

In-Hand Object Pose Tracking via Contact Feedback and GPU-Accelerated Robotic Simulation

Jacky Liang¹, Ankur Handa², Karl Van Wyk², Viktor Makoviychuk², Oliver Kroemer³, Dieter Fox²

Abstract—Tracking the pose of an object while it is being held and manipulated by a robot hand is difficult for vision-based methods due to significant occlusions. Prior works have explored using contact feedback and particle filters to localize in-hand objects. However, they have mostly focused on the static grasp setting and not when the object is in motion, as doing so requires modeling of complex contact dynamics. In this work, we propose using GPU-accelerated parallel robot simulations and derivative-free, sample-based optimizers to track in-hand object poses with contact feedback during manipulation. We use physics simulation as the forward model for robot-object interactions, and the algorithm jointly optimizes for the state and the parameters of the simulations, so they better match with those of the real world. Our method runs in real-time (30Hz) on a single GPU, and it achieves an average point cloud distance error of 6mm in simulation experiments and 13mm in the real-world ones.

I. INTRODUCTION

Performing dexterous manipulation policies benefits from a robust estimate of the pose of the object held in-hand. Despite recent advances in pose estimation and tracking using vision feedback [1]–[3], in-hand object pose tracking still presents a challenge due to significant occlusions. As such, works that require in-hand object poses are currently limited to experiments where the object is mostly visible or rely on multiple cameras [4], or the hand-object transform is fixed or known [5], [6]. To mitigate the issue of visual occlusions, previous works have studied object pose estimation via contacts or tactile feedback, often by using particle filters and knowledge of the object geometry and contact locations. These techniques have been mostly applied to the static-grasp setting, where the object is stationary and in-grasp. Extending these techniques to tracking object poses during in-hand manipulation is difficult, requiring modeling of complex object-hand contact dynamics.

To tackle the problem of in-hand object tracking during robot manipulation, we propose combining a GPU-accelerated, high-fidelity physics simulator [7] as the forward dynamics model with a sample-based optimization framework to track object poses with contacts feedback (Figure 1). First, we initialize a concurrent set of simulations with the initial states of the real robot and the initial pose of the real object, which may be obtained from a vision-based pose registration algorithm assuming the object is not in occlusion in the beginning. The initial poses of the simulated

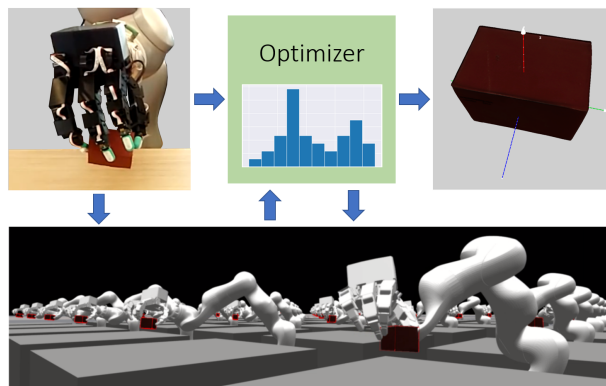


Fig. 1. The proposed in-hand object pose tracking framework. Robot controls are sent to an GPU-accelerated physics simulator that runs many robot simulations in parallel, each with different physics parameters and perturbed object poses. Costs based on observations, such as contact feedback from the real world and from the simulations, are passed to a sample-based derivative-free optimizer that periodically updates the states and parameters of all simulations to better match that of the real world. At any point in time, the pose of the simulation with the lowest cost is chosen as the current object pose estimate.

objects are slightly perturbed and reflect the uncertainty of the vision-based pose registration algorithm. The GPU-accelerated physics simulator runs many concurrent simulations in real-time on a single GPU. As a given policy controls the real robot to approach, grasp, and manipulate the object in-hand, we run the same robot control commands on the simulated robots. We collect observations of the real robot and the simulated robots, which include terms like the magnitude and direction of contacts on the robot hand’s contact sensors. Then, a sample-based optimization algorithm periodically updates the states and parameters of all simulations according to a cost function that captures how well the observations of each simulation matches with those of the real world. In addition, the algorithm also updates simulation parameters, such as mass and friction, to further improve the simulations’ dynamics models of the real world. At any point in time, the object pose estimate is the pose of the robot-object system.

To evaluate the proposed algorithm, we collected a total of 24 in-hand manipulation trajectories with 3 different objects in simulation and in the real world. For experiments, we use the Kuka IIWA7 arm with the 4-finger Wonik Robotics Allegro hand as the end-effector, with each finger outfitted with a SynTouch BioTac contact sensor. Object manipulation trajectories are human demonstrations collected via a hand-tracking teleoperation system. Because we have ground-truth object poses in simulation, we performed detailed ablation studies in simulation experiments to study the properties of

¹Carnegie Mellon University. Work done during internship at NVIDIA. jackyliang@cmu.edu

²NVIDIA Research {ahanda, vmakoviychuk, kvanwyk, dieterf}@nvidia.com

³Carnegie Mellon University. okroemer@andrew.cmu.edu

the proposed algorithm. For real-world experiments, we use a vision-based algorithm to obtain the object pose in the first and last frame of the collected trajectories, where the object is not in occlusion. The pose in the first frame is used to initialize the simulations, and the pose in the last frame is used to evaluate the accuracy of the proposed contact-based algorithm.

