

# MidasTouch: Monte-Carlo inference over distributions across sliding touch

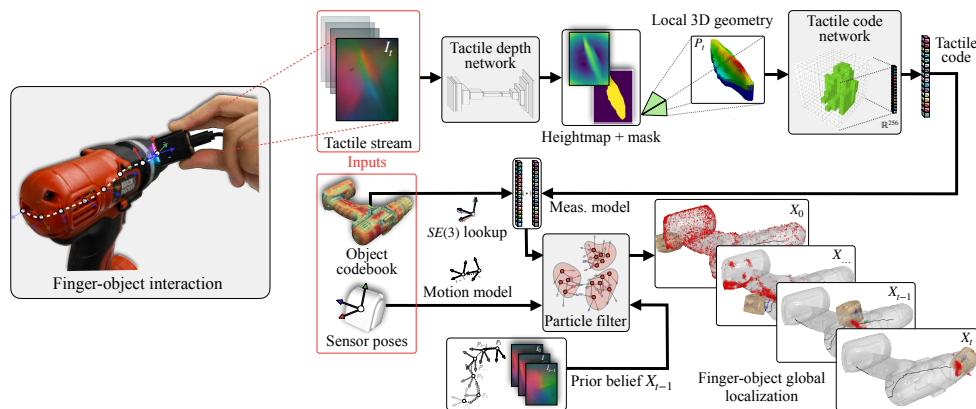
Sudharshan Suresh<sup>1,2</sup>, Zilin Si<sup>1</sup>, Stuart Anderson<sup>2</sup>, Michael Kaess<sup>1</sup>, Mustafa Mukadam<sup>2</sup>

<sup>1</sup>Carnegie Mellon University, <sup>2</sup>Meta AI

<https://suddhu.github.io/midastouch-tactile>

**Abstract:** We present MidasTouch, a tactile perception system for online global localization of a vision-based touch sensor sliding on an object surface. This framework takes in posed tactile images over time, and outputs an evolving distribution of sensor pose on the object’s surface, without the need for visual priors. Our key insight is to estimate local surface geometry with tactile sensing, learn a compact representation for it, and disambiguate these signals over a long time horizon. The backbone of MidasTouch is a Monte-Carlo particle filter, with a measurement model based on a tactile code network learned from tactile simulation. This network, inspired by LIDAR place recognition, compactly summarizes local surface geometries. These generated codes are efficiently compared against a pre-computed tactile codebook per-object, to update the pose distribution. We further release the YCB-Slide dataset of real-world and simulated forceful sliding interactions between a vision-based tactile sensor and standard YCB objects. While single-touch localization can be inherently ambiguous, we can quickly localize our sensor by traversing salient surface geometries.

**Keywords:** Tactile perception, Localization, 3D deep learning



**Figure 1:** MidasTouch performs online global localization of a vision-based touch sensor on an object surface during sliding interactions. Given posed tactile images over time, this system leverages local surface geometry within a nonparametric particle filter to generate an evolving distribution of sensor pose on the object’s surface.

## 1 Introduction

Interactive perception is the Catch-22 [1] of robotics; vision can be used to track objects for downstream contact-rich interactions, but such interactions can occlude visual tracking. In particular, end-effector-object relative positioning is crucial for contact-rich policies like sliding [2–6], in-hand manoeuvres [7–10], finger-gaiting [11], and multi-finger enclosure [12]. Additionally, vision is often affected by object transparency, specularities, and poor scene illumination. With high-dimensional vision-based tactile sensors [13–15], we now have a window into local object interactions. While tactile images from these sensors capture local surface geometry, they lack global context necessary for relative pose tracking. We address this challenge with **MidasTouch**, an online tactile perception system that tracks the evolving pose distribution of a vision-based touch sensor sliding across known objects. We acquire global context by integrating local observations over long horizons.

The form-factor of vision-based tactile sensors has restricted prior methods to small parts [16–18] or local tracking [19, 20]. For everyday objects, we can leverage compositionality [21]; each is made up of local geometries that are spatially fixed relative to one another. Consider a simple mug: made up of a curved body, flat base, rounded handle, and sharp lip. Without global context, a single-touch is ambiguous: a perceived sharp edge could lie anywhere along the lip of the mug. Such a likelihood distribution is spread across the object’s surface and not unimodal, but interaction over long time horizons can disambiguate it. This mirrors haptic *apprehension*, or the exploratory procedures humans perform when presented with a familiar object [22].

We approach this as an analogue to mobile robot Monte-Carlo filtering [23], but instead apply it on the surface manifold of the object. Just as a mobile robot has access to detailed floor plans, odometry, and cameras, manipulators have access to object meshes, end-effector poses, and vision-based touch. This is applicable to environments with known object models like households, warehouses, and factories, further facilitated by large-scale scanned object datasets [24, 25]. While priors from vision can bootstrap our system [26], they are not a prerequisite for global estimation.

In addition to open-sourcing MidasTouch, we release a comprehensive real-world and simulated dataset of sliding DIGIT [15] interactions across standard YCB objects [24] with ground-truth. This serves as both an evaluation of MidasTouch, and as a benchmark currently lacking in the tactile sensing community. Both are accessed from our project [website](#). Specifically, our contributions are:

1. Online particle filtering over posed tactile images for a distribution of finger-object poses,
2. Learned embeddings for vision-based touch using local surface geometry,
3. The YCB-Slide dataset of DIGIT and YCB object interactions for evaluation and benchmarking.

Our framework relies on recent developments in tactile sensing, rendering, and robotics: **(i)** vision-based tactile sensors [13–15, 27–31] like the GelSight and DIGIT have the requisite spatial acuity to discern local geometric features, **(ii)** tactile simulation [32–34] with realistic rendering enables learning tactile observation models and precomputing object-specific interactions with sizeable data, which would be infeasible in the real-world, and **(iii)** learned models for 3D place recognition [35–37] spawned by ubiquitous LIDAR and RGB-D data that can be extended to tactile sensing.

