

Autonomous Exploration for Shape Reconstruction and Measurement via Informative Contact-Guided Planning

Feiyu Zhao, Chenxi Xiao*

Abstract—Coordinate Measuring Machines (CMMs) are widely used for high-precision inspection of industrial parts, particularly in scenarios where visual systems are ineffective or cost-prohibitive. However, conventional CMMs rely on CAD model priors and user-defined probing paths, which limit their applicability and efficiency in measuring freeform parts. To overcome these limitations, we present a fully autonomous, CAD model-free, tactile-based framework that enables dense 3D shape reconstruction to facilitate subsequent measurements. Our approach leverages a dual Gaussian Process Implicit Surface architecture, termed Exploration-Reconstruction GPIS (ER-GPIS), which enables both high-fidelity shape reconstruction and uncertainty estimation on the object’s surface. A hybrid exploration motion planner is then employed to adaptively sample surface geometries by integrating local surface exploration, global exploration, and contact recovery policies for robust shape estimation. Extensive real-world experiments demonstrate that the proposed method effectively reconstructs object geometries across diverse shapes, highlighting its ability to autonomously reconstruct and measure both surfaces and internal features without relying on CAD model priors. Our project webpage is available at <https://aesrm.github.io/>

Index Terms—Force and Tactile Sensing, Reactive and Sensor-Based Planning, Planning under Uncertainty

I. INTRODUCTION

COORDINATE Measuring Machines (CMMs) have long been the golden standard of geometric inspection in modern manufacturing [1]. By physically probing selected points on a part’s surface, CMMs verify dimensional tolerances with exceptional accuracy and repeatability. These systems operate by executing predefined measurement routines derived from CAD models, enabling reliable verification of part geometries, hole positions, flatness, and surface alignments [2]. As a result, CMMs are indispensable in precision-critical domains such as robotics [3], where adherence to dimensional specifications is essential for ensuring performance, interoperability, and compliance to industrial regulations [4].

However, traditional CMM workflows are limited by their reliance on pre-defined paths generated from CAD models [5]. Without access to high quality CAD model, CMMs cannot autonomously determine what or where to measure, rendering additional human efforts for a wide range of scenarios such

Manuscript received: July, 3, 2025; Revised October, 21, 2025; Accepted December, 2, 2025.

This paper was recommended for publication by Editor Ashis Banerjee upon evaluation of the Associate Editor and Reviewers’ comments.

*This research was supported by the Natural Science Foundation of Shanghai (Grant No. 25ZR1402370).

All authors are with the School of Information Science and Technology, ShanghaiTech University, Shanghai, China (*corresponding author, xiaochx@shanghaitech.edu.cn).

Digital Object Identifier (DOI): see top of this page.

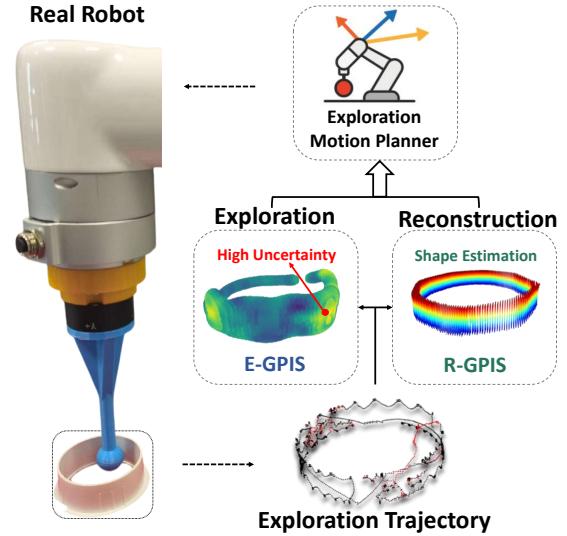


Fig. 1: Overview of the proposed tactile exploration framework. The E-GPIS and R-GPIS models jointly estimate local uncertainty and reconstruct the surface shape. Both outputs guide the Exploration Motion Planner in generating informative actions. The real robot executes these actions, collects contact data, and iteratively refines the reconstruction.

as reverse engineering of legacy components, inspection of worn or damaged parts with missing geometry [6]. Moreover, CMMs are designed for sparse, point-wise verification rather than dense data acquisition. As such, they are fundamentally incapable of reconstructing fine-grained shape [7]. This limitation prevents them from capturing geometric details across freeform shapes, and therefore the usage of CMMs are mostly limited in manual paradigms where human efforts are required.

To bridge this gap, we present a fully autonomous tactile-based framework that models such contact metrology as a joint active exploration and reconstruction process. Instead of relying on predefined paths, our system incrementally probes an unknown object, decides where to explore next, and reconstructs and measures the object surface from contact observations alone (as illustrated in Fig. 1). The core of our framework is a novel dual-Gaussian Process (GP) architecture **ER-GPIS**, which explicitly separates surface estimation from uncertainty modeling in the context of tactile exploration. Unlike previous approaches that use a single Gaussian Process to achieve both objectives [8], our decoupled design improves robustness to intrinsically sparse tactile point clouds while preserving uncertainty estimation capability (Sec. IV).

The proposed ER-GPIS framework relies on incrementally sampling informative contact points from the object surface. To determine the target locations for sampling during explo-

ration, we introduce a hybrid motion planner (Sec. III-B). This planner interactively selects a contact sampling strategy from three policies: local surface sliding to add contact points, global redirection to maximize surface coverage, and contact recovery to maintain reliable interaction. These policies are applied interleavingly to actively guide the end-effector toward informative regions, enabling ER-GPIS to incrementally acquire high-quality contact observations essential for accurate shape reconstruction and uncertainty modeling.

To evaluate the effectiveness of the proposed pipeline, real-world experiments were conducted on a diverse set of test objects, involving both internal cavities and exterior surfaces. The results demonstrate that our system accurately reconstructs 3D geometry using contact data alone and outperforms previous tactile exploration baselines in terms of completeness, accuracy, and robustness. In summary, the key contributions of this work are:

- A CAD-free, vision-free tactile exploration framework that extends contact-based metrology with autonomous, dense 3D shape reconstruction capabilities;
- A dual-Gaussian Process modeling approach (ER-GPIS) that concurrently achieves uncertainty estimation and accurate reconstruction;
- A hybrid tactile exploration planner that achieves both exploration efficiency and robustness;
- Comprehensive experiments for revealing usability and ablation studies.

