

Active Tactile Exploration for Rigid Body Pose and Shape Estimation

Ethan K. Gordon, Bruke Baraki, Hien Bui, Michael Posa¹

Abstract—General robot manipulation requires the handling of previously unseen objects. Learning a physically accurate model at test time can provide significant benefits in data efficiency, predictability, and reuse between tasks. Tactile sensing can compliment vision with its robustness to occlusion, but its temporal sparsity necessitates careful online exploration to maintain data efficiency. Direct contact can also cause an unrestrained object to move, requiring both shape *and* location estimation. In this work, we propose a learning and exploration framework that uses only tactile data to simultaneously determine the shape and location of rigid objects with minimal robot motion. We build on recent advances in contact-rich system identification to formulate a loss function that penalizes physical constraint violation without introducing the numerical stiffness inherent in rigid-body contact. Optimizing this loss, we can learn cuboid and convex polyhedral geometries with less than 10s of randomly collected data after first contact. Our exploration scheme seeks to maximize Expected Information Gain and results in significantly faster learning in both simulated and real-robot experiments. More information can be found at: <https://dairlab.github.io/activetactile>.

I. INTRODUCTION

Robot manipulators in the wild will inevitably come across previously unseen objects and environments. They can benefit from building object models from the limited data available at test time. Such models can be reused across different tasks and systems, and they can make learned manipulation policies more interpretable. In the building of these models, tactile sensing can compliment visual modalities like RGB and depth. Tactile data is robust to darkness, reflections, and occlusions from clutter, the environment, and the object itself. With the advent of inexpensive sensors [1, 2, 3], touch is becoming a more common component of robot manipulation systems [4].

Two challenges inherent to tactile sensing are sparsity and disturbance. By definition, data is only collected about the part of the object in direct physical contact with the sensor, necessitating multiple movements to get a complete picture. Sparsity can be addressed by *active exploration*, where actions are taken to efficiently maximize some information metric, often captured by maintaining a belief distribution over possible object parameterizations [5, 6, 7, 8, 9, 10]. However, reasoning about information gain is complicated by disturbance: every contact has the potential to move the object unless it is heavy or bolted down. Propagating belief through motion is tricky, informing our first key insight: expected information gain (EIG), by foregoing an explicit




Fig. 1. Only tactile data is used to find the pose and geometry of an arbitrary convex object (e.g., parameterized as a 20-vertex polytope as shown here). Actions are chosen to maximize *Expected Information Gain*, related to the info we expect to gain from contact divided by the info previously observed. As parts of the ground-truth object (e.g. a cuboid, shown in gray) are contacted, the corresponding areas of the learned estimate will be assigned a high observed info (blue $\rightarrow$ red).

belief distribution, is a particularly useful metric for active tactile exploration.

Disturbance affects even the model-building itself, as it necessitates learning the object’s pose trajectory in addition to its geometry. That trajectory needs to adhere to the physical constraints of rigid-body motion and contact, which can lead to difficult loss landscapes with both nearly flat and nearly discontinuous regions. Related work in geometry-learning with a known trajectory [11, 12] has shown how a *violation-implicit* loss leads to a landscape more amenable to gradient-based optimization. Our second key insight is that this loss works just as well with the trajectory itself as a decision variable.

In summary, this work presents two contributions.

- 1) A *violation-implicit loss* that enables the simultaneous learning of rigid-body pose and geometry from purely tactile data.
- 2) An active exploration scheme that seeks to maximize *EIG* to enable more efficient learning with fewer robot motions, as highlighted in Fig. 1.

The combination of these contributions yields a scheme that can simultaneously learn the pose and cuboid or convex polytope geometric approximations of arbitrary unknown convex rigid bodies with under 10s of purely tactile data.

II. RELATED WORK

A. Model Building and System Identification

Visual modeling of both environment and object geometry is a mature field with the well-studied paradigms of simultaneous localization and mapping (SLAM) and structure from motion (SfM)[13]. Continuous object representations include radiance fields [14] and Gaussian splatting [15]. Recent work learns the geometry [16] and pose [17] of dynamic objects in video, even in contact-rich environments [18]. For assessing the accuracy of our tactile-only system, we compare against ground truth estimates from FoundationPose [17].

¹All authors are with the General Robotics, Automation, Sensing, and Perception (GRASP) Laboratory, University of Pennsylvania, Philadelphia, PA, USA 19104, {ethankg, bbaraki, xuanhien, posa}@seas.upenn.edu




Fig. 2. Block diagram of the interaction protocol. An estimate of object pose and geometry emerges from tactile and proprioceptive data (top left) by minimizing a violation-implicit loss. In this example with a cuboid parameterization, the top, left, and bottom faces can be directly inferred from data while the right face is unobserved. The expected information gain (EIG) of the next action is optimized using a Gaussian cross-entropy method: robot trajectories are sampled, simulated forward to time horizon H , and scored. Pictured here, $\mathcal{D}_{obs}$ includes contact with left and top, leading to high observed info $\mathcal{I}$ (red), so the upper sample contacting these same sides yields low EIG. The lower sample contacts the right side with low $\mathcal{I}$ (blue), maximizing EIG.

Even without active exploration, tactile sensing has been used in recent work to refine models of both the robot and its environment. Non-optical contact sensors can enable more accurate proprioception in the presence of a known environment [19]. Optical tactile sensors create estimates of local surface geometry that can be used to localize the robot with respect to a known object [20] or vice versa [21].

B. Active Exploration in Robotics

Active exploration has been studied in robotics for decades [22] and involves either the selection of data-gathering actions to maximize information or the curation of data based on information potential. Generally, it includes: (1) some current state of knowledge, (2) an estimate of the information collected by a given action, and (3) an optimization metric quantifying how useful the latter is given the former.

