An Active Extrinsic Contact Sensing for Generalizable Insertion Strategy

Sangwoon Kim and Alberto Rodriguez
{sangwoon, albertor}@mit.edu

Abstract— We propose a framework that actively estimates extrinsic contact with tactile feedback then uses it as a low dimensional input to the peg-in-hole insertion policy. An iSAM graph estimator and an active tactile feedback controller work collaboratively to get an accurate location of the extrinsic contact line. The reduction in the input dimension, from a high-resolution tactile image to a single line, enables the insertion policy to train in simulation. By randomizing the object-hole shape while training the policy in the simulation, the policy can generalize to various object-hole shapes. Lastly, we demonstrate the active extrinsic contact line estimation and test the insertion policy in the real experiment. We show that the proposed method inserts various-shaped objects with a higher success rate and fewer insertion attempts than previous works that use an end-to-end approach.

I. INTRODUCTION

The ability to sense and utilize tactile interaction between fingers and grasped object is a key foundation that enables dexterous manipulation skills [1]. The tactile sensing can be used as a feedback signal to regulate desired tactile configurations. One major problem in tactile-feedback control is localizing and controlling an external contact: a contact between the grasped object and its environment [2], [3]. For example, consider pivoting an unknown object while avoiding slip at the contact between the object and its environment. It requires accurate localization of the contact and an ability to regulate the external contact mode using indirect observation through tactile sensing.

A task where the external contact matters is the peg-in-hole insertion task. A misalignment between the peg and the hole results in contact during the insertion attempt. The contact triggers a tactile signal, which can be used to localize the contact or plan the next insertion attempt [4], [5], [6]. Some of the key challenges in this task are as follows:

- The tactile signal is a partial observation of the state; many different contact configurations can cause the same tactile signal [7].
- The difficulty of modeling the contact mechanics makes it hard for the methods to generalize to unknown objects.
- For learning-based methods, it should have a reasonably high data efficiency.

In this work, we tackle both the contact localization problem and the peg-in-hole insertion task in a combined framework. For the contact localization, we use incremental smoothing and mapping (iSAM) factor graph alongside an

All authors are with Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

active tactile-feedback controller. The factor graph and the feedback controller work collaboratively to estimate the extrinsic contact line. The contact line estimation is then used as an input to the insertion policy, as opposed to the end-to-end approach [6] where the policy takes input as a raw tactile sensor image. Since the input to the policy is a low dimensional representation, the policy training can be done in a simple simulation with randomly shaped objects. Lastly, we demonstrate the contact line estimation and the insertion policy in a real experiment.

II. APPROACH OVERVIEW

The task we solve is a typical peg-in-hole problem. We make several assumptions to implement our framework:

- Object bottom surfaces and hole top surfaces are flat.
- Objects and holes have un-chamfered corners.
- Objects and holes have convex shapes.

The approach is composed of two parts: the active extrinsic contact sensing and the insertion policy. In the active extrinsic contact sensing part, the iSAM graph works collaboratively with the active tactile-feedback controller to estimate the extrinsic contact line. Then, the reinforcement learning (RL) policy takes the estimated contact line as input and computes the action for the insertion.

A. Tactile Module and iSAM Graph

We use GelSlim [8], a vision-based tactile sensor, to capture the deformation image on the robot finger during the insertion (Fig. 1a). The image is passed to the tactile module with a convolutional neural network (CNN) architecture. The tactile module is trained with supervised learning to estimate the relative SE(3) pose between the gripper and the object. The gripper-object relative pose and the robot proprioception data are used in the iSAM graph to estimate the extrinsic contact line.

Fig. 2 shows the iSAM graph. g_i, o_i, and c_i are the variables for gripper pose, object pose, and estimated contact line, respectively. The iSAM enables to add new measurements incrementally and update the estimation in real-time. Each small circle with different colors represents different types of factors:

