

Learning Tactile Models for Factor Graph-based Estimation

Paloma Sodhi^{1,2}, Michael Kaess², Mustafa Mukadam¹, and Stuart Anderson¹

¹Facebook AI Research, ²Carnegie Mellon University

Abstract—We’re interested in the problem of estimating object states from touch during manipulation under occlusions. In this work, we address the problem of estimating object poses from touch during planar pushing. Vision-based tactile sensors provide rich, local image measurements at the point of contact. A single such measurement, however, contains limited information and multiple measurements are needed to infer latent object state. We solve this inference problem using a factor graph. In order to incorporate tactile measurements in the graph, we need local observation models that can map high-dimensional tactile images onto a low-dimensional state space. Prior work has used low-dimensional force measurements or engineered functions to interpret tactile measurements. These methods, however, can be brittle and difficult to scale across objects and sensors. Our key insight is to directly learn tactile observation models that predict the *relative pose* of the sensor given a pair of tactile images. These relative poses can then be incorporated as factors within a factor graph. We propose a two-stage approach: first we learn local tactile observation models supervised with ground truth data, and then integrate these models along with physics and geometric factors within a factor graph optimizer. We demonstrate reliable object tracking using *only tactile* feedback for ~ 150 real-world planar pushing sequences with varying trajectories across three object shapes.

I. INTRODUCTION

We look at the problem of estimating object states such as poses from touch. Consider the example of a robot manipulating an object in-hand: as the object is being manipulated, it is occluded by the fingers. This occlusion renders visual estimation alone insufficient. *Touch* in such cases can provide local, yet precise information about the object state.

The advent of new touch sensors [1–4] has enabled rich, local tactile image-based measurements at the point of contact. However, a single tactile image only reveals limited information that may correspond to multiple object poses, making it hard to directly predict poses from an image. We hence need to be able to reason over multiple such measurements to collapse uncertainty. This results in the inference problem: *Compute the sequence of latent object poses, given a stream of tactile image measurements.*

We solve this inference problem using a factor graph [5, 6] that offers a flexible way to process a stream of measurements while incorporating other priors including physics and geometry. The factor graph relies on having local observation models that can map measurements into the state space. Observation models for high-dimensional tactile measurements are, however, challenging to design. Prior work [7, 8] used low-dimensional force measurements

Work done when Paloma Sodhi interned at Facebook AI Research. psodhi@cs.cmu.edu. Code and supplementary material can be found on <https://psodhi.github.io/tactile-pushing>

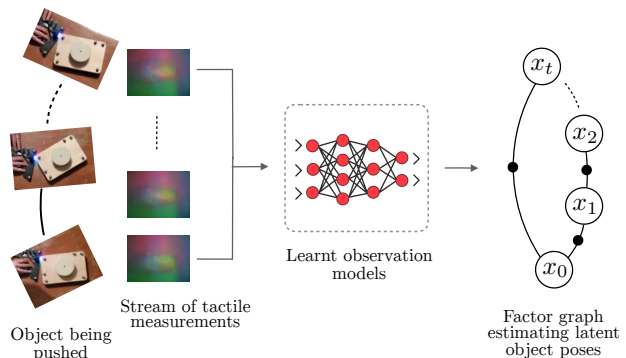


Fig. 1: Estimating latent object poses from a stream of tactile images modeled as an inference over a factor graph. Since design of observation models for such measurements is challenging, we learn these models to be used as factors within the graph.

or engineered functions for high-dimensional tactile force measurements. These, however, can be brittle and difficult to scale across objects and sensors.

In this paper, we propose learning tactile observation models that are incorporated as factors within a factor graph during inference (Fig. 1). The learner takes high-dimensional tactile image measurements and predicts noisy low-dimensional poses that are then integrated by the factor graph optimizer. For the learner, however, accurately predicting object pose directly from a tactile image is typically not possible without additional information such as a tactile map. Our key insight is to instead have the learner predict *relative poses* from image pairs. That is, given a pair of non-sequential tactile images, predict the difference in the pose of the tactile sensor relative to the object in contact. This relative pose information is then used to correct for accumulated drift in the object pose estimate. Observation models trained in this fashion can work for a large class of objects, as long as the local contact surface patches are in distribution to those found in training data.¹

We propose a two-stage approach: we first learn local tactile observation models supervised with ground truth data, and then integrate these along with physics and geometric models as factors within a factor graph optimizer. The tactile observation model is learning-based and integrates with a factor graph. This integration leverages benefits of both learning-based and model-based methods — wherein, (a) our learnt factor is general and more accurate than engineered functions, and (b) integrating this learnt model within a

¹While we do not require a tactile map, we do need a category class of the object shape in the observation model and a geometric model prior in the factor graph. These are much easier to obtain as they do not require tactile sensor readings that cover the full object.

factor graph lets us reason over a stream of measurements while incorporating structural priors in an efficient, real-time manner. Our main contributions are: topsep=0pt

- 1) A novel learnable tactile observation model that integrates as factors within a factor graph.
- 2) Tactile factors that work for multiple object shapes.
- 3) Real-time object tracking during real-world pushing trials using *only tactile* measurements.

II. RELATED WORK

Localization and state estimation are increasingly solved as smoothing problems given the increased accuracy and efficiency over their filtering counterparts [9]. Typically, the smoothing objective is formulated as inference over a factor graph whose variable nodes encode latent states and factor nodes the measurement likelihoods [5, 6, 10]. To incorporate high-dimensional measurements as factors in the graph, one needs an observation model to map between latent states and measurements. These typically are analytic functions e.g. projection geometry for images [11, 12], or scan matching for point clouds [13, 14]. Recent work in visual SLAM has also looked at using such functions on learnt, low-dimensional encodings of the original image measurements [15, 16].

While there exist well-studied models within visual SLAM literature, it is hard to obtain such a-priori models for complex sensors such as vision-based tactile sensors [1, 2] that capture deformations lacking a straightforward metric interpretation. Recent work is seeking to create increasingly accurate simulator models for such sensors [17, 18]. Within tactile manipulation literature, one class of approaches using such sensors make use of tactile images directly as feedback to solve various control tasks such as cable manipulation [19], in-hand marble manipulation [1], box-packing [20]. While efficient for the particular task, it can be difficult to scale across different tasks. An inference module is additionally needed to estimate a common latent state representation like global object poses that can be used for different downstream control and planning tasks.

Prior work on estimating states from touch during manipulation has included filtering methods [21–26], learning-only methods [27, 28], methods utilizing a prior map or tactile simulators [29, 30], and graph-based smoothing methods [7, 8, 31, 32]. Smoothing approaches that model the problem as inference over a factor graph have the benefits of (a) being more accurate than filtering methods, (b) incorporating structural priors unlike learning-only methods, and (c) recovering global object poses from purely local observation models without needing a global map. Moreover, the graph inference objective can be solved in real-time making use of fast, incremental tools [33–35] in the literature.

Of these different methods, the work in [7, 8, 29, 30] is most closely related. Yu et al. [7], Lambert et al. [8] solve for a factor graph-based pose estimation objective but use either low-dimensional force readings or engineered functions for high-dimensional readings from sensors like Biotac. Bauza et al. [29, 30] use different methods for pose estimation but use vision-based tactile sensors most similar

to ours. They solve for a global localization objective with either an offline tactile map [29], or contact image renderings from a simulator [30]. We see this global localization objective complementary to the object state tracking objective achieved with our proposed learnt tactile factors.

III. PROBLEM FORMULATION

We formulate the estimation problem as inference over a factor graph. A factor graph is a bipartite graph with two types of nodes: variables x and factors ϕ . Variable nodes are the latent states to be estimated, and factor nodes encode constraints on these variables such as measurement likelihood functions, or physics, geometric models. Maximum a posteriori (MAP) inference over a factor graph involves maximizing the product of all factor graph potentials, i.e.,

$$\hat{x} = \operatorname{argmax}_x \prod_{i=1}^m \phi_i(x) \quad (1)$$

Under Gaussian noise model assumptions, MAP inference is equivalent to solving a nonlinear least-squares problem [5]. That is, for Gaussian factors $\phi_i(x)$ corrupted by zero-mean, normally distributed noise,

$$\begin{aligned} \phi_i(x) &\propto \exp \left\{ -\frac{1}{2} \|F_i(x)\|_{\Sigma_i}^2 \right\} \\ \Rightarrow \hat{x} &= \operatorname{argmin}_x \frac{1}{2} \sum_{i=1}^m \|F_i(x)\|_{\Sigma_i}^2 \end{aligned} \quad (2)$$

where, $F_i(x)$ are cost functions defined over states x and include measurement likelihoods or priors derived from physical or geometric assumptions. $\|\cdot\|_{\Sigma_i}$ is the Mahalanobis distance with covariance Σ_i .