In Bayesian state estimation, the current state can be represented by a belief distribution. Mutual information (MI) [5] as a metric captures the information gained about this belief distribution by observing some other (e.g. measurement likelihood) distribution, but in dynamic environments this can be intractable to compute exactly [6]. Explicit MI calculation can be avoided by using uncertainty instead, e.g., probing areas where the belief distribution has a high variance [7, 8]. One could also use a heuristic to avoid actions similar to those previously taken [10] or a policy trained to take informative actions [9]. For dynamic objects, propagating belief through a dynamics model that is itself uncertain is tricky: the result is inherently multimodal, and the number of required parameters scales exponentially in the number of moments. Multimodal uncertainty can be approximated by simulating an ensemble of possible models [23].

Another approach is to drop the belief distribution and maximize Fisher Information [14], which only requires a point maximum-likelihood estimate (MLE) and is used extensively in optimal experimental design. Recent work uses this approach to refine a simulator from real world data [24] and plan informative trajectories in contact-rich environments [25]. However, maximizing Fisher information alone does not take into account previously observed information over multiple actions. Ergodic exploration [26, 27] can be

used to plan a sequence of actions such that, while more time is spent in areas likely to have high information, the entire state space is eventually explored. EIG (see Sec. V), derived from Fisher information, explicitly quantifies how much information an action is likely to obtain relative to previously observed information. It has been used in visual model building for active view selection of static objects [14, 15]. This work takes this procedure and implements it for tactile sensing of dynamic objects.

C. Tactile Active Exploration

Multiple works have tackled the problem of tactile active exploration. The most common application is the modeling of an unknown surface using poking or sliding contact from an optical sensor [7, 28, 29], as thoroughly collated by a recent survey [30]. It is also common to estimate the geometry of 3d static objects [8, 10]. When learning both pose and geometry, the latter tends to be treated as a classification problem by maintaining a discrete set of possible object geometries [9, 31, 32]. One work simultaneously estimates the pose and geometry of specifically planar objects [33]. To our knowledge, ours is the first work that learns the pose and geometry (from a continuous class of convex shapes) of a 3D rigid body that is free to move during deployment.

III. PROBLEM FORMULATION

Consider a finite workspace with a rigid known ground plane, unknown convex rigid object, and a robot with known convex end-effectors located at r_t , each fitted with a tactile sensor. They report a contact boolean $c_t \in \{0, 1\}$ and a measured contact normal direction $\hat{n}_{t,m} \in S^2$. The goal is to determine the current pose of the object x_T and its geometry θ in as few actions as possible. In practice, our action space will be a parameterized trajectory described in more detail in Sec. VI-A. Given the possible geometric symmetries in the object, the success metric will be the bidirectional Chamfer distance (bCH) between the estimated $\tilde{S}(\tilde{x}_T, \tilde{\theta})$ and ground-truth $S^*(x_T^*, \theta^*)$ surfaces.

$$C_{AB} = \sum_{a \in A} \min_{b \in B} \|a - b\|_2^2 ; \text{ bCH} = C_{\tilde{S}S^*} + C_{S^*\tilde{S}} \quad (1)$$

The interaction protocol is summarized in Fig. 2.

- 1) **Learning:** (Sec. IV) At current time T , given some observed dataset $\mathcal{D}_{obs} = \bigcup_{t \in [0, T]} \{r_t, c_t, \hat{n}_{t,m}\}$, produce some estimate of the object pose and shape. In practice, this will be the minimum of some loss function $\mathcal{L}$.

$$\tilde{\theta}, \tilde{x}_T = \arg \min_{\theta, x_T} \mathcal{L}(\theta, x_T, \mathcal{D}_{obs}) \quad (2)$$

- 2) Evaluate bCH
- 3) **Exploration:** (Sec. V) Given some known horizon $H > T$, choose a trajectory $r_{T \dots H}$ and collect new data $\mathcal{D}_{acq} = \bigcup_{t \in [T, H]} \{r_t, c_t, \hat{n}_{t,m}\}$.
- 4) Repeat with $\mathcal{D}_{obs} \leftarrow \mathcal{D}_{obs} \cup \mathcal{D}_{acq}$

A. Sensor Model

We model a discrete logistic likelihood for c_t , where the logit $\phi_t(\theta, x_t, r_t)$ is the signed distance between the robot and the object and α is an uncertainty hyperparameter. For a given value of c_t , this likelihood approximates stepwise behavior about $\phi_t = 0$ with a sigmoid logistic function.

$$p(c_t | \theta, x_t, r_t) = \frac{c_t + (1 - c_t)e^{\alpha\phi_t}}{1 + e^{\alpha\phi_t}} \quad (3)$$

If $c_t = 1$, we model our measured normal $\hat{n}_{t,m}$ with a Gaussian¹ likelihood with respect to the expected contact normal $\hat{n}_t(\theta, x_t, r_t)$ with covariance hyperparameter Σ_n .

$$p(\hat{n}_{t,m} | \theta, x_t, r_t) = \frac{c_t \exp(-0.5 \|\hat{n}_{t,m} - \hat{n}_t(\theta, x_t, r_t)\|_{\Sigma_n^{-1}}^2)}{(2\pi)^{3/2} \sqrt{\det \Sigma_n}} + (1 - c_t) \delta(\hat{n}_{t,m}) \quad (4)$$

While the second term with the Dirac delta function δ is included here for completeness, it does not involve any object or action parameters and will be omitted going forward.

