

Interactive Robot-Environment Self-Calibration via Compliant Exploratory Actions

Podshara Chanrungrameekul^{1*}, Kejia Ren^{1*}, Joshua T. Grace², Aaron M. Dollar², Kaiyu Hang¹

Abstract—Calibrating robots into their workspaces is crucial for manipulation tasks. Existing calibration techniques often rely on sensors external to the robot (cameras, laser scanners, etc.) or specialized tools. This reliance complicates the calibration process and increases the costs and time requirements. Furthermore, the associated setup and measurement procedures require significant human intervention, which makes them more challenging to operate. Using the built-in force-torque sensors, which are nowadays a default component in collaborative robots, this work proposes a self-calibration framework where robot-environment spatial relations are automatically estimated through compliant exploratory actions by the robot itself. The self-calibration approach converges, verifies its own accuracy, and terminates upon completion, autonomously purely through interactive exploration of the environment’s geometries. Extensive experiments validate the effectiveness of our self-calibration approach in accurately establishing the robot-environment spatial relationships without the need for additional sensing equipment or any human intervention.

I. INTRODUCTION

Robots are becoming increasingly prevalent in various industrial and household applications, undertaking tasks ranging from pick-and-place to tool manipulation within defined environments. Often, such robot applications require an accurate pose of the robot frame relative to a frame of the workspace to be provided by a calibration procedure before task execution. Robot calibration involves identifying the true geometrical parameters in the robot’s kinematic structure and enhances absolute pose accuracy through software adjustments, avoiding alterations to the robot’s mechanical structure [1]. While certain calibration types, such as hand-eye calibration [2]–[4] or object localization [5], [6], have been previously studied, research on automatic robot-environment calibration for manipulators, a critical aspect of overall robot performance, has been relatively limited.

Traditional robot-environment calibration approaches rely on external measurement devices, which are time-consuming and require significant human intervention in the loop. To automate and streamline this process, various calibration methods have explored sensor modalities like vision, 1D laser sensors, touch-based techniques, tactile sensors, and probes [3], [4], [7]–[11], complicating the calibration and limiting accessibility for users without these resources. When robot physical interactions are needed for calibration, such as




Fig. 1: Our self-calibration framework determines the spatial relationship between the robot and the non-movable, rigid environment through compliant exploratory actions. Visualized in green is the environment’s pose as currently estimated by the robot, and in blue is the ground truth.

probing points in the environment, manual operations are still often required since existing simulation-based interaction and outcome prediction [11]–[13] have not shown to generalize to arbitrary robot-environment setups.

To this end, we propose a framework for robot-environment *self-calibration*, defined as a fully autonomous process for identifying the robot-environment spatial parameters only using the robot’s exploratory actions without any human interventions. We work with setups, as exemplified in Fig. 1, without external sensors or specialized end-effectors, e.g., point contact-based probes. The robot-environment spatial relationship is modeled using a Particle Filter with a distribution of hypothesized poses. Our framework uses the robot’s end-effector to actively interact with the environment through touching and sliding actions to gather informative observations of contacts, refining an inaccurate initial belief iteratively and finally converging the hypotheses towards the true value. The exploratory actions are strategically optimized to collect information that maximally contributes to the self-calibration. Experiments conducted in both simulation and the real world confirm the effectiveness and accuracy of our self-calibration framework. The results show that the proposed framework is able to precisely estimate the robot’s pose relative to the environment, operating in a fully autonomous manner.

*Equal contribution.

¹Department of Computer Science, Rice University, Houston, TX 77005, USA. ²Department of Mechanical Engineering and Material Science, Yale University, New Haven, CT 06511, USA. This work was supported by the US National Science Foundation grant FFR-2133110 and FFR-2132823.

II. RELATED WORK

Traditional calibration: Traditionally, robot-environment calibration involves sensing the actual pose of the robot's end-effector and comparing this information with the poses predicted by the robot-environment kinematic models to acquire calibration data. This process falls into two primary categories: robot calibration utilizing external measurement devices [14]–[16] and robot calibration through the imposition of physical constraints [2]–[4], [17]. External measurement-based robot calibration methods are time-consuming, challenging to operate, and rely on human experts to operate. Many of such calibration approaches [3], [4] possess particular complexities, including the prerequisite of performing hand-eye calibration and robot exterior calibration in advance, with the added requirement for an external chessboard. In this work, we propose a robot-environment self-calibration process to eliminate the need for human intervention, external tools, and complex prerequisites.

Autonomous calibration: Previous works in autonomous calibration have explored multiple sensor modalities, including visions for end-effector tracking [3], [4], [10], 1D laser sensors to gauge distances to the environment [7], and the integration of touch-based techniques using a probe [8], [11] or a poker [9] to establish point contacts. Tactile sensors have been employed to model contact dynamics [13], [18]–[21], and tactile images have been leveraged for richer contact information [5], [6]. A common feature among these approaches is the reliance on specific additional sensors or probes to facilitate calibration. In contrast, our work proposes an approach that allows the robot to self-calibrate *as is* without extra sensors or tools.

Next best action: During an autonomous calibration process, the measurement process iteratively refines the robot's calibration using dynamic interactions between the robot and the environment. Various approaches have been employed to optimize this interaction. While some methods choose actions for calibration exploration by incorporating information gain [11]–[13], others use entropy functions based on particle rejection principles to enhance the decision-making process during calibration [11]. However, these approaches select the next actions by calculating the potential information gain across all hypothesis states, thus limiting the number of samples for the robot's state representation. As an effective alternative, we develop a heuristic algorithm to select the most promising exploratory actions based on an analysis of the environment.