II. RELATED WORKS

Prior works have studied identifying in-hand object-pose with vision only, usually by first segmenting out the robot or human hand in an image before performing pose estimation [8], [9]. However, vision-only approaches degrade in performance for larger occlusions. Another approach is to use tactile feedback to aid object pose estimation. Tactile perception can identify object properties such as materials and pose [10], as well as provide feedback during object manipulation [11]–[13].

For the task of planar pushing where the object is visible, prior works have studied tracking object poses using particle filtering with contacts and vision feedback [14]. The authors of [15] compared a variety of dynamics models and particle filter techniques, and they found that adding noise to applied forces instead of the underlying dynamics yielded more accurate tracking results. One work combined tactile feedback with a vision-based object tracker to track object trajectories during for planar pushing tasks [16], and another applied incremental smoothing and mapping (iSAM) to combine global visual pose estimations with local contact pose readings [17].

For in-hand object pose estimation with tactile feedback, many prior works have explored this problem in the “static-grasp” context, where the robot hand grasps an object and localizes the object pose without moving. These works can be separated into two groups: 1) using point contact locations and 2) using a full tactile map to extract local geometry information around the contacts.

To use contact location feedback for pose estimation, many methods use a variation of Bayesian or particle filtering [18]–[26]. In [27] the authors perform filtering jointly over visual features, hand joint positions, force-torque readings, and binary contact modes. Similar techniques can be applied to pose estimation when the object is not held by the robot hand as well by using force probes [28], [29].

To use tactile maps for pose estimation, earlier works used large, low-resolution tactile arrays to sense contacts in a grid [11], [30], while more recent works use high-resolution tactile sensors mounted on robot finger tips. For example, the algorithm in [31] searches for similar local patches on an object surface to localize the object with respect to the contact location, and the one in [32] fuses GelSight data with a point cloud perceived by a depth sensor before performing pose estimation.

By contrast to the case of static-grasps, our work tackles the more challenging problem of tracking in-hand object pose during manipulation. Prior works have also explored this context. In [33] the authors propose an algorithm that

combines contact locations with Dense Articulated Real-time Tracking (DART) [34], a depth image-based object tracking system, while in [35] the algorithm fuses contact locations with color visual features, joint positions, and force-torque readings. The former algorithm is sensitive to initialization of the object poses, especially when the object appears small in the depth image. The latter work conducted experiments where the objects were mostly visible, so matching visual features alone would give reasonable pose estimates. In addition, neither work explicitly models the dynamics of the robot-object interaction, which limits the type of manipulation tasks during which the object pose can be tracked. To address these challenges, our approach does not assume access to robust visual features during manipulation. Instead, it uses a physics simulator to model both the kinematics and the dynamics of the robot-object system.

III. METHOD

A. Problem Statement

We consider the problem of tracking the pose of an object held in-hand by a robot manipulator during object manipulation. At some time t , let the object pose be $p_t \in SE(3)$. We define a physics dynamics model $s_{t+1} = f(s_t, u_t, \theta)$, where s_t is the state of the world (position and velocities of rigid bodies and of joint angles in articulated bodies), $u_t \in \mathbb{R}^M$ the robot controls (we use desired joint positions as the action space), and $\theta \in \mathbb{R}^N$ the fixed parameters of the simulation (e.g., mass and friction).

For a simulation model f that exactly matches reality given the perfect initializations of p_0 , s_0 , and θ , pose estimation requires only playing back the sequence of actions u_t applied to the robot in the simulation. However, given our imperfect forward model and noisy pose initializations, pose estimation using our method can be improved via observation feedback.

Let D be the number of joints the robot has and L the number of its contact sensors. We define the observation vector o_t as the concatenation of the joint position configuration of the robot $q_t \in \mathbb{R}^D$, the position and rotation of the robot’s contact sensors $P_t^{(l)} \in \mathbb{R}^3, R_t^{(l)} \in SO(3)$ (located on the fingertips), the force vectors of the sensed contacts $c_t^{(l)} \in \mathbb{R}^3$, the unit vector in the direction of the translational slippage on the contact surface $d_t^{(l)} \in \mathbb{R}^2$, and the binary direction of the rotational slippage on the contact surface $r_t^{(l)} \in \{0, 1\}$, where l indexes into the l th contact sensor. The general in-hand pose estimation problem is given the current and past observations $o_{1:t}$, robot controls $u_{1:t}$, and the initial pose p_0 , find the most probable current object pose p_t .

B. Proposed Approach

In this work, we leverage a GPU-accelerated physics simulator as the forward dynamics model to concurrently simulate many robot-object environments to track the in-hand object pose, and we use derivative-free, sample-based optimizers to jointly tune the state and parameters of these simulations to improve tracking performance (Algorithm 1). First, we obtain an estimate of the initial object pose via

a vision-based object pose estimator. We assume this pose estimator can give reliable initial pose estimate p_0 when the robot is not in contact with the object and when the object is not occluded, *i.e.*, before grasping. Then, given the initial object pose estimate and robot configuration, we initialize K concurrent simulations, and at every timestep we copy the real robot actions u_t to all K simulations. Note that the object pose can change when the hand establishes contact, and this will be modeled by the simulator. Let the object pose and the observation of the i th simulation be $p_t^{(i)}$ and $o_t^{(i)}$, the ground truth observations be $o_t^{(\text{gt})}$. Given a cost function \mathcal{C} , we say the current best pose estimate at time t is the pose of the i^* th simulation, $p_t^{(i^*)}$, where the i^* th simulation is the one that incurs the lowest average cost across some past time window T :

$$C_i = \frac{1}{T} \sum_{\Delta t=0}^{T-1} \mathcal{C}(o_{(t-\Delta t)}^{(i)}, o_{(t-\Delta t)}^{(\text{gt})}) \quad (1)$$

$$i^* = \arg \min_i C_i \quad (2)$$

The costs are used to periodically update the simulations and their parameters. This enables better alignment with the real robot-object system.