For our planar pushing setup, states in the graph are the planar object and end-effector poses at every time step $t = 1 \dots T$, i.e. $x_t = [o_t \ e_t]^T$, where $o_t, e_t \in SE(2)$. Factors in the graph incorporate tactile observations $F_{tac}(\cdot)$, quasi-static pushing dynamics $F_{qs}(\cdot)$, geometric constraints $F_{geo}(\cdot)$, and priors on end-effector poses $F_{eff}(\cdot)$.

At every time step, new variables and factors are added to the graph. Writing out Eq. 2 for our setup at time step T ,

$$\begin{aligned} \hat{x}_{1:T} = \operatorname{argmin}_{x_{1:T}} \sum_{t=1}^T \{ &\|F_{qs}(o_{t-1}, o_t, e_{t-1}, e_t)\|_{\Sigma_{qs}}^2 + \\ &\|F_{geo}(o_t, e_t)\|_{\Sigma_{geo}}^2 + \|F_{tac}(o_{t-k}, o_t, e_{t-k}, e_t)\|_{\Sigma_{tac}}^2 + \\ &\|F_{eff}(e_t)\|_{\Sigma_{eff}}^2 \} \end{aligned} \quad (3)$$

Eq. 3 is the optimization objective that we must solve for every time step. Instead of resolving from scratch every time step, we make use of efficient, incremental solvers such as iSAM2 [34] for real-time inference. Individual cost terms in Eq. 3 are described in more detail in Section IV-B.

IV. APPROACH

We present a two-stage approach: we first learn local tactile observation models from ground truth data (Section IV-A), and then integrate these models along with physics and

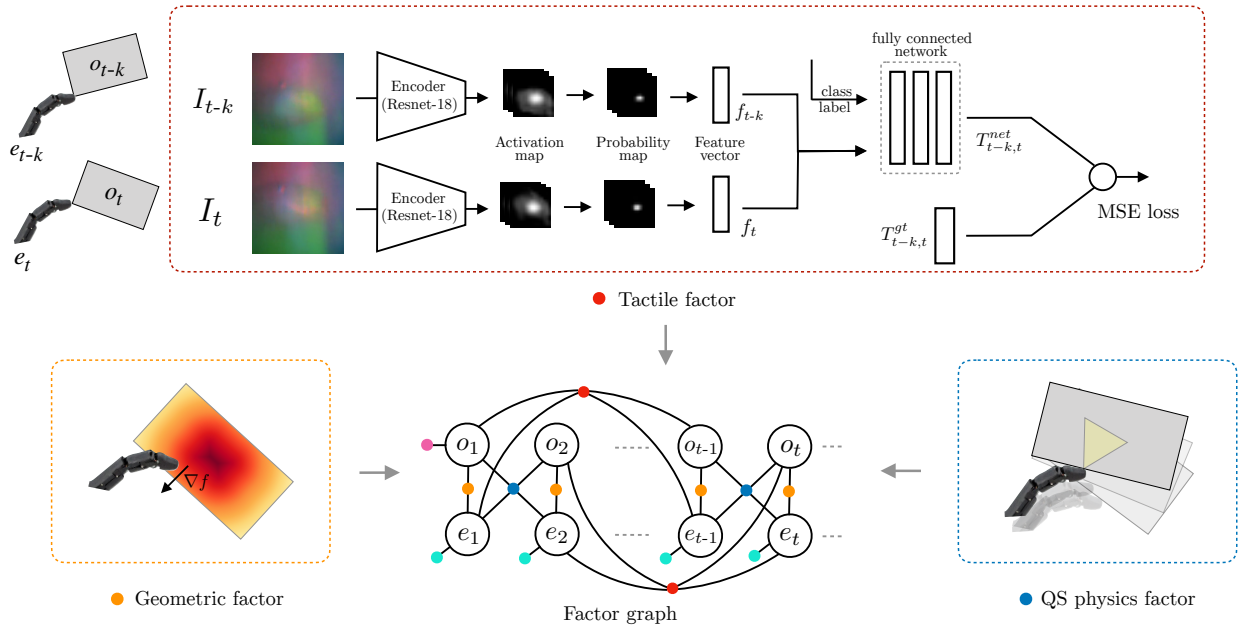


Fig. 2: Overall approach showing the three factors in the factor graph — tactile, physics, geometric. Tactile factor (red) predicts relative sensor poses given a pair of tactile images. Physics factors (blue) model quasi-static pushing. Geometric factors (orange) model object and end-effector intersections. We also include global pose priors as unary factors on end-effector (cyan) and the first object pose (pink).

geometric models as factors within a factor graph (Section IV-B). Fig. 2 illustrates our overall approach showing the three factors, tactile, physics and geometric, being integrated into the factor graph.

A. Tactile observation model

The goal of learning a tactile observation model is to derive the tactile factor cost term $\|F_{tac}(o_{t-k}, o_t, e_{t-k}, e_t)\|_{\Sigma_{tac}}^2$ in Eq. 3 to be used during graph optimization. We do this by predicting a relative transformation and penalizing deviations from this prediction. The relative transformation is that of the sensor (or end-effector) pose relative to the object in contact.

Our learnt tactile observation model consists of a transform prediction network that: Given a pair of non-sequential tactile image inputs $\{I_{t-k}, I_t\}$ at times $\{t-k, t\}$, predicts the relative transformation $T_{t-k,t}^{net}$. This is done by featurizing each image.

Feature learning: For encoding image I_t as feature f_t , we use an auto-encoder with a structural keypoint bottleneck proposed in [1]. It consists of an encoder and decoder using ResNet-18 as the backbone network. The encoder processes image input into a feature map from which K 2D keypoint locations corresponding to maximum feature activations are extracted. At decoding, a Gaussian blob is drawn on an empty feature map for each extracted keypoint. The decoder takes these as inputs and produces a reconstructed image. The auto-encoder is trained in a self-supervised manner with L2 image reconstruction loss along with auxiliary losses optimizing for sparse, non-redundant keypoints.

Transform prediction network: Once we have a trained feature encoder, we use that encoder (with weights fixed) in the transform prediction network to map input image pairs $\{I_{t-k}, I_t\}$ into keypoint feature vectors $\{f_{t-k}, f_t\}$. This is followed by a fully-connected regression network that

predicts a relative 2D transformation $T_{t-k,t}^{net}$ between times $\{t-k, t\}$ (Fig. 2). To make the same transform prediction network work across object classes, we also pass in a one-hot class label vector $c \in \{0, 1\}^{N_c}$. We expand the feature inputs via an outer product with the class vector, $f_{t-k} \otimes c$, $f_t \otimes c$, and pass this expanded input to the fully-connected layers. The label vector c corresponds to a category class of the object shape, we assume we know this for the object being pushed.

The transform prediction network is trained using a mean-squared loss (MSE) against ground truth data $T_{t-k,t}^{gt}$. We make the loss symmetric so that the network learns both regular and inverse transform, that is,

$$MSE = \frac{1}{|\mathcal{S}|} \sum_{(t-k,t) \in \mathcal{S}} \underbrace{(T_{t-k,t}^{gt} - T_{t-k,t}^{net})^2}_{\text{transform error}} + \underbrace{(T_{t,t-k}^{gt} - T_{t,t-k}^{net})^2}_{\text{inverse transform error}} \quad (4)$$

Here, $T_{t-k,t}^{gt}$ is the relative transform between end-effector poses at time steps $\{t-k, t\}$ in object coordinate frame. That is, $T_{t-k,t}^{gt} = [o_t^{-1}e_t]^{-1}[o_{t-k}^{-1}e_{t-k}]$, where $o_t, o_{t-k}, e_t, e_{t-k} \in SE(2)$ are ground truth object and end-effector poses obtained using a motion capture system. \mathcal{S} is the set of all non-sequential image pairs $(t-k, t)$ over a chosen time window $k \in [k_{min}, k_{max}]$.