IV. LEARNING WITH A VIOLATION-IMPLICIT LOSS

A natural step is to seek to minimize the negative log-likelihood of our sensor model. We take the log of the product of Eq. 3 and 4, bringing the discrete selectors c_t and $(1 - c_t)$ out of the logarithm and dropping additive constants.

$$\mathcal{L}_t = \frac{c_t}{2} \|\hat{n}_{t,m} - \hat{n}_t\|_{\Sigma_n^{-1}}^2 + [(c_t - 1)\alpha\phi_t + \log(1 + e^{\alpha\phi_t})] \quad (5)$$

However, via $\hat{n}$ and ϕ , $\mathcal{L}_t$ depends on the object state x_t , so the sum over the trajectory must be optimized under physical constraints.

$$\begin{aligned} \tilde{\theta}, \tilde{x}_T = \arg \min_{\theta, x_T} \sum_t \mathcal{L}_t \text{ s.t. } x_{t+1} &= f_\theta(x_t, r_t, \lambda_t) \\ \lambda_t &= \arg \min_{\lambda} g_\theta(x_t, r_t, \lambda) \text{ s.t. } \lambda \in \mathcal{FC} \end{aligned} \quad (6)$$

Where f_θ are the discrete-time dynamics as a function of contact impulses λ_t . These can themselves be written as the minimizer of a function g_θ [34] representing other physical constraints (e.g. non-penetration) discussed in more detail

¹While a more appropriate likelihood in S^2 may be a univariate Gaussian about the dot product $1 - \hat{n}_{m,t} \cdot \hat{n}_t$ with variance σ^2 , the two are equivalent for diagonal covariance $\Sigma_n = \sigma^2 * I$, though the gradients differ in $\mathbb{R}^3$ before normalization. The 2-norm enables the easy calculation of Eq. 14 that is independent of $\hat{n}_{m,t}$.

below. Under the assumption of Coulomb friction, impulses are constrained to the friction cone $\mathcal{FC} = \{\|\lambda_t^f\|_2 \leq \mu \lambda_t^n\}$, where the total frictional impulse $\|\lambda_t^f\|_2$ is bounded by a constant multiple of the normal impulse λ_t^n , sliding in equality and sticking otherwise.

One could optimize this loss function through shooting and differentiable simulation. With parameter x_0 , compute $x_t = (f_\theta)^t(x_0)$ and perform gradient descent through $(f_\theta)^t$ to find $\tilde{x}_0$. Then $\tilde{x}_T = (f_\theta)^T(\tilde{x}_0)$. As discussed in related work [12], this approach can produce gradients that are near-0 in some regions and near-infinite in others. This is especially true with rigid contact dynamics, longer trajectories, and simulations with a small timestep Δt .

Intuitively, this procedure is analogous to fitting a Heaviside step function $y = Hs(x - \theta)$ using the square-error loss $\sum_{\mathcal{D}} \|y_{\mathcal{D}} - Hs(x_{\mathcal{D}} - \theta)\|_2^2$, capturing error in the y-direction, which has a gradient of 0 almost everywhere and is discontinuous when data lies at the estimate of the parameter θ . An alternative is to minimize the distance between data and the closest point on the function, or “graph distance” $\sum_{\mathcal{D}} \min_x \|(x_{\mathcal{D}}, y_{\mathcal{D}}) - (x, Hs(x - \theta))\|_2^2$. At the cost of a (potentially expensive) inner optimization, the loss gradient is now a finite value (or bounded set) everywhere.

Approximately minimizing graph distance is the essence of the violation-implicit loss, and has been used in geometry learning given a known object trajectory [11]. The analogous procedure in our setting is to bring the near-discontinuous dynamics and physical constraints into the loss function as penalties, and the inner optimization happens over the contact impulses λ .

$$\begin{aligned} \mathcal{L}_v(\theta, x_t, \mathcal{D}) = \sum_t \min_{\lambda_t} \mathcal{L}_t + \|x_{t+1} - f_\theta(x_t, \lambda_t, r_t)\|_2^2 \\ + g_\theta(x_t, \lambda_t, r_t) \text{ s.t. } \lambda_t \in \mathcal{FC} \end{aligned} \quad (7)$$

Without dynamics as a hard constraint, the entire trajectory x_t can be treated as decision variables, reducing the instability of differentiating through a large T at the cost of introducing non-physical local minima in the optimization. In practice, we mitigate the latter using multiple trajectory initializations for the gradient descent. Following [11], our physical loss terms g_θ include the following.

- 1) Complementarity: $\lambda_t^n \phi_t$, i.e., contact force only happens when objects touch.
- 2) Maximum Power Dissipation: $\|v_t^s\|_2^2 \lambda_t^n + (v_t^s)^T \lambda_t^f$, i.e., friction force opposes the contact sliding velocity v_t^s on the boundary of $\mathcal{FC}$
- 3) Non-Penetration: $\min(0, \phi_t)^2$
- 4) Inelasticity (new in this work): $\lambda_t^n \max(v_t^n, 0)$, i.e., the normal force opposes the normal velocity v_t^n and cannot return energy to the object

While [11] was agnostic to elasticity given the known trajectory, since we are estimating an unknown trajectory, we found it useful to assume that all of our collisions would be nearly inelastic. With the above loss terms, our inner optimization loop in λ was a quadratic program (QP) with a second-order cone constraint that could be solved quickly

in Cvxpylayers [35]. Note that, by the envelope theorem, explicitly differentiating through the $\min_\lambda$ is unnecessary, further speeding up the backwards pass.

V. MAXIMIZING EXPECTED INFORMATION GAIN (EIG)

In Sec. V-A, we first review EIG as a popular metric for general active exploration [14, 15, 36]. In Sec. V-B, we detail the specific formulation for our setting: active tactile exploration with a dynamic unknown object and a violation-implicit loss (Sec. IV).

A. Preliminaries

As described in Sec. II-B, active exploration requires some estimate of the information present in observed data $\mathcal{D}_{obs}$, some expectation of information that may be available in future data $\mathcal{D}_{acq}$, and a scalar metric quantifying the utility of that information. A natural step is to consider belief prior $p(\omega|\mathcal{D}_{obs})$ and posterior $p(\omega|\mathcal{D}_{obs}, \mathcal{D}_{acq})$ distributions over model parameters $\omega = \{\theta, x_T\}$. In particular, a common approach is to look at the entropy H of these distributions.

$$H[\omega|\mathcal{D}] := \mathbb{E}_p[-\log p(\omega|\mathcal{D})] \quad (8)$$

A chief difficulty is estimating the entropy of the posterior, as the acquired data itself is a random variable sampled from a likelihood distribution $p(\mathcal{D}_{acq}|\omega, r_t)$.