C. Cost Function

The desired cost function sufficiently correlates with object pose differences during in-hand object manipulation such that a lower cost corresponds to better pose estimations. The cost function we use has the form of:

$$\begin{aligned} \mathcal{C}(o_{(t-\Delta t)}^{(i)}, o_{(t-\Delta t)}^{(\text{gt})}) &= w_1 \|q_t^{(i)} - q_t^{(\text{gt})}\|_2 \\ &+ \sum_{l=1}^L (w_2 \|P_t^{(i,l)} - P_t^{(\text{gt},l)}\|_2 + w_3 |\Delta(R_t^{(i,l)}, R_t^{(\text{gt},l)})| \\ &+ w_4 (1 - \alpha_{(i,l)}) + w_5 \alpha_{(i,l)} |\Delta M(c_t^{(i,l)}, c_t^{(\text{gt},l)})| \\ &+ w_6 \alpha_{(i,l)} |\Delta \phi(c_t^{(i,l)}, c_t^{(\text{gt},l)})| \\ &+ w_7 (1 - \beta_{(i,l)}) + w_8 \beta_{(i,l)} |\Delta \phi(d_t^{(i,l)}, d_t^{(\text{gt},l)})| \\ &+ w_9 (1 - \gamma_{(i,l)}) + w_{10} \gamma_{(i,l)} |r_t^{(i,l)} - r_t^{(\text{gt},l)}| \end{aligned} \quad (3)$$

For the first term in the cost function, comparing q_t 's between the simulated and real-world robots is useful even if they share the same u_t , because q_t can be different depending on the collision constraints imposed by the current pose of the object in contact with the robot hand, which might make it physically impossible for a joint to reach a commanded target angle.

A contact sensor is in contact if its force magnitude is greater than a threshold. $\alpha_{(i,l)}$ is 1 when the binary contact state of the l th contact sensor of the i th simulation agrees with that of the real contact sensor and 0 otherwise. Similarly, $\beta_{(i,l)}$ is 1 when the l th contact sensor of the i th simulation agrees with the real contact sensor in whether or not the sensor is undergoing translational slippage, 0 otherwise; $\gamma_{(i,l)}$ is the same but for rotational slippage.

For any two vectors, $\Delta M(\cdot, \cdot)$ gives the difference of their magnitudes, and $\Delta \phi(\cdot, \cdot)$ gives the angle between them. For

any two rotations R_a and R_b , $\Delta(R_a, R_b)$ gives the angle of the axis-angle representation of $R_a^{-1}R_b$.

The weights of the cost terms, w_i 's, are determined such that the corresponding mean magnitude of each term is roughly normalized to 1.

D. Addressing Uncertainty and the Sim-to-Real Gap

There are two sources of uncertainty regarding object pose estimation via simulation: 1) the initial pose estimation p_0 from the vision-based pose estimator is noisy and 2) there is a mismatch between the simulated and real-world dynamics, partly caused by imperfect modeling and partly caused by the unknown real-world physics parameters θ .

To address the first issue of initial pose uncertainty, 1) we perturb the initial pose estimations across the different simulations by sampling from a distribution centered around the vision-based estimated pose $p_0^{(i)} \sim \mathcal{N}(p_0, \Sigma_p)$, and 2) we increase the number of simulations K ($K = 40$ in our experiments). If K is arbitrarily large, then it is with high probability that the true initial pose will be sufficiently represented in the set of simulations, and a well-designed cost function will select the correct simulation with the correct pose. To perform sampling over initial object poses, we sample the translation and rotation separately. Translation is sampled from an isotropic normal distribution, while rotation is sampled by drawing zero-mean, isotropic tangent vectors in $so(3)$ and then applying it to the mean rotation [36].

To address the second issue of mismatch between simulated and real-world physics (the ‘‘sim-to-real’’ gap), we propose using derivative-free, sample-based optimization algorithms to tune θ during pose tracking. Specifically, after every T time steps, we pass the average costs of all simulations during this window along with the simulation state and parameters to a given optimizer. The optimizer determines the next set of simulations with their own updated parameters. The simulations in the next set are sampled from simulations from the current set, with some added perturbations to the simulation parameters and object pose. Such exploration maintains the diversity of the simulations, preventing them from getting stuck in sub-optimal simulation parameters or states due to noisy observations.

Although it is desirable to have $\theta^{(i^*)}$ converge to the true $\theta^{(\text{gt})}$, this is not necessary to achieve good pose estimation. In addition, due to differences in simulated and real-world dynamics, we do not expect the optimal θ for reducing \mathcal{C} to be their corresponding real-world values.