## 2 Related work

**Tactile pose inference:** Binary contact sensors have been used in conjunction with particle filters for global estimation of robot hands relative to simple geometries [38–43]. These methods require a large amount of touches, but serve as a touchstone for Monte-Carlo methods that have seen great success in mobile robotics [23, 44]. With tactile arrays, local patch measurements are integrated for planar estimation [45, 46] or for 3D alignment [47]. The synergy of vision and touch gives much needed global context [48, 49] but is outside our current work’s scope.

With vision-based touch, Li et al. [16] show planar small-part localization by directly computing the homography on GelSight heightmaps. Relative pose-tracking has also been learnt directly from tactile images either via recurrent [50], or auto-encoder [15, 19] networks. Additionally, local tracking has been explored with online factor-graph optimization [19, 20, 51]. The aforementioned approaches are unimodal and rely on good pose initializations; MidasTouch is nonparametric and requires no such knowledge. Additionally, Kelestemur et al. [52] show category-level object pose estimation with a parallel jaw gripper.

Closely related is Tac2Pose [17, 18], which performs pose estimation of small-parts with the GelSlim [14]. It produces a distribution of object poses learned in tactile simulation from contact shapes. Similarly, Gao et al. [53] estimates contact location on objects through vision-based touch and audio with a small, discrete set of measurements. We differentiate our work in both context and approach: **(i)** we filter over long time horizons rather than single/multi touch predictions, **(ii)** tactile embeddings are learned from local surface geometry rather than images, and **(iii)** we consider objects considerably larger than the robot finger.

**Place recognition for touch:** A compact representation for tactile images enables easy frame-to-frame or frame-to-model tracking. Inspired by the computer vision community, a popular interme-

diate is a learned embedding from either RGB [15, 19] or binary contact masks [17, 18]. Realistic tactile simulators, like TACTO [32], enable training these models to the scale and generality of arbitrary real-world interactions. For example, Bauza et al. [18] learn object-specific embeddings to match 2D contact shapes.

For contact-rich manipulation, local 3D geometry is a natural candidate, and iterative closest point (ICP) has shown promise for frame-to-frame tracking [20, 54]. For vision-based touch, 3D geometry is obtained either via photometric stereo [13, 55, 56], or image-to-heightmap models [20, 31, 57–59]. However, ICP is only suitable for local tracking and not global localization: it is sensitive to initialization and is intractable to scale as a sampling-based measurement model.

Succinctly, what is an accurate and efficient similarity metric for tactile geometry? With the ubiquity of LIDAR/RGB-D data, matching unordered point sets is vital for place recognition in the SLAM community [36, 37]. Following the seminal PointNet [60], Choy et al. [61] later developed efficient and expressive sparse 3D convolution. This backbone was used to aggregate local point features to global point-cloud embeddings for LIDAR place-recognition [37]. We show tactile geometries can be compacted into codes with the same architecture, for easy frame-to-model queries.

### 3 Problem formulation

For a vision-based tactile sensor dynamically sliding along a known object’s surface, our goal is to track the distribution of 6D sensor pose  $\mathbf{x}_t \in SE(3)$  in the object-centric frame. Such pose context is useful for downstream planning and control, for instance manipulating an object in-hand [7]. The sensor is affixed on a robot finger and we have access to the end-effector pose, but its relative position and orientation on the object’s surface is unknown. At each timestep  $t$ , our measurements are a tactile image, and noisy sensor pose in the robot’s reference-frame,  $\mathbf{z}_t = \{\mathbf{I}_t \in \mathbb{R}^{240 \times 320 \times 3}, \mathbf{p}_t \in SE(3)\}$ . For simplicity, we assume the object is stationary and sliding is a single continuous contact interaction. The dimensions of the objects are much larger than the sensor’s contact area. We assume an uninformative initial prior for  $\mathbf{x}_t$ , but later show the benefit of visual priors.

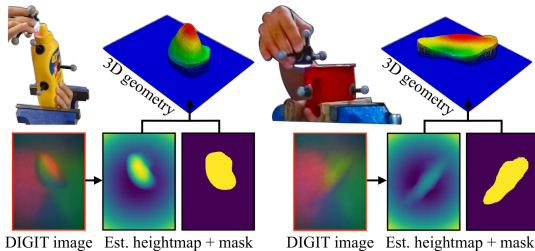
### 4 Global localization during sliding touch

MidasTouch comprises three distinct modules, as illustrated in Figure 1: a tactile depth network (TDN: Section 4.1), tactile code network (TCN: Section 4.2), and particle filter (Section 4.3). At a high-level, the TDN first converts a tactile image  $\mathbf{I}_t$  into its local 3D geometry  $\mathbf{P}_t$  via a learned observation model. The 3D information is then condensed into a tactile code  $\mathbf{E}_t$  by the TCN through a sparse 3D convolution network. Finally, the downstream particle filter uses these learned codes in its measurement model, and outputs a sensor pose distribution that evolves over time. Throughout this work, we use 10 YCB objects with diverse geometries in our tests: sugar\_box, tomato\_soup\_can, mustard\_bottle, bleach\_cleanser, mug, power\_drill, scissors, adjustable\_wrench, hammer, and baseball. In Appendix G, we further demonstrate the method on small parts.

#### 4.1 Tactile depth network (TDN)

The tactile depth network learns the inverse sensor model to recover local 3D geometry from a tactile image. We adapt a fully-convolutional residual network [59, 62], and train it in a supervised fashion to predict local heightmaps from tactile images.

**Network and training:** This network uses a ResNet-50 backbone, with a series of up-projection blocks. With the optical tactile simulator TACTO [32], we render a large collection of DIGIT images with ground-truth heightmaps. The images are from simulated interaction with 40 YCB [24] object meshes, while holding out all test objects. The images are rendered at 5000 different poses per object: randomizing for contact point, orientation, and indentation depth. For sim2real transfer, we



**Figure 2:** Real-world DIGIT images from interactions in the YCB-Slide dataset (more in Appendix A). Our tactile depth network (Section 4.1) is trained in simulation, and predicts 3D geometry given input tactile image.

calibrate TACTO with real-world data from multiple independent DIGITs, and apply random lighting augmentations (refer Appendix A). Per object, the split of train-validation-test is 4000:500:500.