II. RELATED WORKS

A. Coordinate Measuring Machines

Coordinate Measuring Machines measure object geometries and are generally classified as contact or non-contact systems. Non-contact CMMs employ sensors such as cameras, offering fast and efficient measurements, but they are costly and struggle with reflective or occluded features [9]. Contact CMMs use mechanical probes to achieve precise, low-cost measurements that are robust to surface variations [10]. Such systems are more commonly used, but rely on manually defined trajectories [11], and often result in sparse data. To automate this process, recent research has leveraged CAD models for feature identification and motion planning [12], optimizing probing paths for efficiency and collision avoidance [13], [14]. However, these approaches are not suitable when CAD data are unavailable or inaccurate [2], leaving manual measurement still the most common approach. To address this limitation, we propose a CAD-independent tactile exploration framework that reconstructs object surfaces and enables adaptive measurement through real-time contact feedback in uncertain environments.

B. 3D Reconstruction

Reconstructing 3D shape has been one central focus in robot perception and computer vision. Traditional methods rely on explicit representations such as point clouds or voxel grids [15], [16], which require a large number of points or camera views to achieve fine detail. Recent research has shifted toward implicit representations, such as NeRF [17] and 3D

Gaussian Splatting [18] can generate photorealistic, novel-view renderings. However, since these methods were originally designed for visual rendering. Recent research has combined touch, such as creating multimodal representations [19], and refining contact relationship to enhance geometric fidelity [20], [21]. Signed Distance Functions (SDFs) offer another powerful form of implicit representation [22]. They are continuous, memory-efficient, and guarantee watertight surfaces, making them particularly suitable for applications involving contact and physical interaction. Consequently, SDFs have become a common choice for representing shape information derived from tactile sensing, such as via GP regressors [8], [23], [24].

C. Active Tactile Exploration

Active tactile exploration is a process in which a robot leverages tactile sensing to proactively probe and infer object properties. This exploration serves multiple purposes. First, it enables the robot to inspect object geometry in regions where visual data are unavailable due to occlusions [23], [25], transparency [26], or reflectivity [20], [27], providing an alternative means to recover object shape. Beyond geometric understanding, tactile exploration allows the acquisition of additional object properties, such as estimating the center of mass [28], assessing hardness [29], and facilitating manipulation skill learning [30], thereby extending perception to multimodal information beyond vision. This ability to actively acquire information is essential for robots to develop a comprehensive understanding of their surroundings.

D. Uncertainty-based Tactile Exploration for Reconstruction

One research direction for reconstruction leveraging active tactile exploration focuses on determining the next-best probing location to maximize information gain. This process first requires quantifying uncertainty about the object's shape. Common approaches employ GP [31], [32] and ergodic metrics [33], [34] to represent and evaluate such uncertainty. Once an uncertainty map is established, various planning strategies can be applied to guide the exploration. These range from greedy methods [35], which myopically select the next best touch point, to short-horizon planners based on optimization [8], [36], and long-horizon planners that optimize an entire sequence of future actions [37]. More recently, reinforcement learning has been explored to predict next-best actions [38]. Despite these advancements, most existing studies have been validated only on very limited set of objects. Our work extends these methods toward practical applications and demonstrates improved robustness through comprehensive experiments.

III. METHODOLOGY

To enable autonomous shape reconstruction and measurement, we present a unified framework via active tactile exploration. The proposed approach integrates two main components: (1) a dual-Gaussian Process model, **ER-GPIS**, which simultaneously estimates object shape and associated uncertainty, and (2) a hybrid Exploration Motion Planner that leverages the estimated shape and uncertainty to guide exploration movements. The overall pipeline of the proposed system is illustrated in Fig. 2.

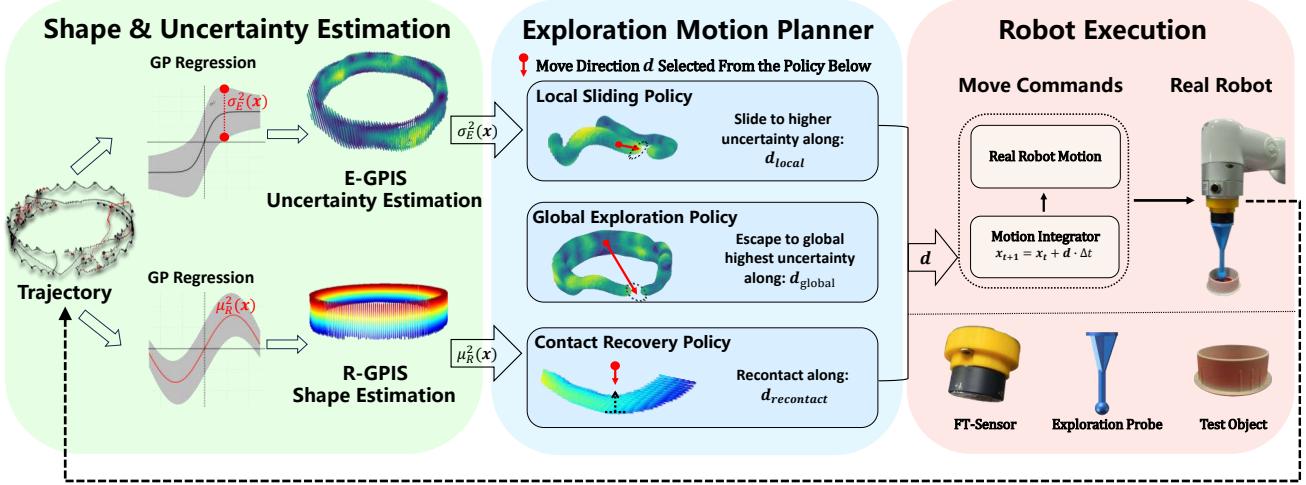


Fig. 2: Overview of the proposed technical pipeline. The system consists of three interconnected modules: **(1) Shape & Uncertainty Estimation**: Two Gaussian Process models are employed. *E-GPIS* estimates uncertainty $\sigma_E^2(x)$ to guide exploration, while the *R-GPIS* estimates $\mu_R^2(x)$ for shape reconstruction. **(2) Exploration Motion Planner**: The motion direction d is calculated from either of three policies to guide exploration: **Local Sliding Policy** slides toward higher local uncertainty d_{local} ; **Global Exploration Policy** redirects toward highest global uncertainty along d_{global} ; **Contact Recovery Policy** restores contact along surface normals $d_{\text{recontact}}$. **(3) Robot Execution**: The robot system includes an FT-sensor, and an exploration probe. The selected direction d is executed on this robot system via a low-level controller.