- **Gripper Prior** takes the robot proprioception measurement as a prior on the gripper pose.
- **Object Prior** imposes the prior belief on object pose.
- **Gelslim Deformation** takes the gripper-object relative pose from the tactile module and impose the relation between $g_i, g_{i+1}, o_i, o_{i+1}$.
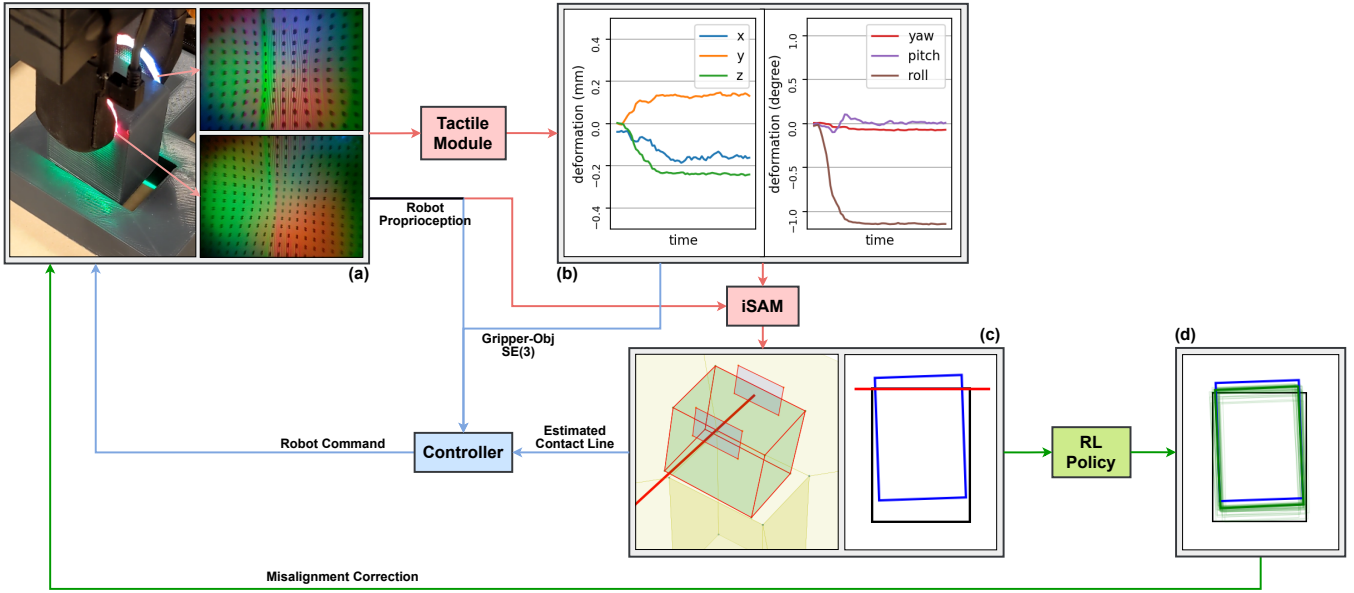
Fig. 1. Approach Overview. iSAM factor graph (red) and active tactile-feedback controller (blue) runs collaboratively to estimate extrinsic contact line. RL policy (green) takes the estimated extrinsic contact line as input and computes the next action. (a) Scene of insertion attempt and the tactile image captured by GelSlim finger. (b) Gripper-Object relative pose computed by the tactile module. (c) 3D and top view of the extrinsic contact line estimation. The bold red line is the current estimate of the extrinsic contact line. (d) Actor output. The blue rectangle is the current pose and the green rectangles are the candidate poses for the next insertion attempt.
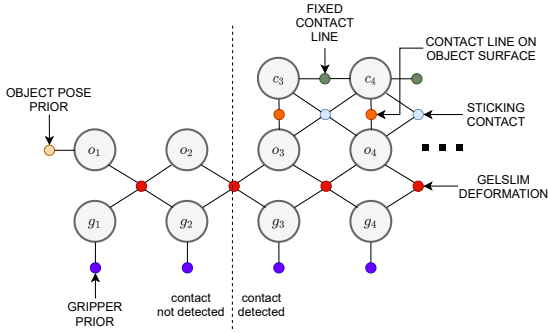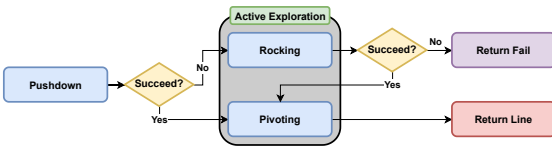

Fig. 2. iSAM graph

Fig. 3. Active Tactile-Feedback Controller

- **Contact Line on Object Surface** constrains the estimated contact line to be on the object's bottom surface.
- **Fixed Contact Line** constrains the estimated contact line to be stationary over time.
- **Sticking Contact** constrains the relative translation between the object and the contact line to be fixed, which corresponds to the sticking contact.

B. Active Tactile-Feedback Controller

Fig. 3 shows the schematics of the active tactile-feedback controller. Its logic comprises the passive push-down and the active exploration with rocking and pivoting. During push-down phase, it executes a proportional control on gripper-object relative deformation. The vertical deformation is set to a non-zero value to ensure the object contacts the environment with sufficient normal force. Other components are set to zero.

If the object does not tilt enough so the iSAM graph fails to estimate a contact line with enough confidence, it enters the rocking phase. In the rocking phase, the robot is commanded to follow a cone-like trajectory while maintaining deformation smaller than a set threshold. If the iSAM graph fails to find a contact line even after the rocking, the controller stops and the estimator returns a failure output.

If the iSAM graph succeeds in estimating an extrinsic contact line either in push-down or rocking, the controller enters the pivoting phase to help the iSAM graph to gain more confidence. The controller tries to pivot the object around the extrinsic contact line. It is assumed that the tactile deformation will remain constant during the pivoting if the object pivots without slipping and also maintaining a constant contact force. The controller commands the robot to rotate around the current contact line estimate while the components other than the rotation axis are controlled by a proportional control. The proportional gains are exerted to the direction where it tries to keep the tactile deformation constant.

