

Tactile Probabilistic Contact Dynamics Estimation of Unknown Objects

Jinhoo Kim^{*1}, Yifan Zhu^{*2}, and Aaron Dollar²

Abstract—We study the problem of rapidly identifying contact dynamics of unknown objects in partially known environments. The key innovation of our method is a novel formulation of the contact dynamics estimation problem as the joint estimation of contact geometries and physical parameters. We leverage DeepSDF, a compact and expressive neural-network-based geometry representation over a distribution of geometries, and adopt a particle filter to estimate both the geometries in contact and the physical parameters. In addition, we couple the estimator with an active exploration strategy that plans information-gathering moves to further expedite online estimation. Through simulation and physical experiments, we show that our method estimates accurate contact dynamics with fewer than 30 exploration moves for unknown objects touching partially known environments.

I. INTRODUCTION

While robot manipulation technologies have advanced rapidly, deploying robots that can robustly perform manipulation with external contact in novel environments such as assembly and cooking remains a significant challenge. Model-free methods require environments that closely resemble or are identical to the training environment and degrade in the unfamiliar settings. Meanwhile, model-based approaches require dynamics models that are challenging to quickly acquire in unstructured environments where visual occlusions and sensor noises are prevalent.

In this work, we aim to rapidly identify an accurate contact dynamics model for a grasped unknown rigid object in a partially known rigid environment based exclusively on tactile measurements, and iteratively refine the model during interaction as shown in Fig. 1. Previous works have focused on explicitly estimating the contact locations and types [1], [2] or obtaining a linear complementarity model [3]. However, explicitly determining contact positions and types is very challenging in real-world scenarios with complex object shapes, where there are abundant contact modes that also change frequently. In addition, a linear complementary model is a local approximation that inevitably introduces approximation errors especially when the local contact geometries are highly nonlinear.

Instead, we contribute a novel formulation of the contact dynamics estimation problem as the joint estimation of contact geometries and physical parameters, and our proposed method quickly captures contact dynamics in the wild with




Fig. 1. Our work considers estimating the contact dynamics of unknown planar rigid objects rigidly attached to a robot with an unknown rigid transform in a fully rigid partially known environment. Our method uses a particle filter with a compact representation of object geometry to quickly estimate the contact dynamics accurately based on tactile measurement only.

few assumptions. Essential to our method is a compact and expressive geometric representation leveraging DeepSDF [4]. This yields a representation of the object geometry using a learned continuous signed distance function (SDF) representation based on neural networks. The compact geometric representation allows us then to adopt a particle filter, which has shown robustness for tasks involving nonlinear and discontinuous contact dynamics [5], [6]. Our representation ensures that the parameter space remains low dimensional, avoiding the curse of dimensionality of sampling-based methods, and enables the object parameters and geometry to be jointly estimated in a particle filter. In addition to the estimator, the quality of online samples also plays a key role in efficient estimation. We augment the particle filter with an active exploration strategy based on information theory that plans exploration actions with maximum expected information gain.

We evaluate our method on unknown objects in an environment of unknown friction with a flat surface and a vertical wall of unknown height and position. In both simulation and physical experiments, our estimator quickly estimates the contact dynamics with high accuracy. In simulation, the estimator shows less than 4 N of contact force prediction errors on new testing trajectories with ground truth force magnitudes going up to 25 N, while in physical experiments, it predicts with less than 0.5 N of error with 10 N of ground truth force magnitudes. This is achieved after fewer than 30 exploration actions. In addition, in simulation, we show that the active exploration strategy reduced wrench prediction error by more than 10% compared to random action selection.

*: Denotes equal contributions.

¹J. Kim is with the Department of Mechanical Engineering, ETH Zurich, Zurich, Switzerland. Work done as a visiting scholar at Yale University. kimjin@ethz.edu

²Y. Zhu and A. Dollar are with the Department of Mechanical Engineering and Materials Science, Yale University, New Haven, United States. {yifan.zhu, aaron.dollar}@yale.edu