A. Shape and Uncertainty Estimation via ER-GPIS

Unlike visual observations with dense point clouds, tactile sensing provides inherently sparse contact points, resulting in limited prior information for surface reconstruction. To address this sparsity, the GPIS algorithm [39] has previously been adopted, which regresses the Signed Distance Function (SDF) by leveraging spatial smoothness assumptions encoded in kernel functions. This provides prior that enables inference of the global object shape from only a few contact locations. Specifically, the GPIS model is characterized by a GP prior:

$$f(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')) \quad (1)$$

where $m(\mathbf{x})$ is a prior mean function, and $k(\mathbf{x}, \mathbf{x}') : \mathbb{R}^3 \times \mathbb{R}^3 \rightarrow \mathbb{R}$ is a covariance (kernel) function. Here, \mathbf{x} and \mathbf{x}' denote input locations in the domain \mathbb{R}^3 . The posterior distribution $f(\mathbf{x})$ is given by:

$$p(f(\mathbf{x}) | \mathbf{X}, \mathbf{y}) \sim \mathcal{N}(\mu(\mathbf{x}), \sigma^2(\mathbf{x})) \quad (2)$$

with

$$\mu(\mathbf{x}) = m(\mathbf{x}) + \mathbf{k}^T(\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1}(\mathbf{y} - m(\mathbf{X})) \quad (3)$$

$$\sigma^2(\mathbf{x}) = k(\mathbf{x}, \mathbf{x}) - \mathbf{k}^T(\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1}\mathbf{k} \quad (4)$$

Here, $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_N]^T$ and $\mathbf{y} = [y_1, \dots, y_N]^T$ are the observed inputs and outputs. The kernel matrix $\mathbf{K} \in \mathbb{R}^{N \times N}$ has entries $\mathbf{K}_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$, and $\mathbf{k} = [k(\mathbf{x}, \mathbf{x}_1), \dots, k(\mathbf{x}, \mathbf{x}_N)]^T$ is the kernel vector between the test point \mathbf{x} and training inputs. Conditioned on the prior mean function $m(\mathbf{x})$ that encodes function values in regions without observations, the GPIS model implicitly represents the object surface as the zero level set of the posterior mean:

$$\mathcal{S} = \{\mathbf{x} \in \mathbb{R}^3 \mid \mu(\mathbf{x}) = 0\}. \quad (5)$$

The posterior variance $\sigma^2(\mathbf{x})$ quantifies the uncertainty associated with the prediction.

However, faithfully regressing SDF values requires information from both inside and outside the object, while interior regions are physically inaccessible to touch. A common practice is to augment the training set with auxiliary off-surface samples treated as pseudo-observations with prescribed SDF targets [39]; but it has been reported that placing such off-surface points in inaccurate locations requires caution to avoid biasing the model [35]. Driess *et al.* [8] addressed this by setting the prior mean to $m = 1$, encoding the assumption that most of the space is empty and thus avoiding the need for artificial points and normals. Nevertheless, we found this approach is still sensitive to hyperparameters and prone to reconstruction instability. More recently, uGPIS [40] achieved robust shape reconstruction without internal points or surface normals, but their method lacks reliable uncertainty estimation, which is essential for exploration guidance.

To overcome these limitations, we propose the Exploration-Reconstruction GPIS (ER-GPIS) framework, which addresses the aforementioned issues by achieving shape reconstruction and uncertainty modeling using two separate Gaussian Process models. This design allows ER-GPIS to leverage the robustness of uGPIS while overcoming the limitations of both uGPIS and conventional GPIS: it eliminates the need for normal-based off-surface points and simultaneously provides reliable uncertainty estimation. Specifically, ER-GPIS consists of two components: Exploration GPIS (E-GPIS), which computes uncertainty for motion planning, and Reconstruction GPIS (R-GPIS) for robust surface reconstruction, as detailed below.

1) *Exploration GPIS (E-GPIS)*: To estimate uncertainty that highlights informative regions for next-best touch, we construct an E-GPIS model that predicts spatial uncertainty using a dataset of contact and dummy points overlapping with

the probe, denoted as \mathcal{D}_E . Specifically, $\mathcal{D}_E = \mathcal{D}_c \cup \mathcal{D}_d$, where \mathcal{D}_c contains the observed contact locations with 0 SDF values, and \mathcal{D}_d consists of known void locations without contact (with SDF value 1). The E-GPIS computes the predictive variance $\sigma_E^2(\mathbf{x})$ at unvisited locations via Eq. 2, generating a spatial uncertainty map that guides the exploration planner (see Sec. III-B). We set the prior mean function $m(\mathbf{x}) = 1$, following [8], which encodes the assumption that regions far from detected contacts are empty.

2) *Reconstruction GPIS (R-GPIS)*: To robustly reconstruct object shape, ER-GPIS incorporates a second Gaussian Process model to regress occupancy probability for shape estimation, following [40]. Our method overcomes uGPIS's inability to accurately estimate uncertainty. To achieve this, a global dataset $\mathcal{D}_R = \mathcal{D}_c \cup \mathcal{D}_a$ is constructed, containing contact points \mathcal{D}_c labeled with occupancy probability 1, along with additional auxiliary points \mathcal{D}_a uniformly sampled in void regions labeled as 0. The R-GPIS model is then trained on this binary occupancy dataset \mathcal{D}_R , producing a posterior mean field $\mu_R(\mathbf{x})$ representing the predicted occupancy probability. For R-GPIS, we set the prior mean function $m(\mathbf{x}) = 0$ to indicate zero occupancy when far from known contact points.

To obtain the reconstructed object shape $\mathcal{S}_{\text{recon}}$ from the posterior mean field, we apply a bounded threshold based on the predicted mean values $\mu_R(\mathbf{x})$ at the contact points. Since all contact points are assumed to lie on the object surface, we first compute the $\mu_R(\mathbf{x})$ for all points in the contact set \mathcal{D}_c , and define the lower and upper bounds $[\mu_{\min}, \mu_{\max}]$ as the 5th and 95th percentiles of $\mu_R(\mathbf{x})$ for all $\mathbf{x} \in \mathcal{D}_c$, respectively. The final reconstruction is then defined as:

$$\mathcal{S}_{\text{recon}} = \{\mathbf{x} \in \mathbb{R}^3 \mid \mu_{\min} \leq \mu_R(\mathbf{x}) \leq \mu_{\max}\} \quad (6)$$

B. Exploration Motion Planner

To plan motion for tactile exploration, we propose a hybrid motion planner that autonomously selects among three motion policies based on the current context:

- 1) *Local Sliding Policy* guides the robot to maximize the utility of collected data via sliding motion;
- 2) *Global Exploration Policy* prevents the robot from getting trapped in local optima by redirecting it toward regions of higher uncertainty;
- 3) *Contact Recovery Policy* re-establishes contact with the surface when contact is lost.