To optimize the parameters of the K simulations and make their simulated states more closely track that of the real world, we evaluate three derivative-free, sample-based optimizers:

1) *Weighted Resampling (WRS)*: WRS forms a probability mass function (PMF) over the existing simulation states $s^{(1:K)}$ and samples K times with replacement from that distribution to form the next set of simulations. To form the PMF, WRS applies softmax over the simulation costs:

$$P(i) = \frac{\exp -\lambda(C_i - \min_j C_j)}{\sum_{i=1}^K \exp -\lambda(C_i - \min_j C_j)} \quad (4)$$

Here, λ is a temperature hyperparameter that determines the sharpness of the distribution. After resampling, we perform exploration on all simulations by perturbing 1) the simulation parameters θ and 2) the object pose.

Simulation parameters are perturbed by sampling from an isotropic normal distribution around the previous parameters: $\theta_{\tau+1}^{(i)} \sim \mathcal{N}(\theta_{\tau}^{(i)}, \Sigma_{\theta})$, where Σ_{θ} is predefined. The subscript τ denotes the optimizer update step (after τ update steps the simulation has ran for a total of τT time steps).

For object pose perturbation, adding noise to the pose directly while the object is held in-hand is impractical; most delta poses would result in mesh penetration and are hence invalid. This issue was noted in [15], [37], and like those works we perturb the objects by applying perturbation forces to the object in each simulation, with each force sampled from a zero-mean isotropic normal distribution $v^{(i)} \sim \mathcal{N}(0, \Sigma_v)$.

2) *Relative Entropy Policy Search (REPS)*: Unlike [38], which also uses REPS [39] to tune simulation parameters to address the sim-to-real gap, we use a sample-based variant of REPS that computes weights for each simulation and samples from a distribution formed by the softmax of those weights. Whereas WRS uses a fixed λ parameter to shape the distribution, REPS solves for an adaptive temperature parameter η that best improves the performance of the overall distribution subject to ϵ , a constraint on the KL-divergence between the old and updated sample distributions.

To use REPS, we reformulate the costs as rewards by setting $R_i = \max_j C_j + \min_j C_j - C_i$. We compute η at every step by optimizing the dual function $g(\eta)$, and then we use η to form the PMF:

$$\eta^* = \arg \min_{\eta} \eta \epsilon + \eta \log \frac{1}{K} \sum_{i=1}^K \exp \frac{R_i}{\eta} \quad (5)$$

$$P(i) = \frac{\exp \frac{R_i}{\eta^*}}{\sum_{j=1}^K \exp \frac{R_j}{\eta^*}} \quad (6)$$

After resampling, every simulation is perturbed in the same manner as in WRS.

3) *Population-Based Optimization (PBO)*: Inspired by Population-Based Training (PBT) [40], this algorithm first ranks all simulations by their average costs and finds the top K_{best} simulations with the lowest costs. Then, it 1) exploits by replacing the remaining $K - K_{\text{best}}$ simulations with copies of the K_{best} ones, sampled with replacement, and 2) explores by perturbing the K_{best} simulations in the same way as WRS.

PBO effectively uses a shaped cost that depends only on the relative ordering of the simulation costs and not their magnitudes, potentially making the optimizer more robust to noisy costs.

E. Hyperparameters

Each of the proposed optimizers has a distribution-shaping hyperparameter used to balance exploration with exploitation. There are 5 additional hyperparameters for our proposed framework:

- T - the time steps the algorithm waits for every update.

- K - the number of concurrent simulations.
- θ_0 - the initial normal distribution over simulation parameters.
- Σ_p - the diagonal covariance matrix for the normal distribution over initial pose perturbation.
- Σ_{θ} and Σ_v - the diagonal covariances of normal distributions of perturbations used for exploration.

A larger K is generally better than a smaller K , with the caveat that the resulting simulation is slower and may not be practical in application. Σ_p should be large enough such that the actual initial pose is well represented in the initial pose distribution. However, K should be increased with a larger Σ_p and the covariance of θ_0 to ensure that the density of the samples is high enough to capture a wider distribution.

We note two additional trade-offs with these hyperparameters. One is the exploration-exploitation trade-off in the context of optimizing for θ , and the other is the trade-off between optimizing for θ and for $p_t^{(i^*)}$. Making Σ_{θ} or Σ_v wider will increase the speed at which the set of simulation parameters “move,” and the optimizer will explore more than it exploits. Increasing T improves the optimization for θ as the optimizer has more samples to evaluate each simulation. However, updating the simulation parameters too slowly might lead to drift in pose estimation if the least-cost simulation is sufficiently different from the real world, potentially leading to divergent behavior. The worst-case divergent behavior occurs when force perturbation or some simulation parameters lead to an irrecoverable configuration, where the object falls out of the hand or brings the object into a pose such that small force perturbations cannot bring it back to the correct pose. It is acceptable if a few samples become divergent. Their costs will be high, so they will be discarded and replaced by ones that are not divergent during optimizer updates.

There are connections between our approach and previous works that use particle filtering. However, prior works were mostly applied in the static-grasp setting where the forward model of the particle filter is a constant model. Instead, we track the object during manipulation and use a physics simulator as the forward model. In addition to tracking the object pose, the proposed algorithm also identifies the context of the forward model by tuning the simulation parameters θ , which are not affected by the forward or observation models. We focus on optimizers based on discrete samples and not continuous distributions, such as the popular Covariance Matrix Adaptation Evolution Strategy (CMA-ES) [41], because we cannot easily sample in-hand object poses from some distribution due to complicated mesh-penetration constraints imposed by contacts.