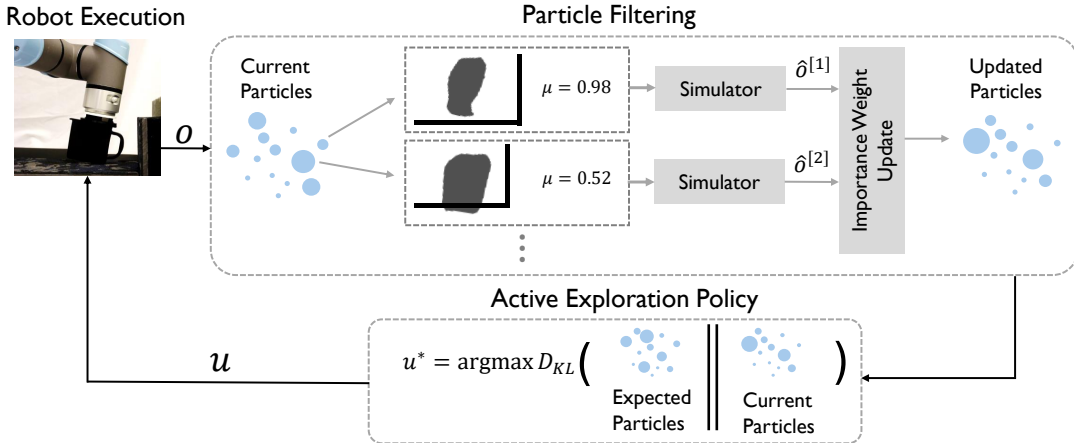


Fig. 2. Overview of the proposed estimation process. The robot manipulates an unknown object in a rigid environment with unknown surface height and wall position. Our method leverages DeepSDF [4] (a compact learned SDF) as a representation of object geometry allowing joint estimation of geometry and physical parameters using a particle filter. We then employ an active exploration strategy to select actions which give the maximum expected information gain measured in terms of the Kullback–Leibler (KL) divergence.

II. RELATED WORK

A. Contact Dynamics Estimation for Tactile Measurements

Many existing works aim to explicitly estimate the contact type (e.g. point vs. line contact) and contact information (e.g. location of point contact) from tactile measurements. These works adopt different estimation strategies, including direct equation fitting for narrow classes of contacts [1], maximum a posterior (MAP) estimation via factor graphs [2], [7], [8], and particle filters [9], [10], [11]. Some of these works assume unknown object geometry and a partially known environment, and estimate contact dynamics by explicitly considering the contact type, parameters, and the coefficient of friction [1], [2], [7], [8], [9]. While these methods are effective when estimating limited contact types or modes, they fail to scale to cases with multiple contacts of different contact types or when the contact changes rapidly in realistic manipulation tasks. Other works [10], [11] assume that object and environment geometries are known. The most similar prior work to this was presented by Pankert and Hutter [11], where a particle filter is used to estimate the in-hand transform of an insertion object and the world transform of a target box of known geometries. The key difference is our method estimates these transforms along with the geometries and physical parameters with minimal prior knowledge, as is essential for deploying robots in novel unknown environments.

Instead of directly estimating the contact parameters, another line of work aims to directly learn local dynamics, modeled as a linear complementarity system [12], [3]. In these works, the matrices of a linear complementarity system are fitted from observation data. Compared to these, our work estimates the geometries in contact, and avoids approximation errors from linearization.

B. Particle Filtering for Contact Dynamics

Particle filtering is a non-parametric filtering approach that can support multimodal probability distributions using




Fig. 3. Experiment Setup for both simulation (left) and physical world (right). In simulation, the environment used is a flat surface with an unknown height and a wall with an unknown position (Wall). The objects tested are 2D slices of YCB objects [16]. In physical experiments, the Mustard_Bottle, Mug, and Lemon are tested in the Wall environment. Contact wrenches are measured by the internal end-effector force torque sensor (F/T) of the UR5e robot arm.

particles. This property is highly valuable for rigid body contact estimation as the problem is often multi-modal. Many existing works have adopted particle filters for estimation problems involving contact or nonlinear dynamics [13], [14], [5], [15], [6]. Our main contribution lies in the novel formulation of the contact estimation problem as nonlinear multimodal filtering of both geometry and physical parameters.

III. PROBLEM DEFINITION

We consider quasi-static manipulation of planar rigid objects rigidly attached to the robot with an unknown rigid transform, in a fully rigid environment. In this work, we assume the object geometry is unknown and the environment geometry is partially known, e.g. a flat surface with an unknown height and a wall located at an unknown position. Our goal is to estimate the probabilistic dynamics function $p(x_{l+1}, w_{l+1} | x_l, u_l)$ at time l , given past observations of the object poses $x_1, \dots, x_l \in SE(2)$, contact wrenches $w_1, \dots, w_l \in \mathbb{R}^3$, and position commands of the impedance controller $u_1, \dots, u_l \in SE(2)$. For convenience of notation, we also denote observation $o_l = [x_l, w_l]$.