B. Factor graph optimization

Once we have the learnt local tactile observation model, we integrate it along with physics and geometric models as factors within a factor graph. The factor graph optimizer then solves for the joint objective in Eq. 3. Here we look at each of the cost terms in Eq. 3 in more detail.

Tactile factor: Measurements from tactile sensor images are incorporated as relative tactile factors denoted by the cost $\|F_{tac}(\cdot)\|_{\Sigma}^2$ in Eq. 3. Our relative tactile factor is a quadratic cost penalizing deviation from a predicted value, that is,

$$\|F_{tac}(o_{t-k}, o_t, e_{t-k}, e_t)\|_{\Sigma_{tac}}^2 := \|T_{t-k,t}^{graph} \ominus T_{t-k,t}^{net}\|_{\Sigma_{tac}}^2 \quad (5)$$

where, $T_{t-k,t}^{net}$ is the *predicted* relative transform from the transform prediction network that takes as inputs tactile image measurement at time steps $\{t-k, t\}$. $T_{t-k,t}^{graph}$ is the *estimated* relative transform using current variable estimates in the graph. \ominus denotes difference between two manifold elements. $k \in [k_{min}, k_{max}]$ is the time step window over which these relative tactile factors are added. We choose this to be some subset of the training window set $|S|$ in Eq. 4.

$T_{t-k,t}^{graph}$ computed using graph variable estimates $\{o_{t-k}, e_{t-k}, o_t, e_t\} \in SE(2)$, and $T_{t-k,t}^{net}$ computed as transform prediction network output can be expressed as,

$$\begin{aligned} T_{t-k,t}^{graph} &= [o_{t-k}^{-1} e_{t-k}]^{-1} [o_t^{-1} e_t] \\ T_{t-k,t}^{net} &\leftarrow \phi^{net}(I_{t-k}, I_t) \end{aligned} \quad (6)$$

Quasi-static physics factor: To model object dynamics as it is pushed, we use a quasi-static physics model. The quasi-static approximation assumes end-effector trajectories executed with negligible acceleration, i.e. the applied pushing force is just large enough to overcome friction without imparting an acceleration [36].

We use the velocity-only quasi-static model from [37] that uses a convex polynomial to approximate the limit surface for force-motion mapping. For sticking contact, the contact point pushing velocity must lie inside the motion cone, resulting in the following relationship between object and contact point velocities:

$$\begin{aligned} DV &= V_p \\ \Rightarrow \begin{bmatrix} 1 & 0 & -p_y \\ 0 & 1 & p_x \\ -p_y & p_x & -c^2 \end{bmatrix} \begin{bmatrix} v_x \\ v_y \\ \omega \end{bmatrix} &= \begin{bmatrix} v_{px} \\ v_{py} \\ 0 \end{bmatrix} \end{aligned} \quad (7)$$

where, V is the object twist, V_p is the contact point velocity, and $c = \tau_{max}/f_{max}$ is a hyper-parameter dependent on pressure distribution of the object [36]. We calculate this value assuming pressure to be distributed either uniformly or at the corners/edges of the objects.

The dynamics in Eq. 7 is incorporated as the quadratic quasi-static cost term $\|F_{qs}(\cdot)\|_{\Sigma}^2$ in Eq. 3. Expanding this,

$$\|F_{qs}(o_{t-1}, o_t, e_{t-1}, e_t)\|_{\Sigma_{qs}}^2 := \|DV - V_p\|_{\Sigma_{qs}}^2 \quad (8)$$

Object twist V is computed using object poses o_{t-1}, o_t and contact point velocity V_p is computed using end-effector contact point estimate $p_t = [e_t^x \ e_t^y]$. That is,

$$V = \begin{bmatrix} v_x \\ v_y \\ \omega \end{bmatrix} = \mathcal{R}(o_t) \begin{bmatrix} (o_t^x - o_{t-1}^x)/dt \\ (o_t^y - o_{t-1}^y)/dt \\ (o_t^\theta - o_{t-1}^\theta)/dt \end{bmatrix} \quad (9)$$

$$V_p = \begin{bmatrix} v_{px} \\ v_{py} \\ 0 \end{bmatrix} = \mathcal{R}(o_t) \begin{bmatrix} (e_t^x - e_{t-1}^x)/dt \\ (e_t^y - e_{t-1}^y)/dt \\ 0 \end{bmatrix} \quad (10)$$

where, $\mathcal{R}(o_t)$ rotates object and contact point velocities into current object frame.

Geometric factor: We would like to add a geometric constraint to ensure that the contact point lies on the object surface. We do so as an intersection cost between the end-effector and the object. This is similar to the obstacle avoidance cost in [38, 39] but instead of a one-sided contact function, we use a two-sided function to penalize the contact point from lying on either side of the object surface.

To compute this intersection cost, we transform end-effector contact point $p_t = [e_t^x \ e_t^y] \in \mathbb{R}^2$ into current object frame as $o_t^{-1} p_t$, and look up its distance value in a precomputed 2D signed distance field map $\xi^{sdf}(\cdot)$ of the object centered around the object. This is incorporated as the quadratic geometric factor cost term $\|F_{geo}(\cdot)\|_{\Sigma}^2$ in Eq. 3. Expanding this,

$$\|F_{geo}(o_t, e_t)\|_{\Sigma_{geo}}^2 := \|\xi^{sdf}(o_t^{-1} p_t)\|_{\Sigma_{geo}}^2 \quad (11)$$

End-effector priors: Finally we also model uncertainty about end-effector locations as unary pose priors on the end-effector variables. These priors currently come from motion capture readings with added noise, but for a robot end-effector, these would instead come from the robot kinematics.

The pose priors are incorporated as the quadratic end-effector factor cost term $\|F_{eff}(\cdot)\|_{\Sigma}^2$ in Eq. 3. Expanding,

$$\|F_{eff}(e_t)\|_{\Sigma_{eff}}^2 := \|e_t \ominus \tilde{e}_t^{mc}\|_{\Sigma_{eff}}^2 \quad (12)$$

where, $\tilde{e}_t^{mc} = e_t^{mc} \oplus \mathcal{N}(0, \Sigma_{eff})$ are poses from the motion capture system with added Gaussian noise.

V. RESULTS AND EVALUATION

We evaluate our approach qualitatively and quantitatively on a number of real-world planar pushing trials where the pose of an object is unknown and must be estimated. We compare against a set of baselines on metrics like learning errors, estimation accuracy and runtime performance. Learnt tactile factors are trained using PyTorch [40]. These, along




Fig. 3: Experimental setup (a) shows overall planar pushing setup with motion capture (b) shows closeups of the Digit tactile sensor on an end-effector and the end-effector pushing a rectangular object.




Fig. 4: Contact within an image is detected if % of significantly ($>4\sigma$) different pixels exceed a threshold.

with engineered physics and geometric factors, are incorporated within GTSAM C++ library [41]. We use the iSAM2 [34] solver for efficient, incremental optimization.

A. Experimental setup

Fig. 3(a) shows the overall experimental setup for the pushing trials. We use an OptiTrack motion capture system to record ground truth object and end-effector poses for training and evaluating tactile observation models. Fig. 3(b) shows a closeup of the object, end-effector and the Digit tactile sensor mounted on an end-effector. The Digit sensor provides high-dimensional RGB images of the local deformation at the contact point [1].

B. Tactile factor learning

We now look at performance of the first stage of our approach, i.e. learning tactile observation models.

Contact detection: The first step is to detect contact in tactile images. We do so by first subtracting a mean no-contact image from the current image. If % of significantly ($>4\sigma$) different pixels exceeds a threshold, then contact is declared. We found this simple method works reliably on different pushing trials.

Keypoint features: Fig. 5 shows results for $K=1$ keypoint features learnt using the auto-encoder network described in Sec IV-A. The tactile image shows an elliptical contact patch where the curvature varies with local surface geometry of the object. The learnt keypoint features are able to track the patch center over time.