IV. EXPERIMENTS

We evaluate the performance of our proposed approach with both simulation and the real-world experiments using an Allegro Hand mounted on a Kuka IIWA7 robot arm. In-hand object manipulation trajectories are first collected with a hand-tracking teleoperation system, and we evaluate pose estimation errors by running our proposed algorithms offline

Algorithm 1 Sample-Based Optimization for In-Hand Object Pose Tracking via Contact Feedback

Input: \mathcal{C} , T , K , s_0 , p_0 , Σ_p , θ_0 , Σ_θ , Σ_v , opt

- 1: Initialize K simulations with s_0 , $p_0^{(i)} \sim \mathcal{N}(p_0, \Sigma_p)$, and $\theta_0 \sim \theta_0$
 - 2: Initialize index of current best pose estimate $i^* \leftarrow 0$
 - 3: **while** TRUE **do**
 - 4: **for** $t \in \{1, \dots, T\}$ **do**
 - 5: Obtain u_t and $o_t^{(gt)}$ from real robot
 - 6: **for** $i \in \{1, \dots, K\}$ **do**
 - 7: Step simulation $s_{t+1}^{(i)} = f(s_t^{(i)}, u_t, \theta_\tau^{(i)})$
 - 8: Obtain simulation observation $o_t^{(i)}$
 - 9: Compute average cost across time window:

$$C_i = \frac{1}{t} \sum_{\Delta t=0}^{t-1} \mathcal{C}(o_{t-\Delta t}^{(i)}, o_{t-\Delta t}^{(gt)})$$
 - 10: **end for**
 - 11: Update index of best pose estimate:

$$i^* \leftarrow \arg \min_i C_i$$
 - 12: **end for**
 - 13: Update simulations according to optimizer:

$$s_{\tau+1}^{(1:K)}, \theta_{\tau+1}^{(1:K)} \leftarrow opt(s^{(1:K)}, \theta_\tau^{(1:K)}, C_{1:K}, \Sigma_\theta, \Sigma_v)$$
 - 14: **end while**
-

against the collected trajectories. These trajectories start and end with the object not in the hand and not in occlusion. Because we have access to ground truth object poses in simulation experiments, we perform detailed ablation studies in simulation to study the effects of different hyperparameters on algorithm performance. While we can compare pose estimation errors during the entire trajectory in simulation experiments, this is not possible in the real world. For real-world experiments, we use PoseRBPf [3], a recent RGB-D, particle-filter based pose estimation algorithm to obtain initial and final object poses. We treat these initial and final object poses as ground truth and compare the final pose with the one predicted by our proposed algorithm.

We mount the 4-finger 16-DoF Allegro hand on the 7-Dof Kuka IIWA7 robot arm. To obtain contact feedback in the real world, we attach SynTouch BioTac sensors to each of the finger tips. While BioTac sensors does not explicitly give force or slippage information, past works have studied how to extract such information from the sensor’s raw electrode readings to predict contact force [42], [43], slip direction [44], and grasp stability [45]–[47]. For real-world experiments, we use the trained model from [43] to estimate force vectors c_t , but currently we do not estimate slippage from the BioTac sensors, so the cost function in real-world experiments do not contain the slippage terms. Simulations were conducted on a computer with an Nvidia GTX 1080 Ti GPU, Intel i7-8700K CPU, and 16GB of memory.

We use 3 objects from the Yale-Columbia-Berkeley (YCB) objects dataset (the spam can, foam brick, and toy banana), with models, textures, and point clouds obtained from the dataset published in [1]. These objects were chosen because they fit the size of the Allegro hand and are light enough so that robust precision grasps can be formed (we emptied the spam can to reduce its weight).




Fig. 2. Pose tracking error comparison across different optimizers for initial pose noise levels in all simulation experiments. The length of the black vertical lines denote 1 standard deviation. Optimizer methods generally have lower mean and variance of ADD, but their relative ordering varies depends on the amount of initial pose noise. REPS and PBO achieve the best ADD performance in the medium noise case - 5.8mm and 5.9mm respectively.

For each object, in both simulation and real-world experiments, we give 2 demonstrations of 2 types of manipulation trajectories: 1) pick and place with finger-grasp and in-hand object rotation, and 2) the same but with finger tips breaking and re-establishing contact during the grasp (finger gaing). This gives a total of 24 trajectories for analysis for both simulation and real-world experiments. In both trajectory types the object undergoes translational and rotational slippage from both inertial forces and push-contacts with the table. Each trajectory lasts about a minute. Given that we can run the pose estimation algorithm at about 30Hz, we obtain a total of about 2k frames per trajectory.

The teleoperation system is described in detail in a concurrent work under review. The input to the system is a point cloud of the hand of the human demonstrator. Then, a neural network based on PointNet++ [48] maps the point cloud to an estimate of the hand’s pose relative to the camera as well as the joint angles of the hand. These estimates along with an articulated hand model [49] and the original point cloud are then given to DART, which performs tracking by refining upon the neural network estimates. Finally, to perform kinematic retargetting, we solve an optimization problem that finds the Allegro hand joint angles that result in finger tip poses close to those of the human hand.

In addition to our proposed optimizers (WRS, REPS, PBO), we also evaluate the following two baselines: Open Loop (OLP) and Identity (EYE). OLP tracks the object pose with 1 simulation. EYE is initialized with a set of noisy initial poses and always picks the pose of the lowest-cost simulation, but it does not perform any resampling or optimizer updates.