**Image to 3D:** At each timestep, we pass the  $240 \times 320$  RGB tactile image  $\mathbf{I}_t$  through the network to get heightmap  $\mathbf{H}_t$ . To remove non-contact areas we compute a mask  $\mathbf{C}_t$  by depth thresholding, and get  $\hat{\mathbf{H}}_t = \mathbf{H}_t \odot \mathbf{C}_t$ . Finally, the heightmap is reprojected to 3D via the camera’s known perspective projection model  $\hat{\mathbf{H}}_t \mapsto \mathbf{P}_t$ . Evaluations on the test set show heightmap RMSE of 0.135 mm with respect to ground-truth. Figure 2 shows examples of the TDN on real-world DIGIT interactions.

#### 4.2 Tactile code network (TCN)

This network summarizes large, unordered point clouds of local geometry into a low-dimensional embedding space, or code. If two sensor measurements are nearby in pose-space, they observe similar geometries and therefore, their codes will also be nearby in embedding-space. However, the inverse is not necessarily true: measurements from two opposite corners of a cube may have similar codes, but their corresponding sensor poses are dissimilar. This is an inherent challenge of tactile localization, the so-called *contact non-uniqueness* highlighted by Bauza et al. [18]. In our work getting the most-likely modes of this distribution is sufficient since the downstream particle filter can then disambiguate them temporally.

An analog in the SLAM community is the loop closure problem from 3D data. LIDAR place recognition modules, such as PointNetVLAD [36] and MinkLoc3D [37], use metric learning to learn point cloud similarity. While tactile images differ from natural images, the geometries from LIDAR and tactile sensors (normalized for scale) are similar. The MinkLoc3D architecture, based on MinkowskiNet [61], performs state-of-the-art point cloud retrieval through sparse 3D convolutions.

**Network and training:** The architecture comprises of three components: (i) Voxelization of  $\mathbf{P}_t$  into a quantized sparse tensor  $\hat{\mathbf{P}}_t = \{\langle \hat{x}_i, \hat{y}_i, \hat{z}_i, 0/1 \rangle\}$ , (ii) Feature pyramid network [63] for per-voxel local features  $\hat{\mathbf{P}}_t^f = \{\langle \hat{x}_i, \hat{y}_i, \hat{z}_i, \hat{\mathbf{f}}_i \in \mathbb{R}^{256} \rangle\}$ , and (iii) Generalized-mean pooling for point cloud embedding vector  $\mathbf{E}_t \in \mathbb{R}^{256}$ . We use the pre-trained weights, learned from four comprehensive LIDAR datasets [37], and fine-tune them with TACTO data collected over the 40 YCB objects. This is trained with a contrastive triplet loss, and we accumulate positive and negative local 3D geometries through pose-supervision.



**Figure 3:** Tactile codebook per object visualized as a spectral colorspace map using t-SNE [64]. Each codebook comprises of 50k densely sampled poses with their corresponding 256-dimensional tactile code. Similar hues denote sensor poses that elicit similar tactile codes. We can clearly delineate local geometric features: edges (*sugar\_box*), ridges (*power\_drill*), corners (*scissors*), and complex texture (*baseball*).

**Tactile codebook:** Once we have a code  $\mathbf{E}_t$ , we would like to efficiently compare this with a dense set of contacts. Inspired by prior work [18, 26], we build a tactile codebook comprising of  $M = 50k$  randomized sensor poses on each object’s mesh. We evenly sample these points and normals on the mesh with random orientations and indentations (refer Appendix A). We feed these poses into TACTO to generate a dense set of tactile images for each object.

We pass the generated images through the TDN + TCN to get a codebook for the specific object  $o$ :  $\mathcal{C}_o = \{\langle \mathbf{p}^{[m]}, \mathbf{E}^{[m]} \rangle\}_{m=1}^M$ , where  $\mathbf{p}^{[m]}$  are  $SE(3)$  sensor poses in codebook and  $\mathbf{E}^{[m]}$  are corresponding tactile codes. Figure 3 shows a t-SNE visualization of the codebooks, a colorspace representation of local geometric similarity. We observe regions with similar geometries have identical hues and the traversal of our sensor between these geometries will give us valuable measurement signals.

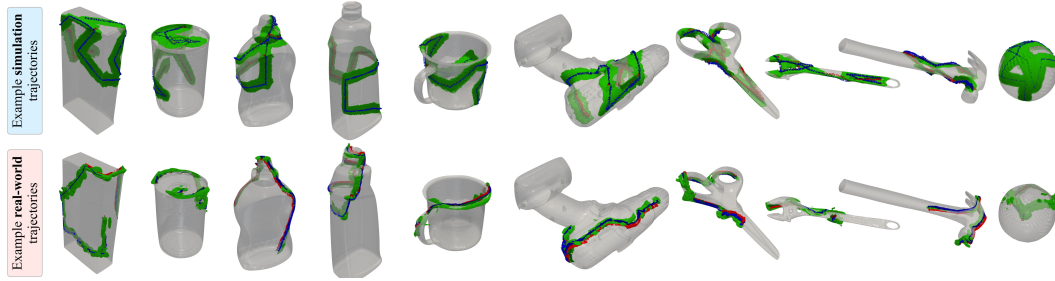
We build a KD-Tree with 6-element vectors:  $\{\langle \mathbf{p}_{\text{trans}}^{[m]}, \alpha \log(\mathbf{p}_{\text{rot}}^{[m]}) \rangle\}_{m=1}^M$  for nearest-neighbor search. Here,  $\log(\cdot)$  is the  $SO(3)$  logarithm map obtained via Theseus [65], and  $\alpha = 0.01$  is the rotation scaling factor. Thus, given a candidate pose  $\mathbf{x}_t^{[i]}$  we do not have to render and obtain



**Hypothesis clustering:** Alongside a distribution of poses, downstream tasks may need distinct pose hypotheses. To achieve this, we hierarchically cluster particles in  $\mathbb{R}^3$  using DB-SCAN [68]. We average the cluster positions and quaternions [69] to get a hypothesis set  $h_t : \{\mathbf{x}_t^1 \dots \mathbf{x}_t^H\}$ .