For a Gaussian, the entropy can be expressed in terms of the log-determinant of the covariance: $H[\mathcal{N}(\mu, \Sigma)] = 1/2 \log \det \Sigma + \text{const}$. This is why it is common to use the (co)variance of the belief as a heuristic for high-information regions. But for dynamic objects, propagating a Gaussian prior through time to get that covariance is difficult. The result is fundamentally multi-modal, and just fitting a Gaussian can be completely non-physical. However, it is possible to fit a Gaussian to a second-order approximation of an arbitrary distribution about *any modal peak*. In that case, the effective covariance is equal to the Hessian of the log of the distribution about that peak ([36], Prop. 3.5).

$$H[\omega|\mathcal{D}] \approx \frac{1}{2} \log \det \nabla_\omega^2 (-\log(p(\omega|\mathcal{D})))|_{\tilde{\omega}} + c \quad (9)$$

We now have a local approximation of entropy that only requires a local maximum a posteriori (MAP) estimate $\tilde{\omega}$ of the model parameters. Given an uninformative prior (i.e. $\text{Cov}(\omega|\emptyset)^{-1} \rightarrow 0$), the Hessian of the posterior equals the Hessian of the likelihood, and $\tilde{\omega}$ becomes a local MLE.

$$\begin{aligned} (9) &= \frac{1}{2} \log \det \nabla_\omega^2 (-\log(p(\mathcal{D}|\omega)))|_{\tilde{\omega}} + c \\ &:= \frac{1}{2} \log \det \mathcal{I}(\mathcal{D}, \tilde{\omega}) + c \end{aligned} \quad (10)$$

The Hessian of the negative log-likelihood about an MLE is called *observed information*. Fisher information $\mathcal{F}(\mathcal{D}, \tilde{\omega}) = \mathbb{E}_{p(\mathcal{D})} \mathcal{I}(\mathcal{D}, \tilde{\omega})$ is defined as the expected value of the observed information from future data. EIG, which is defined to be the difference in entropy between the belief prior and the belief posterior, can be computed without explicitly

maintaining either of those distributions by using Fisher and observed information.

$$\begin{aligned} EIG &:= H[\omega|\mathcal{D}_{obs}] - H[\omega|\mathcal{D}_{obs} \cup \mathcal{D}_{acq}] \\ &\propto \log \det(\mathcal{F}(\mathcal{D}_{acq})\mathcal{I}(\mathcal{D}_{obs})^{-1} + \mathbf{I}) \end{aligned} \quad (11)$$

For a full derivation we direct the reader to [36], Sec. 3-5.

B. Computing EIG

In our setting, the negative log-likelihood is our loss $-\log p(\omega|\mathcal{D}) = \mathcal{L}_t$ (Eq. 5), therefore, we can calculate observed information as our loss Hessian at the MLE $\tilde{\theta}, \tilde{x}_{0...T}$ computed by our learning scheme.

$$\mathcal{I}(\mathcal{D}_{obs}) = \nabla_{\theta, x_T}^2 \sum_{t=0}^T \mathcal{L}_t(\tilde{\theta}, \tilde{x}_t, r_t) \quad (12)$$

Dividing $\mathcal{L}_t$ into two components ($\hat{n}_{m,t}$ and c_t), the result for each is a term proportional to the Hessian of the logit ($\nabla_\omega^2(\phi_t, \hat{n}_t)$) and a term proportional to the outer product of the gradient of the logit ($\nabla_\omega(\phi_t, \hat{n}_t)\nabla_\omega(\phi_t, \hat{n}_t)^T$). The former can be difficult to calculate. Fortunately, its multiplicand depends on the measurements and is 0 in expectation (noting that $\mathbb{E}[\hat{n}_{m,t}] = \hat{n}_t$ and $\mathbb{E}[c_t] = (1 + e^{\alpha\phi_t})^{-1}$).

$$\begin{aligned} \mathbb{E}_{\hat{n}_{m,t}}[\nabla_\omega^2(\hat{n}_t \cdot (\hat{n}_{m,t} - \hat{n}_t))] &= \nabla_\omega^2 \hat{n}_t \cdot \mathbb{E}[(\hat{n}_{m,t} - \hat{n}_t)] \\ \nabla_\omega^2(\phi_t) \mathbb{E}_{c_t}[(c_t - 1) + e^{\alpha\phi_t}(1 + e^{\alpha\phi_t})^{-1}] &= 0 \text{ (identically)} \end{aligned} \quad (13)$$

While observed information is *not* taken in expectation, we omit the above terms, assuming that their contribution will tend towards negligibility as more data is collected.

$$\begin{aligned} \mathcal{I}(\mathcal{D}_{obs}) &\approx \sum_{t=0}^T \alpha^2 \left(\nabla_\omega \phi_t \left(\frac{e^{\alpha\phi_t}}{(1 + e^{\alpha\phi_t})^2} \right) \nabla_\omega \phi_t^T \right) \\ &\quad + c_t \nabla_\omega \hat{n}_t \Sigma_n^{-1} \nabla_\omega \hat{n}_t^T \end{aligned} \quad (14)$$

While the gradients can be evaluated at the MLE trajectory $\tilde{x}_{0...T}$, they must be taken with respect to the parameters we actually care about: $\tilde{\theta}, \tilde{x}_T$. Frustratingly, this necessitates quantifying the sensitivity of *past* measurements to the *current* state. Even with a differentiable simulation, contact dynamics cannot in general be simulated backwards (multiple starting states can lead to the same end state). One could consider taking the dynamics Jacobian $\frac{dx_t}{dx_{t-1}}$ and directly inverting it to try and approximate this sensitivity, but in our experience this leads to severe numerical instability. For this work, we found success computing the sensitivity at each time step with respect to only that time step's state and summing the results, e.g., $\sum_t \nabla_{x_t} \phi_t \nabla_{x_t} \phi_t^T$. This is equivalent to assuming $\frac{dx_t}{dx_{t-1}} = \mathbf{I}$. While sufficient for the sliding object trajectories that emerged in our experiments, we do not believe this approach is likely to work for more chaotic disturbances, which is an avenue for future work.