Similar to previous works [1]–[3], we use Average Distance Deviation (ADD) [50] as the evaluation metric. ADD computes the average distance between corresponding points in the object point cloud situated at the ground truth pose and at the predicted pose. Unlike [1]–[3], we do not use its symmetric variant, ADD-S, which does not penalize pose differences across object symmetries (*e.g.* for poses of a sphere that share the same translation, any rotation difference gives 0 error). This is desirable for resolving visual ambiguities for pose registration but not for tracking.

A. Simulation Experiments

For simulation experiments we build upon our previous work in GPU-Accelerated robotics simulation [51]. The arm




Fig. 3. Ablation studies in simulations across different initial pose noise levels. Exploring too little or too much (how wide the Σ_θ and Σ_v are.) typically result in higher mean and variance of ADD. In all other experiments we use the “med” level of exploration. Increasing the number of simulations tend to reduce mean and variance of ADD. Contacts feedback help reduce the mean and variance of ADD, especially as noise level increases. While using slip detection helps in the case of low and medium initial pose noise, the advantage disappears when the noise is high.

and hand in the simulation is controlled via a joint-angle PD controller, and we tuned the controller’s gains so that the joint angle step responses are similar to those of the real robot. To speed up simulation, we simplify the collision meshes of the robot and objects. This is done first by applying TetWild [52] which gives a mesh with triangles that are more equilateral, then with Quadric Edge Collapse Decimation¹. Each simulation generates at most 200 contacts during manipulation, and we run $K = 40$ simulations at 30Hz.

We performed simulation experiments with varying amounts of initial pose noise. Three levels were tested: “Low” has a translation standard deviation of 1mm and a rotation standard deviation of 0.01 radians. “Med” is 5mm and 0.1 radians, and “High” is 10mm 1 radian.

See Figure 2 for a comparison of the optimizers on tracking in-hand object poses across all the simulation trajectories. ADD increases as the initial pose error increases, and the mean ADD for the optimizer-based methods tends to be lower. While EYE sometimes achieves comparable mean ADD with the optimizer methods, the latter ones generally have much smaller error variance and max error. This result is expected as the optimizers focus the distribution of simulations towards better performing ones over time. In the medium noise case, REPS and PBO achieve the best ADD with a mean of 5.8mm and 5.9mm respectively.

See Figure 3 for results of ablation studies in simulation performed over the hyperparameters governing exploration distance (how much simulations are perturbed), the number of parallel simulations, and whether or not contact and slip detection feedback is used in the cost function.

B. Real-World Experiments

We evaluate our algorithm on real-world trajectories similar to those collected in simulation. We use PoseRBPF to register the object pose in the first and last frames of a trajectory. The initial pose estimate is used to initialize the simulations, while the last one is used to evaluate the accuracy of our contacts-based pose tracking algorithm. Unlike simulation experiments, the real-world experiments initialize the object pose by sampling from the distribution over object poses from PoseRBPF, so the initial pose samples correspond to the uncertainties of the vision-based pose estimation algorithm.




Fig. 4. Real-world object pose tracking experiments across 3 objects. The adaptive optimizers generally perform better than the baselines, although all methods struggle on the banana. Due to the banana’s long moment arm, precision grasps with the fingertips tend to be very unstable, and the banana undergoes significant rotational slippage or can slip out of hand in ways that are sensitive to initial conditions. The best ADD achieved with Foam is 14.1mm by PBO, and with Spam is 12.2mm by REPS.

See Figure 4 for real-world experiment results. The ADDs are higher than those from simulation experiments. This is due to both that the real-world dynamics is more dissimilar with simulations than are simulations with different parameters, and that real-world observations are noisier than those in simulations. We observe that no optimizer is able to track the toy banana for the real-world data. The object’s long moment arm and low friction coefficient makes its slippage behavior difficult to model precisely. This is a failure mode of our algorithm, where if all of the simulations become divergent (e.g. the banana rotates in the wrong direction, or falls out of hand), then the algorithm cannot recover in subsequent optimizer updates. The best ADD achieved with Foam is 14.1mm by PBO, and with Spam is 12.2mm by REPS.

V. CONCLUSION

We introduce a sample-based optimization algorithm for tracking in-hand object poses during manipulation via contact feedback and GPU-accelerated robotic simulation. The parallel simulations concurrently maintains many beliefs about the real world and model object pose changes that are caused by complex contact dynamics. The optimization algorithm tunes simulation parameters during object pose tracking to further improve tracking performance. In future work, we plan to integrate contact sensing with vision-based pose tracking in-the-loop.

ACKNOWLEDGMENT

The authors thank Renato Gasoto, Miles Macaklin, Tony Scudiero, Jonathan Tremblay, Stan Birchfield, Qian Wan, and Mike Skolones for their help with GPU-accelerated robotic simulations, Xinke Deng and Arsalan Mousavian for PoseRBPF, and Balakumar Sundaralingam and Tucker Hermans for BioTac sensing. This work is in part supported by the NSF Graduate Research Fellowship Program Grant No. DGE 1745016 and the Office of Naval Research Grant No. N00014-18-1-2775.