III. PROBLEM FORMULATION

This work addresses a self-calibration problem for estimating the pose of a robot manipulator in an uncalibrated environment without the requirements of any external sensors such as cameras. Specifically, we denote the fixed world frame of the static environment as $X^s \in SE(3)$, the base frame of the robot as $X^b \in SE(3)$, and the body frame of the robot as $X^e \in SE(3)$ rigidly attached on the robot's end-effector. Given the geometric model of the static environment expressed in the fixed world frame, the estimation goal is to

obtain the pose of the robot, $X \in SE(3)$, relative to the environment (i.e., the transformation from X^s to X^b).

We approach the self-calibration stated above through Particle Filter-based probabilistic estimation. It requires the robot manipulator to actively interact with the environment by touching and sliding to obtain informative observations of contacts between the robot and the environment. By updating the particle beliefs based on the information obtained from observations and sequential importance resampling (SIR), the Particle Filter can approximate the distribution of the system's hidden state (i.e., the robot's pose X) and iteratively refine it based on new observations.

A. Particle Filter-based Self-Calibration

At each time step t of the iterative process of Particle Filtering, we denote the set of M importance samples as $\mathcal{X}_t = \{X_t^1, \dots, X_t^m, \dots, X_t^M\}$, called particles for simplicity. The value of each particle is a hypothesized pose of the robot, i.e., $X_t^m \in SE(3)$. Based on the observational data and the actions taken by the robot, the likelihood of having a given particle $X_t^m \in \mathcal{X}_t$ is given by the posterior:

$$X_t^m \sim p(X_t | Z_{1:t}, U_{1:t}) \quad (1)$$

where $X_t \in SE(3)$ is the robot's pose to be estimated at time t ; $Z_{1:t}$ and $U_{1:t}$ are the observations and robot controls until time t respectively.

However, explicitly deriving the true posterior distribution above in real-world scenarios is often impractical. As such, Particle Filters rely on a proposal distribution π to iteratively approximate the real distribution:

$$\pi(X_t) \sim p(X_t | X_{t-1}, U_{t-1})\pi(X_{t-1}) \quad (2)$$

With a set of particles \mathcal{X}_t sampled from π , the estimated value of the robot's pose, given by the expected value of X_t , can be approximated through importance sampling:

$$\mathbb{E}[X_t] = \int p(X_t)X_t dX_t \approx \sum_{m=1}^M X_t^m w_t^m \quad (3)$$

where $w_t^m \propto p(X_t^m)/\pi(X_t^m)$ is the normalized weight associated with each particle.

Assuming the system is Markovian, the weight w_t^m can be recursively computed using Bayes' Theorem:

$$w_t^m = \frac{p(Z_t | X_t^m)p(X_t^m | X_{t-1}^m, U_t)}{\pi(X_t^m | X_{t-1}^m, U_t, Z_t)} \quad (4)$$

where $p(Z_t | X_t^m)$ is the observation model, and $p(X_t^m | X_{t-1}^m, U_t)$ is the motion model, which describes how the system transitions under the control input U_t .

B. Motion Model and Observation Representation

In this work, we use the robot's touching and sliding actions to interact with the uncalibrated environment. Since the environment is static and the robot manipulator is fixed in the environment, the robot's sliding action U_t will have no impact on the robot's pose relative to the environment.

Algorithm 1 Self-Calibration using Particle Filter

Input: Initial guess of the robot’s pose $X_0 \in SE(3)$.

Output: Estimated pose of the robot $X_t \in SE(3)$.

```

1:  $t \leftarrow 0$ 
2: while not terminated do                                ▷ Sec. VI-A
3:    $t \leftarrow t + 1$ 
4:    $\mathcal{X}_t \leftarrow \emptyset$                                 ▷ Set of Particles Initialized as Empty
5:    $U_t \leftarrow \text{ACTIONSELECTION}(U_{1:t-1})$            ▷ Sec. V-B
6:    $Z_t \leftarrow \text{EXECUTE}(U_t)$                          ▷ Observation from FT Sensor
7:   for  $m = 1, \dots, M$  do
8:      $X_t^m \sim \mathcal{N}(X_{t-1}^m, \sigma_t \mathbb{I})$            ▷ Motion Model, Eq. (5)
9:      $w_t^m \leftarrow \text{Eq. (7)}$                          ▷ Weight Update in Sec. IV-B
10:  end for
11:   $\text{NORMALIZE}(\{w_t^1, \dots, w_t^M\})$ 
12:  for  $m = 1, \dots, M$  do
13:    Sample  $i$  with probability  $\propto w_t^i$                  ▷ Resampling
14:     $\mathcal{X}_t.append(X_t^i)$                                ▷ New Particles
15:  end for
16: end while
17: return  $X_t \leftarrow \sum_{m=1}^M X_t^m w_t^m$            ▷ Eq. (3)

```

As such, we model the motion model with a multivariate Gaussian distribution independent of the robot action U_t :

$$p(X_t^m | X_{t-1}^m, U_t) \sim \mathcal{N}(X_{t-1}^m, \sigma_t \mathbb{I}) \quad (5)$$

where σ_t scales the covariance of the Gaussian to simulate system uncertainties. This motion model is equivalent to perturbing each particle in the previous time step with random Gaussian noise.

After the robot executes the action U_t at time t , it will observe Z_t with the readings of its internal joint encoders and force-torque (FT) sensor. We represent this observation as a set of J_t discrete contact events while the robot is executing the sliding action U_t , that is, $Z_t = \{(q_t^1, c_t^1), \dots, (q_t^j, c_t^j), \dots, (q_t^{J_t}, c_t^{J_t})\}$. Each contact event $(q_t^j, c_t^j) \in Z_t$ consists of a robot configuration (e.g., joint angles) $q_t^j \in \mathbb{R}^N$ where N is the number of DoFs of the robot, and a binary value $c_t^j \in \{0, 1\}$ indicating whether the robot is in collision with the environment.