The above three policies constitute a state machine. This section details the implementation of each policy as follows.

1) *Local Sliding Policy*: Compared to the discrete sampling used in conventional CMMs, sliding motion improves sampling efficiency by generating a continuous stream of contact points. This sliding motion is implemented using a hybrid force-position control scheme that maintains compliant contact while guiding movement toward regions of high predicted uncertainty. Specifically, the hybrid controller slides the probe along the tangential direction of the object surface to acquire informative data, while simultaneously applying force along the surface normal to maintain a user-specified target contact force f_{target} . To enable this, we first estimate the surface normal \mathbf{n}_f from the sensed contact force $\mathbf{f} = \frac{\mathbf{f}}{\|\mathbf{f}\|}$. The tangential

direction at the contact point is then defined using a projection onto the local tangent plane, which constrains the sliding direction using a projection matrix $\mathbf{P}_{\mathbf{n}_f}$:

$$\mathbf{P}_{\mathbf{n}_f}(\mathbf{x}) = \mathbf{I} - \mathbf{n}_f \mathbf{n}_f^T \quad (7)$$

Next, to guide the robot toward unexplored regions, we define a utility function that encourages movement toward areas of high uncertainty while penalizing actions that exceed the workspace boundaries. This function is defined as:

$$a(\mathbf{x}) = \sigma_E^2(\mathbf{x}) - \psi_z(x_z) \quad (8)$$

Here, $\sigma_E^2(\mathbf{x})$ represents the uncertainty predicted by E-GPIS. The term $\psi_z(x_z)$ aims to constrain the motion within the workspace. In our implementation, we regulate the vertical position of the probe using the following equation:

$$\psi_z(x_z) = \begin{cases} \lambda(x_z - z_{\max})^2, & x_z > z_{\max} \\ \lambda(z_{\min} - x_z)^2, & x_z < z_{\min} \\ 0, & \text{otherwise} \end{cases} \quad (9)$$

To maximize the accumulated utility function $\int a(x)$ along the sliding motion path, robot actions are determined by combining two components: a lateral term $\frac{\mathbf{P}_{\mathbf{n}_f}(\mathbf{x}) \nabla a(\mathbf{x})}{\|\mathbf{P}_{\mathbf{n}_f}(\mathbf{x}) \nabla a(\mathbf{x})\|}$ which drives the end-effector toward informative regions, and a normal term aligned with \mathbf{n}_f , which maintains compliant contact with the surface. These components are combined to produce the final direction for local sliding motion:

$$\mathbf{d}_{\text{local}} = \alpha \cdot \frac{\mathbf{P}_{\mathbf{n}_f}(\mathbf{x}) \nabla a(\mathbf{x})}{\|\mathbf{P}_{\mathbf{n}_f}(\mathbf{x}) \nabla a(\mathbf{x})\|} + \gamma \cdot (f_c - f_{\text{target}}) \cdot \mathbf{n}_f \quad (10)$$

where $f_c = \mathbf{f} \cdot \mathbf{n}_f$ is the contact force projected along the estimated surface normal, and α, γ are weights that balance the influence of the exploration and compliant interaction components, respectively.

2) *Global Exploration Policy*: The utility function (Eq. 8) is non-convex, and therefore the control law (Eq. 10) may become trapped in local optima. To address this, *Global Exploration Policy* is triggered when the maximum displacement $\max(\|\mathbf{x}_i - \mathbf{x}_0\|)$ from the initial probe position \mathbf{x}_0 (at the beginning of the current iteration) falls below a threshold δ .

During global exploration, the probe is temporarily detached from the object surface and redirected to a new target location on the reconstructed manifold by R-GPIS: $\mathbf{x} \in \mathcal{S}_{\text{recon}}$. This location is selected as the one with the highest utility value:

$$\mathbf{x}_{\text{global}} = \arg \max_{\mathbf{x} \in \mathcal{S}_{\text{recon}}} a(\mathbf{x}) \quad (11)$$

Then, the exploratory direction is computed as an unit vector from the current probe position \mathbf{x}_t to the selected target:

$$\mathbf{d}_{\text{global}} = \frac{\mathbf{x}_{\text{global}} - \mathbf{x}_t}{\|\mathbf{x}_{\text{global}} - \mathbf{x}_t\|} \quad (12)$$

Note that in order to move to $\mathbf{x}_{\text{global}}$, the probe follows different strategies depending on the type of object geometry. For internal cavities, it directly moves along $\mathbf{d}_{\text{global}}$ until the next contact is made. For exterior surfaces, the probe first rises to a collision-free height z_{\max} , then moves laterally along $\mathbf{d}_{\text{global}}$ to a position above the $\mathbf{x}_{\text{global}}$, and finally descends vertically until contact is re-established. This strategy enables the robot to escape local optima in exploration and resume collection in more informative regions, thereby improving overall efficiency for surface coverage.

3) *Contact Recovery Policy*: When sliding along object surfaces, the sliding policy may fail to maintain stable contact, especially near high-curvature or discontinued surface regions. To address this, we introduce a *Contact Recovery Policy* that helps the probe re-establish contact after detachment.

The recontact direction is estimated from the R-GPIS model as the normalized gradient of the predictive mean. Theoretically, this guides the probe back toward the object, as the occupancy probability increases along this direction. The resulting recontact direction is given by:

$$\mathbf{d}_{\text{recontact}} = \hat{\mathbf{n}}(\mathbf{x}) = \frac{\nabla \mu_R(\mathbf{x})}{\|\nabla \mu_R(\mathbf{x})\|} \quad (13)$$

IV. EXPERIMENTS

We evaluated the effectiveness of the proposed approach in shape reconstruction and measurement tasks using a diverse set of objects. First, we describe the hardware setup in Sec. IV-A. Then, we conduct comprehensive experiments to answer the following key questions:

- **Q1**: Can our method autonomously explore unknown object surfaces without geometric priors (Sec. IV-B)?
- **Q2**: Can the framework accurately reconstruct object shapes using only contact feedback to support geometric measurements (Sec. IV-C)?
- **Q3**: Does the full system achieve more robust and complete reconstructions compared to its ablated variants and prior tactile exploration methods (Sec. IV-D)?

