based control barrier functions,” *IEEE Robotics and Automation Letters*, vol. 9, no. 9, pp. 7731–7738, 2024.
- [13] M. A. Khan, T. Ibuki, and A. Chatterjee, “Gaussian control barrier functions: Non-parametric paradigm to safety,” *IEEE Access*, vol. 10, pp. 99823–99836, 2022.
- [14] W. Hashimoto, K. Hashimoto, A. Wachi, X. Shen, M. Kishida, and S. Takai, “Bayesian meta-learning on control barrier functions with data from on-board sensors,” *arXiv preprint arXiv:2308.05306*, 2023.
- [15] S. Keyumarsi, M. W. S. Atman, and A. Gusrialdi, “Lidar-based online control barrier function synthesis for safe navigation in unknown environments,” *IEEE Robotics and Automation Letters*, vol. 9, no. 2, pp. 1043–1050, 2024.
- [16] H. Oleynikova, A. Millane, Z. Taylor, E. Galceran, J. Nieto, and R. Siegwart, “Signed distance fields: A natural representation for both mapping and planning,” in *RSS 2016 workshop: geometry and beyond-representations, physics, and scene understanding for robotics*, University of Michigan, 2016.
- [17] K. Long, C. Qian, J. Cortés, and N. Atanasov, “Learning barrier functions with memory for robust safe navigation,” *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 4931–4938, 2021.
- [18] O. Williams and A. Fitzgibbon, “Gaussian process implicit surfaces,” in *Gaussian Processes in Practice*, 2006.
- [19] L. Wu, K. M. B. Lee, L. Liu, and T. Vidal-Calleja, “Faithful euclidean distance field from log-gaussian process implicit surfaces,” *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 2461–2468, 2021.
- [20] B. Lee, C. Zhang, Z. Huang, and D. D. Lee, “Online continuous mapping using gaussian process implicit surfaces,” in *IEEE International Conference on Robotics and Automation*, pp. 6884–6890, 2019.
- [21] S. Dragiev, M. Toussaint, and M. Gienger, “Gaussian process implicit surfaces for shape estimation and grasping,” in *IEEE International Conference on Robotics and Automation*, pp. 2845–2850, 2011.
- [22] G. Z. Gandler, C. H. Ek, M. Björkman, R. Stolkin, and Y. Bekiroglu, “Object shape estimation and modeling, based on sparse gaussian process implicit surfaces, combining visual data and tactile exploration,” *Robotics and Autonomous Systems*, vol. 126, p. 103433, 2020.
- [23] 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, 2016.
- [24] C. K. Williams and C. E. Rasmussen, *Gaussian processes for machine learning*, vol. 2. MIT press Cambridge, MA, 2006.
- [25] S. Osher, R. Fedkiw, and K. Piechor, “Level set methods and dynamic implicit surfaces,” *Appl. Mech. Rev.*, vol. 57, no. 3, pp. B15–B15, 2004.
- [26] M. P. Do Carmo, *Differential geometry of curves and surfaces: revised and updated second edition*. Courier Dover Publications, 2016.
- [27] J. e. a. Gawlikowski, “A survey of uncertainty in deep neural networks,” *arXiv preprint arXiv:2107.03342*, 2021.
- [28] M. Khan, T. Ibuki, and A. Chatterjee, “Safety uncertainty in control barrier functions using gaussian processes,” in *IEEE International Conference on Robotics and Automation*, pp. 6003–6009, 2021.
- [29] T. P. Minka, “Expectation propagation for approximate bayesian inference,” in *Proceedings of the 17th conference on Uncertainty in artificial intelligence*, pp. 362–369, 2001.
- [30] E. Snelson and Z. Ghahramani, “Sparse gaussian processes using pseudo-inputs,” in *Advances in Neural Information Processing Systems*, pp. 1257–1264, MIT Press, 2006.
- [31] M. Titsias, “Variational learning of inducing variables in sparse gaussian processes,” in *Artificial Intelligence and Statistics*, 2009.
- [32] C. E. Rasmussen and H. Nickisch, “Gaussian processes for machine learning (gpml) toolbox,” *Journal of machine learning research*, vol. 11, no. Nov, pp. 3011–3015, 2010.
- [33] M. L. Stein, *Interpolation of spatial data: some theory for kriging*. Springer Science & Business Media, 1999.
- [34] D. W. Mellinger, “Trajectory generation and control for quadrotors,” 2012.
- [35] B. Xu and K. Sreenath, “Safe teleoperation of dynamic uavs through control barrier functions,” in *2018 IEEE International Conference on Robotics and Automation*, pp. 7848–7855, 2018.
- [36] Q. Nguyen and K. Sreenath, “Exponential control barrier functions for enforcing high relative-degree safety-critical constraints,” in *American Control Conference*, pp. 322–328, IEEE, 2016.