The implementation of Particle Filters for robot self-calibration is detailed in Alg. 1. Given an initial guess of the robot’s pose $X_0 \in SE(3)$, the robot needs to actively determine the sliding action U_t that it will perform to collect observations, as will be detailed in Sec. V. The observations will then be used to evaluate and update the weights of particles for resampling, as described in Sec. IV. By repeating this procedure, the robot will iteratively refine the approximated distribution of the robot’s pose X_t until the estimated robot’s pose converges close to its ground truth, based on sets of particles and observations that will be introduced in Sec. VI.

IV. UPDATE RULE OF PARTICLE WEIGHTS

A critical factor for the success of Particle Filtering is the particle weight assignment based on the observations. A robust weight evaluation can greatly facilitate the efficient resampling of particles. As discussed in Sec. III-B, we observe binary contact events detected by the internal FT sensor of the robot. Hence, the observation model $p(Z_t |$




Fig. 2: The geometry of the robot’s end-effector is approximated as a point cloud \mathcal{P}^e . The grid represents a voxelized cache SDF $\Phi_{\mathcal{O}}$ where each voxel’s color corresponds to the distance to the surface boundary $\partial\mathcal{O}$ of the environment. The left figure shows a non-contact scenario where the signed distance $d_t^{m,j}$ between the end-effector and the environment surface is positive; the right figure shows a penetration scenario where the signed distance is less than a threshold, $d_t^{m,j} < -\delta_p$.

X_t^m) is the probability of contact modeled with the hypothesized distance between the robot’s end-effector and the environment.

In Sec. IV-A, we implement an efficient evaluation of the hypothesized distance between the robot’s end-effector and the environment by using a Signed Distance Field (SDF). Then, based on this hypothesized distance, the rule for updating the particle weights w_t^m is detailed in Sec. IV-B.

A. Geometric Representations of Robot and Environment

We consider the scenario where the robot manipulator operates in a static environment. We denote the geometric model of the environment by $\mathcal{O} \subset \mathbb{R}^3$ and its surface boundary by $\partial\mathcal{O} \subset \mathcal{O}$, both expressed in the fixed world frame X^s . The computation of the distance between two exact geometric models in the 3D space can be expensive, especially for complex object shapes. In Particle Filtering, such computation needs to be performed for each particle, which limits the algorithm’s speed. To this end, the geometric model of the environment is approximated by a volumetric grid. We represent the geometry of the robot’s end-effector by a finite point cloud \mathcal{P}^e , containing a set of L pre-sampled points on the surface of the end-effector denoted by $\mathcal{P}^e = \{p_1, \dots, p_l, \dots, p_L\}$. The position of each point $p_l \in \mathbb{R}^3$ is given in the body frame of the robot X^e attached to the end-effector. Then, the distance between the end-effector and the environment can be approximated by the distance between the nearest point in \mathcal{P}^e to the volumetric model of the environment. To further speed up the distance computation, we generate a discrete Signed Distance Field (SDF) [22] with a resolution of Δ_{sdf} for the environment, as illustrated in Fig. 2. The SDF of the environment \mathcal{O} is mathematically represented by a function $\Phi_{\mathcal{O}} : \mathbb{R}^3 \rightarrow \mathbb{R}$, which returns the signed distance to the surface boundary $\partial\mathcal{O}$ of the environment given a query point in the world frame X^s .

With the approximate geometric representations, distance computations can be efficiently performed. Given a hypothesized robot’s pose (i.e., a particle value) X_t^m and an observed joint configuration of the robot $q_t^j \in \mathbb{R}^N$, the

hypothesized distance between the robot’s end-effector and the environment $d_t^{m,j}$ can be quickly obtained by querying the discrete SDF:

$$d_t^{m,j} = \min_{p_l \in \mathcal{P}^e} \Phi_{\mathcal{O}} \left(X_t^m \cdot \Gamma(q_t^j) \cdot p_l \right) \quad (6)$$

where $\Gamma : \mathbb{R}^N \rightarrow SE(3)$ is the forward kinematics of the robot that computes the transformation from the base frame X^b to the end-effector’s body frame X^e of the robot given a joint configuration; $X_t^m \cdot \Gamma(q_t^j) \cdot p_l$ is the transformation chain that transforms the coordinates of a point $p_l \in \mathcal{P}^e$ (on the surface of the end-effector) in the end-effector’s body frame X^e to the fixed world frame X^s , where X_t^m , $\Gamma(q_t^j)$ and p_l are all given by their homogeneous representations.

B. Particle Weight Evaluation

At each time step t of Particle Filters, the particle weight w_t^m is updated based on the observation Z_t consisting of J_t pairs of joint configuration and detected contact (q_t^j, c_t^j) . When a contact is detected, i.e., $c_t^j = 1$, we use a Gaussian with a preset variance σ_P^2 to model the likelihood of observing this contact. Specifically, if the hypothesized distance between the end-effector and the environment $d_t^{m,j}$ for particle X_t^m is close to zero based on the calculation in Eq. (6), the likelihood of observing a contact for this particle is large, and we assign a large weight to this particle. In contrast, if $d_t^{m,j}$ is large, which conflicts with the observed contact, we will assign a small value to the particle weight.

