

## MIT Open Access Articles

*Tactile Dexterity: Manipulation Primitives with Tactile Feedback*

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

**Citation:** Hogan, Francois R, Ballester, Jose, Dong, Siyuan and Rodriguez, Alberto. 2020. "Tactile Dexterity: Manipulation Primitives with Tactile Feedback." Proceedings - IEEE International Conference on Robotics and Automation.

**Published Version:** 10.1109/ICRA40945.2020.9196976

**Publisher:** Institute of Electrical and Electronics Engineers (IEEE)

**Permanent Link:** <https://hdl.handle.net/1721.1/139614>


**Version:** Original manuscript: author's manuscript prior to formal peer review

**Terms of use:** <http://creativecommons.org/licenses/by-nc-sa/4.0/>



# Tactile Dexterity: Manipulation Primitives with Tactile Feedback

Francois R. Hogan, Jose Ballester, Siyuan Dong, and Alberto Rodriguez  
Massachusetts Institute of Technology — Department of Mechanical Engineering  
[https://mcube.mit.edu/research/tactile\\_dexterity.html](https://mcube.mit.edu/research/tactile_dexterity.html)



**Abstract**—This paper develops closed-loop tactile controllers for dexterous robotic manipulation with a dual-palm robotic system. *Tactile dexterity* is an approach to dexterous manipulation that plans for robot/object interactions that render interpretable tactile information for control. We divide the role of tactile control into two goals: 1) control the contact state between the end-effector and the object (contact/no-contact, stick/slip) by regulating the stability of planned contact configurations and monitoring undesired slip events; and 2) control the object state by tactile-based tracking and iterative replanning of the object and robot trajectories.

Key to this formulation is the decomposition of manipulation plans into sequences of manipulation primitives with simple mechanics and efficient planners. We consider the scenario of manipulating an object from an initial pose to a target pose on a flat surface while correcting for external perturbations and uncertainty in the initial pose of the object. We experimentally validate the approach with an ABB YuMi dual-arm robot and demonstrate the ability of the tactile controller to react to external perturbations.

## I. INTRODUCTION

This paper studies the use of tactile sensing for dexterous robotic manipulation, or *tactile dexterity*. Despite the evidence that humans heavily depend on the sense of touch to manipulate objects [1], robots still rely mostly on visual feedback. The vision-based approach has been effective for tasks such as pick-and-place [2] but it presents fundamental limitations to accomplish dexterous manipulation tasks that depend on more accurate and controlled contact interactions, such as object reorientation, object insertion, or almost any kind of object use.

The mechanics of object manipulation are driven primarily by the relative motions and forces at the frictional interfaces between object, end-effectors and environment. As a result, we believe that tactile sensors, with their ability to localize

contact geometry, detect contact motion, and infer contact forces, should be at the center of our manipulation plans.

Our approach to *tactile dexterity*, described in Sec. III, is based on planning for robot/object interactions that render interpretable tactile information for control. At its core, this is an approach to robotic manipulation that puts tactile feedback at the center to bypass some of its common caveats. Tactile information is by nature local and, in general, not sufficient to fully describe the state of a manipulated system [3]. Furthermore, the design of feedback policies for systems undergoing physical contact is challenging, even when full state feedback is available [4], [5].

Given these challenges, the key question becomes: How do we structure manipulation planning so that tactile feedback is relevant and convenient? Our approach to tactile dexterity:

- **Targets contact interactions** that produce tactile imprints rich in geometric features (for estimation); and that define dynamic systems with simple mechanics and efficient closed-loop policies (for control). In Sec. IV we describe four such manipulation primitives: *grasp*, *push*, *pivot*, and *pull*.
- **Partitions the role of tactile control** into controlling the *contact state* between end-effector and object; and controlling the *object state* in its environment. We describe a formulation for both in Sec. V.

We present a first study of this approach for the scenario of a dual-arm robot equipped with high-resolution tactile-sensing palms (based on GelSlim [6]) as end-effectors, and tasked with manipulating an object on a table-top from an arbitrary initial pose to an arbitrary target pose. An offline graph-search task planner, described in Sec. VI, sequences manipulation primitives, which are executed in a closed-loop fashion by the robot. The focus of the experiments described in Sec. VII is to evaluate the robustness of the system to external perturbations and to uncertainty in the initial pose of the object.