A. System Setup and Task Description

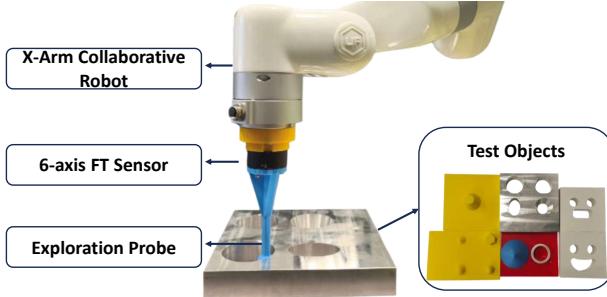


Fig. 3: Real-world experimental setup. The system consists of a 6-DoF X-Arm robotic manipulator, a force–torque sensor, a probing tool, and a diverse set of test objects for evaluation.

Our hardware configuration is shown in Fig. 3. The system consists of an X-Arm 6 collaborative robot equipped with a γ45 force–torque (FT) sensor (Decent Inc.) mounted at the end effector. A probe with a spherical tip is used to interact with object surfaces during exploration. We compensate for the probe’s spherical geometry through post-processing. First, the extracted shape S_{recon} is represented as a point cloud. Then, R-GPIS predicts per point normal as the gradient of the occupancy function. The final shape is obtained by shifting the points outward along the reversed normal direction by the probe radius (10mm). Finally, an outlier removal process is applied following [41].

Since the X-Arm does not support joint-level force control, we implement a hybrid position-based controller at the task

space. At each step, the motion command is computed according to the selected exploration policy. The position is then updated as $\Delta \mathbf{x}_t = \mathbf{d} \cdot \Delta t$, where $\mathbf{d} \in \{\mathbf{d}_{\text{local}}, \mathbf{d}_{\text{global}}, \mathbf{d}_{\text{recontact}}\}$ is the motion direction computed by the Exploration Motion Planner (Sec. III-B), and Δt is the control time step.

For all experiments, the control step size was set to $\Delta t = 0.001$. The Exploration Motion Planner was configured with fixed parameters: $\alpha = 0.4$, $\gamma = 1.0$, $\lambda = 1000$, and $f_{\text{target}} = 1.0\text{N}$. The parameters α and γ were empirically tuned to achieve smooth motion by balancing normal and tangential force components, while λ was found to be robust to parameter variations. The target force f_{target} was chosen to ensure stable yet minimally intrusive contact. For OBJ.1–4 (cavity-type objects), the system executed up to 120 ER-GPIS updates, with motion constrained along the cavity depth. For OBJ.5–8 (surface-type objects), up to 400 updates were allowed, with motion constrained along the object height.

B. Effectiveness of the Tactile Exploration Policy

To evaluate whether the proposed tactile exploration policy can autonomously guide the robot to explore unknown object surfaces without geometric priors (Q1), we conducted experiments on two types of structures, as shown in Fig. 4-1. The test objects include four internal cavities (a cylindrical hole, an inclined ring, an irregular cavity, and an emoji-like structure) and four exterior surfaces (a conical frustum, a hemispherical bump, a smooth dome, and a polygonal surface). The corresponding exploration trajectories are shown in Fig. 4-2. For all objects, we observed trajectories from sliding motion (black lines) and global redirection to underexplored regions with high uncertainty (red lines).

During exploration, the evolution of uncertainty acquired by the exploration policy is shown in Fig. 5. Overall, the sliding policy reduces uncertainty over time. When the rate of uncertainty reduction slows down (indicating that the utility function is trapped in local optima, as seen where the curve in Fig. 5 flattens), the policy redirects the probe toward regions with higher uncertainty, causing an increase in the uncertainty metric. This mechanism enables the policy to escape local optima and guide the robot toward unexplored regions, which is essential for achieving complete and robust surface coverage. Finally, the trajectory covers most regions of the object surface (Fig. 4-2). These results confirm the proposed method supports model-free tactile exploration across diverse surface profiles (Q1).

C. Accuracy of Reconstruction and Measurement

Next, we evaluate whether the proposed framework can accurately reconstruct object geometry (Q2). This requires effective exploration (Q1), and also a high-fidelity reconstruction algorithm capable of handling sparse tactile observations.

To assess reconstruction quality, we measure accuracy using both qualitative visualization (Fig. 4-3) and three quantitative metrics. First, the Chamfer Distance (CD) measures global shape similarity between the reconstructed point cloud and the ground truth (downsampled to 2000 points from the mesh) and is sensitive to incomplete regions. Second, the

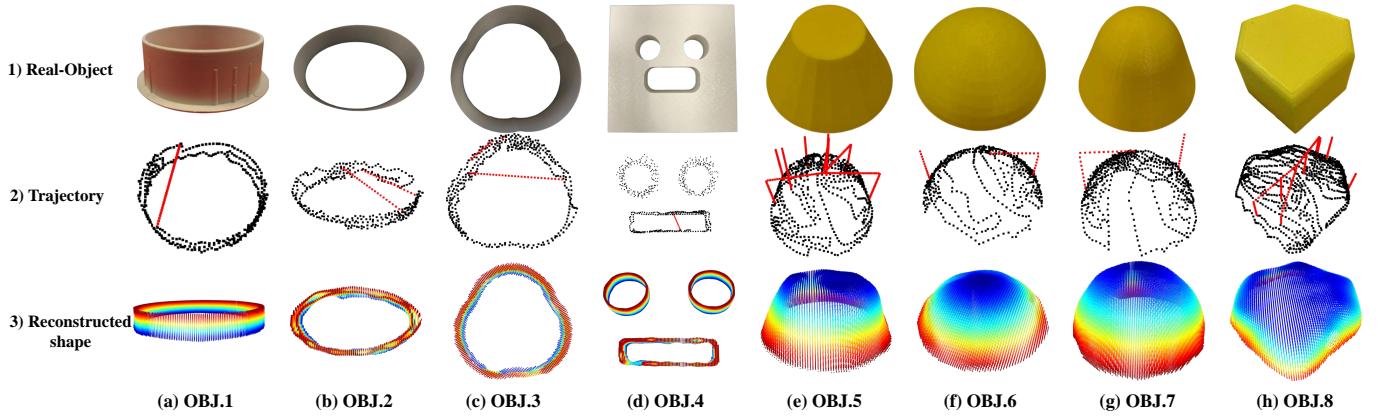


Fig. 4: Visualization of eight test cases (OBJ.1–OBJ.8): real objects (top), exploration trajectories (middle), and reconstructed shapes (bottom). OBJ.1–OBJ.4 correspond to internal cavities (cylindrical, inclined, irregular, and emoji-shaped), while OBJ.5–OBJ.8 represent exterior surfaces (conical, hemispherical, domed, and polygonal). In the middle row, black traces indicate local sliding motions along contacted surfaces, and red segments or points denote global-policy redirections toward high-uncertainty regions. The bottom row presents reconstructions generated by the proposed R-GPIS using tactile contacts only, where point colors represent height along the z -axis.