When no contact is observed, $c_t^j = 0$, we introduce a rejection mechanism for particles significantly inconsistent with the observation. If the hypothesized distance $d_t^{m,j}$ associated with a particle is less than a penetration distance threshold $-\delta_P$, as shown in Fig. 2 (right), it indicates the robot is predicted to greatly penetrate the environment. Then, the weight of this particle will be multiplied by an arbitrarily small value ϵ . The weight of each particle w_t^m is then aggregated as the product of all observed contact events in Z_t stemming from a sliding action U_t . The updated rule for particle weights assignment under different cases discussed above is given by:

$$w_t^m = \prod_{j=1}^{J_t} \begin{cases} \frac{1}{\sqrt{2\pi\sigma_P^2}} \exp\left(-\frac{|d_t^{m,j}|^2}{2\sigma_P^2}\right) & c_t^j = 1 \\ \epsilon & c_t^j = 0 \wedge d_t^{m,j} < -\delta_P \\ 1 & \text{otherwise} \end{cases} \quad (7)$$

After the particle weights are updated, they are normalized by $w_t^m = w_t^m / \sum_{m=1}^M w_t^m$ (Line 11 of Alg. 1). As such, we can use the weights as the sampling probabilities to resample particles.

V. ROBOT COMPLIANT ACTION SELECTION

Each robot action entails a time cost; hence, it is imperative to reduce the number of actions needed for accurate self-calibration and to ensure that every action captures a maximal amount of information. To this end, we use sliding actions of the robot to gather as much observation




Fig. 3: (a) An illustrative plot of a sliding action U_t . The blue dot is the reference contact location r_t on the environment surface, and the black arrow represents the normal vector \hat{n}_t at the contact location. The recorded observations for contact and non-contact are shown by the red and green dots, respectively. (b) Examples of convex segmentation for objects in Fig. 5. (c) Examples of a real robot executing two different sliding actions.

information as possible, as described in Sec. V-A. In Sec. V-B, we develop a strategy to select the most promising sliding action from multiple candidates heuristically.

A. Compliant Sliding Action

The robot will inevitably have unexpected collisions with the environment while interacting with it through sliding actions due to the inaccuracy of the self-calibration and real-world uncertainties, such as control errors. In response to such uncertainties, enabling robot compliance becomes imperative to facilitate safe interactions and maintain continuous contact with the uncalibrated environment.

Specifically, a sliding action U_t is associated with a reference contact (r_t, \hat{n}_t) on the environment surface, which includes a contact location $r_t \in \mathbb{R}^3$ and its corresponding normal vector $\hat{n}_t \in \mathbb{R}^3$, both expressed in the fixed world frame X^s . To execute the sliding action U_t , the robot moves close to and above the reference location r_t . Following this, to ensure safe interactions, the robot utilizes Cartesian impedance control to guide its end-effector to gently touch the environment surface and detect contact.

Upon contact detection by the FT sensors of the robot, the robot transitions into a sliding motion perpendicular to the surface normal vector \hat{n}_t , moving along the surface while maintaining contact, as shown in Fig. 3 (a). The sliding continues until the robot loses contact with the environment or exceeds a pre-defined maximum sliding distance. Throughout the sliding action, the robot collects contact events $(q_t^j, c_t^j) \in Z_t$ at fixed distance intervals, regardless of the contact status. Fig. 3 (c) showcases two different sliding actions executed by a real robot.

B. Action Optimization

To further maximize the information gained from a single robot action, we propose a strategic approach to optimize the action selection to enhance estimation efficiency. The objects in the environment are divided into convex segments using CoACD [23], as shown in Fig. 3 (b). Each segment

captures the local geometric information of the environment and serves as a focal point guiding the estimation process to concentrate on different parts of the environment. For each segment, we pre-sample a set of D possible contact candidates $\mathcal{D} = \{(r_1, \hat{n}_1), \dots, (r_d, \hat{n}_d), \dots, (r_D, \hat{n}_D)\}$. Each contact candidate (r_d, \hat{n}_d) consists of a contact location $r_d \in \mathbb{R}^3$ on the surface of the segment and a normal vector $\hat{n}_d \in \mathbb{R}^3$ at the contact location. As discussed in Sec. V-A, each robot sliding action U_t is associated with a reference contact (r_t, \hat{n}_t) , selected from the pre-sampled contact candidates \mathcal{D} of all segments.

When the robot selects a sliding action U_t at time t to optimize the information gathering, it will retrieve all the previous actions $U_{1:t-1}$ that the robot has executed. The robot will strategically select an action that explores different geometric features of the environment unsampled in the previous steps. Specifically, the robot first randomly picks a segment of the environment. For each contact candidate $(r_d, \hat{n}_d) \in \mathcal{D}$ in this segment, we define a local sparsity $\rho_{\mathcal{D}}(d)$, quantified as the angular difference between its normal vector \hat{n}_d and the normal vectors of the reference contacts associated with all previous actions, i.e., $\hat{n}_{1:t-1}$:

$$\rho_{\mathcal{D}}(d) = \min_{\tau=1, \dots, t-1} \cos^{-1} \hat{n}_d \cdot \hat{n}_{\tau} \quad (8)$$

The strategy is to select the contact candidate with the maximum local sparsity as the reference contact for U_t , i.e., $(r_t, \hat{n}_t) = (r_{d^*}, \hat{n}_{d^*}) \in \mathcal{D}$ where d^* is calculated by

$$d^* = \arg \max_{d=1, \dots, D} \rho_{\mathcal{D}}(d) \quad (9)$$

The robot then moves its end-effector to the selected location r_{d^*} , uses Cartesian impedance control to touch the environment surface, and slides its end-effector in a direction perpendicular to \hat{n}_{d^*} to execute U_t , as described in Sec. V-A. In cases where the selected contact candidate results in an infeasible action due to the reachability of the robot, it will attempt the next best candidate until a feasible one is found. Alternatively, there is a small probability of switching to another segment to explore alternate contact options and select a reference contact from the candidates on the new segment. This adaptive strategy ensures the viability of the selected contact and the sliding action associated with this contact.