Tactile model performance: Table I compares mean-squared losses (Eq. 4) on the validation dataset for different choices of the transform prediction network described in Sec.




Fig. 5: Learnt keypoint features in blue tracking the local contact patch in tactile images for (top) rect, (middle) disc and (bottom) ellip objects.

TABLE I: Tactile Model Performance (Validation loss)

Model complexity	Datasets			
	Disc	Rect	Ellip	Combined
const	5.5e-3	15e-3	6.7e-3	9.6e-3
linear with eng-feat	5.1e-3	6.7e-3	4.2e-3	5.9e-3
linear with learnt-feat	1.0e-3	1.1e-3	1.7e-3	1.5e-3
nonlinear with eng-feat	3.0e-3	5.2e-3	3.3e-3	4.2e-3
nonlinear with learnt-feat	1.0e-3	1.7e-3	3.0e-3	2.4e-3

IV-A. We use a total dataset of ~ 3000 tactile images per object shape with a 50-50 train-test split. For training the transform prediction network, we create a dataset of pairwise images over a chosen time window $(t-k, t)$ with $k \in [10, 40]$. Table I, *const* is a zeroth order model that predicts mean relative transform of the training dataset. This is equivalent to using only contact detection information irrespective of contact patch locations in the images. *eng-feat* and *learnt-feat* represent engineered vs learnt keypoint features. For engineered features, we find a least-squares ellipse fit to detected contours in the image, as a generalization to points/line features used in related work [42]. *linear* and *nonlinear* refer to using linear or nonlinear ReLU activations in the fully-connected layers. We see that models using *learnt-feat* have lower losses over *eng-feat*. We also see *linear* models have lower losses than *const*. The *nonlinear* models don't show a significant improvement over the *linear* models. This is because there is a strong approximately linear relationship between the keypoint feature motions and relative rotation of the sensor about the contact point.

C. Factor graph optimization

We now look at the final task performance of estimating object poses using learnt tactile factors along with physics/geometric factors. For all runs, we use the same linear with learnt-feat tactile model network from Table I trained on the combined datasets. The tactile model network is conditioned on category label c of the object shape being used. c corresponds to different clusters for local curvatures — disc, ellip, rect-corners and rect-edges. We keep the covariance parameters same across all runs in the graph, i.e. $\Sigma_{tac} = \text{diag}(1, 1, 1e-5)$, $\Sigma_{qs} = \text{diag}(1e-3, 1e-3, 1e-3)$, $\Sigma_{geo} = \text{diag}(1e-2)$, $\Sigma_{eff} = \text{diag}(1e-5, 1e-5, 1e-5)$. First pose of the object is assumed to be known, and added as a unary prior to graph.

Qualitative tracking performance: Fig. 6(a) shows that physics and geometric factors alone cause object poses to drift over time. Fig. 6(b)-(d) shows tracking with different tactile models: *const*, *learnt*, and *oracle*. *const* predicts a constant relative transform value for each tactile factor. This approach is equivalent to using only a contact detector. While it improves over *no tactile*, it is unable to correct object rotations relative to contact point leading to drifting object pose estimates. This effect is most pronounced in the rect trajectories, where pushing along a corner cause large object rotations about the contact point. *learnt* is our proposed method using the transform prediction network, while *oracle* provides ground truth relative transforms. We see that the

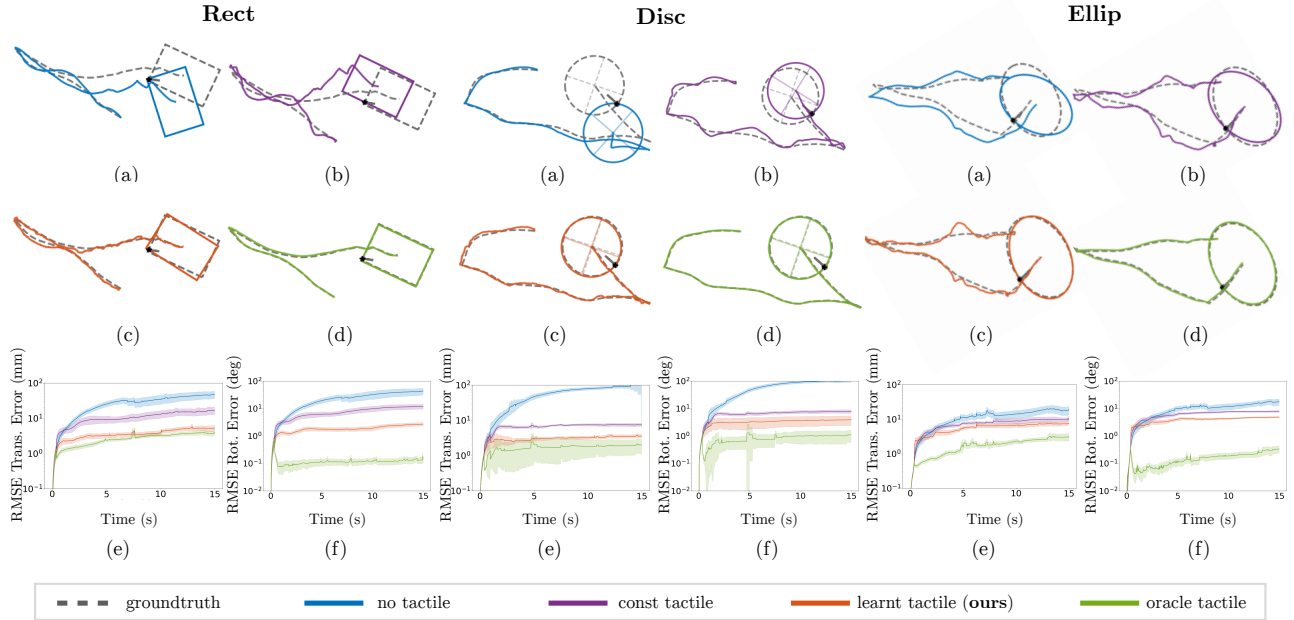


Fig. 6: Object pose tracking errors over time. (a)-(d) show qualitative tracking results, (e), (f) show quantitative mean-variance plots (semi-log scale) for 50 pushing trials each with rect, disc, ellip objects.

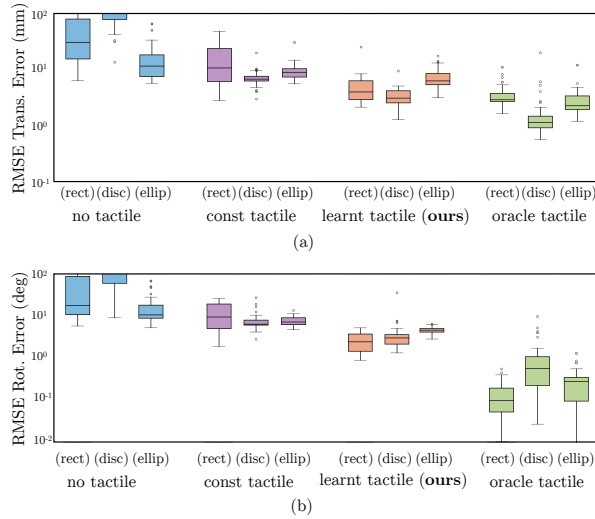


Fig. 7: Final estimation error box plots showing (a) translation and (b) rotation error values (semi-log scale).




Fig. 8: Graph optimizer runtime

learnt tactile model recovers object poses close to their true trajectory, and matches oracle performance closely.

Quantitative tracking performance: Fig. 6(e),(f) show mean-variance plots of RMSE rotational and translation object pose errors over time. Errors are computed over 50 pushing sequences each for rect, disc, ellip objects. The sequences have varying pushing trajectories making and

breaking contact typically 2-3 times. Fig. 7 additionally shows summary statistics of the RMSE translation and rotational errors at the final time step $T = 15s$. The learnt model performance is closest to oracle tactile performance recovering true object poses up to $\sim 10mm$ translational errors and $\sim 5^\circ$ rotational errors.

Runtime performance: Finally, Fig. 8 shows runtime per iteration of the graph optimizer. Runtime stays relatively constant with new measurements and priors added every step.