## 5 The YCB-Slide dataset

To evaluate MidasTouch, and enable further research in tactile sensing, we introduce our YCB-Slide dataset (more in Appendix D). It comprises of DIGIT sliding interactions on the 10 YCB objects from Section 4. We envision this can contribute towards efforts in tactile localization, mapping, object understanding, and learning dynamics models. We provide access to DIGIT images, sensor poses, ground-truth mesh models, and ground-truth heightmaps + contact masks (simulation only).



**Figure 5:** Example sliding trajectories from simulated and real trials on the 10 YCB objects. Overlaid in green are the local 3D geometries captured by the tactile sensor, and the contact poses as RGB coordinate axes.

**Simulated interactions:** We simulate sliding across objects using TACTO, mimicking realistic interaction sequences. The pose sequences are geodesic-paths of fixed length  $L = 0.5\text{m}$ , connecting random waypoints on the mesh. We corrupt sensor poses with zero-mean Gaussian noise  $\sigma_{\text{trans}} = 0.5\text{mm}$ ,  $\sigma_{\text{rot}} = 1^\circ$ . We record five trajectories per-object; 50 interactions in total. A representative sample of the trajectories, along with their accrued tactile geometries are shown in Figure 5.

**Real-world interactions:** In the real-world, we perform sliding experiments through handheld operation of the DIGIT. We keep each YCB object stationary with a heavy-duty bench vise, and slide along the surface and record 60s trajectories at 30Hz. We use an OptiTrack system for timesynced sensor poses, with 8 cameras tracking the reflective markers. We affix six markers on the DIGIT and four on each test object. The canonical object pose is adjusted to agree with the ground-truth mesh models. Minor misalignment of sensor poses from human error are rectified by postprocessing the trajectories to lie on the object surface. We record five logs per-object, for a total of 50 sliding interactions. Our experimental setup and dataset are illustrated in Figures 5 and 6.

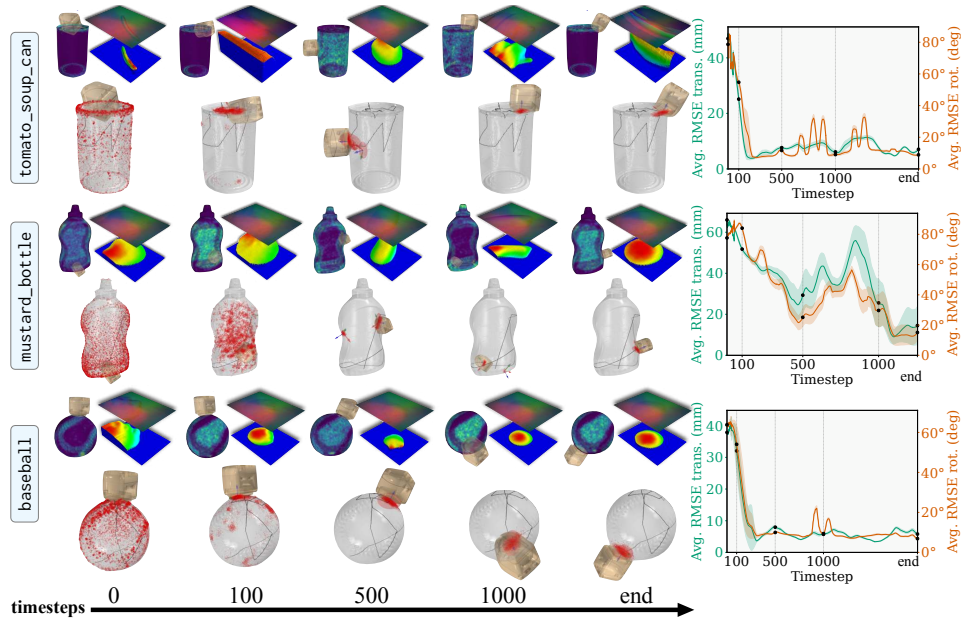


**Figure 6:** Real-world sliding trials in the YCB-Slide dataset. Inset is an example tactile image from the interactions, capturing the local geometry of the object.

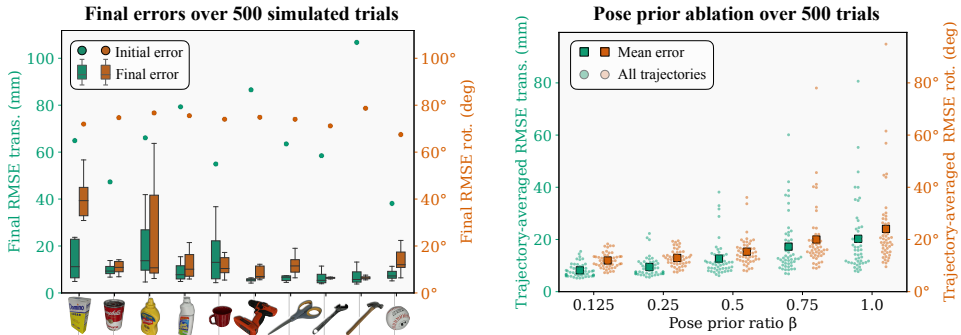
## 6 Experimental results

**Simulation:** We evaluate MidasTouch over the 50 trajectories collected in Section 5. As the filter is non-deterministic, each trajectory is run 10 times for averaged results over 500 trials. Figure 7 shows qualitative results for three representative trajectories. At each timestep, we visualize the pose distribution and plot the average particle RMSE with respect to the ground-truth pose. We see convergence to the most-likely pose hypothesis over the sliding interactions. Upon convergence, the hypothesis clustering step sets an averaged pose for each cluster: visualized as the pose axes with uncertainty ellipses. We also highlight the current tactile image, surface geometry, and comparison with the tactile codebook. These trials are initialized with pose prior  $[\sigma_{\text{trans}}, \sigma_{\text{rot}}]$  from Section 4.3.