VI. CONVERGENCE CRITERIA AND VERIFICATION OF THE ESTIMATION

We assume the ground truth of the robot's pose is not obtainable, and no other external sensors are available to verify the quality of the estimation during the self-calibration. Thus, we need an automatic mechanism to self-measure the estimation convergence and end the self-calibration when needed. In Sec. VI-A, we define criteria to measure the confidence of the estimation based on the statistical state of the particles and implement an adaptive strategy to control the spreading of the particles under different conditions to avoid overconfidence in a potentially false estimation.

A. Confidence and Convergence Criteria

We define three binary criteria to measure the confidence of the estimation, based on the set of particles \mathcal{X}_t at the current step or all particles in the previous steps $\mathcal{X}_{1:t}$:

1) *Particle Confidence*: $\mathcal{C}(\mathcal{X}_t) \in \{0, 1\}$ is determined by examining the variance of the particles' values:

$$\mathcal{C}(\mathcal{X}_t) = \begin{cases} 1 & \sigma(\mathcal{X}_t) < \theta_V \\ 0 & \text{otherwise} \end{cases} \quad (10)$$

where the dimension-wise variance $\sigma(\mathcal{X}_t) \in \mathbb{R}^6$ independently evaluates the variance of all six dimensions of \mathcal{X}_t , after converting the representation of each $X_t^m \in SE(3)$ in the set \mathcal{X}_t to a six-dimensional vector (i.e., x, y, z, roll, pitch, and yaw). When the variance is below a threshold $\theta_V \in \mathbb{R}^6$, indicating all particles tend to converge to the same estimate, the Particle Filter estimator is confident about the estimation.

2) *Particle Stability*: $\mathcal{S}(\mathcal{X}_{1:t}) \in \{0, 1\}$ is evaluated based on the particle values of all previous steps $\mathcal{X}_{1:t}$:

$$\mathcal{S}(\mathcal{X}_{1:t}) = \begin{cases} 1 & \text{ran}_M < \epsilon_M \wedge \text{ran}_V < \epsilon_V \\ 0 & \text{otherwise} \end{cases} \quad (11)$$

where $\text{ran}_M \in \mathbb{R}^6$ is the dimension-wise range of $X_{t-h:t}$, the estimates (i.e., the weighted average of particle values) over the past h iterations, and $\text{ran}_V \in \mathbb{R}^6$ is the range of $\{\sigma(\mathcal{X}_{t-h}), \dots, \sigma(\mathcal{X}_t)\}$, the dimension-wise variances of the particle values in the past h iterations.

In other words, this ensures both the weighted average and the variance of the particle values do not change much during the past h steps. If their ranges ran_M and ran_V are smaller than their respective thresholds $\epsilon_V, \epsilon_M \in \mathbb{R}^6$, the estimation is considered consistent over the time window, meaning the particles are stable.

3) *Particle Consistency*: $\mathcal{V}(\mathcal{X}_t, Z_t) \in \{0, 1\}$ is used to evaluate whether the estimation is consistent with the most recent observation Z_t . Assuming the estimated robot pose determined by the weighted average of particle values X_t is correct, for each observed joint configuration q_t^j , we calculate an estimated distance $d_t^{E,j}$ between the robot's end-effector and the environment using Eq. (6) with X_t instead of X_t^m . For all observed events in Z_t , if the estimated distance $d_t^{E,j}$ is not consistent with the detected contact c_t^j , thresholded by a distance δ_E , then the Particle Consistency criterion will be considered unsatisfied. This criterion is met only when the estimated distance $d_t^{E,j}$ is consistent with the detected contact for all observed events $(q_t^j, c_t^j) \in Z_t$:

$$\mathcal{V}(\mathcal{X}_t, Z_t) = \begin{cases} 0 & \exists (q_t^j, c_t^j) \in Z_t, \\ & \text{s.t. } (c_t^j = 1 \wedge d_t^{E,j} > \delta_E) \\ & \text{or } (c_t^j = 0 \wedge d_t^{E,j} \leq -\delta_E) \\ 1 & \text{otherwise} \end{cases} \quad (12)$$

$$d_t^{E,j} = \min_{p_i \in \mathcal{P}^e} \Phi_{\mathcal{O}} \left(X_t \cdot \Gamma(q_t^j) \cdot p_i \right)$$

When all three criteria are met for five consecutive iterations, the Particle Filter will terminate and return the weighted average of the current particle values as the final estimation for the robot's pose.

B. Adaptive Particle Spreading

As the robot's pose is time-invariant in the motion model of the system, as introduced in Sec. III-B, the motion model is a random Gaussian perturbation scaled by σ_t . When σ_t is large, the particles will disperse between iterations, whereas when σ_t is small, the particles will become more stable and more easily trapped at a local minimum. Therefore, to better deal with scenarios where the particles distribute differently, we adaptively adjust the scale of the variance in the motion model, i.e., the value of σ_t :

$$\sigma_t = \begin{cases} \alpha\sigma_{t-1} & \mathcal{C}(\mathcal{X}_t) \wedge \neg\mathcal{V}(\mathcal{X}_t, Z_t) \\ \beta\sigma_{t-1} & \mathcal{S}(\mathcal{X}_{1:t}) \wedge \mathcal{V}(\mathcal{X}_t, Z_t) \\ \sigma_{t-1} & \text{otherwise} \end{cases} \quad (13)$$

where these criteria were defined in Sec. VI-A; $\alpha > 1$ and $\beta < 1$ are scaling factors.