This work was supported by the Amazon Research Awards (ARA) and the Toyota Research Institute (TRI).

Corresponding author: Francois Hogan <fhogan@mit.edu>

Project website: [https://mcube.mit.edu/research/tactile\\_dexterity.html](https://mcube.mit.edu/research/tactile_dexterity.html)

Project video: <https://youtu.be/f59FoS-hV7c>

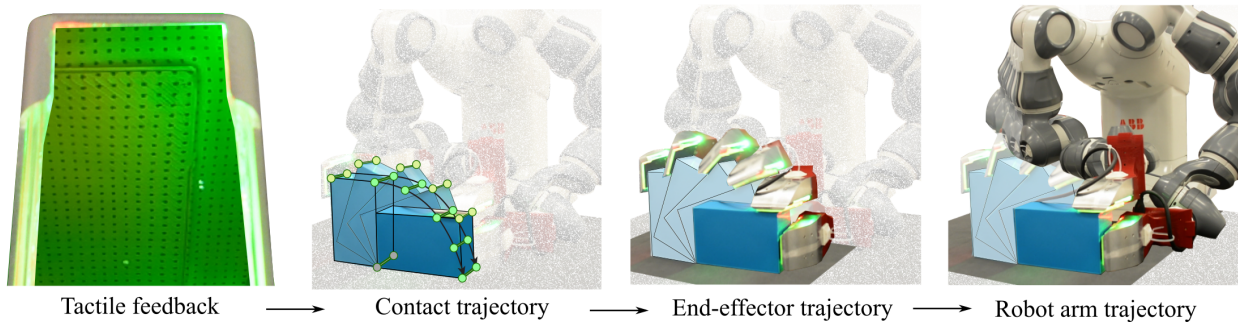


Fig. 1. **Planning to pivot a box.** Pivoting a box involves a hierarchy of planning problems. Tactile sensors allow us to observe and control contact points. The trajectory of contact points is determined by the end-effector pose trajectory, which in turn is generated by the joint trajectory of the robot arm. We propose an approach to dexterous manipulation that structures the planning problem in a way that when unexpected behavior of the contact points is observed with the tactile sensors, the robot can replan the motion of the contact points, palms, and arms in real-time.

## II. RELATED WORK

This section outlines previous relevant research in exploiting tactile sensing for manipulation control.

**Slip Control** Veiga [7] learns a predictor of slippage of an object in a grasp. This model is used in [8] to maintain stable multi-fingered grasps under external perturbations, where each finger acts independently to enforce sticking contact. Dong [9] develops an incipient slip detection algorithm with a GelSlim vision-based tactile sensor [6], and uses it to design a closed-loop tactile controller that maintains stable grasps in prehensile experiments such as a cap into a bottle with a parallel-jaw gripper. Li [10] develops an impedance controller that adapts the grasp on an object when it is predicted to be unstable. These studies focus on controlling stick/slip interactions but do not explicitly control the trajectory of the manipulated object.

**Pose Control** Tian [11] trains a deep convolutional neural network to predict the motion of a ball rolled on the ground using a tactile finger directly in tactile space. This model is controlled using sample-based MPC. A drawback of this approach is its need for large quantities of real-world data, which would be challenging to collect for the rich palm/objects interactions considered in this paper.

Li [12] shows that localized object features can be exploited with tactile sensing to recover an accurate estimate of its pose. This strategy has been shown to be effective at performing challenging manipulation tasks such as part insertion with small tolerances. Izatt [13] fuses tactile and visual perception by interpreting tactile imprints as local 3D pointclouds within a Kalman filter framework for object pose estimation. More recently, Bauza [14] develops a tactile based pose estimation algorithm that exploits a high resolution tactile map of the object to localize tactile imprints.