Figure 8 [left] accumulates quantitative metrics over the 500 simulated trials. First, we show the final pose errors across all trials compared against the initial error. Overall, the averaged final pose errors are  $0.74\text{cm}$  and  $9.43^\circ$ ; the per-object errors roughly correlate to the single-touch errors from Figure 4. While the error drops significantly for most trials, it fluctuates depending on each trajectory’s salient geometries (or lack thereof). Larger pose-errors can be at-

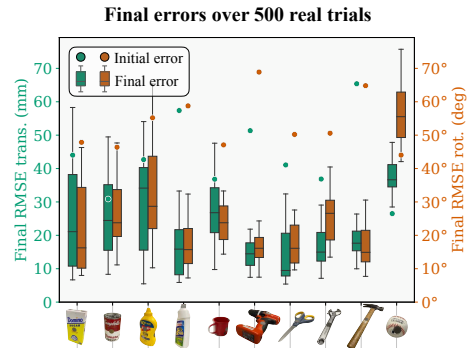


**Figure 7:** Snapshots of simulated sliding results on three YCB objects. For each row: **[top]** the tactile images, local geometries, and heatmap of pose likelihood with respect to the tactile codebook, **[bottom]** pose distribution evolving over time, and converging to the most-likely hypothesis after encountering salient geometries, **[right]** average translation and rotation RMSE of the distribution over time with variance over 10 trials.

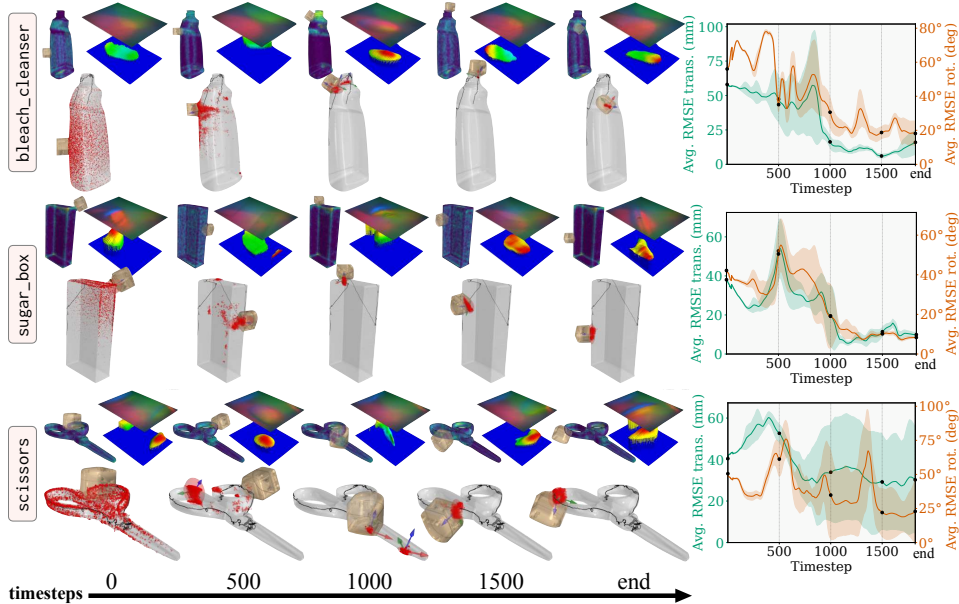


**Figure 8:** **[left]** Boxplot of final trajectory error over 500 simulation trials. For each object, we plot the averaged initial and final RMSE for the particle set. We observe better convergence for tools with salient geometries, as opposed to symmetric objects. **[right]** Ablation over initial pose uncertainty, to show lower average trajectory error with better visual priors. With a weaker initialization ( $\beta = 1.0$ ), outliers in pose-error are more prevalent. tributed to tracking multiple modes that equally explain the same sliding sequence (please refer to supplementary video). We consider this a benefit of our multi-modal framework, and is especially prevalent for symmetric objects like the sugar\_box, mustard\_bottle, and mug. We perform ablations over pose prior (Figure 8 [right]), initializing each trial with uncertainty  $\beta \times [\sigma_{\text{trans}}, \sigma_{\text{rot}}]$ . This serves as a proxy for visual-priors; better initializations lead to fewer candidate modes. The full-stack operates, on average, at the rate of 10Hz.

**Real-world:** We run 500 similar trials for the real-world data from Section 5 with  $\beta = 0.5$ , which serves as a bound for a reasonable prior estimate available from vision. The real-world tactile heightmaps are noisier than simulation, thus a tighter initialization prevents particle depletion. In Figure 10 we present three representative trajectories that show the hypotheses converge to the ground-truth pose. The accumulated



**Figure 9:** Boxplot of final error over 500 real-world trials from the 50 YCB-Slide trajectories.



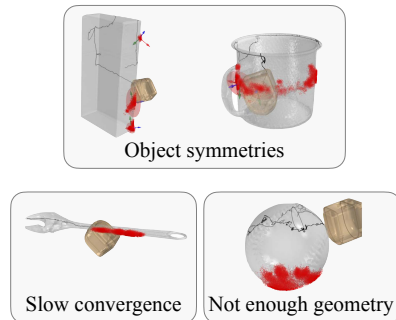
**Figure 10:** Snapshots of real-world sliding results on three YCB objects. For each row: **[top]** the tactile images, local geometries, and heatmap of pose likelihood with respect to the tactile codebook, **[bottom]** pose distribution evolving over time, and converging to the most-likely hypothesis after encountering salient geometries, **[right]** average translation and rotation RMSE of the distribution over time with variance over 10 trials.

statistics in Figure 9 shows reduced final error across all objects, except baseball. The averaged final pose errors are higher than the simulation results: 1.97cm and 21.48°. Once again, we see tools and intricate objects are the easiest to localize on while symmetric objects are the hardest. An anomaly is the baseball, on which we fail to localize in all trials. We attribute this to a lack of distinct geometry detected from real tactile images, effectively meaning we are trying to localize on a featureless sphere. These failure modes are highlighted in Figure 11. We present best-hypothesis error metrics and further qualitative results in Appendix E.