Fisher information can be similarly calculated. For a given action $r_{T...H}$, we can simulate forward to compute $\tilde{x}_{T...H}$. We note that this future trajectory will be, by construction, a local MLE. The results of simulation will perfectly obey

physical constraints and simulated measurements are identical to their modeled counterparts, so $\nabla_{x_{T...H}} \mathcal{L}_{T...H} = 0$.

$$\mathcal{F}(\mathcal{D}_{acq}) = \mathbb{E}_{c_t, \hat{n}_{m,t}} \left[\nabla_{\omega}^2 \sum_{t=T}^H \mathcal{L}_t(\tilde{\theta}, \tilde{x}_t, r_t) \right] \quad (15)$$

The computation is identical to Eq. 14, with the approximation becoming equality as Fisher information is taken in expectation. The same note from the previous paragraph applies for gradient computation with respect to $\tilde{x}_T$. Differentiable calculation of ϕ and $\hat{n}$ (itself a QP) were implemented with Jaxopt [37]. With both Fisher and observed information, we can directly compute EIG with Eq. 11.

VI. EXPERIMENT SETUP

A. Action Space

As mentioned in Sec. III, our action space is a parameterized robot trajectory. For each of our robot’s 2 end-effectors, we execute a 1s straight line trajectory that approaches our guess of the object’s centroid from a fixed radius. To avoid sensor collisions, each finger is restricted to the +X or -X half-space (bounded from below by the ground). Therefore, our final space is composed of 2, 2D approach angles ($S^2 \otimes S^2$). In practice, we use the “right-ascension” (a rotation about +Z from the +/-X-axis) and “declination” (a rotation towards the +Z axis). On the real robot, both rotations are limited to $[-\pi/4, \pi/4]$ such that all collisions are likely to happen with the tactile sensor’s surface, though in simulation we relax the bounds to $[-2\pi/5, 2\pi/5]$.

EIG is optimized using an iterative Gaussian cross-entropy method. Initially, a set of n_{samp} actions are sampled uniformly. We then construct a new sampling distribution by sorting the actions by EIG and taking a Gaussian fit of only the best $n_{best} < n_{samp}$. After the last iteration, the final action is chosen uniformly randomly from the final n_{best} .

For the computation of bCH for evaluation only, FoundationPose [17] visually computes the ground-truth object pose at 30Hz. As we do not have any size prior for the object, we ignore the initial uniform sweep of the workspace and only begin counting actions at the moment of first contact. In practice, this means randomly selecting an action for one finger centered on the ground truth position.

B. Hardware Setup

Our hardware setup is summarized in Fig. 3. We utilize a Trifinger robot [38] modified into a dual-finger configuration 180-degrees apart. Each finger is instrumented with a Densetact [3], which streams sensor images at 30Hz. Tactile data acquisition leveraged the correlation between observed fingertip motion primitives from each frame and applied forces (normal, shear, and torsional). Contact motion was computed using the optical flow algorithm [39] between the current frame and a reference frame with no contact. We perform a Helmholtz-Hodge decomposition [40] of the resulting vector field to separate out the irrotational component. Empirically, we find the surface deformation to be correlated with the scalar potential $\psi(u, v)$ of this field. The center-of-pressure (CoP) can be measured by taking a weighted




Fig. 3. Hardware setup: a Trifinger [38] modified to have 2 fingers directly opposing each other, each instrumented with a Densetact [3] sensor. Densetact images (green) can be used to locate the center-of-pressure (yellow), which in turn can be projected onto the nominal model to compute the measured normal vector $\hat{n}_{m,t}$ (pink). A single RealSense running FoundationPose (red) is used for evaluation *only* to calculate bCH.

average $\sum_{u,v} \psi(u, v) \cdot (u, v)$, and the contact normal can be imputed from a model of the sensor’s surface. For the contact boolean, we place a threshold on ψ near the CoP.

C. Object Parameterization

The object pose in $SE(3)$ is parameterized as a quaternion and a displacement. Since the former is normalized before every computation, for the calculation of Fisher and observed information we omit the scalar component.

In this work we consider 2 convex object parameterizations. The cuboid is parameterized in $\mathbb{R}_{>0}^3$ by its side lengths. The convex polytope is parameterized by the convex hull of its vertices in $\mathbb{R}^{3 \times n_v}$, where in our experiments $n_v = 20$. One benefit of EIG is that the object can be reparameterized at any time. During gradient descent, if any point moves internally, off of the convex hull, it becomes effectively invisible to both the loss function and the final bCH calculation, permanently reducing the expressiveness of the polytope parameterization. To avoid this, we can replace it with a random point on the convex hull. We also add a regularizer that encourages the polytope to have at least 3 vertices on the ground at the end of the trajectory, as the violation-implicit loss otherwise does not penalize instability.

VII. RESULTS

A. Simulation Results

We first evaluated our method in a Drake [41] simulation running at 1kHz (with data still reported at 30Hz). For the cuboid parameterization, our ground truth was a box with a randomized shape between 3 and 6cm, with our initial guess randomly sized between 1 and 3cm (in practice, a smaller initial guess lead to more stable learning). For the polytope parameterization, we used Platonic solids of fixed size (2cm bounding radius) as the ground-truth, and our




Fig. 4. bCH between the ground truth and estimated shapes for cuboid (top) and convex polytope (bottom) parameterizations in simulation. Shaded region is the 95% confidence interval for the mean of a normal distribution ($N = 10$). One outlier trial for random (marked with $\times$) does not factor into pictured 95% CI.

initial guess was always 20 random vertices sampled from the 1cm sphere. Each trial, we initialized both the ground truth and our initial guess to a random pose within 2cm of the center of the workspace. After the initial random action that made contact with the ground truth, we took 5 (for cuboid) or 8 (for polytope) subsequent actions, recording bCH after each one. The results are summarized in Fig. 4.