A key idea in the above framework is the synergistic interaction between the controller and the iSAM estimator. A better contact line estimation from the iSAM graph helps the controller to pivot the object with less slipping. On the other direction, a better pivoting with a consistent pivoting axis helps the iSAM graph to get more accurate contact line

TABLE I

Success Rate and average number of insertion attempts for various object-hole pairs.

	Object (blue) & Hole (black) Shape	◯	◯	⬡	⬡	◯	◯	▭
Active iSAM-RL	Success Rate (%)	100	97	99	100	95	95	97
	Avg. Attempt #	1.94	1.98	2.40	2.94	2.61	3.04	4.13
RL-end2end (Dong [6])	Success Rate (%)	97	-	97	-	98	-	90
	Avg. Attempt #	2.96	-	3.83	-	2.34	-	5.42
SL-end2end (Dong [5])	Success Rate (%)	85	-	70	-	94	-	15
	Avg. Attempt #	3.04	-	3.34	-	2.60	-	3.83

estimate. The interaction acts as a stable attractor that enables to find an accurate contact line estimate with the simple proportional controller.

C. RL Policy

The RL policy takes the estimated extrinsic contact line from the iSAM as the input and computes the SE(2) pose correction for the next insertion attempt. Since the input to the policy is a low dimensional representation (a single line), it is easy to simulate the policy; given two random shape polygons, each representing the object and the hole, and an SE(2) misalignment between the two, we can find a contact line that the object can pivot around.

The policy training is done purely in simulation. We vary the object and hole shape on every training episode so the policy can be generalized to unseen object and hole shapes. The random shape is generated by randomly scattering points and drawing a convex hull around the points.

We use twin delayed deep deterministic policy gradient (TD3) [9] with recurrent neural networks to train the policy. The use of recurrent networks enables the policy to deal with partial observability by considering the previous history of observations and actions.

III. Result

Table I shows the success rate and the average number of insertion attempts for various object-hole pairs. Each object-hole pair has 2.25 mm of clearance. Results from previous works [6], [5] are also shown for comparison. Note that [6], [5] used the demonstrated object-hole pair as training cases, while our method used it as testing cases. The proposed method performed similarly or better for each object-hole pair. Especially for the rectangle object-hole, the proposed method outperformed [6] by a 7% margin in success rate with 1.3 less number of average attempts.

IV. CONCLUSIONS

In this short paper, we proposed a framework that estimates extrinsic contact then uses it as a low dimensional representation of the contact configuration. In the estimation part, an iSAM graph and an active tactile feedback controller worked collaboratively to get an accurate estimate of the contact line. In the policy part, by reducing the input dimension from a high-resolution image to a single line, we enabled the

RL policy to train in simulation. Also, by randomizing the object-hole shape during training, the RL policy was able to generalize to various object-hole shapes. As the future direction of this framework, we plan to extend this approach to more general manipulation skills.

REFERENCES

[1] Howe, Robert D. "Tactile sensing and control of robotic manipulation." *Advanced Robotics* 8.3 (1993): 245-261.

[2] Rodriguez, Alberto. "The unstable queen: Uncertainty, mechanics, and tactile feedback." *Science Robotics* 6.54 (2021).

[3] Ma, Daolin, Siyuan Dong, and Alberto Rodriguez. "Extrinsic Contact Sensing with Relative-Motion Tracking from Distributed Tactile Measurements." *2021 IEEE/RSJ International Conference on Robotics and Automation (ICRA)*. IEEE, 2021.

[4] Yu, Kuan-Ting, and Alberto Rodriguez. "Realtime state estimation with tactile and visual sensing for inserting a suction-held object." *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018.

[5] Dong, Siyuan, and Alberto Rodriguez. "Tactile-based insertion for dense box-packing." *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2019.

[6] Dong, Siyuan, et al. "Tactile-RL for Insertion: Generalization to Objects of Unknown Geometry." *2021 IEEE/RSJ International Conference on Robotics and Automation (ICRA)*. IEEE, 2021.

[7] Newman, Wyatt S., Yonghong Zhao, and Y-H. Pao. "Interpretation of force and moment signals for compliant peg-in-hole assembly." *Proceedings 2001 ICRA. IEEE International Conference on Robotics and Automation (Cat. No. 01CH37164)*. Vol. 1. IEEE, 2001.

[8] Taylor, Ian, Siyuan Dong, and Alberto Rodriguez. "GelSlim3. 0: High-Resolution Measurement of Shape, Force and Slip in a Compact Tactile-Sensing Finger." *arXiv preprint arXiv:2103.12269* (2021).

[9] Fujimoto, Scott, Herke Hoof, and David Meger. "Addressing function approximation error in actor-critic methods." *International Conference on Machine Learning. PMLR*, 2018.