¹<https://help.sketchfab.com/hc/en-us/articles/205852789-MeshLab-Decimating-a-model>

REFERENCES

- [1] Y. Xiang, T. Schmidt, V. Narayanan, and D. Fox, "PoseCNN: A convolutional neural network for 6d object pose estimation in cluttered scenes," *RSS*, 2018.
- [2] Y. Li, G. Wang, X. Ji, Y. Xiang, and D. Fox, "DeepIM: Deep iterative matching for 6d pose estimation," in *ECCV*, 2018.
- [3] X. Deng, A. Mousavian, Y. Xiang, F. Xia, T. Bretl, and D. Fox, "PoseRBPF: A rao-blackwellized particle filter for 6d object pose tracking," *RSS*, 2019.
- [4] M. Andrychowicz, B. Baker, M. Chociej, R. Jozefowicz, B. McGrew, J. Pachocki, A. Petron, M. Plappert, G. Powell, A. Ray, *et al.*, "Learning dexterous in-hand manipulation," *arXiv preprint arXiv:1808.00177*, 2018.
- [5] A. Rajeswaran, V. Kumar, A. Gupta, G. Vezzani, J. Schulman, E. Todorov, and S. Levine, "Learning complex dexterous manipulation with deep reinforcement learning and demonstrations," *RSS*, 2018.
- [6] H. Zhu, A. Gupta, A. Rajeswaran, S. Levine, and V. Kumar, "Dexterous manipulation with deep reinforcement learning: Efficient, general, and low-cost," *ICRA*, 2019.
- [7] M. Macklin, K. Erleben, M. Müller, N. Chentanez, S. Jeschke, and V. Makovychuk, "Non-smooth newton methods for deformable multi-body dynamics," *arXiv preprint arXiv:1907.04587*, 2019.
- [8] C. Choi, J. Del Preto, and D. Rus, "Using vision for pre-and post-grasping object localization for soft hands," in *International Symposium on Experimental Robotics*. Springer, 2016.
- [9] M. Kovic, D. Kragic, and J. Bohg, "Learning to estimate pose and shape of hand-held objects from rgb images," *arXiv preprint arXiv:1903.03340*, 2019.
- [10] S. Luo, J. Bimbo, R. Dahiya, and H. Liu, "Robotic tactile perception of object properties: A review," *Mechatronics*, 2017.
- [11] Y. Chebotar, O. Kroemer, and J. Peters, "Learning robot tactile sensing for object manipulation," in *IROS*, 2014.
- [12] R. Li, R. Platt, W. Yuan, A. ten Pas, N. Roscup, M. A. Srinivasan, and E. Adelson, "Localization and manipulation of small parts using gelsight tactile sensing," in *IROS*, 2014.
- [13] M. A. Lee, Y. Zhu, P. Zachares, M. Tan, K. Srinivasan, S. Savarese, L. Fei-Fei, A. Garg, and J. Bohg, "Making sense of vision and touch: Learning multimodal representations for contact-rich tasks," *arXiv preprint arXiv:1907.13098*, 2019.
- [14] 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*, 2015.
- [15] S. Li, S. Lyu, J. Trinkle, and W. Burgard, "A comparative study of contact models for contact-aware state estimation," in *IROS*, 2015.
- [16] A. Lambert, M. Mukadam, B. Sundaralingam, N. Ratliff, B. Boots, and D. Fox, "Joint inference of kinematic and force trajectories with visuo-tactile sensing," *arXiv preprint arXiv:1903.03699*, 2019.
- [17] K.-T. Yu and A. Rodriguez, "Realtime state estimation with tactile and visual sensing. application to planar manipulation," in *ICRA*, 2018.
- [18] C. Corcoran and R. Platt, "A measurement model for tracking hand-object state during dexterous manipulation," in *ICRA*, 2010.
- [19] R. Platt, F. Permenter, and J. Pfeiffer, "Using bayesian filtering to localize flexible materials during manipulation," *IEEE Transactions on Robotics*, 2011.
- [20] M. Chalon, J. Reinecke, and M. Pfanne, "Online in-hand object localization," in *IROS*, 2013.
- [21] L. Zhang, S. Lyu, and J. Trinkle, "A dynamic bayesian approach to real-time estimation and filtering in grasp acquisition," in *ICRA*, 2013.
- [22] S. Javdani, M. Klingensmith, J. A. Bagnell, N. S. Pollard, and S. S. Srinivasa, "Efficient touch based localization through submodularity," in *ICRA*, 2013.
- [23] J. Bimbo, P. Kormushev, K. Althoefer, and H. Liu, "Global estimation of an objects pose using tactile sensing," *Advanced Robotics*, 2015.
- [24] G. Vezzani, U. Pattacini, G. Battistelli, L. Chisci, and L. Natale, "Memory unscented particle filter for 6-dof tactile localization," *IEEE Transactions on Robotics*, 2017.
- [25] D. Álvarez, M. A. Roa, and L. Moreno, "Tactile-based in-hand object pose estimation," in *Iberian Robotics conference*. Springer, 2017.
- [26] Y. Ding, J. Bonse, R. Andre, and U. Thomas, "In-hand grasping pose estimation using particle filters in combination with haptic rendering models," *International Journal of Humanoid Robotics*, 2018.
- [27] P. Hebert, N. Hudson, J. Ma, and J. Burdick, "Fusion of stereo vision, force-torque, and joint sensors for estimation of in-hand object location," in *ICRA*, 2011.