IV. METHOD

We propose a method to efficiently capture the contact dynamics via the estimation of contact geometries and physical

parameters with minimal assumptions. We take advantage of the well-established quasi-static rigid body contact dynamics and solve the problem of estimating the contact dynamics via the geometry and physical parameters θ of a quasi-static rigid body simulator. The dynamics function is then represented as $p(x_{l+1}, w_{l+1}|x_l, u_l) = \int p(x_{l+1}, w_{l+1}|x_l, u_l, \theta) \text{bel}(\theta) d\theta$, where $p(x_{l+1}, w_{l+1}|x_l, u_l, \theta)$ is the simulator and $\text{bel}(\theta)$ is the current belief of θ , which is recursively updated over time using particle filter. In the case of a deterministic simulator, we simply denote the simulator as $(x_{l+1}, w_{l+1}) = f_\theta(x_l, u_l)$. Summarized in Fig. 2, our method first adopts a compact and expressive geometric representation in the form of a learned continuous Signed Distance Function (SDF) representation of the distribution of geometries. Then, treating contact dynamics estimation as a nonlinear multimodal filtering problem on the unknown parameters of a quasi-static rigid body simulator, we adopt a particle filter for efficient estimation. Finally, we couple our estimator with an active exploration strategy based on information theory to collect samples with high information. In this section, we will first discuss the geometric representation, then detail the quasi-static rigid body simulator, present our estimation algorithm based on particle filtering, and finally the active exploration algorithm.

A. Geometry Representation

A key ingredient of our method is the choice of the geometric representation. Representations such as meshes or point clouds, while extremely flexible, contain too many parameters for any practical objects, making the estimation intractable. Instead, we take advantage of an object geometry prior by utilizing DeepSDF [4], which is a signed distance function based on a neural network that can represent a broad class of shapes in a single latent vector. We denote it as $g(p, z) : \mathbb{R}^2 \times \mathbb{R}^{d_z} \rightarrow \mathbb{R}$, which takes in the query 2D position of a point in the object frame $p \in \mathbb{R}^2$, and the latent geometry vector $z \in [-1, 1]^{d_z}$, and outputs the signed distance $d \in \mathbb{R}$ at the query position from the object surface (positive for being outside of the object, and negative for being inside). We use the trained latent vector z and the scaling parameter s to represent the object geometry.

For experiments, we assume that the unknown objects are drawn from the same distribution that the DeepSDF is trained on and we have partial knowledge of the terrain geometry, e.g. flat surface with unknown height. However, partial knowledge is not necessarily a limitation of our method since a similar compact geometric representation of the terrain to the object can be adopted and we intend to explore this in future work. As we will demonstrate in the Results section, we use $d_z = 2$ and achieved accurate estimation across all objects tested.

B. Quasi-Static Contact Dynamics

In quasi-static simulation, we assume that the robot moves slowly such that the Coriolis forces and accelerations can be ignored. We also assume the rigid bodies are always in force equilibrium, which is a mild assumption for many robot

manipulation tasks. In addition, the robot uses a Cartesian-space impedance controller for commanding the end-effector poses with known impedance K . Note that due to the quasi-static assumption, there is no longer a concept of velocity, and the system is driven by position commands of the impedance controller of the robot. We follow the simulator proposed by Pang et al. [17], and briefly present the key aspects of the simulator below.

Denoting the state of the actuated and un-actuated rigid objects in the environment with subscripts a and un , at each time step, the quasistatic rigid body simulator solves a linear complementarity problem (LCP) at time l :

Find v_{l+1} , s.t.:

$$J_{c,u}^T f_{c,un} + h\tau_{un} = 0$$

$$J_{c,a}^T f_{c,a} + h\tau_a + hK(x^* - x_l - hv_{l+1}) = 0$$

$$0 \leq \phi_l + hf_c^T v_{l+1} \perp f_c \geq 0$$

Terrain constraints.

Where v_{l+1} is the velocity at the next step and h is the time step such that $x_{l+1} = x_l + hv_{l+1}$. Note that quasi-static systems do not actually have velocity, this term here is just used for simulating the objects forward in time. J_c is the contact Jacobian, f_c is the frictional contact force that depends on the coefficient of friction μ , and τ are external forces including gravity and constraint forces such as those that constrain a rigid object that is a fixed terrain to be static. ϕ is the signed distance function. The expression $hK(x^* - x_l - h * v_{l+1})$ is the impulse applied by the impedance controller of stiffness K over h . Note that we use the minus sign $-$ for both positional and angular differences for rotation matrices. In our problem, we consider robots with a single rigidly attached actuated object interacting with static terrain. This LCP can be expressed as the KKT condition of a quadratic program [17], and we solve it efficiently with OSQP [18].