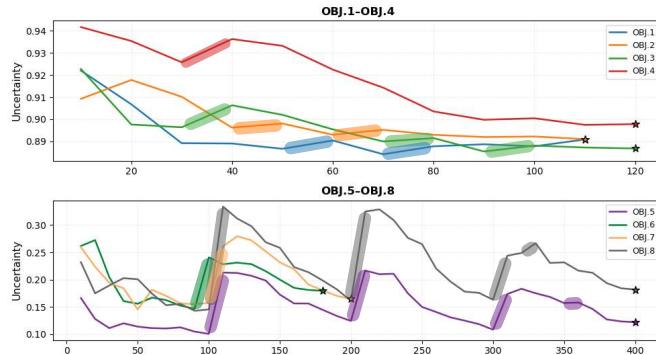


Fig. 5: Evolution of the estimated uncertainty during exploration on OBJ.1–OBJ.4 (top) and OBJ.5–OBJ.8 (bottom). The vertical axis denotes the mean predictive variance of ER-GPIS at the probe’s location, where higher values indicate faster information acquisition, and a rise reflects the activation of the global exploration policy. The horizontal axis represents the exploration steps.

Root Mean Square Deviation (RMSD) reflects local surface accuracy using one-sided point-to-mesh distances. Third, to evaluate whether the measured shape dimensions are biased, we compute the diameter error of the circumscribed circle between the measured data and the ground truth.

The results in Fig. 4-3 and Table I demonstrates that our method achieves high-fidelity shape reconstruction. All reconstructed shapes exhibit globally consistent topology and geometrically plausible surfaces. Quantitatively, the average resulting Chamfer Distance (CD) is 5.270 mm^2 , the average resulting Root Mean Square Deviation (RMSD) is 2.046 mm , and the average resulting diameter error is 2.865 mm . Notably, performance varies with surface complexity: for instance, OBJ.1, OBJ.6 and OBJ.7 (cylindrical, hemispherical and domed) achieve the highest accuracy due to their uniform curvature and consistent contact normals. In contrast, larger errors observed on OBJ.3 and OBJ.8 (irregular and polyg-

onal) reflect challenges introduced by discontinuous edges, tilted surfaces, and reduced contact coverage. Despite slight inaccuracies, all reconstructions are complete and successful. This confirms that our method enables reliable reconstruction and measurement across both internal and external geometries using tactile data alone without requiring CAD priors (Q2).

D. Ablation Studies

To examine whether the full integration of our exploration policy outperforms ablated baselines (Q3), we conduct ablation studies using OBJ.1. This object is selected for comparison because it is the simplest object that all methods can reconstruct to a comparable extent. Other objects may present challenges for baselines, resulting in large variance or failures.

Comparison Experiments. Comparisons were conducted under the following configurations: (1) our full framework with all components enabled; (2) an ablated variant without the global exploration policy, but with ER-GPIS and the contact recovery policy; (3) without the contact recovery policy, but with ER-GPIS and the global exploration policy; (4) without both the global exploration and contact recovery policies, but with ER-GPIS; (5) a minimal baseline using a single GPIS for both exploration and reconstruction, consistent with previous works [8]; and (6) the same baseline as (5) but augmented with our global exploration strategy. The results are summarized in Table II. Our full pipeline consistently achieves the best overall performance, with the removal of any introduced component leading to degraded accuracy and efficiency. Besides, we observed our approach outperforms the baseline method in reconstruction quality, though the wall-clock exploration time is slightly longer due to usage of global exploration policy (6 min 18 sec vs. 5 min 34 sec) (Q3).

Compare with RL-based approach. Furthermore, we evaluated our method against the RL-based *Actexplore* [38] using a 3D-printed replica of its demonstration object. Our method achieved near-complete surface coverage within 25 min and

TABLE I: Quantitative evaluation of reconstruction accuracy, runtime, and peak memory usage across all test objects.

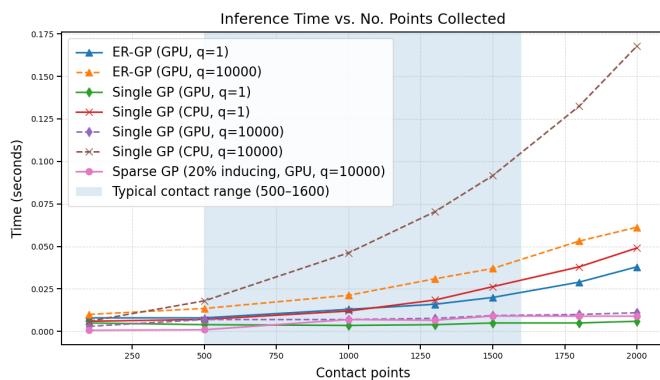
Object	OBJ.1 (cylindrical)	OBJ.2 (inclined)	OBJ.3 (irregular)	OBJ.4 (emoji)	OBJ.5 (conical)	OBJ.6 (hemispherical)	OBJ.7 (domed)	OBJ.8 (polygonal)	Average
CD (mm^2)	3.02	4.50	4.35	8.53	7.54	2.86	3.01	8.80	5.270
RMSD (mm)	0.85	1.10	5.17	2.81	1.94	1.19	1.22	2.09	2.046
Dia. Err. (mm)	3.81	2.10	4.89	5.29	2.12	0.97	1.08	2.66	2.865
Run Time (min)	6m18s	6m03s	6m15s	8m39s	25m00s	14m19s	18m15s	21m39s	–
GPU Memory (GB)	3.6	3.4	3.5	4.1	12.1	7.4	7.5	8.1	–

TABLE II: Compare with ablated variants and baseline [8].

Variant	CD \downarrow	RMSD \downarrow	Dia. Err. \downarrow
Benchmark (Ours)	3.02	0.85	3.81
w/o Global Policy	5.91	2.32	4.65
w/o Recontact Policy	7.78	2.85	5.14
w/o Global + Recontact Policy	19.92	5.41	–
Driess <i>et al.</i> [8]	12.45	2.49	14.18
Driess <i>et al.</i> [8] + Global Policy	11.02	2.35	11.47

Note: CD is measured in mm^2 ; RMSD and Dia. Err. are in mm .