## 7 Conclusion and limitations

In this work, we demonstrate finger-object global localization from posed tactile images. This online method outputs a pose-distribution on the object’s surface that converges over time as the sensor traverses salient surface geometries. Specifically, our system is the first to learn tactile embeddings from local 3D geometry, and disambiguate them with a nonparametric particle filter. Our experiments demonstrate the surprising effectiveness of pairing learned tactile perception with Monte-Carlo methods to resolve distribution ambiguities.

**Limitations:** The current formulation is limited to a moving sensor relative to a fixed-pose object (or vice-versa). In future work, we wish to incorporate a dynamic object in our motion model through (i) visual measurements [26], and/or (ii) local in-hand tracking [20]. Our method currently does not work in scenarios where we lack ground-truth object models. Along with reconstructing objects [70], we would need to build the tactile codebook on-the-fly. For future in-hand manipulation tasks, it is also necessary to scale MidasTouch to multi-contact configurations. Our on-surface assumption leads to poor behavior when we break contact with the object, and the effects of shear on the tactile image sequence have been ignored [71]. Object compositionality doesn’t hold for deformable or articulated objects, so modeling these properties is an interesting future direction [72]. With a differentiable filter [73] we can fine-tune the system end-to-end for more robust real-world performance.



**Figure 11:** Failure modes in the real-world: (i) it may be hard to converge to the true hypothesis for objects with symmetries, (ii) slow convergence of the filter can lead to large pose uncertainty, (iii) lack of discernible geometry can result in drift from the true mode.



## Acknowledgments

We thank Wenzhen Yuan, Ming-Fang Chang, Wei Dong, Maria Bauza Villalonga, and Antonia Bronars for insightful discussions. We are grateful towards the CMU AI Maker Space and Greg Armstrong for facilitating the collection of the YCB-Slide dataset. The authors acknowledge funding from Meta AI, and this work was partially carried out while Sudharshan Suresh interned at Meta AI.

## References

- [1] J. Heller. *Catch-22: a novel*, volume 4. Simon and Schuster, 1961.
- [2] J. Shi, J. Z. Woodruff, P. B. Umbanhowar, and K. M. Lynch. Dynamic in-hand sliding manipulation. *IEEE Trans. on Robotics (TRO)*, 33(4):778–795, 2017.
- [3] N. Chavan-Dafle, R. Holladay, and A. Rodriguez. In-hand manipulation via motion cones. *Proc. Robotics: Science and Systems (RSS)*, 2018.
- [4] J. Shi, H. Weng, and K. M. Lynch. In-hand sliding regrasp with spring-sliding compliance and external constraints. *IEEE Access*, 8:88729–88744, 2020.
- [5] Y. She, S. Wang, S. Dong, N. Sunil, A. Rodriguez, and E. Adelson. Cable manipulation with a tactile-reactive gripper. *Intl. J. of Robotics Research (IJRR)*, 40(12-14):1385–1401, 2021.
- [6] J. Kerr, H. Huang, A. Wilcox, R. Hoque, J. Ichnowski, R. Calandra, and K. Goldberg. Learning self-supervised representations from vision and touch for active sliding perception of deformable surfaces. *arXiv preprint arXiv:2209.13042*, 2022.
- [7] O. 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. *Intl. J. of Robotics Research (IJRR)*, 39(1):3–20, 2020.
- [8] T. Chen, J. Xu, and P. Agrawal. A system for general in-hand object re-orientation. In *Proc. Conf. on Robot Learning (CoRL)*, pages 297–307. PMLR, 2022.
- [9] A. Handa, A. Allshire, V. Makoviychuk, A. Petrenko, R. Singh, J. Liu, D. Makoviichuk, K. Van Wyk, A. Zhurkevich, B. Sundaralingam, Y. Narang, J.-F. Lafleche, D. Fox, and G. State. DeXtreme: Transfer of agile in-hand manipulation from simulation to reality. *arXiv*, 2022.
- [10] H. Qi, A. Kumar, R. Calandra, Y. Ma, and J. Malik. In-Hand Object Rotation via Rapid Motor Adaptation. In *Conference on Robot Learning (CoRL)*, 2022.
- [11] B. Sundaralingam and T. Hermans. Geometric in-hand regrasp planning: Alternating optimization of finger gaits and in-grasp manipulation. In *Proc. IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 231–238. IEEE, 2018.
- [12] S. Brahmabhatt, A. Handa, J. Hays, and D. Fox. Contactgrasp: Functional multi-finger grasp synthesis from contact. In *Proc. IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 2386–2393. IEEE, 2019.
- [13] W. Yuan, S. Dong, and E. H. Adelson. GelSight: High-resolution robot tactile sensors for estimating geometry and force. *Sensors*, 17(12):2762, 2017.
- [14] E. Donlon, S. Dong, M. Liu, J. Li, E. Adelson, and A. Rodriguez. GelSlim: A high-resolution, compact, robust, and calibrated tactile-sensing finger. In *Proc. IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 1927–1934. IEEE, 2018.
- [15] M. Lambeta, P.-W. Chou, S. Tian, B. Yang, B. Maloon, V. R. Most, D. Stroud, R. Santos, A. Byagowi, G. Kammerer, et al. DIGIT: A novel design for a low-cost compact high-resolution tactile sensor with application to in-hand manipulation. *IEEE Robotics and Automation Letters (RA-L)*, 5(3):3838–3845, 2020.