For the cuboid parameterization, Fig. 6 (top) shows an example 5-action learning trajectory. Observed information tends to jump quickly, as with only 3 geometry parameters a single action with low-variance data hitting the corners is sufficient to learn all of them with high certainty. While uniform random exploration does reach 5mm within 5 actions (5s of data), EIG exploration consistently reaches $< 4\text{mm}$ in the same time. At time step 5 this represents a significant improvement (two-tailed t-test with independent samples, $p < 0.05$ with a Bonferroni correction of 5 for all time steps, $N = 10$).

For the polytope parameterization, Fig. 6 (middle) shows 5 actions (excluding actions 3, 6, and 7 for brevity) of an example 8-action learning trajectory. Both EIG and Random exploration reach a bCH of 5mm consistently within 8 actions (8s of data). While not rising to the same statistical significance, a clear difference can still be seen between EIG and Random, with the latter requiring 3 extra actions for the 95% confidence interval to drop below 1cm. Note that this interval for Random already omits a catastrophic outlier, where the centroid of the guess was completely outside of the ground truth object, leading to 2 actions without contact before they re-aligned. Combined with an observable reduction in the variance of bCH across all actions, we can confidently say that the data supports the claim that EIG-




Fig. 5. bCH between the ground truth and estimated shapes on the real robot. Shaded region is the 95% confidence interval for the mean of a normal distribution ($N = 6$).




Fig. 6. Example learning trajectories for (from top to bottom): a simulated cuboid ground truth with a cuboid parameterization, a simulated regular octahedron ground truth with a 20-vertex polytope parameterization, and a real cuboid ground truth with a cuboid parameterization. Colors represent per-vertex (for the polytope) or total observed information $\log \det \mathcal{I}(\mathcal{D}_{obs})$. Each image is re-centered on the ground truth, which moved 1-2cm over the course of the trial.

maximizing exploration provides a more consistent learning trajectory. At the same time, even Random exploration still achieved consistent results with 5-8s of tactile data.

B. Real Robot Results

On our real Trifinger and Densetact setup, we used a cuboid ground truth and parameterization. The ground truth object had dimensions $5.8 \times 4.9 \times 5.3\text{cm}$ and weighing $\sim 0.5\text{kg}$. The object generally moved 1 – 2cm from its starting pose during each trial. Initialization was randomized as in simulation, with 5 actions performed after first contact. Fig. 6(bottom) shows an example 5-action learning trajectory.

The results are summarized in Fig. 5. While the overall absolute bCH is worse than simulation, a large portion can be attributed to 3-5mm of proprioceptive error in each end-effector. We do still see a significant improvement in EIG over Random exploration at time step 5 (two-tailed t-test with independent samples, $p < 0.05$ with a Bonferroni correction of 5 for all time steps, $N = 6$).

VIII. DISCUSSION

We have demonstrated that we can identify the geometry and pose of a continuous class of unknown dynamic convex rigid bodies with only 5-8s of tactile contact data. Furthermore, we have shown that our active exploration

procedure, by maximizing expected information gain, can lead to significantly faster and more consistent performance.

This procedure is not without limitations. Further hardware improvements can increase performance, e.g., the Denstact struggled to detect light touches, thus limiting our experiments to moderately heavy objects. More generally, EIG as a local metric can perform poorly if the MLE is a non-optimal local minimum. Our action space, which relied on the estimated object pose, was similarly vulnerable. This required multiple initializations to mitigate, which on top of the increasing dataset size over the course of the experiment slowed down training. Also, as discussed in Sec. V-B, computing EIG with respect to x_T requires quantifying the sensitivity of past measurements to future states, which is unavoidably poorly posed for stiff contact dynamics. Subject to Coulomb friction, no amount of smoothing can change the fact that, if the object is at rest ($z_T = 0$), it could have always been there or just arrived. While an identity dynamics Jacobian was sufficient for this work, further investigation into this issue is necessary.

With that said, this success can still be expanded to more fully demonstrate the model-building capability of tactile sensing. For example, while the current approach requires a convex object and end-effector to define a unique collision point, it should be possible to model non-convex objects as a union of convex shapes, further extending the space of possible geometry representations. Additionally, this work can be combined with a vision-based model-building solution [14, 18] to create a robust multi-modal system. This is a step towards the goal of a robot that can use all of the sensors at its disposal to model previously-unseen arbitrary objects in its environment throughout its entire deployment lifetime.

ACKNOWLEDGMENTS

This work was supported by an NSF CAREER Award under Grant No. FRR-2238480 and the RAI Institute. The authors extend their gratitude to Hrishikesh Sathyanarayan and Ian Abraham from the Intelligent Autonomy Lab at Yale University for all the insightful discussions early in the conception of this work.