**Dexterous Manipulation** Erdman [15] first studied the use of robotic palms as a way to achieve rich manipulation skills by exploiting the mobility of robotic manipulators. Researchers have investigated robotic palms for scooping [16], [17], tilting [18], grasping [19], [20], and collaborative manipulation [21], [22]. This paper draws inspiration from [23]–[25] by structuring complex manipulation behavior as

a combination of simpler manipulation primitives.

## III. APPROACH

This paper develops closed-loop tactile controllers that enable robust manipulation behavior, where the robot can react to external object perturbations. We propose an approach that divides the role of tactile control into:

- **Contact state control.** Enforce a desired contact formation (contact/no-contact, stick/slip) between the end effector and the object.
- **Object state control.** Control the motion of the object using a contact based state estimator for tactile tracking and iterative replanning of the trajectories of the object, end-effector and robot.

This paper considers the scenario of manipulating an object from an initial pose to a target pose on a table top. We consider a robotic platform that is 1) dexterous, where both robotic palms are controlled independently, and 2) tactile sensorized, where each palm is equipped with high-resolution tactile sensors (based on GelSlim technology [6]). We formulate the manipulation problem as a sequencing of manipulation primitives, where each primitive is designed to have a prescribed contact interaction between the robotic palms, the object, and the environment. Fig. 2 shows the four manipulation primitive considered in this paper: *grasp*, *push*, *pivot*, and *pull*. This structuring of the manipulation problem into simpler behaviors gives us the freedom to design interactions for which we understand the mechanics, are able to interpret the tactile information, and can develop effective planning algorithms.

## IV. MECHANICS OF MANIPULATION PRIMITIVES

This section describes the mechanics of the four manipulation primitives in Fig. 2. To model the palm-object-environment interaction we assume:

- Known geometry of object, robot, and environment.
- Known coefficients of friction.
- Rigid-body interaction.
- Coulomb’s friction law.
- Quasi-static interaction.
- Surface contacts with uniform pressure distribution.




Fig. 2. **Manipulation primitives.** a) *grasp*: the palms align as a parallel-jaw gripper to grasp an object. b) *push*: one palm contacts the object laterally to manipulate it within the plane. c) *pivot*: the object is rotated about an edge on the table by both palms. d) *pull*: one palm presses down on the top of the object to slide it within the support plane.

These simplifications will help in designing fast trajectory planning algorithms in Sec. VI. The unmodeled aspects of the interactions (non-uniform pressure distributions, inertial forces, deformation of contact, etc.) are addressed by designing closed-loop tactile controllers in Section V that react to undesired contact events such as slip.

Assuming quasi-static interactions, force equilibrium dictates that contact forces on the object (applied by the end-effector or environment) are balanced by:

$$\sum_{i=1}^C \mathbf{G}_i(\mathbf{q})^\top \mathbf{w}_i = \mathbf{w}_{ext}, \quad (1)$$

where  $\mathbf{q} = [\mathbf{p}_o^\top \quad \mathbf{p}_{p,l}^\top \quad \mathbf{p}_{p,r}^\top]^\top$  is the concatenation of the object pose and the left/right palm poses,  $\mathbf{w}_i = [\mathbf{c}_i^\top \quad \boldsymbol{\tau}_i^\top]^\top$  is the applied wrench on the object by the  $i^{\text{th}}$  contact in the contact frame,  $\mathbf{w}_{ext}$  is the external force applied by gravity in the world frame,  $\mathbf{G}_i$  is a grasp matrix transforming the coordinates of a contact wrench from the contact frame to the world frame [26], and  $C$  is the number of contacts.

Contact forces are constrained within the friction cone in accordance to Coulomb’s frictional law. Denoting the normal and tangential components of the contact force as  $\mathbf{c}_i = [f_{n,i} \quad \mathbf{f}_{t,i}^\top]^\top$ , we express Coulomb’s frictional law as:

$$f_{n,i} \geq 0 \quad (2)$$

$$|\mathbf{f}_{t,i}| \leq \mu |f_{n,i}|. \quad (3)$$

In the case of point contact interactions (*grasp* and *pivot*), the contact is unable to sustain frictional moments, implying  $\boldsymbol{\tau}_i = \mathbf{0}$ . For contacts modeled using surface contacts (*push* and *pull*), the surface is able to resist a certain amount of frictional moment. We model surface contacts using the limit surface [27], which describes the set of forces and moments that can be transmitted through the contact. In practice, we make use of the ellipsoidal approximation to the limit surface introduced in [28] that gives a simple analytical representation of the limit surface. A detailed description of the mechanics of similar manipulation primitives can be found in [5] for non-prehensile pushing, [29] for prehensile pushing, and [30] for pivoting.