VI. DISCUSSION

We presented a factor graph based inference approach for estimating object poses from touch using vision-based tactile sensors. We proposed learning tactile observation models that directly integrate as factors within the graph. We demonstrated that our method is able to reliably track object poses for over 150 real-world planar pushing trials using tactile measurements alone.

As future work, in the tactile observation model, we would like to learn to model a distribution over relative poses instead of only the mean-squared error values. This should improve performance in cases where the relative pose uncertainty is asymmetric or varies significantly between contact episodes. We would also like to learn richer feature descriptors that describe contact patch geometry in addition to the patch centers. This should allow us to not have to condition on the object category label, and instead capture this information with the feature descriptor itself. Finally, to make these tactile factors work on more complex manipulation tasks, different physics priors will need to be incorporated.

ACKNOWLEDGEMENTS

We would like to thank the DIGIT team, in particular R. Calandra, P.W. Chou, and M. Lambeta for support with the sensor, software and insightful discussions. We would also like to thank Z. Dong, D. Gandhi, and anonymous reviewers for helpful feedback and suggestions.

REFERENCES

- [1] 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*, vol. 5, no. 3, pp. 3838–3845, 2020.
- [2] W. Yuan, S. Dong, and E. H. Adelson, “GelSight: High-resolution robot tactile sensors for estimating geometry and force,” *Sensors*, vol. 17, no. 12, p. 2762, 2017.
- [3] 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 *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2018.
- [4] A. Yamaguchi and C. G. Atkeson, “Combining finger vision and optical tactile sensing: Reducing and handling errors while cutting vegetables,” in *IEEE Intl. Conf. on Humanoid Robots (Humanoids)*, pp. 1045–1051, 2016.
- [5] F. Dellaert and M. Kaess, “Factor graphs for robot perception,” *Foundations and Trends in Robotics*, vol. 6, no. 1-2, pp. 1–139, 2017.
- [6] F. Dellaert, “Factor graphs: Exploiting structure in robotics,” *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 4, 2020.
- [7] K.-T. Yu and A. Rodriguez, “Realtime state estimation with tactile and visual sensing. application to planar manipulation,” in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pp. 7778–7785, 2018.
- [8] A. S. Lambert, M. Mukadam, B. Sundaralingam, N. Ratliff, B. Boots, and D. Fox, “Joint inference of kinematic and force trajectories with visuo-tactile sensing,” in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pp. 3165–3171, 2019.
- [9] C. Cadena, L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I. Reid, and J. J. Leonard, “Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age,” *IEEE Trans. Robotics*, vol. 32, no. 6, pp. 1309–1332, 2016.
- [10] D. M. Rosen, K. J. Doherty, A. T. Espinoza, and J. J. Leonard, “Advances in inference and representation for simultaneous localization and mapping,” *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 4, 2021.
- [11] R. Mur-Artal and J. D. Tardós, “ORB-SLAM2: An open-source SLAM system for monocular, stereo, and RGB-D cameras,” *IEEE Trans. Robotics*, vol. 33, no. 5, pp. 1255–1262, 2017.
- [12] J. Engel, T. Schöps, and D. Cremers, “LSD-SLAM: Large-scale direct monocular SLAM,” in *Eur. Conf. on Computer Vision (ECCV)*, pp. 834–849, Springer, 2014.
- [13] W. Dong, J. Park, Y. Yang, and M. Kaess, “GPU accelerated robust scene reconstruction,” in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pp. 7863–7870, IEEE, 2019.
- [14] P. V. Teixeira, M. Kaess, F. S. Hover, and J. J. Leonard, “Underwater inspection using sonar-based volumetric submaps,” in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pp. 4288–4295, 2016.
- [15] J. Czarnowski, T. Laidlow, R. Clark, and A. J. Davison, “DeepFactors: Real-time probabilistic dense monocular SLAM,” *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 721–728, 2020.
- [16] M. Bloesch, J. Czarnowski, R. Clark, S. Leutenegger, and A. J. Davison, “CodeSLAM—learning a compact, optimisable representation for dense visual slam,” in *Proc. IEEE Int. Conf. Computer Vision and Pattern Recognition*, pp. 2560–2568, 2018.
- [17] S. Wang, M. Lambeta, P-W. Chou, and R. Calandra, “Tacto: A fast, flexible and open-source simulator for high-resolution vision-based tactile sensors,” *arXiv preprint arXiv:2012.08456*, 2020.
- [18] A. Agarwal, T. Man, and W. Yuan, “Simulation of vision-based tactile sensors using physics based rendering,” *arXiv preprint arXiv:2012.13184*, 2020.
- [19] Y. She, S. Wang, S. Dong, N. Sunil, A. Rodriguez, and E. Adelson, “Cable manipulation with a tactile-reactive gripper,” in *Robotics: Science and Systems (RSS)*, 2020.
- [20] S. Dong and A. Rodriguez, “Tactile-based insertion for dense box-packing,” *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2019.
- [21] G. Izatt, G. Mirano, E. Adelson, and R. Tedrake, “Tracking objects with point clouds from vision and touch,” in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pp. 4000–4007, 2017.
- [22] B. Saund, S. Chen, and R. Simmons, “Touch based localization of parts for high precision manufacturing,” in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pp. 378 – 385, 2017.
- [23] F. M. Behbahani, G. Singla-Buxarraisi, and A. A. Faisal, “Haptic SLAM: An ideal observer model for bayesian inference of object shape and hand pose from contact dynamics,” in *International Conference on Human Haptic Sensing and Touch Enabled Computer Applications*, pp. 146–157, 2016.
- [24] M. C. Koval, N. S. Pollard, and S. S. Srinivasa, “Pose estimation for planar contact manipulation with manifold particle filters,” *Intl. J. of Robotics Research*, vol. 34, no. 7, pp. 922–945, 2015.
- [25] Z. Pezzementi, C. Reyda, and G. D. Hager, “Object mapping, recognition, and localization from tactile geometry,” in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pp. 5942–5948, 2011.
- [26] Z. Pezzementi, E. Plaku, C. Reyda, and G. D. Hager, “Tactile-object recognition from appearance information,” *IEEE Trans. Robotics*, vol. 27, no. 3, pp. 473–487, 2011.
- [27] B. Sundaralingam, A. S. Lambert, A. Handa, B. Boots, T. Hermans, S. Birchfield, N. Ratliff, and D. Fox, “Robust learning of tactile force estimation through robot interaction,” in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2019.
- [28] 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 *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pp. 3988–3993, 2014.
- [29] M. Bauza, O. Canal, and A. Rodriguez, “Tactile mapping and localization from high-resolution tactile imprints,” in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pp. 3811–3817, 2019.
- [30] M. Bauza, E. Valls, B. Lim, T. Sechopoulos, and A. Rodriguez, “Tactile object pose estimation from the first touch with geometric contact rendering,” *Conference on Robot Learning (CoRL)*, 2020.
- [31] K.-T. Yu and A. Rodriguez, “Realtime state estimation with tactile and visual sensing for inserting a suction-held object,” in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pp. 1628–1635, 2018.
- [32] 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,” *arXiv preprint arXiv:2011.07044*, 2020.
- [33] M. Kaess, A. Ranganathan, and F. Dellaert, “iSAM: Incremental smoothing and mapping,” *IEEE Trans. Robotics*, vol. 24, no. 6, pp. 1365–1378, 2008.
- [34] M. Kaess, H. Johannsson, R. Roberts, V. Ila, J. Leonard, and F. Dellaert, “iSAM2: Incremental smoothing and mapping using the Bayes tree,” *Intl. J. of Robotics Research*, vol. 31, pp. 216–235, Feb. 2012.
- [35] P. Sodhi, S. Choudhury, J. G. Mangelson, and M. Kaess, “ICS: Incremental constrained smoothing for state estimation,” in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2020.
- [36] K. M. Lynch, H. Maekawa, and K. Tanie, “Manipulation and active sensing by pushing using tactile feedback,” in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, vol. 1, 1992.
- [37] J. Zhou, J. A. Bagnell, and M. T. Mason, “A fast stochastic contact model for planar pushing and grasping: Theory and experimental validation,” in *Robotics: Science and Systems (RSS)*, 2017.
- [38] M. Mukadam, J. Dong, X. Yan, F. Dellaert, and B. Boots, “Continuous-time Gaussian process motion planning via probabilistic inference,” *The International Journal of Robotics Research*, vol. 37, no. 11, pp. 1319–1340, 2018.
- [39] N. Ratliff, M. Zucker, J. A. Bagnell, and S. Srinivasa, “CHOMP: Gradient optimization techniques for efficient motion planning,” in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pp. 489–494, 2009.
- [40] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, “PyTorch: An imperative style, high-performance deep learning library,” in *neurips*, pp. 8024–8035, 2019.
- [41] F. Dellaert, “Factor graphs and GTSAM: A hands-on introduction,” tech. rep., Georgia Institute of Technology, 2012.
- [42] F. R. Hogan, J. Ballester, S. Dong, and A. Rodriguez, “Tactile dexterity: Manipulation primitives with tactile feedback,” *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2020.