For scenarios where the particles have confidence in their estimation, yet the estimation conflicts with the recent observation data, i.e., $\mathcal{C}(\mathcal{X}_t) \wedge \neg\mathcal{V}(\mathcal{X}_t, Z_t)$, the particles are likely trapped at a local minimum. To address this, we introduce larger divergence to the particle distribution by increasing σ_t , allowing for broader exploration and preventing confinement of the estimation to local minima.

On the contrary, if particles remain stable and the current estimation aligns with the observation (i.e., $\mathcal{S}(\mathcal{X}_{1:t}) \wedge \mathcal{V}(\mathcal{X}_t, Z_t)$), it indicates that σ_t should be decreased to densify the particles, to enhance estimation confidence by reducing the size of the hypothesized distribution. This scaling approach ensures adaptability in scenarios where the estimation conflicts with the observations and fine-tunes the particle distribution for improved estimation accuracy and robustness.

VII. EXPERIMENTAL EVALUATION

In this study, we thoroughly evaluate the proposed framework for robot self-calibration and investigate its performance in both simulated and real-world scenarios. The experiments are conducted in the Gazebo simulator [24] and on a real Franka Emika Panda robot arm. In the real world, the system monitors the robot joint's forces and torques, which trigger collision detection when the values surpass a certain threshold. Throughout the experiments, the robot is equipped with a 3D-printed end-effector cover that mirrors the shape of a typical Franka gripper and serves as a protective measure to prevent damage to the robot.

A. Experiments in Simulation

As defined in Sec. III, we task our framework to self-calibrate the robot's pose with touching and sliding actions. We simulate an initial transformation error as $X - X_0 \sim \mathcal{U}(-e, e)$ where \mathcal{U} is a uniform distribution and e indicates the interval size of the distribution. The initial translational errors are simulated to be within ± 15 cm in the X-, Y-, and Z-axis, while the initial rotational errors are within ± 0.15 rad (8.6°). The number of sampled particles is fixed at $M = 100,000$. We conduct five experiments with different initial X_0 values to self-calibrate the robot against each of

the environmental objects, including a table, a shelf, a rigid object with the shape of a beanbag, a nightstand, and a winerack, as illustrated in Fig. 5. For each object, we sample a total of 1,000 contact candidates.

There are two major features to be highlighted in the experiments: 1) Without any tactile sensors for determining the contact location, our framework relies on accurate distance calculations between the hypothesized end-effector and the environment, which is crucial for achieving precise self-calibration. 2) After each measurement, an efficient method is required to filter out incorrect particles based on observations. Our framework effectively addresses these challenges, as demonstrated in the experiments detailed below. The self-calibration accuracy is evaluated by comparing the final translational and rotational errors, calculated as the difference between the expected value of X_t from Eq. (3) and the ground truth X . Due to the similarity between translational and rotational errors, here we focus on the experimental results of translational errors. Fig. 4 presents average translational errors for each experimental setup, with error bars indicating the standard deviation.

1) *Geometric Representation Analysis*: This experiment investigates the influence of geometric representation resolution on self-calibration accuracy reflected by the final translation error. We show that self-calibration improves its accuracy with more accurate representations of the environment and the robot's end-effector.

We evaluate the voxel grid size of the signed distance field (i.e., Δ_{sdf} in Sec. V-B) at 0.2, 0.3, 0.5, 1, and 2 cm, all with an end-effector point cloud size $L = 600$. The voxel size affects the precision of distance calculations between the end-effector and the surrounding environment, directly impacting the particle updates during self-calibration. The results in Fig. 4(a) show a direct correlation between voxel size and calibration accuracy. Specifically, smaller voxel sizes consistently lead to improved calibration outcomes. For the voxel grid size of 0.2 cm, the average final translational error is 0.53 ± 0.22 cm across all objects, highlighting the importance of finer voxel resolutions in self-calibration.

Different end-effector's point cloud sizes (i.e., L in Sec. V-A), including 60, 100, 600, 1000, and 2000, are tested for end-effector representation with a fixed SDF Resolution of 0.2 cm. Objects with diverse contact characteristics show varying impacts on calibration accuracy, as demonstrated in Fig. 4(b). For example, objects like the winerack and the shelf, which often make point contacts, are more sensitive to changes in the size of the point cloud. In contrast, objects with more flat and larger surfaces, like tables and nightstands, demonstrate similar accuracies when different L values are set. Increasing the point cloud size generally leads to not only increased calibration accuracy but also a rise in average computational time for particle evaluation. As reported in Tab. I, the average computational time increases from 2.7s to 12s as L increases from 60 to 2000. This indicates a noticeable trade-off between computation time and calibration accuracy.




Fig. 4: Self-calibration performance is evaluated in simulation with 100,000 particles. The reported data illustrate the effects of: (a) SDF resolution, (b) end-effector point cloud size L , and (c) σ_P (a particle weight evaluation parameter). The results, averaged over 5 experimental runs for each environmental object, showcase the final translational error (cm), with error bars representing the standard deviation.



Fig. 5: Environmental objects used in experiments: Baxter (real), table (sim/real), shelf (sim/real), beanbag (sim), nightstand (sim), winerack (sim). The green overlays are the objects' mesh models, displayed relative to the robot's base frame based on the robot's self-calibrated pose.