- [16] 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 *Proc. IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 3988–3993. IEEE, 2014.
- [17] M. Bauza, E. Valls, B. Lim, T. Sechopoulos, and A. Rodriguez. Tactile object pose estimation from the first touch with geometric contact rendering. In *Proc. Conf. on Robot Learning, CoRL*, 2020.
- [18] M. Bauza, A. Bronars, and A. Rodriguez. Tac2Pose: Tactile object pose estimation from the first touch. *arXiv preprint arXiv:2204.11701*, 2022.
- [19] P. Sodhi, M. Kaess, M. Mukadam, and S. Anderson. Learning tactile models for factor graph-based estimation. In *Proc. IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 13686–13692. IEEE, 2021.
- [20] P. Sodhi, M. Kaess, M. Mukadam, and S. Anderson. Patchgraph: In-hand tactile tracking with learned surface normals. In *Proc. IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2022.
- [21] J. Hawkins, M. Lewis, M. Klukas, S. Purdy, and S. Ahmad. A framework for intelligence and cortical function based on grid cells in the neocortex. *Frontiers in neural circuits*, 12:121, 2019.
- [22] S. J. Lederman and R. L. Klatzky. Hand movements: A window into haptic object recognition. *Cognitive psychology*, 19(3):342–368, 1987.
- [23] F. Dellaert, D. Fox, W. Burgard, and S. Thrun. Monte Carlo localization for mobile robots. In *Proc. IEEE Intl. Conf. on Robotics and Automation (ICRA)*, volume 2, pages 1322–1328. IEEE, 1999.
- [24] B. Calli, A. Singh, J. Bruce, A. Walsman, K. Konolige, S. Srinivasa, P. Abbeel, and A. M. Dollar. Yale-CMU-Berkeley dataset for robotic manipulation research. *Intl. J. of Robotics Research (IJRR)*, 36(3):261–268, 2017.
- [25] L. Downs, A. Francis, N. Koenig, B. Kinman, R. Hickman, K. Reymann, T. B. McHugh, and V. Vanhoucke. Google scanned objects: A high-quality dataset of 3D scanned household items. *Proc. IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2022.
- [26] X. Deng, A. Mousavian, Y. Xiang, F. Xia, T. Bretl, and D. Fox. PoseRBPF: A Rao-Blackwellized particle filter for 6D object pose tracking. *IEEE Trans. on Robotics (TRO)*, 37(5):1328–1342, 2021.
- [27] A. Yamaguchi and C. G. Atkeson. Combining finger vision and optical tactile sensing: Reducing and handling errors while cutting vegetables. In *Proc. IEEE-RAS Intl. Conf. on Humanoid Robots (Humanoids)*, pages 1045–1051. IEEE, 2016.
- [28] B. Ward-Cherrier, N. Pestell, L. Cramphorn, B. Winstone, M. E. Giannaccini, J. Rossiter, and N. F. Lepora. The TacTip family: Soft optical tactile sensors with 3D-printed biomimetic morphologies. *Soft robotics*, 5(2):216–227, 2018.
- [29] A. Alspach, K. Hashimoto, N. Kuppawamy, and R. Tedrake. Soft-bubble: A highly compliant dense geometry tactile sensor for robot manipulation. In *Proc. IEEE Intl. Conf. on Soft Robotics (RoboSoft)*, pages 597–604. IEEE, 2019.
- [30] A. Padmanabha, F. Ebert, S. Tian, R. Calandra, C. Finn, and S. Levine. OmniTact: A multi-directional high-resolution touch sensor. In *Proc. IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 618–624. IEEE, 2020.
- [31] S. Wang, Y. She, B. Romero, and E. Adelson. GelSight Wedge: Measuring high-resolution 3D contact geometry with a compact robot finger. In *Proc. IEEE Intl. Conf. on Robotics and Automation (ICRA)*. IEEE, 2021.

- [32] S. Wang, M. M. Lambeta, P.-W. Chou, and R. Calandra. TACTO: A fast, flexible, and open-source simulator for high-resolution vision-based tactile sensors. *IEEE Robotics and Automation Letters (RA-L)*, 2022.
- [33] A. Agarwal, T. Man, and W. Yuan. Simulation of vision-based tactile sensors using physics based rendering. In *Proc. IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 1–7. IEEE, 2021.
- [34] Z. Si and W. Yuan. Taxim: An example-based simulation model for gelsight tactile sensors. *IEEE Robotics and Automation Letters (RA-L)*, 2022.
- [35] C. Choy, J. Park, and V. Koltun. Fully convolutional geometric features. In *Proc. Intl. Conf. on Computer Vision (ICCV)*, pages 8958–8966, 2019.
- [36] M. A. Uy and G. H. Lee. PointNetVLAD: Deep point cloud based retrieval for large-scale place recognition. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 4470–4479, 2018.
- [37] J. Komorowski. MinkLoc3D: Point cloud based large-scale place recognition. In *Proc. Winter Conf. on Applications of Computer Vision (WACV)*, pages 1790–1799, 2021.
- [38] C. Corcoran and R. Platt. Tracking object pose and shape during robot manipulation based on tactile information. In *Proc. IEEE Intl. Conf. on Robotics and Automation (ICRA)*, volume 2. Citeseer, 2010.
- [39] A. Petrovskaya and O. Khatib. Global localization of objects via touch. *IEEE Trans. on Robotics (TRO)*, 27(3):569–585, 2011.
- [40] L. Zhang and J. C. Trinkle. The application of particle filtering to grasping acquisition with visual occlusion and tactile sensing. In *2012 IEEE International Conference on Robotics and Automation*, pages 3805–3812. IEEE, 2012.
- [41] M. Chalon, J. Reinecke, and M. Pfanne. Online in-hand object localization. In *Proc. IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 2977–2984. IEEE, 2013.
- [42] B. Saund, S. Chen, and R. Simmons. Touch based localization of parts for high precision manufacturing. In *Proc. IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 378–385. IEEE, 2017.
- [43] 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 *Proc. IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 4673–4680. IEEE, 2017.
- [44] S. Thrun. Particle filters in robotics. In *UAI*, volume 2, pages 511–518. Citeseer, 2002.
- [45] Z. Pezzementi, C. Reyda, and G. D. Hager. Object mapping, recognition, and localization from tactile geometry. In *2011 IEEE International Conference on Robotics and Automation*, pages 5942–5948. IEEE, 2011.
- [46] S. Luo, W. Mou, K. Althoefer, and H. Liu. Localizing the object contact through matching tactile features with visual map. In *Proc. IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 3903–3908. IEEE, 2015.
- [47] 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 (RA-L)*, 1(1): 570–577, 2016.
- [48] K.-T. Yu and A. Rodriguez. Realtime state estimation with tactile and visual sensing. application to planar manipulation. In *Proc. IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 7778–7785. IEEE, 2018.