## V. TACTILE CONTROL

In this section, we describe tactile-based controllers that give robust behavior to the manipulation primitives. We use GelSlim [6], an optical-based tactile sensor that renders high-resolution images of the contact surface geometry and strain field, as shown in Fig. 5, and in the video <https://youtu.be/f59FoS-hV7c>.

A primary goal of the tactile controller is to enforce a desired contact state between the palms and the object. For example, during the pivot maneuver in Fig. 2(c), we want both palms to maintain sticking contacts with the corners of the object. By monitoring incipient slippage, we design a controller that engages when necessary to regulate the applied forces on the object to prevent further slip. We refer to this as *Contact State Control* (Sec. V-A). This is illustrated in Fig. 3, where the palms rotate and apply additional normal force on the object in reaction to a slip event.

An important consequence of undesired slippage is that the position of the object has deviated from the desired trajectory. To address this, we design a tactile-based state estimator, illustrated in Fig. 5, that tracks in real-time local features on the object and uses the updated object pose estimate to continuously replan object, palm, and robot arm trajectories for the manipulation primitive. We refer to this as *Object State Control* (Sec. V-B)

### A. Contact State Control

Each primitive assumes a particular contact formation between the palms and the object. This assumption is likely to be broken as unmodeled perturbations are applied on the system and cause undesired slippage. We design a contact state controller that acts to enforce the planned contact modes by reacting to a binary incipient slip signal  $s_i \in \{0, 1\}$  at contact  $i$ , based on [9].

Coulomb friction states that slippage occurs when the contact force lies on the boundary of the friction cone, as shown in Fig. 4. Given an undesired slippage signal, we find local palm configuration and force adjustments that result in a more stable sticking solution (i.e. contact forces nearer to the center of the friction cone). The stability margin  $\phi$  in Fig. 4 is defined as the shortest distance from the contact force to the friction cone boundary and evaluates the risk of slipping of a particular contact. The goal of the controller is to find robot




Fig. 3. **Contact state control.** How should the robot react to undesirable slippage? We design a model based tactile controller that determine locally optimal robot adjustments to recover from contact state deviations.

adjustments that either maximize (for sticking) or minimize (for sliding) the stability margin of a contact to enforce a desired contact mode. In this paper we develop primitives that restrict to sticking contacts between the gripper and the object.

Given the slippage signal  $s_i \in \{0, 1\}$  at each contact and the current robot pose configuration  $\mathbf{q}_p = [\mathbf{p}_{p,l}^T \mathbf{p}_{p,r}^T]^T$ , we search for a robot palm pose adjustment  $\Delta\mathbf{q}_p$  and contact force  $\mathbf{w}_i$  that maximize the stability margin of a particular contact  $i$  (indicated by the weight  $\beta_i$ ) by solving:

$$\begin{aligned} & \max_{\Delta\mathbf{q}_p, \mathbf{w}_i} \beta_i \phi_i \\ \text{s.t.} \quad & \sum_{i=1}^C \mathbf{G}_i(\mathbf{q}_p + \Delta\mathbf{q}_p)^T \mathbf{w}_i = \mathbf{w}_{ext} \\ & f_{n,i} \geq 0 \\ & \|\mathbf{f}_{t,i}\|_2 \leq \mu |f_{n,i}| \end{aligned} \quad (4)$$

with  $\mathbf{G}_i$ ,  $\mathbf{q}_p$ ,  $\mathbf{w}_{ext}$  as defined in Section IV and where  $\phi$  is the stability margin shown in Fig. 4. The notation  $\mathbf{q}_p + \Delta\mathbf{q}_p$  is abused here for simplicity, where we ensure that  $\mathbf{q}_p + \Delta\mathbf{q}_p$  belongs to  $SE(3)$  and does not violate the kinematics (i.e. maintaining contact and non-penetration) of the contact formation. For example, in Fig. 2(c), the palms can press harder and/or pivot about the contact edge. The parameter  $\beta_i$  is a weight used to biased the effort of the regulating behavior to contact  $i$  based on the slip signal  $s_i$ .

