Active Tactile Pose and Shape Estimation of Highly Dynamic Objects

Ethan K. Gordon, Michael Posa[1]

***Abstract*— Learning a physically accurate object model at test time can provide significant benefits in predictability and reuse between tasks. Tactile sensing compliments 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 our main conference work [1], we introduced an active learning and exploration framework that uses only tactile data to simultaneously determine the shape and location of rigid objects with minimal robot motion. This abstract proposes an extension to that work, addressing the fact that information quantification in dynamic systems requires reasoning about the likelihood of *past* measurements given a *current* state. Though marginalization and Markov Chain Monte Carlo (MCMC) trajectory sampling, we can avoid the difficulties of poorly-conditioned "backwards" simulation and nearly-discontinuous contact dynamics. The resulting information matrix in principle should be more robust to highly-dynamic trajectories such as tumbling.**

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. In this task, tactile sensing, with its robustness to darkness, reflections, and occlusions, can compliment visual modalities. With the advent of inexpensive sensors [2, 3, 4], touch is becoming a more common component of robot manipulation systems [5].

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 [6, 7, 8, 9, 10, 11]. 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, motivating the use of expected information gain (EIG), which can forego the use of an explicit belief distribution.

Our most recent work [1] introduces both (1) a *violation-implicit loss* that avoids differentiating through nearly-discontinuous rigid-body contact dynamics to enable the simultaneous leaning of rigid-body pose and geometry, and (2) a framework for computing EIG, a function of both observed and expected future information, in the dynamic setting. The

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

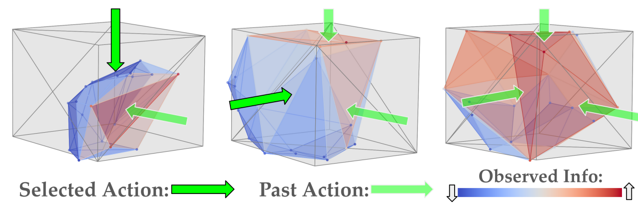

Fig. 1. Only tactile data is used to find the pose and geometry of an arbitrary convex object (e.g., a 20-vertex polytope as shown here). [1]

computation of EIG involved the gradient of the likelihood of *past* measurements given *current* (or future) states, which at face value requires the gradient of a backwards simulation. This was made tractable by setting the dynamics Jacobian to the identity matrix. While successful for minimally dynamic trajectories, this approach is likely to fail for more chaotic behaviors such as tumbling. This abstract proposes a modified computation based on marginalization of the likelihood and Monte Carlo trajectory sampling to compute EIG, avoiding both backwards simulation and the need for this simplification.

II. 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 measurements m_t, and we assume a known per-timestep measurement likelihood $p(m_t|\theta, x_t)$. The goal is to determine the current pose of the object x_T and its geometry θ in as few actions as possible. The interaction protocol is summarized as follows.

1) **Learning:** At current time T, given some observed dataset $m_{t<T}$, produce some estimate of the object pose and shape. In practice, this will be the minimum of some loss function $\mathcal{L}$, with the entire trajectory $\tau = x_{t<T}$ optimized simultaneously.

$$\tilde{\theta}, \tilde{x}_T, \tilde{\tau} = \arg\min_{\theta, x_T, \tau} \mathcal{L}(\theta, x_T, \tau, r_t, m_{t<T}) \quad (1)$$

2) **Exploration:** Given some known horizon $H > T$, choose a trajectory $r_{T\ldots H}$ and collect new data $m_{T\ldots H}$. Repeat with $T \leftarrow H$.

III. Preliminaries

Active exploration often involves a scalar quantifying the information present in past data and expected in future data. A common approach is to consider the entropy H belief prior $p(\omega|m_{t<T})$ and posterior $p(\omega|m_{t<H})$ distributions over model parameters $\omega = \{\theta, x_T\}$, where $H[\omega|m_{[t]}] := \mathbb{E}_p[-\log p(\omega|m_{[t]})]$. For a Gaussian, the entropy can be

expressed in terms of the log-determinant of the covariance: $H[\mathcal{N}(\mu, \Sigma)] = 1/2 \log \det \Sigma + const$. 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 the Hessian of the log of the posterior (or, with an uninformed prior, the likelihood) about that peak. This in expectation is equal to the variance of the gradient around the peak.

$$\begin{aligned} H[\omega | m_{[t]}] &\approx \frac{1}{2} \log \det Var \left(\nabla_\omega \log(p(m_{[t]}|\omega)|_{\tilde{\omega}} \right) \\ &:= \frac{1}{2} \log \det \mathcal{I}(m_{[t]}, \tilde{\omega}) \end{aligned} \tag{2}$$

The empirical gradient variance from past measurements is called *observed information*. Fisher information $\mathcal{F}(r_{T \cdots H}, \tilde{\omega}) = \mathbb{E}_{p(m_{[t]}|r_{[t]})} \mathcal{I}(m_{[t]}, \tilde{\omega})$ is the expected variance given unknown future measurements. EIG for a robot trajectory can be computed by using Fisher and observed information without explicitly maintaining the posterior.

$$EIG(r_{T \cdots H}) :\propto \log \det(\mathcal{F}(r_{T \cdots H}) \mathcal{I}(m_{t<T})^{-1} + \mathbf{I}) \tag{3}$$

For a full derivation we direct the reader to [12], Sec. 3-5.