- [49] A. N. Chaudhury, T. Man, W. Yuan, and C. G. Atkeson. Using collocated vision and tactile sensors for visual servoing and localization. *IEEE Robotics and Automation Letters (RA-L)*, 7(2):3427–3434, 2022.
- [50] S. Dong and A. Rodriguez. Tactile-based insertion for dense box-packing. In *Proc. IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 7953–7960. IEEE, 2019.
- [51] S. Kim and A. Rodriguez. Active extrinsic contact sensing: Application to general peg-in-hole insertion. In *Proc. IEEE Intl. Conf. on Robotics and Automation (ICRA)*. IEEE, 2022.
- [52] T. Kelestemur, R. Platt, and T. Padir. Tactile pose estimation and policy learning for unknown object manipulation. *arXiv preprint arXiv:2203.10685*, 2022.
- [53] R. Gao, Z. Si, Y.-Y. Chang, S. Clarke, J. Bohg, L. Fei-Fei, W. Yuan, and J. Wu. ObjectFolder 2.0: A multisensory object dataset for Sim2Real transfer. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [54] N. Kuppaswamy, A. Castro, C. Phillips-Grafflin, A. Alspach, and R. Tedrake. Fast model-based contact patch and pose estimation for highly deformable dense-geometry tactile sensors. *IEEE Robotics and Automation Letters (RA-L)*, 5(2):1811–1818, 2019.
- [55] M. K. Johnson and E. H. Adelson. Retrographic sensing for the measurement of surface texture and shape. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 1070–1077, 2009. doi:10.1109/CVPR.2009.5206534.
- [56] M. K. Johnson, F. Cole, A. Raj, and E. H. Adelson. Microgeometry capture using an elastomeric sensor. *ACM Trans. Graph.*, 30(4), July 2011. ISSN 0730-0301.
- [57] M. Bauza, O. Canal, and A. Rodriguez. Tactile mapping and localization from high-resolution tactile imprints. In *Proc. IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 3811–3817. IEEE, 2019.
- [58] R. Ambrus, V. Guizilini, N. Kuppaswamy, A. Beaulieu, A. Gaidon, and A. Alspach. Monocular depth estimation for soft visuotactile sensors. In *Proc. IEEE Intl. Conf. on Soft Robotics (RoboSoft)*, 2021.
- [59] S. Suresh, Z. Si, J. G. Mangelson, W. Yuan, and M. Kaess. ShapeMap 3-D: Efficient shape mapping through dense touch and vision. In *Proc. IEEE Intl. Conf. on Robotics and Automation (ICRA)*, Philadelphia, PA, USA, May 2022.
- [60] C. R. Qi, H. Su, K. Mo, and L. J. Guibas. PointNet: Deep learning on point sets for 3D classification and segmentation. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 652–660, 2017.
- [61] C. Choy, J. Gwak, and S. Savarese. 4D spatio-temporal ConvNets: Minkowski convolutional neural networks. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 3075–3084, 2019.
- [62] I. Laina, C. Rupprecht, V. Belagiannis, F. Tombari, and N. Navab. Deeper depth prediction with fully convolutional residual networks. In *Proc. Intl. Conf. on 3D Vision (3DV)*, pages 239–248. IEEE, 2016.
- [63] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie. Feature pyramid networks for object detection. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 2117–2125, 2017.
- [64] L. Van der Maaten and G. Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.

- [65] L. Pineda, T. Fan, M. Monge, S. Venkataraman, P. Sodhi, R. T. Chen, J. Ortiz, D. DeTone, A. Wang, S. Anderson, J. Dong, B. Amos, and M. Mukadam. Theseus: A Library for Differentiable Nonlinear Optimization. *Advances in Neural Information Processing Systems*, 2022.
- [66] W. Yuan, R. Li, M. A. Srinivasan, and E. H. Adelson. Measurement of shear and slip with a GelSight tactile sensor. In *Proc. IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 304–311. IEEE, 2015.
- [67] S. Thrun. Probabilistic robotics. *Communications of the ACM*, 45(3):52–57, 2002.
- [68] E. Schubert, J. Sander, M. Ester, H. P. Kriegel, and X. Xu. DBSCAN revisited, revisited: why and how you should (still) use DBSCAN. *ACM Transactions on Database Systems (TODS)*, 42(3):1–21, 2017.
- [69] F. L. Markley, Y. Cheng, J. L. Crassidis, and Y. Oshman. Averaging quaternions. *Journal of Guidance, Control, and Dynamics*, 30(4):1193–1197, 2007.
- [70] S. Suresh, M. Bauza, K.-T. Yu, J. G. Mangelson, A. Rodriguez, and M. Kaess. Tactile SLAM: Real-time inference of shape and pose from planar pushing. In *Proc. IEEE Intl. Conf. on Robotics and Automation (ICRA)*, May 2021.
- [71] K. Aquilina, D. A. Barton, and N. F. Lepora. Shear-invariant sliding contact perception with a soft tactile sensor. In *Proc. IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 4283–4289. IEEE, 2019.
- [72] Y. Wi, P. Florence, A. Zeng, and N. Fazeli. VIRDO: Visio-tactile implicit representations of deformable objects. In *Proc. IEEE Intl. Conf. on Robotics and Automation (ICRA)*. IEEE, 2022.
- [73] R. Jonschkowski, D. Rastogi, and O. Brock. Differentiable particle filters: End-to-end learning with algorithmic priors. In *Proc. Robotics: Science and Systems (RSS)*, 2018.
- [74] McMaster-Carr. <https://www.mcmaster.com/>. Accessed: 2022-08-25.