The optimization program in (4) is non-convex because  $\phi$  is nonlinear and because the constraint associated with static equilibrium is bilinear. In [31], the surrogate stability margin  $\alpha$  illustrated in Fig. 4 is proposed as a convex approximation to  $\phi$ . We relax then the optimization program using the surrogate margin  $\alpha$  and linearizing the static equilibrium equation (1) about the current robot configuration  $\mathbf{q}_p^*$  and




Fig. 4. The stability margin  $\phi$  evaluates how close a contact is to the slipping boundary. The goal of the contact state controller is to maximize/minimize the stability margin to encourage/discourage slippage.

contact forces  $f_{n,i}^*$ ,  $\mathbf{f}_{t,i}^*$  for computational efficiency. The linearized contact state controller becomes:

$$\begin{aligned} & \max_{\Delta\mathbf{q}_p, \Delta\mathbf{w}_i, \alpha_i} \beta_i \alpha_i \\ \text{s.t.} \quad & \sum_{i=1}^C \left( \frac{\partial (G_i \mathbf{w}_i)}{\partial \mathbf{q}_p} \Big|_* \Delta\mathbf{q}_p + \frac{\partial (G_i \mathbf{w}_i)}{\partial \mathbf{w}_i} \Big|_* \Delta\mathbf{w}_i \right) = \mathbf{w}_{ext} \\ & f_{n,i} \geq 0 \\ & \alpha_i \geq 0 \\ & \|\mathbf{f}_{t,i}\|_2 \leq \mu f_{n,i} \\ & \|\mathbf{f}_{t,i}\|_2 \leq \mu |f_{n,i} - \alpha_i|, \end{aligned} \quad (5)$$

where the symbol  $(\cdot)^*$  is used to evaluate a term at the nominal configuration. Note that the surrogate margin  $\alpha$  is not invariant to scale, i.e., we could artificially increase the margin simply by scaling up the magnitude of all the forces. To regulate this behavior, in practice, we add bound constraints to the allowable  $\Delta\mathbf{q}$  and  $\Delta\mathbf{w}$ . The optimization program in (5) takes the form of a convex quadratic program under a polyhedral approximation to the friction cone [32].

## B. Object State Control

The contact state controller in Sec. V-A regulates the applied forces on the object to enforce a desired contact mode. This tactile controller reacts to fight against external perturbations but does not have the ability to change the planned trajectory of the object in response to a perturbation. To address this, we design an object state controller running in parallel tasked with replanning object/palm trajectories to drive the object to its target location.

In this paper we track two types of features: points (corners of the object) and lines (edges of the object). We formulate the tactile object state estimator as an optimization program that updates the pose  $\mathbf{p}_o$  of the object to satisfy the geometric constraints associated with the tactile features, as shown in Fig. 5. We quantify the error between the previous and the updated pose estimate using the distance  $d_{TS}(\mathbf{p}_o, \mathbf{p}_o^*)$ , where  $d_{TS}$  is defined as the weighted sum of the Euclidean metric in  $\mathbb{R}^3$  and the great circle angle metric in  $SO(3)$  for the respective components [33]. We enforce that detected lines are collinear with their associated edge on the object mesh and that the sensed points are coincident with the object corner. In addition to the detected geometric constraint, we constrain the estimated object pose to satisfy the geometric




Fig. 5. **Tactile object localization.** By localizing descriptive object geometric features (lines, points), we can track the pose of the object.

constraints consistent with the current manipulation primitive. For example, for the `pull` primitive, we enforce that the bottom surface of the object is in contact and aligned with the table top.

The estimated object pose is used to update the nominal robot palm pose trajectory which allows the robot to adapt to local object perturbations. This is further described in Sec. VI-B.