“–” indicates measurement failure due to unclosed reconstruction.

**Fig. 6:** Inference time comparison of different GP regressors on GPU and CPU, q denotes No. points paralleled for query.

successfully reconstructed the object’s shape, whereas *Actexplore* covered only local regions even after 40min and failed to produce a usable reconstruction. We attribute this advantage to the efficiency of sliding motion, and our optimization-based approach, which does not suffer from generalization issues when encountering new objects.

Computational Costs. We analyzed the computational cost of ER-GPIS, implemented using a GPU-accelerated Gaussian Process framework [42]. The model’s inference and memory complexity are both $\mathcal{O}(N^2)$, similar to other GPIS-based methods [8]. Although maintaining two GP models increases resource usage, the cost remains manageable. Fig. 6 compares the inference time of ER-GP with other GP variants on both GPU and CPU. For contact point sets up to 2,000 (typical for most objects), ER-GP achieves inference times on the order of tens of milliseconds for batched evaluations of 10,000 query points. Although it is slightly slower than our modified GPU-based GP implementation (following [8]), it remains significantly faster than conventional CPU-based approaches. Furthermore, employing Sparse GP methods [43], [44] could offer additional scalability at the expense of accuracy. Considering the communication latency between the computing unit and robot hardware, the overall planning loop operates at approximately 8–16 Hz.

V. DISCUSSIONS

The proposed approach enables robust shape exploration, reconstruction, and measurement using a robotic system. Despite its effectiveness, several technical limitations were observed. First, the spatial resolution of the reconstructed surface is fundamentally limited by the hardware. Our system was implemented and evaluated on a collaborative robotic arm whose motion precision does not yet match that of high-end metrology-grade commercial CMMs. This choice was motivated by the closed-source nature of most industrial CMMs, which do not permit algorithmic customization.

Second, the reconstruction accuracy is constrained by the expressive ability of the GPIS model. Specifically, the kernel-based representation struggles to capture sharp and high-curvature regions. As a result, the model tends to smooth out fine geometric details such as edges and sharp corners. This limitation can be partially addressed by collecting more contact data in high-curvature regions. Future work will explore post-processing strategies to enhance geometric sharpness [45] and investigate the use of visuo-tactile sensors for richer local geometric information to improve reconstruction accuracy, such as following the perception setting of [38], and will further seek to reduce runtime through algorithmic and implementation optimizations so as to better adapt the framework to industrial scenarios.

VI. CONCLUSIONS

This paper presents a fully autonomous, tactile-based framework that enables dense 3D shape acquisition and reconstruction for industrial part measurement. The proposed system integrates three key components: a dual Gaussian Process framework (ER-GPIS) for simultaneous shape reconstruction and uncertainty estimation, coupled with a hybrid exploratory motion planning strategy that integrates sliding motion, global exploration, and contact recovery. The framework is validated on a real robot using a diverse set of objects featuring both internal cavities and exterior surface profiles. The system achieved accurate reconstructions with high coverage, reporting low reconstruction errors and reliable dimensional measurements across various geometries. These results demonstrate the framework’s potential to enable model-free, contact-based inspection in scenarios where conventional coordinate measuring machines cannot perform automated measurements and where vision-based methods fail under challenging visual conditions (e.g., occlusion, transparency, or poor illumination). Future work will focus on enhancing the shape expressiveness of ER-GPIS, integrating visuo-tactile sensors, and deploying the system on metrology-grade CMMs to further validate its measurement precision.

REFERENCES

[1] S. Sadaoui and N. Phan, "Touch probe measurement in dimensional metrology: A review," *International Journal of Automotive and Mechanical Engineering*, vol. 18, no. 2, pp. 8647–8657, 2021.

[2] K.-C. Fan and M. C. Leu, "Intelligent planning of cad-directed inspection for coordinate measuring machines," *Computer Integrated Manufacturing Systems*, vol. 11, no. 1-2, pp. 43–51, 1998.

[3] M. R. Driels, W. Swayze, and S. Potter, "Full-pose calibration of a robot manipulator using a coordinate-measuring machine," *The International Journal of Advanced Manufacturing Technology*, vol. 8, no. 1, pp. 34–41, 1993.

[4] J. A. Sladek, "Coordinate metrology," *Springer Tracts in Mechanical Engineering*, Jan 2016. [Online]. Available: <https://doi.org/10.1007/978-3-662-48465-4>

[5] S. N. Spitz, "Dimensional inspection planning for coordinate measuring machines," Ph.D. dissertation, University of Southern California, 1999.

[6] Z. Geng and B. Bidanda, "Review of reverse engineering systems – current state of the art," *Virtual and Physical Prototyping*, vol. 12, pp. 161–172, Mar 2017. [Online]. Available: <https://doi.org/10.1080/17452759.2017.1302787>

[7] S. Terzić, J. Dedić, D. Lazarević, Ž. Šarkočević, and B. Nedić, "Machining contact and non-contact inspection technologies in industrial application," *Journal of Production Engineering*, pp. 55–60, 2018.

[8] D. Driess, P. Englert, and M. Toussaint, "Active learning with query paths for tactile object shape exploration," in *2017 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, 2017, pp. 65–72.

[9] S. Catalucci, A. Thompson, S. Piano, D. T. Branson III, and R. Leach, "Optical metrology for digital manufacturing: a review," *The International Journal of Advanced Manufacturing Technology*, vol. 120, no. 7, pp. 4271–4290, 2022.

[10] S. Yan, J. Shi, G. Li, C. Hao, Y. Wang, H. Yu, and W. Zhou, "Advances in aeroengine cooling hole measurement: A comprehensive review," *Sensors*, vol. 24, no. 7, p. 2152, 2024.

[11] Z. Luthuli, "Traceability of measurements in optical coordinate measuring machines," Ph.D. dissertation, Stellenbosch: Stellenbosch University, 2020.

[12] E. Abouel Nasr, A. Al-Ahmari, A. A. Khan, S. H. Mian, O. Abdulhameed, and A. Kamrani, "Integrated system for automation of process, fixture and inspection planning," *Journal of the Brazilian Society of Mechanical Sciences and Engineering*, vol. 42, no. 1, p. 52, 2020.

[13] Y. Liu, W. Zhao, R. Sun, and X. Yue, "Optimal path planning for automated dimensional inspection of free-form surfaces," *Journal of Manufacturing Systems*, vol. 56, pp. 84–92, 2020.