IV. Proposed EIG Computation Methods

As we would like to quantify the information we have about the state of the object at the current time T, the key challenge with computing EIG is the efficient computation of the gradient $\nabla_\omega \log(p(m_t|\{\theta, x_T\})) \forall t$.

A. Baseline: Identity Jacobian

As proposed in [1], we can consider a direct computation of the gradient by imagining the state $x_t = g^{T-t}(x_T)$ as being generated by a backwards dynamics function $x_{t-1} = g(x_t)$. By the chain rule:

$$\nabla_\omega \log(p(m_t|\omega)) = \nabla_{x_t} \log(p(m_t|\{\theta, x_t\}) \cdot \nabla_{x_T} g^{T-t}(x_T) \tag{4}$$

Since g is in general not well-defined for Coulomb-frictional contact, we set $\nabla g = \mathbf{I}$, the identity matrix. This reduces the gradient to just that of the measurement model $p(m_t|\theta, x_t)$, which is in our control and well-defined.

Besides ease of computation, the benefit of this approach is that we can define $p(m_t|\theta, x_t)$ such that the variance of its gradient, in expectation, is independent of the specific m_t, obviating the need to explicitly sample potential measurements for the computation of $\mathbb{E}_{p(m_t|r_t)}$ in $\mathcal{F}$. For example, a Gaussian likelihood has this property. The downside is that it neglects the dynamics between states, which can lead to significantly incorrect information calculations, particularly for chaotic motion.

B. Baseline: Differentiable Simulation

While g is not well defined, the forward dynamics function $x_t = f(x_{t-1})$ is. Therefore, for observed information, we can utilize the learned trajectory $\tilde{\tau}$ and compute the gradient with respect to $t = 0$:

$$\mathcal{I}(\omega_0) = Var \left(\nabla_\omega \log(p(m_t|\{\theta, x_t\}) \cdot \nabla_{x_0} f^t(x_t)|_{\tilde{\tau}} \right) \tag{5}$$

Since EIG involves the inverse of observed information, we can "bring forward" observed information to be with respect to time $t = T$ again using the forward Jacobian.

$$\mathcal{I}(\omega_T)^{-1} = \nabla_{\omega_0} f^T(x_0) \mathcal{I}(\omega_0)^{-1} \left(\nabla_{\omega_0} f^T(x_0) \right)^\intercal \tag{6}$$

This step is not necessary for expected information. The potential downside with this approach is that the gradients of f may be poorly conditioned numerically (being nearly-0 or nearly-∞). This is discussed in previous work [13, 14] and motivated the use of the *violation-implicit* loss in [1].

C. Marginalization and MCMC Sampling

Our key insight is that we can marginalize the measurement likelihood over trajectories τ:

$$\nabla_{x_T} p(m_{t<T}|x_T) \approx^s \nabla_{x_T} \sum_{\tau \sim U} p(m_{t<T}|\tau) p(\tau|x_T) \tag{7}$$

Where the marginalization integral is approximated by a summation over uniform trajectory samples, and $p(m_{t<T}|\tau) = \prod_t p(m_t|x_t)$. Under the assumption that the trajectory likelihood is in the exponential family, $p(\tau|x_T) \propto \exp(h(\tau, x_T))$, we can write $\nabla_{x_T} \log p(m_{t<T}|x_T)$ in terms of the softmax S over uniformly sampled τ:

$$= (S_{\tau \sim U}(\log p(m_t|\tau) + h(\tau, x_T)) - S_{\tau \sim U}(h)) \cdot \nabla_{x_T} h \tag{8}$$

Taking uniform trajectory samples is intractable. However, by instead taking N samples using MCMC, the h drops from inside the softmax.

$$\approx^s \left(S_{\tau \sim MC}(\log p(m_t|\tau)) - \frac{1}{N} \right) \cdot \nabla_{x_T} h(\tau, x_T) \tag{9}$$

At the cost of sampling, and by using the *violation-implicit* loss formulation for h, we avoid both backwards simulation and the need for ∇f. Additionally, sampling can be made more efficient by starting at the known likely trajectory $\tilde{\tau}$.

One additional potential upside of this approach is that it can also be applied to $\mathcal{F}$ to get the expected information with respect to the *end state* x_H. These trajectories will in principle make the state more certain in the future, rather than only reasoning about the present.

V. Next Steps

We propose running the experiments as-described in [1], Sec. 6, with the three EIG computation methods listed above. Our hypothesis is that for highly dynamic trajectories, our proposed method will lead to faster learning outcomes compared to the baselines for the reasons listed here.

Acknowledgments: This work was supported by an NSF CAREER Award (FRR-2238480) and the RAI Institute.

References

[1] E. K. Gordon, B. Baraki, H. Bui, and M. Posa, "Active tactile exploration for rigid body pose and shape estimation," 2025. [Online]. Available: https://arxiv.org/abs/2510.13595

[2] 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

[3] 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

[4] 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

[5] 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

[6] J. Hejna, S. Mirchandani, A. Balakrishna, A. Xie, A. Wahid, J. Tompson, 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

[7] 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

[8] 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

[9] 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

[10] 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

[11] 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

[12] 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=UVDAKQANOW

[13] 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

[14] 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