L	60	100	600	1000	2000
Time(s)	2.7 ± 0.3	2.9 ± 0.1	5.5 ± 0.2	7.8 ± 0.8	12.4 ± 0.9

TABLE I: Average Computation Time (s) for particle evaluation with different sizes of the point cloud (L).

2) *Particle Evaluation Analysis*: This experiment focuses on particle weight assignment, governed by variance σ_P detailed in Sec. IV-B. Different values of σ_P (0.001, 0.002, 0.005, 0.01, and 0.02) are investigated, as reported in Fig. 4(c). A smaller σ_P imposes a stricter criterion for particle rejection and makes the filter more similar to using binary contact representation, leading to suboptimal calibration and particle starvation. On the other hand, a larger σ_P relaxes rejection criteria, keeping particles with a more erroneous hypothesis state and decreasing calibration accuracy. We observe a U-shaped graph with the lowest average final translational error of 0.53 cm at $\sigma_P = 0.005$.

B. Real Robot Experiments

In real-world deployments, calibration accuracy may vary both due to errors in the environment CAD models and measurement inaccuracies. Here, we validate the feasibility of the simulation-tested parameters in the real world.

We use the optimal parameters obtained from simulation experiments, where $\Delta_{sdf} = 0.2$ cm, $L = 600$, and $\sigma_P = 0.005$. The task involves self-calibrating the robot w.r.t. a table and a shelf in simulation and the real world, as depicted in Fig. 5. Each experiment is repeated 10 times. An example run is presented in Fig. 6, where we show the translational and rotational errors compared to the ground truth, the progression of the robot base's particle set \mathcal{X}_t , and the estimation overlays during the calibration process. The complete results are summarized in Tab. II.

The quality of actions, particularly the average number of contacts J_t across all actions, significantly influences

calibration accuracy. Calibrating a real-world shelf, disturbed by physically noisy touches against bolts, ridges, and narrow surfaces, results in fewer contacts per action (just under 4) compared to simulation (above 5). As a result, the final translational error in the real-world experiments is 0.72 cm, compared to 0.53 cm in the simulation. In contrast, calibrating a table without these variations in both simulation and the real world yields similar contact numbers, leading to comparable final translation and rotational errors.

We further see that self-calibration demonstrates adaptability to different geometric situations. Calibrating the robot with respect to the table requires fewer actions for the particles to gain confidence and stability due to its relatively simple geometry. The required number of actions is reduced from 30 in the real world to 20 in simulation, while the translational error remains around 0.5 cm. Conversely, when calibrating a shelf, the translational error decreases from 0.7cm (real-world) to 0.53cm (simulation), while the average number of actions remains at 30.

Furthermore, we conduct the evaluation in a real-world application by applying the self-calibration framework w.r.t. a Baxter robot, as shown in Fig. 5. Given the absence of ground-truth information, a visual assessment is conducted. The results demonstrate a high self-calibration accuracy, underscoring the potential of the proposed framework in real-world robotic applications.

	Translational Error(cm)	Rotational Error(10^{-2} rad)	# Actions	# Contact
Table(Sim)	0.51 ± 0.17	0.75 ± 0.69	21.8 ± 5.8	6.0 ± 3.0
Table(Real)	0.53 ± 0.14	1.10 ± 0.41	17.2 ± 6.5	6.1 ± 3.0
Shelf(Sim)	0.53 ± 0.23	0.50 ± 0.20	30.5 ± 8.3	5.3 ± 2.9
Shelf(Real)	0.72 ± 0.30	1.49 ± 0.70	29.8 ± 9.4	3.9 ± 1.0

TABLE II: Experimental results for self-calibration using Table and Shelf from Fig. 5 both in simulation and real environment.

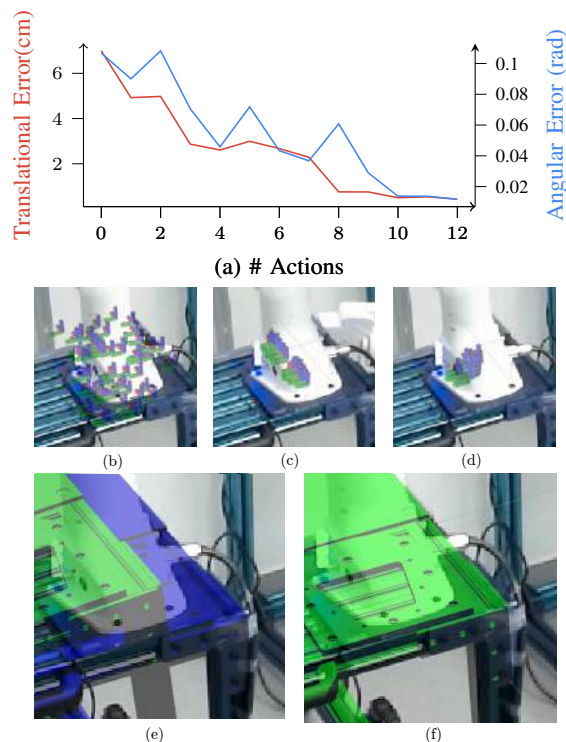


Fig. 6: (a) Translational (red) and rotational (blue) errors during the self-calibration. (b) initial hypothesis set \mathcal{X}_1 , (c) hypothesis set \mathcal{X}_τ updated based on the observation from touching and sliding action, (d) final hypothesis set \mathcal{X}_t . The green models overlay the initial guess (e) and the final estimation (f).