REFERENCES

- [1] M. Lambeta, P.-W. Chou, S. Tian, B. Yang, B. Maloon, V. R. Most, D. Stroud, R. Santos, A. Byagowi, G. Kammerer, D. Jayaraman, and R. Calandra, "DIGIT: A Novel Design for a Low-Cost Compact High-Resolution Tactile Sensor with Application to In-Hand Manipulation," *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pp. 3838–3845, July 2020, arXiv:2005.14679 [cs]. [Online]. Available: <http://arxiv.org/abs/2005.14679>
- [2] M. K. Johnson, F. Cole, A. Raj, and E. H. Adelson, "Microgeometry capture using an elastomeric sensor," in *ACM SIGGRAPH 2011 papers*, ser. SIGGRAPH '11. New York, NY, USA: Association for Computing Machinery, July 2011, pp. 1–8. [Online]. Available: <https://doi.org/10.1145/1964921.1964941>
- [3] W. K. Do, B. Jurewicz, and M. K. III, "DenseTact 2.0: Optical Tactile Sensor for Shape and Force Reconstruction," Mar. 2023, arXiv:2209.10122 [cs]. [Online]. Available: <http://arxiv.org/abs/2209.10122>
- [4] A. I. Weinberg, A. Shirizly, O. Azulay, and A. Sintov, "Survey of Learning-based Approaches for Robotic In-Hand Manipulation," Oct. 2024, arXiv:2401.07915 [cs]. [Online]. Available: <http://arxiv.org/abs/2401.07915>
- [5] J. Hejna, S. Mirchandani, A. Balakrishna, A. Xie, A. Wahid, J. Thompson, P. Sanketi, D. Shah, C. Devin, and D. Sadigh, "Robot Data Curation with Mutual Information Estimators," Feb. 2025, arXiv:2502.08623 [cs]. [Online]. Available: <http://arxiv.org/abs/2502.08623>
- [6] T. Schneider, B. Belousov, G. Chalvatzaki, D. Romeres, D. K. Jha, and J. Peters, "Active Exploration for Robotic Manipulation," Oct. 2022, arXiv:2210.12806 [cs]. [Online]. Available: <http://arxiv.org/abs/2210.12806>
- [7] Z. Yi, R. Calandra, F. Veiga, H. van Hoof, T. Hermans, Y. Zhang, and J. Peters, "Active tactile object exploration with Gaussian processes," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct. 2016, pp. 4925–4930, iSSN: 2153-0866. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/7759723>
- [8] D. Driess, P. Englert, and M. Toussaint, "Active learning with query paths for tactile object shape exploration," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Sept. 2017, pp. 65–72, iSSN: 2153-0866. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/8202139>
- [9] M. M. Zhang, N. Atanasov, and K. Daniilidis, "Active end-effector pose selection for tactile object recognition through Monte Carlo tree search," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Sept. 2017, pp. 3258–3265, iSSN: 2153-0866. [Online]. Available: <https://ieeexplore.ieee.org/document/8206161>
- [10] A.-H. Shahidzadeh, S. J. Yoo, P. Mantripragada, C. D. Singh, C. Fermüller, and Y. Aloimonos, "AcTExplore: Active Tactile Exploration on Unknown Objects," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, May 2024, pp. 3411–3418. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/10611667>
- [11] S. Pfrommer, M. Halm, and M. Posa, "ContactNets: Learning Discontinuous Contact Dynamics with Smooth, Implicit Representations," Nov. 2020, arXiv:2009.11193 [cs]. [Online]. Available: <http://arxiv.org/abs/2009.11193>
- [12] B. Bianchini, M. Halm, N. Matni, and M. Posa, "Generalization Bounded Implicit Learning of Nearly Discontinuous Functions," in *Proceedings of The 4th Annual Learning for Dynamics and Control Conference*. PMLR, May 2022, pp. 1112–1124, iSSN: 2640-3498. [Online]. Available: <https://proceedings.mlr.press/v168/bianchini22a.html>
- [13] O. Özyeşil, V. Voroninski, R. Basri, and A. Singer, "A survey of structure from motion," *Acta Numerica*, vol. 26, pp. 305–364, May 2017. [Online]. Available: <https://www.cambridge.org/core/journals/acta-numerica/article/survey-of-structure-from-motion/C4B2E7BB10BC2C11AF71BC80B584D378>
- [14] W. Jiang, B. Lei, and K. Daniilidis, "FisherRF: Active View Selection and Uncertainty Quantification for Radiance Fields using Fisher Information," Dec. 2024, arXiv:2311.17874 [cs]. [Online]. Available: <http://arxiv.org/abs/2311.17874>
- [15] W. Jiang, B. Lei, K. Ashton, and K. Daniilidis, "AG-SLAM: Active Gaussian Splatting SLAM," Oct. 2024, arXiv:2410.17422 [cs] version: 1. [Online]. Available: <http://arxiv.org/abs/2410.17422>
- [16] B. Wen, J. Tremblay, V. Blukis, S. Tyree, T. Muller, A. Evans, D. Fox, J. Kautz, and S. Birchfield, "BundleSDF: Neural 6-DoF Tracking and 3D Reconstruction of Unknown Objects," Mar. 2023, arXiv:2303.14158 [cs]. [Online]. Available: <http://arxiv.org/abs/2303.14158>
- [17] B. Wen, W. Yang, J. Kautz, and S. Birchfield, "FoundationPose: Unified 6D Pose Estimation and Tracking of Novel Objects," in *2024 IEEE/CVF Conference on*

Computer Vision and Pattern Recognition (CVPR). Seattle, WA, USA: IEEE, June 2024, pp. 17 868–17 879. [Online]. Available: <https://ieeexplore.ieee.org/document/10655554/>