APPENDIX

Here we look at some additional results for object pose tracking during planar pushing tasks. Fig. 9 shows the three different objects being pushed by the Digit tactile sensor mounted on an end-effector. Objects are pushed over 50 varying sequences for each object. We perform following variations across the pushing sequences:

- Trajectories are varied as combinations of straight-line and curved pushes.
- Each sequence makes and breaks contact typically 2-3 times.
- Object is pushed at different locations, e.g. pushing at both edges and corners for the rect object.

Figs. 10, 11 show qualitative object tracking results for the rect object, and Fig. 12 for the other objects: disc and ellip over these varying sequences. In the *no tactile* case, use

of physics and geometric factors alone cause object poses to drift over time. The last three columns show object poses with different tactile models: *const*, *learnt*, *oracle*. *const* predicts a constant relative transform for each tactile factor. This approach is equivalent to using only a contact detector, i.e. using measurements that a standard non image-based contact sensor would give. While it improves over *no tactile*, it is unable to correct object rotations relative to contact point leading to drifting object pose estimates. This effect is most pronounced in the rect trajectories, where pushing along a corner cause large object rotations about the contact point. *learnt* is our proposed method using the transform prediction network, while *oracle* provides ground truth relative transforms. We see that the *learnt* tactile model recovers object poses close to their true trajectory, and matches oracle performance closely.




Fig. 9: Different objects with dimensions used in planar pushing trials.