C. Contact Dynamics Estimation via Particle Filtering

We treat dynamics estimation as a nonlinear multimodal filtering problem on the geometry and physical parameters θ . Therefore, θ is the concatenation of the latent SDF shape vector, the size scale, the pose of the object with respect to the end-effector, the environment shape parameter (such as the position of an unknown wall), and the coefficient of friction μ . Given the nonlinear and discontinuous nature of contact dynamics, we adopt particle filter, a non-parametric filtering approach that has shown good performance for problems involving nonlinear dynamics [5], [15], [19]. In particle filters, we denote the particle set with M particles at time l as: $\Theta_l := \theta_l^{[1]}, \dots, \theta_l^{[M]}$ and the associated weights at time l as: $\Omega_l := \omega_l^{[1]}, \dots, \omega_l^{[M]}$. The belief of the state at time l , or in our case $\text{Bel}(\Theta_l) \approx \sum_i^M \omega_l^i \delta(\theta_l - \theta_l^i)$, is represented as a set of particles θ_l^i and the associated weights ω_l^i . The essence of the particle algorithm is the same as the Bayes filter but with the beliefs represented as particles. This allows the particles to represent multi-modal distributions, which is very common for contact dynamics estimation.

Algorithm 1: Particle Filtering Step

Input: Current Particle Set Θ_l and Their Associated Importance Weights Ω_l, O_l, U_l, R, r

- 1 **for** $i \leftarrow 1$ **to** M **do**
- 2 $\omega_{l+1}^{[i]} \leftarrow 1$
- 3 **for** $j \leftarrow 1$ **to** $N - 1$ **do**
- 4 $o, o', u \leftarrow O_l[j], O_l[j + 1], U_l[j]$
- 5 $\hat{o} = f_{\theta^{[i]}}(o, u)$ \triangleright Predict Observation
- 6 $\omega_{l+1}^{[i]} \leftarrow \omega_{l+1}^{[i]} \omega_l^{[i]} \mathcal{N}(\hat{o}; o', R)$ \triangleright Update Weights
- 7 $\Omega_{l+1} \leftarrow$ Weight Normalization
- 8 **if** $n_{eff}(\Theta_l) \leq \frac{M}{2}$ **then**
- 9 $\Theta_{l+1} \sim Bel(\Theta_l), \Omega_l = \frac{1}{M}$ \triangleright Resample
- 10 $\Theta_{l+1} \leftarrow \Theta_{l+1} + \beta, \beta \sim \mathcal{N}(0, rVar(\Theta_l))$ \triangleright Roughening
- 11 **return** $\Theta_{l+1}, \Omega_{l+1}$

Our particle filtering algorithm is summarized in Alg. 1. Given the current belief of particles, it first predicts the observation at time l based on the previous observation and the previous command input. Then, the importance weight of each particle is calculated by the probability density function of a Gaussian centered at the observation with a fixed diagonal covariance matrix $R \in \mathbb{R}^{d_\theta}$ [20], after which the weights are normalized. One challenge in using tactile feedback for contact dynamics estimation is that a single measurement is not very informative. To overcome this challenge, instead of only using the most recent observation and action, we use a history of N recent observations and actions (denoted as O_l and U_l) to update $Bel(\Theta_l)$, which was first introduced in [21]. Here, we note that particles are not updated through a process model as the particles represent the fixed geometry and physical parameters. Instead, we use the roughening method [22] by adding artificial noise after resampling to prevent particle depletion, which happens when a small number of particles dominate the distribution. Artificial noise is sampled from a zero-mean Gaussian with variance scaled by a roughness r relative to the particle variance. To also mitigate particle depletion due to frequent resampling [20], we only resample when the effective sample size $n_{eff} = \sum_i^M \frac{1}{\omega_i^2}$ is $\leq M/2$.

D. Active Exploration

Instead of executing random exploration moves to collect information, we adopt an active exploration strategy based on information theory. We follow the active learning approach for particle filter proposed by Hauser [23]. We choose an action that would maximize the expected information gain (EIG), where the information gain is defined as the Kullback–Leibler (KL) divergence between the belief of θ before and after an update from the observation o . Specifically, we choose an action u^* from a set of action candidates U

Algorithm 2: Expected Information Gain

Input: Action u, Θ_l, Ω_l

- 1 $G \leftarrow 0$ \triangleright Initialize EIG
- 2 **for** $i \leftarrow 1$ **to** M **do**
- 3 $\hat{o}_{l+1}^{[i]} = f_{\theta^{[i]}}(x_l, u_l)$ \triangleright Simulate the Observation
- 4 $\bar{\Theta}_l, \bar{\Omega}_l \leftarrow$ Particle Filter Step with Alg. 1
- 5 $G \leftarrow G + w_l^{[i]} D_{KL}(\bar{\Theta}_l, \bar{\Omega}_l \| \Theta_l, \Omega_l)$ \triangleright Calculate Information Gain
- 6 **return** G

according to:

$$u^* = \operatorname{argmax}_{u \in U} \operatorname{EIG}(u) = \int_o D_{KL}(Bel(\theta|o) \| Bel(\theta)) P(o|\theta, u) do$$

$$D_{KL}(Bel(\theta|o) \| Bel(\theta)) = \int_\theta Bel(\theta|o) \log \frac{Bel(\theta|o)}{Bel(\theta)} d\theta$$

To calculate EIG for a single action, we take advantage of the particles representation of $Bel(\theta)$. Shown in Alg. 2, for each particle, we simulate the observation with the action u , perform a particle filter step (without resampling) based on the observation, and calculate the KL divergence for the updated belief. Then EIG is the weighted sum of these information gains. Note that the KL divergence can be easily calculated by directly using the weights before and after the importance weight updates. This is an $O(NM^2)$ operation that requires many calls to the simulator where N is the number of actions [23]. Therefore, we randomly downsample $M/5$ particles and weights for EIG calculation.