- [18] B. Bianchini, M. Zhu, M. Sun, B. Jiang, C. J. Taylor, and M. Posa, “Vysics: Object Reconstruction Under Occlusion by Fusing Vision and Contact-Rich Physics,” Apr. 2025, arXiv:2504.18719 [cs]. [Online]. Available: <http://arxiv.org/abs/2504.18719>
- [19] M. C. Koval, N. S. Pollard, and S. S. Srinivasa, “Pose estimation for planar contact manipulation with manifold particle filters,” *The International Journal of Robotics Research*, vol. 34, no. 7, pp. 922–945, June 2015. [Online]. Available: <http://journals.sagepub.com/doi/10.1177/0278364915571007>
- [20] S. Suresh, Z. Si, S. Anderson, M. Kaess, and M. Mukadam, “MidasTouch: Monte-Carlo inference over distributions across sliding touch,” in *Proceedings of The 6th Conference on Robot Learning*. PMLR, Mar. 2023, pp. 319–331, iSSN: 2640-3498. [Online]. Available: <https://proceedings.mlr.press/v205/suresh23a.html>
- [21] P. Sodhi, M. Kaess, M. Mukadam, and S. Anderson, “PatchGraph: In-hand tactile tracking with learned surface normals,” in *2022 International Conference on Robotics and Automation (ICRA)*, May 2022, pp. 2164–2170. [Online]. Available: <https://ieeexplore.ieee.org/document/9811953>
- [22] R. Bajcsy, “Active perception,” *Proceedings of the IEEE*, vol. 76, no. 8, pp. 966–1005, Aug. 1988. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/5968>
- [23] D. Pathak, D. Gandhi, and A. Gupta, “Self-Supervised Exploration via Disagreement,” June 2019, arXiv:1906.04161 [cs, stat]. [Online]. Available: <http://arxiv.org/abs/1906.04161>
- [24] M. Memmel, A. Wagenmaker, C. Zhu, P. Yin, D. Fox, and A. Gupta, “ASID: Active Exploration for System Identification in Robotic Manipulation,” June 2024, arXiv:2404.12308 [cs]. [Online]. Available: <http://arxiv.org/abs/2404.12308>
- [25] H. Sathyanarayan and I. Abraham, “Behavior Synthesis via Contact-Aware Fisher Information Maximization,” May 2025, arXiv:2505.12214 [cs]. [Online]. Available: <http://arxiv.org/abs/2505.12214>
- [26] I. Abraham, A. Mavrommati, and T. Murphey, “Data-Driven Measurement Models for Active Localization in Sparse Environments,” in *Robotics: Science and Systems XIV*. Robotics: Science and Systems Foundation, June 2018. [Online]. Available: <http://www.roboticsproceedings.org/rss14/p45.pdf>
- [27] I. Abraham, A. Prabhakar, M. J. Z. Hartmann, and T. D. Murphey, “Ergodic Exploration using Binary Sensing for Non-Parametric Shape Estimation,” *IEEE Robotics and Automation Letters*, vol. 2, no. 2, pp. 827–834, Apr. 2017, arXiv:1709.01560 [cs]. [Online]. Available: <http://arxiv.org/abs/1709.01560>
- [28] J. Hu, S. Cui, S. Wang, R. Wang, and Y. Wang, “Active shape reconstruction using a novel visuotactile palm sensor,” *Biomimetic Intelligence and Robotics*, vol. 4, no. 3, p. 100167, Sept. 2024. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2667379724000251>
- [29] J. Ketchum, A. Prabhakar, and T. D. Murphey, “Active Exploration for Real-Time Haptic Training,” May 2024, arXiv:2405.11776 [cs]. [Online]. Available: <http://arxiv.org/abs/2405.11776>
- [30] T. Schneider, G. Duret, C. d. Farias, R. Calandra, L. Chen, and J. Peters, “Tactile MNIST: Benchmarking Active Tactile Perception,” June 2025, arXiv:2506.06361 [cs]. [Online]. Available: <http://arxiv.org/abs/2506.06361>
- [31] J. Kim, Y. Zhu, and A. Dollar, “Tactile Probabilistic Contact Dynamics Estimation of Unknown Objects,” Sept. 2024, arXiv:2409.17470 [cs]. [Online]. Available: <http://arxiv.org/abs/2409.17470>
- [32] J. Xu, H. Lin, S. Song, and M. Ciocarlie, “TANDEM3D: Active Tactile Exploration for 3D Object Recognition,” Feb. 2023, arXiv:2209.08772 [cs]. [Online]. Available: <http://arxiv.org/abs/2209.08772>
- [33] 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 *2021 IEEE International Conference on Robotics and Automation (ICRA)*, May 2021, pp. 11 322–11 328, iSSN: 2577-087X. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/9562060>
- [34] M. Anitescu, “Optimization-based simulation of nonsmooth rigid multibody dynamics,” *Mathematical Programming*, vol. 105, no. 1, pp. 113–143, Jan. 2006. [Online]. Available: <https://doi.org/10.1007/s10107-005-0590-7>
- [35] A. Agrawal, B. Amos, S. Barratt, S. Boyd, S. Diamond, and Z. Kolter, “Differentiable convex optimization layers,” in *Advances in Neural Information Processing Systems*, 2019.
- [36] A. Kirsch and Y. Gal, “Unifying Approaches in Active Learning and Active Sampling via Fisher Information and Information-Theoretic Quantities,” *Transactions on Machine Learning Research*, 2022. [Online]. Available: <https://openreview.net/pdf?id=UVDKQANOW>
- [37] M. Blondel, Q. Berthet, M. Cuturi, R. Frostig, S. Hoyer, F. Llinares-López, F. Pedregosa, and J.-P. Vert, “Efficient and modular implicit differentiation,” *arXiv preprint arXiv:2105.15183*, 2021.
- [38] M. Wüthrich, F. Widmaier, F. Grimmering, J. Akpo, S. Joshi, V. Agrawal, B. Hammoud, M. Khadiv, M. Bogdanovic, V. Berenz, J. Viereck, M. Naveau, L. Righetti, B. Schölkopf, and S. Bauer, “TriFinger: An Open-Source Robot for Learning Dexterity,” Jan. 2021, arXiv:2008.03596 [cs]. [Online]. Available: <http://arxiv.org/abs/2008.03596>
- [39] Z. Teed and J. Deng, “RAFT: Recurrent All-Pairs Field Transforms for Optical Flow,” Aug. 2020, arXiv:2003.12039 [cs]. [Online]. Available: <http://arxiv.org/abs/2003.12039>
- [40] Y. Zhang, Z. Kan, Y. Yang, Y. A. Tse, and M. Y. Wang, “Effective estimation of contact force and torque for vision-based tactile sensors with helmholtz–hodge decomposition,” *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 4094–4101, 2019.
- [41] R. Tedrake, *Drake Software*. Course Notes for MIT 6.421, 2024. [Online]. Available: <https://manipulation.csail.mit.edu/drake.html>