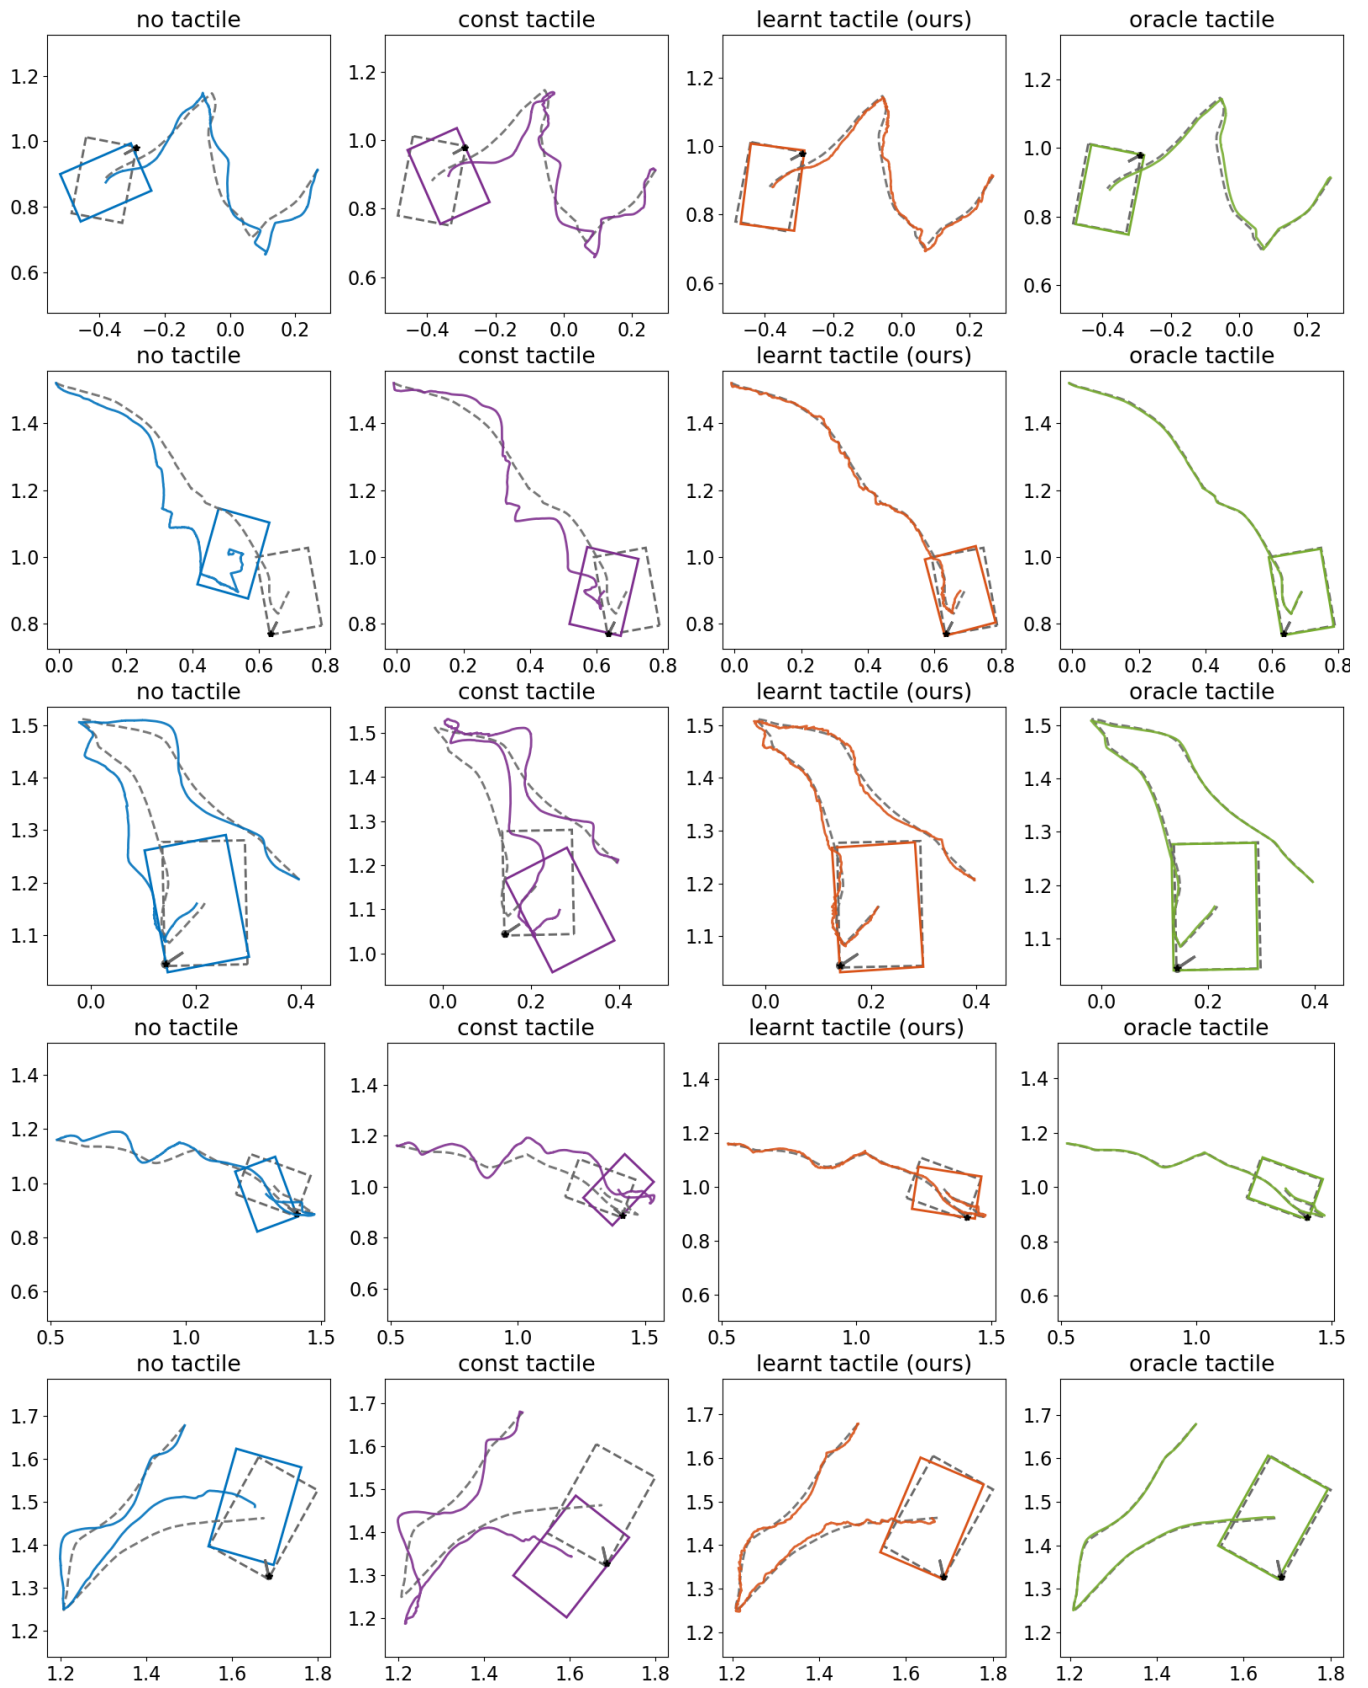


Fig. 10: Rect object pose tracking over time (15s).

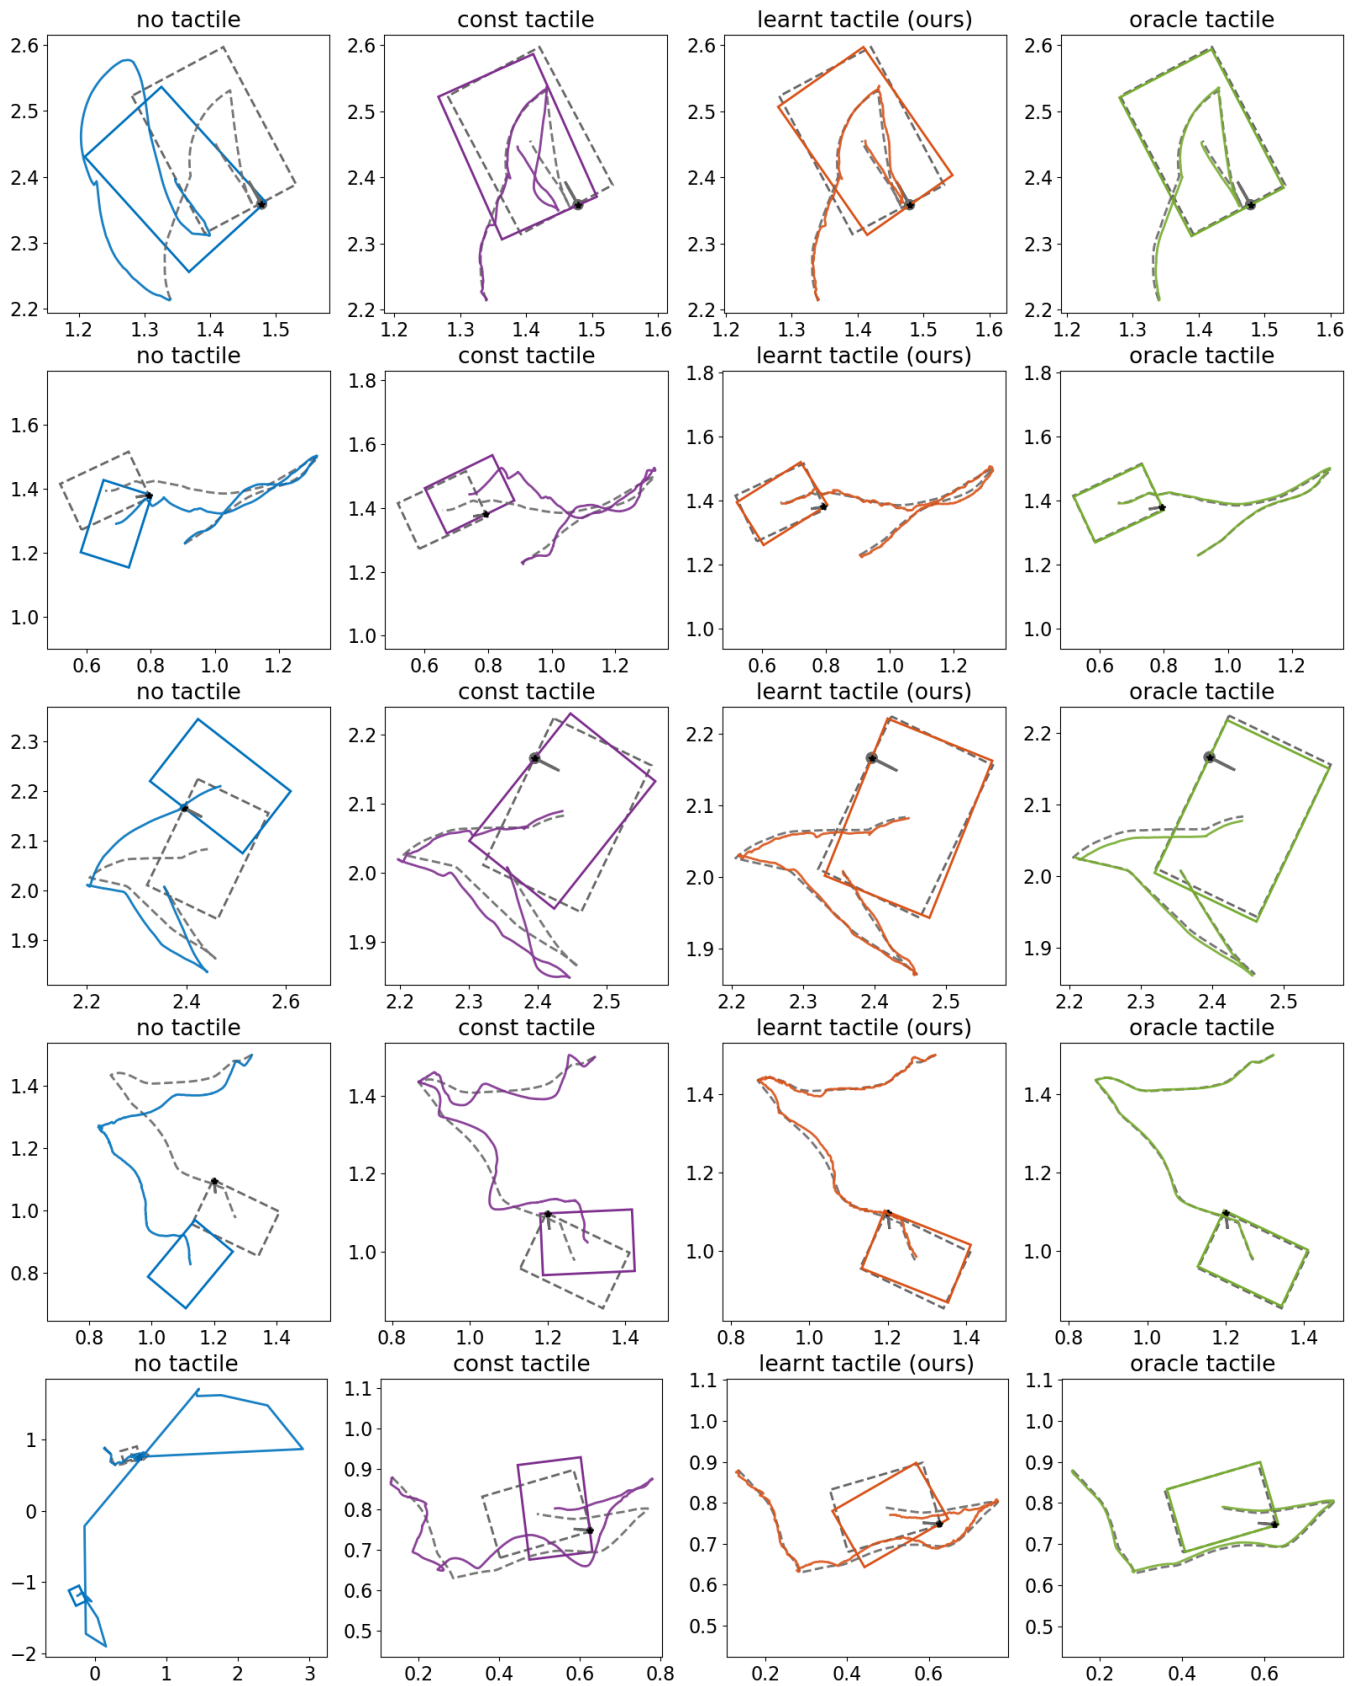


Fig. 11: Rect object pose tracking over time (15s).