V. EXPERIMENTS AND RESULTS

We evaluate our estimation pipeline in both simulation and physical experiments. In this section, we first present the results of simulated experiments and then discuss physical experiments. We train the DeepSDF function on the 2D cross-sections of selected 21 objects from the YCB dataset [16] to represent the object geometry, where the latent geometry vector dimension $d_z = 2$. For both simulation and physical experiments, we set the impedance for the controller as $K = [100 \text{ N/m}, 100 \text{ N/m}, 50 \text{ Nm}]$. Note that the controller is very stiff in rotation and compliant in x - and z -axis. In addition, each action changes the target pose to approximately translate 1 cm or rotate 5° .

A. Simulated Experiments

Shown in Fig. 3, our environment in simulation is a flat ground with unknown height g_h and a vertical wall located at an unknown position p_w . The parameter θ is a 9-dimensional vector, including the DeepSDF latent vector z , the object pose relative to the end-effector frame $T_{o,ee}$, size scale s , and the surface coefficient of friction μ . We assume that we know the possible range of the unknown parameters, where the unknown parameters are selected from

Environment	Exploration Strategy	$F_x(N)$	$F_z(N)$	$\tau(Nm)$	$x(mm)$	$z(mm)$
Sim	Random	3.71 ± 3.92	3.23 ± 4.11	0.260 ± 0.619	7.33 ± 10.1	6.00 ± 4.52
	Active	3.66 ± 2.24	2.81 ± 2.14	0.212 ± 0.182	7.72 ± 10.0	5.64 ± 4.05
	Expert	3.14 ± 1.70	2.68 ± 2.74	0.213 ± 0.140	5.87 ± 9.49	5.40 ± 4.05
Real	Active Expert	0.490 ± 0.688	0.212 ± 1.378	0.111 ± 0.235	5.88 ± 4.15	1.90 ± 13.4

TABLE I

TESTING RESULTS FOR SIMULATION AND REAL EXPERIMENTS, AVERAGED OVER 11 OBJECTS ON 3 TESTING TRAJECTORIES AND 3 OBJECTS ON 1 TESTING TRAJECTORY, RESPECTIVELY.

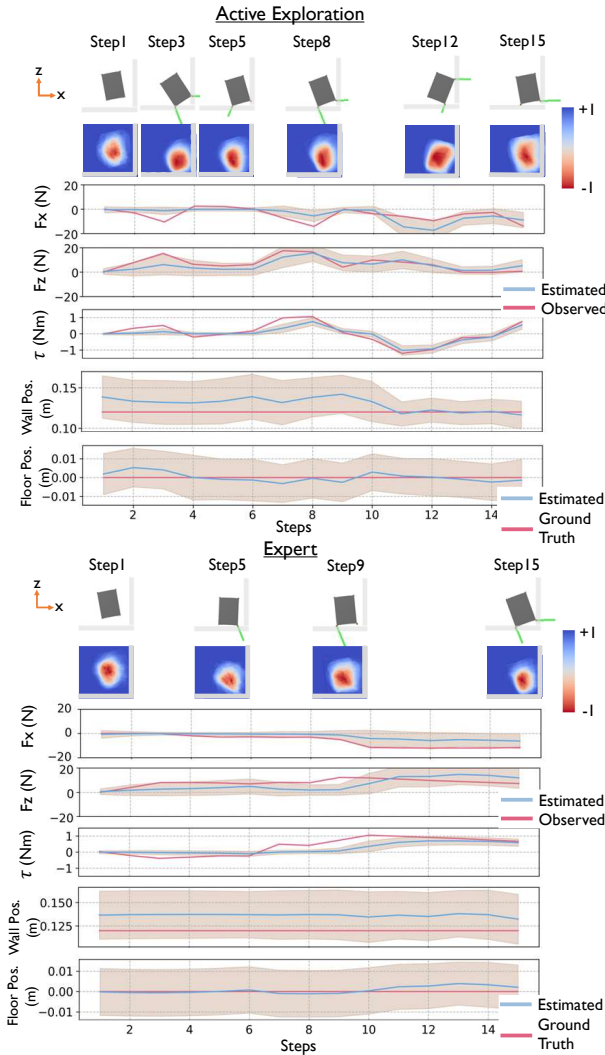


Fig. 4. Estimation progress for the Master_Chef.Can object for Active exploration and Expert exploration policies. The estimations of wrench, wall position, and floor position during the exploration trajectory are plotted with standard deviation. The trajectories and the current distribution of geometries are also visualized, with contact forces shown as green lines. The geometry distribution is the weighted indicator function (+1 outside of the object and -1 inside the object) for the top 100 particles in Θ with the largest weights.