VIII. CONCLUSION

This work proposes an exploratory action-based probabilistic estimation framework for self-calibrating the spatial relationship between the robot and its environment. Our approach gathers contact information and iteratively updates particle beliefs using exploratory touching and sliding actions. These actions are strategically chosen to enhance observation informativeness. An automatic verification mechanism is developed to ensure estimation quality and allows prompt termination when appropriate. With extensive simulation and real-world experiments, we demonstrate the accuracy and effectiveness of our self-calibration framework in multiple different environments. Our particle update mechanism effectively converges hypothesized distribution towards the ground truth and consistently provides accurate calibration. In future work, we will extend our approach to self-calibrating more sensors, e.g., cameras, purely using the robot's exploratory actions without any human intervention or extra sensors.

REFERENCES

- [1] Z. Roth, B. Mooring, and B. Ravani, "An overview of robot calibration," *IEEE Journal on Robotics and Automation*, vol. 3, no. 5, pp. 377–385, 1987.
- [2] M. Qiu, "A practical and robust method for hand-eye calibration," *IFAC Proceedings Volumes*, vol. 28, no. 16, pp. 275–280, 1995.
- [3] Y. Meng and H. Zhuang, "Autonomous robot calibration using vision technology," *Robotics and Computer-Integrated Manufacturing*, vol. 23, no. 4, pp. 436–446, 2007.

- [4] G. Du and P. Zhang, "Online robot calibration based on vision measurement," *Robotics and Computer-Integrated Manufacturing*, vol. 29, no. 6, pp. 484–492, 2013.
- [5] M. B. Villalonga, A. Rodríguez, B. Lim, E. Valls, and T. Sechopoulos, "Tactile object pose estimation from the first touch with geometric contact rendering," in *Conference on Robot Learning*. PMLR, 2021, pp. 1015–1029.
- [6] M. Bauza, A. Bronars, and A. Rodríguez, "Tac2pose: Tactile object pose estimation from the first touch," *The International Journal of Robotics Research*, vol. 42, no. 13, pp. 1185–1209, 2023.
- [7] Y. Jiang, L. Yu, H. Jia, H. Zhao, and H. Xia, "Absolute positioning accuracy improvement in an industrial robot," *Sensors*, vol. 20, no. 16, 2020.
- [8] A. Petrovskaya and O. Khatib, "Global localization of objects via touch," *IEEE Transactions on Robotics*, vol. 27, no. 3, pp. 569–585, 2011.
- [9] A. Sipos and N. Fazeli, "Simultaneous contact location and object pose estimation using proprioception and tactile feedback," in *IEEE International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 3233–3240.
- [10] K. Stepanova, J. Rozlivek, F. Puciow, P. Krsek, T. Pajdla, and M. Hoffmann, "Automatic self-contained calibration of an industrial dual-arm robot with cameras using self-contact, planar constraints, and self-observation," *Robotics and Computer-Integrated Manufacturing*, vol. 73, p. 102250, 2022.
- [11] B. Saund, S. Chen, and R. Simmons, "Touch based localization of parts for high precision manufacturing," in *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 378–385.
- [12] P. K. Murali, A. Dutta, M. Gentner, E. Burdet, R. Dahiya, and M. Kaboli, "Active visuo-tactile interactive robotic perception for accurate object pose estimation in dense clutter," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 4686–4693, 2022.
- [13] P. Hebert, T. Howard, N. Hudson, J. Ma, and J. W. Burdick, "The next best touch for model-based localization," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2013, pp. 99–106.
- [14] 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, pp. 34–41, 1993.
- [15] M. To and P. Webb, "An improved kinematic model for calibration of serial robots having closed-chain mechanisms," *Robotica*, vol. 30, no. 6, pp. 963–971, 2012.
- [16] H. Wang, S. Shen, and X. Lu, "A screw axis identification method for serial robot calibration based on the poe model," *Industrial Robot: An International Journal*, vol. 39, no. 2, pp. 146–153, 2012.
- [17] C. Gong, J. Yuan, and J. Ni, "A self-calibration method for robotic measurement system," *Journal of Manufacturing Science and Engineering*, vol. 122, no. 1, pp. 174–181, 2000.
- [18] M. C. Koval, M. Klingensmith, S. S. Srinivasa, N. S. Pollard, and M. Kaess, "The manifold particle filter for state estimation on high-dimensional implicit manifolds," in *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 4673–4680.
- [19] M. C. Koval, N. S. Pollard, and S. S. Srinivasa, "Pre-and post-contact policy decomposition for planar contact manipulation under uncertainty," *The International Journal of Robotics Research*, vol. 35, no. 1-3, pp. 244–264, 2016.
- [20] S. Javdani, M. Klingensmith, J. A. Bagnell, N. S. Pollard, and S. S. Srinivasa, "Efficient touch based localization through submodularity," in *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2013, pp. 1828–1835.
- [21] M. C. Koval, N. S. Pollard, and S. S. Srinivasa, "Pose estimation for planar contact manipulation with manifold particle filters," *The International Journal of Robotics Research*, vol. 34, no. 7, pp. 922–945, 2015.
- [22] P. F. Felzenszwalb and D. P. Huttenlocher, "Distance transforms of sampled functions," *Theory of Computing*, vol. 8, no. 19, pp. 415–428, 2012.
- [23] X. Wei, M. Liu, Z. Ling, and H. Su, "Approximate convex decomposition for 3d meshes with collision-aware concavity and tree search," *ACM Transactions on Graphics (TOG)*, vol. 41, no. 4, pp. 1–18, 2022.
- [24] N. Koenig and A. Howard, "Design and use paradigms for gazebo, an open-source multi-robot simulator," in *IEEE International Conference on Intelligent Robots and Systems (IROS)*, vol. 3. IEEE, 2004, pp. 2149–2154.