[14] S. M. Stojadinović and V. D. Majstorović, *An intelligent inspection planning system for prismatic parts on CMMs*. Springer, 2019.

[15] B. Curless and M. Levoy, "A volumetric method for building complex models from range images," in *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, 1996, pp. 303–312.

[16] M. Kazhdan, M. Bolitho, and H. Hoppe, "Poisson surface reconstruction," in *Proceedings of the fourth Eurographics symposium on Geometry processing*, vol. 7, no. 4, 2006.

[17] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, "Nerf: Representing scenes as neural radiance fields for view synthesis," *Communications of the ACM*, vol. 65, no. 1, pp. 99–106, 2021.

[18] B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis, "3d gaussian splatting for real-time radiance field rendering," *ACM Trans. Graph.*, vol. 42, no. 4, pp. 139–1, 2023.

[19] Y. Dou, F. Yang, Y. Liu, A. Loquercio, and A. Owens, "Tactile-augmented radiance fields," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 26529–26539.

[20] A. Swann, M. Strong, W. K. Do, G. S. Camps, M. Schwager, and M. Kennedy, "Touch-gs: Visual-tactile supervised 3d gaussian splatting," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2024, pp. 10511–10518.

[21] I. Fang, K. Shi, X. He, S. Tan, Y. Wang, H. Zhao, H.-J. Huang, W. Yuan, C. Feng, and J. Zhang, "Fusionsense: Bridging common sense, vision, and touch for robust sparse-view reconstruction," 2024. [Online]. Available: <https://arxiv.org/abs/2410.08282>

[22] J. J. Park, P. Florence, J. Straub, R. Newcombe, and S. Lovegrove, "Deepsdf: Learning continuous signed distance functions for shape representation," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 165–174.

[23] S. Suresh, Z. Si, J. G. Mangelson, W. Yuan, and M. Kaess, "Shapemap 3-d: Efficient shape mapping through dense touch and vision," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 7073–7080.

[24] D. Driess, D. Hennes, and M. Toussaint, "Active multi-contact continuous tactile exploration with gaussian process differential entropy," in *2019 International Conference on Robotics and Automation (ICRA)*, 2019, pp. 7844–7850.

[25] E. Smith, R. Calandra, A. Romero, G. Gkioxari, D. Meger, J. Malik, and M. Drozdzal, "3d shape reconstruction from vision and touch," *Advances in Neural Information Processing Systems*, vol. 33, pp. 14193–14206, 2020.

[26] P. K. Murali, B. Porr, and M. Kaboli, "Touch if it's transparent! actor: Active tactile-based category-level transparent object reconstruction," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023, pp. 10792–10799.

[27] W. Yuan, S. Dong, and E. H. Adelson, "Gelsight: High-resolution robot tactile sensors for estimating geometry and force," *Sensors*, vol. 17, no. 12, p. 2762, 2017.

[28] S. McGovern, H. Mao, and J. Xiao, "Learning to estimate centers of mass of arbitrary objects," *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 1848–1853, 2019. [Online]. Available: <https://api.semanticscholar.org/CorpusID:210971047>

[29] W. Yuan, C. Zhu, A. Owens, M. A. Srinivasan, and E. H. Adelson, "Shape-independent hardness estimation using deep learning and a gelsight tactile sensor," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 951–958.

[30] P. Lin, Y. Huang, W. Li, J. Ma, C. Xiao, and Z. Jiao, "Pp-tac: Paper picking using tactile feedback in dexterous robotic hands," *arXiv preprint arXiv:2504.16649*, 2025.

[31] F. Khadivar, K. Yao, X. Gao, and A. Billard, "Online active and dynamic object shape exploration with a multi-fingered robotic hand," *Robotics and Autonomous Systems*, vol. 166, p. 104461, 2023.

[32] F. Luan, C. Wang, Z. Wang, J. Yue, Y. Zhou, and B. He, "A practical tactile exploration with enhanced sampling for fast shape estimation," *Measurement*, vol. 252, p. 117350, 2025.

[33] C. Bilaloglu, T. Löw, and S. Calinon, "Whole-body ergodic exploration with a manipulator using diffusion," *IEEE Robotics and Automation Letters*, vol. 8, no. 12, pp. 8581–8587, 2023.

[34] C. Bilaloglu, T. Löw, and S. Calinon, "Tactile ergodic coverage on curved surfaces," *IEEE Transactions on Robotics*, 2025.

[35] Z. Yi, R. Calandra, F. Veiga, H. van Hoof, T. Hermans, Y. Zhang, and J. Peters, "Active tactile object exploration with gaussian processes," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2016, pp. 4925–4930.

[36] C. Rosales, F. Spinelli, M. Gabiccini, C. Zito, and J. L. Wyatt, "Gpatlas-rrt: A local tactile exploration planner for recovering the shape of novel objects," *International Journal of Humanoid Robotics*, vol. 15, no. 01, p. 1850014, 2018.

[37] C. Xiao and J. Wachs, "Nonmyopic informative path planning based on global kriging variance minimization," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 1768–1775, 2022.

[38] A.-H. Shahidzadeh, S. J. Yoo, P. Mantripragada, C. D. Singh, C. Fermüller, and Y. Aloimonos, "Actexplore: Active tactile exploration on unknown objects," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 3411–3418.

[39] O. Williams and A. Fitzgibbon, "Gaussian process implicit surfaces," in *Gaussian Processes in Practice*, 2006.

[40] C. Xiao, N. Madapana, and J. Wachs, "Fingers see things differently (fist-d): An object aware visualization and manipulation framework based on tactile observations," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 4249–4256, 2021.

[41] Open3D Team. Point cloud outlier removal. Open3D. Open3D Documentation.

[42] J. Gardner, G. Pleiss, K. Q. Weinberger, D. Bindel, and A. G. Wilson, "Gpytorch: Blackbox matrix-matrix gaussian process inference with gpu acceleration," *Advances in neural information processing systems*, vol. 31, 2018.

[43] F. Leibfried, V. Dutordoir, S. John, and N. Durrande, "A tutorial on sparse gaussian processes and variational inference," *arXiv preprint arXiv:2012.13962*, 2020.

[44] M. K. Titsias, "Variational model selection for sparse gaussian process regression," *Report, University of Manchester, UK*, 2009.

[45] Y. Liu, A. Obukhov, J. D. Wegner, and K. Schindler, "Point2cad: Reverse engineering cad models from 3d point clouds," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 3763–3772.