a uniform distribution of $g_h \sim \mathcal{U}[-0.02m, 0.02m]$, $p_w \sim \mathcal{U}[0.09m, 0.18m]$, $z \sim \mathcal{U}[-1.0, -1.0], [1.0, 1.0]$, $T_{o,ee} \sim \mathcal{U}[-0.02m, -0.02m, -0.2rad], [0.02m, 0.02m, 0.2rad]$, $s \sim \mathcal{U}[0.8, 1.2]$, and $\mu \sim \mathcal{U}[0.1, 0.9]$.

As shown in Fig. 3, the simulation experiments




Fig. 5. Evaluation of the dynamics model estimated by Active exploration, and Expert exploration policies for the Master_Chef.Can object. The estimations of wrenches are plotted with standard deviation. The evaluation trajectory is also visualized at the top.

use 11 planar objects which are cross sections of objects from the YCB dataset. For the particle filter, we use $M = 5,000$ particles, the observation noise parameters of $[R_{F_x}, R_{F_y}, R_\tau, R_x, R_z, R_\phi] = [30N, 30N, 0.3Nm, 0.0001m, 0.0001m, 0.002rad]$, roughness of $r = 0.3$, and memory length $N = 5$. For each object in each environment, we compare three different exploration strategies, examples of which are shown in Fig. 4. The first is a baseline exploration strategy that adds random movements to a basic exploration trajectory (Random). This basic exploration trajectory simply commands the robot to make contact with the floor and move to the right. A zero mean Gaussian noise with variance $[0.03m, 0.03m, 0.25rad]$ is added to the position command. The second is our active exploration strategy which adds exploration actions to the same basic exploration trajectory. The action set is 27 position commands that is a uniform grid with range $[\pm 0.03m, \pm 0.03m, \pm 0.025rad]$. Note that due to limitations in computation, our active exploration strategy only plans one step ahead. Therefore, this is a local strategy and it requires the basic exploration trajectory as guidance to avoid getting stuck. We hope to alleviate the computation requirement of our active exploration strategy in the future. The last is an “expert” exploration trajectory which is tuned by the authors

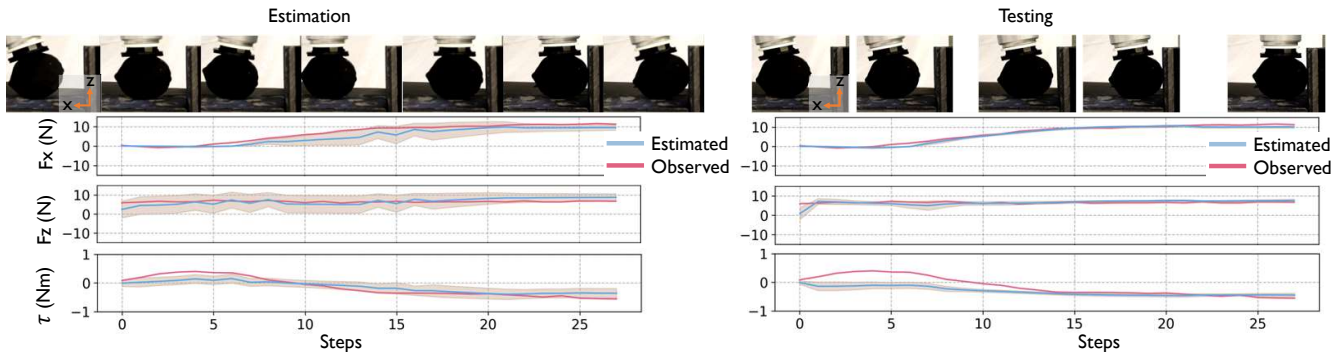


Fig. 6. Example trial of the Lemon object in physical experiments. Left: The estimation progress by our estimator, with an expert exploration trajectory. Right: Testing of the estimated dynamics model on a different trajectory, where both the estimated wrench and the standard deviation are shown.

based on the ground truth wall position. We use the same expert strategy across all the objects. We run each exploration strategy for 15 time steps.