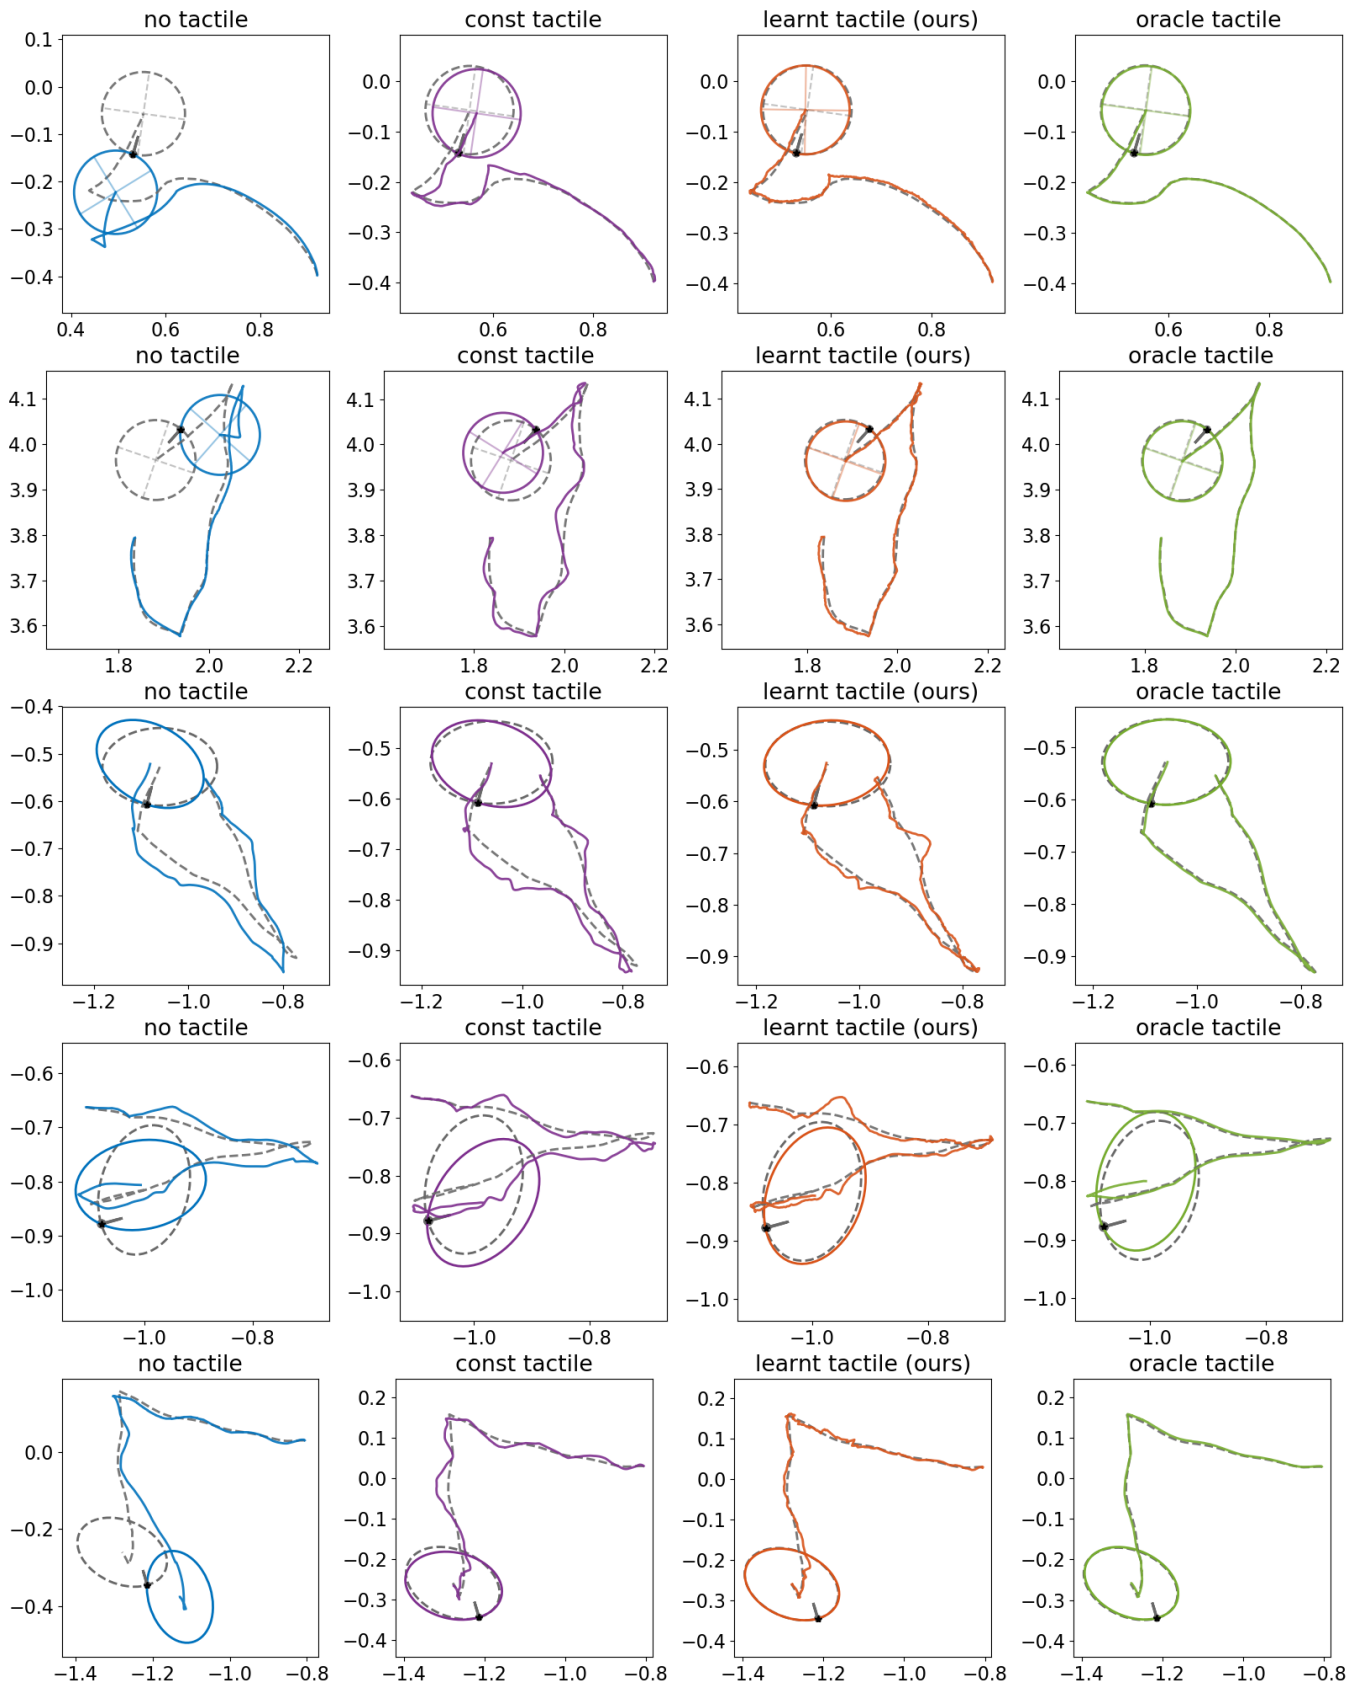


Fig. 12: Disc, Ellip object pose tracking over time (15s).