- [28] A. Petrovskaya and O. Khatib, "Global localization of objects via touch," *IEEE Transactions on Robotics*, 2011.
- [29] B. Saund, S. Chen, and R. Simmons, "Touch based localization of parts for high precision manufacturing," in *ICRA*, 2017.
- [30] Z. Pezzementi, C. Reyda, and G. D. Hager, "Object mapping, recognition, and localization from tactile geometry," in *ICRA*, 2011.
- [31] J. Bimbo, S. Luo, K. Althoefer, and H. Liu, "In-hand object pose estimation using covariance-based tactile to geometry matching," *IEEE Robotics and Automation Letters*, 2016.
- [32] G. Izatt, G. Mirano, E. Adelson, and R. Tedrake, "Tracking objects with point clouds from vision and touch," in *ICRA*, 2017.
- [33] T. Schmidt, K. Hertkorn, R. Newcombe, Z. Marton, M. Suppa, and D. Fox, "Depth-based tracking with physical constraints for robot manipulation," in *ICRA*, 2015.
- [34] T. Schmidt, R. A. Newcombe, and D. Fox, "DART: Dense articulated real-time tracking," in *RSS*, 2014.
- [35] M. Pfanne, M. Chalon, F. Stulp, and A. Albu-Schäffer, "Fusing joint measurements and visual features for in-hand object pose estimation," *IEEE Robotics and Automation Letters*, 2018.
- [36] E. Eade, "Lie groups for 2d and 3d transformations," *URL http://ethaneade.com/lie.pdf, revised Dec*, 2013.
- [37] D. J. Duff, T. Mörwald, R. Stolkin, and J. Wyatt, "Physical simulation for monocular 3d model based tracking," in *ICRA*, 2011.
- [38] Y. Chebotar, A. Handa, V. Makovychuk, M. Macklin, J. Issac, N. Ratliff, and D. Fox, "Closing the sim-to-real loop: Adapting simulation randomization with real world experience," *ICRA*, 2019.
- [39] J. Peters, K. Mulling, and Y. Altun, "Relative entropy policy search," in *Twenty-Fourth AAAI Conference on Artificial Intelligence*, 2010.
- [40] M. Jaderberg, V. Dalibard, S. Osindero, W. M. Czarnecki, J. Donahue, A. Razavi, O. Vinyals, T. Green, I. Dunning, K. Simonyan, *et al.*, "Population based training of neural networks," *arXiv preprint arXiv:1711.09846*, 2017.
- [41] N. Hansen, S. D. Müller, and P. Koumoutsakos, "Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (cma-es)," *Evolutionary computation*, vol. 11, no. 1, pp. 1–18, 2003.
- [42] K. Hang, M. Li, J. A. Stork, Y. Bekiroglu, F. T. Pokorny, A. Billard, and D. Kragic, "Hierarchical fingertip space: A unified framework for grasp planning and in-hand grasp adaptation," *IEEE Transactions on robotics*, vol. 32, no. 4, pp. 960–972, 2016.
- [43] B. Sundaralingam, A. Handa, B. Boots, T. Hermans, S. Birchfield, N. Ratliff, D. Fox, *et al.*, "Robust learning of tactile force estimation through robot interaction," *ICRA*, 2019.
- [44] M. A. Abd, I. J. Gonzalez, T. C. Colestock, B. A. Kent, and E. D. Eneberg, "Direction of slip detection for adaptive grasp force control with a dexterous robotic hand," in *2018 IEEE/ASME International Conference on Advanced Intelligent Mechatronics (AIM)*. IEEE, 2018, pp. 21–27.
- [45] Z. Su, K. Hausman, Y. Chebotar, A. Molchanov, G. E. Loeb, G. S. Sukhatme, and S. Schaal, "Force estimation and slip detection/classification for grip control using a biomimetic tactile sensor," in *Humanoids*. IEEE, 2015, pp. 297–303.
- [46] Y. Chebotar, K. Hausman, Z. Su, A. Molchanov, O. Kroemer, G. Sukhatme, and S. Schaal, "Bigs: Biotac grasp stability dataset," in *ICRA 2016 Workshop on Grasping and Manipulation Datasets*, 2016.
- [47] F. Veiga, B. B. Edin, and J. Peters, "In-hand object stabilization by independent finger control," *arXiv preprint arXiv:1806.05031*, 2018.
- [48] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, "Pointnet++: Deep hierarchical feature learning on point sets in a metric space," in *Advances in neural information processing systems*, 2017, pp. 5099–5108.
- [49] Y. Hasson, G. Varol, D. Tzionas, I. Kalevatykh, M. J. Black, I. Laptev, and C. Schmid, "Learning joint reconstruction of hands and manipulated objects," in *CVPR*, 2019.
- [50] S. Hinterstoisser, V. Lepetit, S. Ilic, S. Holzer, G. Bradski, K. Konolige, and N. Navab, "Model based training, detection and pose estimation of texture-less 3d objects in heavily cluttered scenes," in *Asian conference on computer vision*. Springer, 2012, pp. 548–562.
- [51] J. Liang, V. Makovychuk, A. Handa, N. Chentanez, M. Macklin, and D. Fox, "GPU-accelerated robotic simulation for distributed reinforcement learning," *CoRL*, 2018.
- [52] Y. Hu, Q. Zhou, X. Gao, A. Jacobson, D. Zorin, and D. Panozzo, "Tetrahedral meshing in the wild," *ACM Trans. Graph.*, vol. 37, no. 4, pp. 60:1–60:14, July 2018. [Online]. Available: <http://doi.acm.org/10.1145/3197517.3201353>