Once the estimation is completed, we fix the particles and test the estimated dynamics model on three different testing trajectories in the same partially known environment for 30 time steps. One example is shown in Fig. 5. Presented in Table I, we report the quantitative results, averaged over 11 objects on 3 testing trajectories. The metric we used are the mean absolute error (MAE) between the ground truth and the weighted mean predictions of the top 100 particles with the largest weights. We report MAE for wrench predictions and next pose predictions except for rotation. This is because we use a very stiff controller in rotation and the error is minuscule. Overall, active exploration outperforms Random, but is worse than the Expert trajectory. We think that this is mainly due to the short horizon and the small action set we are using due to computation limit. Meanwhile, we observe that for the particular example shown in Fig. 5, active exploration outperforms Expert. This shows that a fixed exploration strategy, as is done for Expert, is not necessarily good for all objects, demonstrating the value of active exploration. As shown in Fig. 4, active learning approach uses both the bottom and top right part of the object to make a contact with the wall to simultaneously estimate unknown object and environment geometry. As a result, from time step 11, the wall position prediction quickly converges to the ground truth position. In terms of computation, it takes less than 1 s to do the estimation step and about 5 s to perform an active exploration step on a computer with an Intel i9-13900KF CPU, 64 GB of RAM, and an NVIDIA GeForce RTX 4090 GPU.

B. Physical Experiments

As shown in Fig. 3, we use a UR5e robot arm with 3D printed objects mounted on the end effector to perform the task with three objects in the Wall environment. We implemented an impedance controller on the UR5e robot. As the UR5e does not offer torque control, we approximate impedance control by using the position controller and the end-effector F/T sensor. However, this requires the environment to be deformable. We used gym tiles that deform about 1-3 mm with a 10 N contact force in our experiments. Despite

this violation of the rigid body assumption, we show that our estimator still performs accurate estimation.

We use the same hyperparameters for the particle filter, except for the observation noise parameter $R = [20 \text{ N}, 20 \text{ N}, 0.5 \text{ Nm}, 0.1 \text{ m}, 0.1 \text{ m}, 0.1 \text{ rad}]$. Here, the initial distribution over the unknown parameters are the following: $g_h \sim \mathcal{U}[-0.02 \text{ m}, 0.02 \text{ m}]$, $p_w \sim \mathcal{U}[-0.90 \text{ m}, -0.80 \text{ m}]$, $z \sim \mathcal{U}[[-1.0, -1.0], [1.0, 1.0]]$, $T_{o,ee} \sim \mathcal{U}[[-0.02 \text{ m}, -0.02 \text{ m}, -0.2 \text{ rad}], [0.02 \text{ m}, 0.02 \text{ m}, 0.2 \text{ rad}]]$, $s \sim \mathcal{U}[0.8, 1.2]$, and $\mu \sim \mathcal{U}[0.5, 1.6]$. To showcase the best performance of our estimator, instead of using the active learning strategy, we adopt an Expert exploration trajectory for these experiments. We report the quantitative results in Table I and show the estimation and testing trajectories for Lemon in Fig. 6. The estimation of Mug is also shown in Fig. 1. Despite the only information given in this case is that the shape of the environment is flat with a wall, our estimator is able to quickly estimate the contact dynamics, achieving less than 0.5 N of force prediction error where the ground truth magnitude goes up to 10 N. We believe our method shows great promise towards contact dynamics estimation in the open world.

VI. DISCUSSION AND CONCLUSION

In this work, we present a method to quickly estimate accurate contact dynamics for unknown objects in a partially known environment. Through both simulated and physical experiments, we demonstrate the accuracy of our estimator. We also show the effectiveness of our active exploration approach in the simulated experiments. One requirement of our method is the presence of a good geometry prior. We believe that with the abundance of 3D geometry data and the development of 3D large vision foundation models, such a requirement is not an obstacle.

There are a number of future directions we would like to pursue. First, we would like to extend this to 3D and lift the restrictions for a partially known environment. We hope to also investigate techniques to improve the sample efficiency of particle filters. Next, we would like to improve the computation efficiency of our active exploration strategy. Additionally, we want to explore training a reinforcement learning agent to learn an exploration strategy that best suits

our estimation pipeline. Finally, we would like to adopt our estimator for downstream manipulation tasks.

REFERENCES

- [1] N. Doshi, O. Taylor, and A. Rodriguez, "Manipulation of unknown objects via contact configuration regulation," in *ICRA*, 2022.
- [2] O. Taylor, N. Doshi, and A. Rodriguez, "Object manipulation through contact configuration regulation: multiple and intermittent contacts," 2023.
- [3] W. Jin and M. Posa, "Task-driven hybrid model reduction for dexterous manipulation," *IEEE Transactions on Robotics (TRO)*, vol. 40, pp. 1774–1794, 2024. [Online]. Available: <https://ieeexplore.ieee.org/document/10415517>
- [4] 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.
- [5] Z. Jiang, W. Zhou, H. Li, Y. Mo, W. Ni, and Q. Huang, "A new kind of accurate calibration method for robotic kinematic parameters based on the extended kalman and particle filter algorithm," *IEEE Transactions on Industrial Electronics*, vol. 65, no. 4, pp. 3337–3345, 2018.
- [6] G. Vezzani, U. Pattacini, G. Battistelli, L. Chisci, and L. Natale, "Memory unscented particle filter for 6-dof tactile localization," *Trans. Rob.*, vol. 33, no. 5, p. 1139–1155, oct 2017. [Online]. Available: <https://doi.org/10.1109/TRO.2017.2707092>
- [7] S. Kim and A. Rodriguez, "Active extrinsic contact sensing: Application to general peg-in-hole insertion," 2022.
- [8] S. Kim, D. K. Jha, D. Romeres, P. Patre, and A. Rodriguez, "Simultaneous tactile estimation and control of extrinsic contact," 2023.
- [9] K. Gadeyne, T. Lefebvre, and H. Bruyninckx, "Bayesian hybrid model-state estimation applied to simultaneous contact formation recognition and geometrical parameter estimation," *The International Journal of Robotics Research*, vol. 24, no. 8, pp. 615–630, 2005. [Online]. Available: <https://doi.org/10.1177/0278364905056196>
- [10] K. Hertkorn, M. A. Roa, C. Preusche, C. Borst, and G. Hirzinger, "Identification of contact formations: Resolving ambiguous force torque information," in *2012 IEEE International Conference on Robotics and Automation*. IEEE, 2012, pp. 3278–3284.
- [11] M. H. Johannes Pankert, "Learning contact-based state estimation for assembly tasks," in *IROS*, 2023.
- [12] W. Jin, A. Aydinoglu, M. Halm, and M. Posa, "Learning linear complementarity systems," in *Proceedings of The 4th Annual Learning for Dynamics and Control Conference*, ser. Proceedings of Machine Learning Research, R. Firoozi, N. Mehr, E. Yel, R. Antonova, J. Bohg, M. Schwager, and M. Kochenderfer, Eds., vol. 168. PMLR, 23–24 Jun 2022, pp. 1137–1149. [Online]. Available: <https://proceedings.mlr.press/v168/jin22a.html>
- [13] M. C. Koval, M. R. Dogar, N. S. Pollard, and S. S. Srinivasa, "Pose estimation for contact manipulation with manifold particle filters," in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2013, pp. 4541–4548.
- [14] 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 *2017 IEEE International Conference on Robotics and Automation (ICRA)*, 2017, pp. 4673–4680.
- [15] E. N. Chatzi and A. W. Smyth, "The unscented kalman filter and particle filter methods for nonlinear structural system identification with non-collocated heterogeneous sensing," *Structural Control and Health Monitoring*, 2009. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/stc.290>
- [16] B. Calli, A. Walsman, A. Singh, S. Srinivasa, P. Abbeel, and A. M. Dollar, "Benchmarking in manipulation research: Using the yale-cmu-berkeley object and model set," *IEEE Robotics and Automation Magazine*, vol. 22, no. 3, pp. 36–52, 2015.
- [17] T. Pang and R. Tedrake, "A convex quasistatic time-stepping scheme for rigid multibody systems with contact and friction," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 6614–6620.
- [18] B. Stellato, G. Banjac, P. Goulart, A. Bemporad, and S. Boyd, "OSQP: an operator splitting solver for quadratic programs," *Mathematical Programming Computation*, vol. 12, no. 4, pp. 637–672, 2020. [Online]. Available: <https://doi.org/10.1007/s12532-020-00179-2>
- [19] K. Hang, W. G. Bircher, A. S. Morgan, and A. M. Dollar, "Manipulation for self-identification, and self-identification for better manipulation," *Science robotics*, vol. 6, no. 54, p. eabe1321, 2021.
- [20] J. Elfving, E. Torta, and R. van de Molengraft, "Particle filters: A hands-on tutorial," p. 438, 2021.

- [21] G. Vezzani, U. Pattacini, G. Battistelli, L. Chisci, and L. Natale, "Memory unscented particle filter for 6-dof tactile localization," *IEEE Transactions on Robotics*, vol. 33, no. 5, pp. 1139–1155, 2017.
- [22] T. Li, T. P. Sattar, Q. Han, and S. Sun, "Roughening methods to prevent sample impoverishment in the particle phd filter," in *Proceedings of the 16th International Conference on Information Fusion*, 2013, pp. 17–22.
- [23] K. Hauser, "Randomized belief-space replanning in partially-observable continuous spaces," in *Algorithmic Foundations of Robotics IX: Selected Contributions of the Ninth International Workshop on the Algorithmic Foundations of Robotics*. Springer, 2010, pp. 193–209.