## VI. PLANNING

We formulate the manipulation planning problem as a sequencing of manipulation primitives, as described in Sec. IV. This section briefly describes an approach for how to 1) sequence the manipulation primitives to achieve the task and 2) plan robot trajectories within a manipulation primitive a to achieve a desired object transformation.

### A. Search for Primitive Sequence

We formulate the search for the primitive sequence as a graph search problem. We adapt the regrasp graphs developed in [30], [34] to include a broader set of manipulation primitives. The nodes of the manipulation graph represent possible object stable placements and the edges represent manipulation primitive actions transforming the object from a stable placement to another. We use Dijkstra’s algorithm to search for the shortest path within the constructed graph achieving the desired pose to pose reconfiguration. In the experiment in the title figure, the primitive sequence `pull`→`pivot`→`push` is found to achieve the desired object transformations. We refer the reader to [34] and the project website [https://mcube.mit.edu/research/tactile\\_dexterity.html](https://mcube.mit.edu/research/tactile_dexterity.html) for more details on the graph search planning framework.

### B. Trajectory Planning

After the graph search planner has determined a sequence of primitives, we compute robot/object trajectories independently for each primitive. We rely on primitive specific planners for which we require that can regenerate trajectories at high frequency to exploit the timely estimated of object pose from the tactile object tracker from Section V-B.

For `grasp` and `pull`, we plan object motions that kinematically stick to the palms. For these, we directly interpolate

between the initial and final poses of the object to determine the palm pose trajectory. For `pivot`, we compute an interpolated object trajectory between the initial and final poses about a specified center of rotation. Solving (1), we find the palm pose relative to the object that maintains a sticking interaction at all contacts. For `push`, we use a Dubins’ car planner that computes the time-optimal trajectory connecting the initial and final object configurations with a single push based on [35]. We consider pushes on all sides of the object and execute the trajectory with the shortest path.

## VII. RESULTS

We validate our approach to *tactile dexterity* on an ABB YuMi dual-arm robot with real experiments visualized in <https://youtu.be/f59FoS-hV7c> that evaluate the ability of the tactile controller to handle external perturbations on a table top manipulation task.

The title figure shows snapshots of an experiment where the robot manipulates an object from an initial pose  $\mathbf{q}_0 = [0.3, -0.2, 0.07, 0.38, 0.60, 0.60, 0.38]^T$  to a target pose  $\mathbf{q}_f = [0.45, 0.3, 0.045, 0.0, 0.71, 0.0, 0.71]^T$ . To achieve the task, the robot follows the sequence: `pull` the object to the middle of the table, `pivot` the object to its target placement, and `push` it to its target location. The initial `pull` primitive is necessary to move the object to a location that allows the robot to perform a pivot manoeuvre with well defined inverse kinematics and that avoids collisions with the environment.

In Fig. 6, we evaluate the closed-loop performance of the tactile controller in Section V on individual primitives. We illustrate the regulation task of maintaining an object in a stationary pose for the `pull` and `grasp` primitives. The regulation task allows to better visualize the reactive capabilities of the controller without loss of generality. In each experiment, we apply two successive impulsive forces on the object and evaluate the stabilizing capabilities of the tactile controller. Figure 6 plots the mean and standard deviation of the error between the object’s desired and measured poses for 5 consecutive experiments. In both cases, the controller quickly reacts to the disturbance by 1) detecting slippage events at the contact interfaces and 2) tracking the pose of the object using the detected object edge in the tactile signal. First, the applied normal force is increased following (5) in reaction to the detected slippage at the contact interface. Second, the robot replans its trajectory from updates on the object state and this quickly returns to its nominal pose. For evaluation purposes, we track the ground truth pose of the object using an Apriltag visual marker [36].

## VIII. DISCUSSION

This paper introduces *tactile dexterity*, an approach to dexterous manipulation that exploits tactile sensing for reactive control. By structuring the manipulation problem as a sequence of manipulation primitives that render interpretable tactile information, we enable tactile object state estimation and tactile object control, yielding robust manipulation behavior. This requires targeting primitive interactions to those

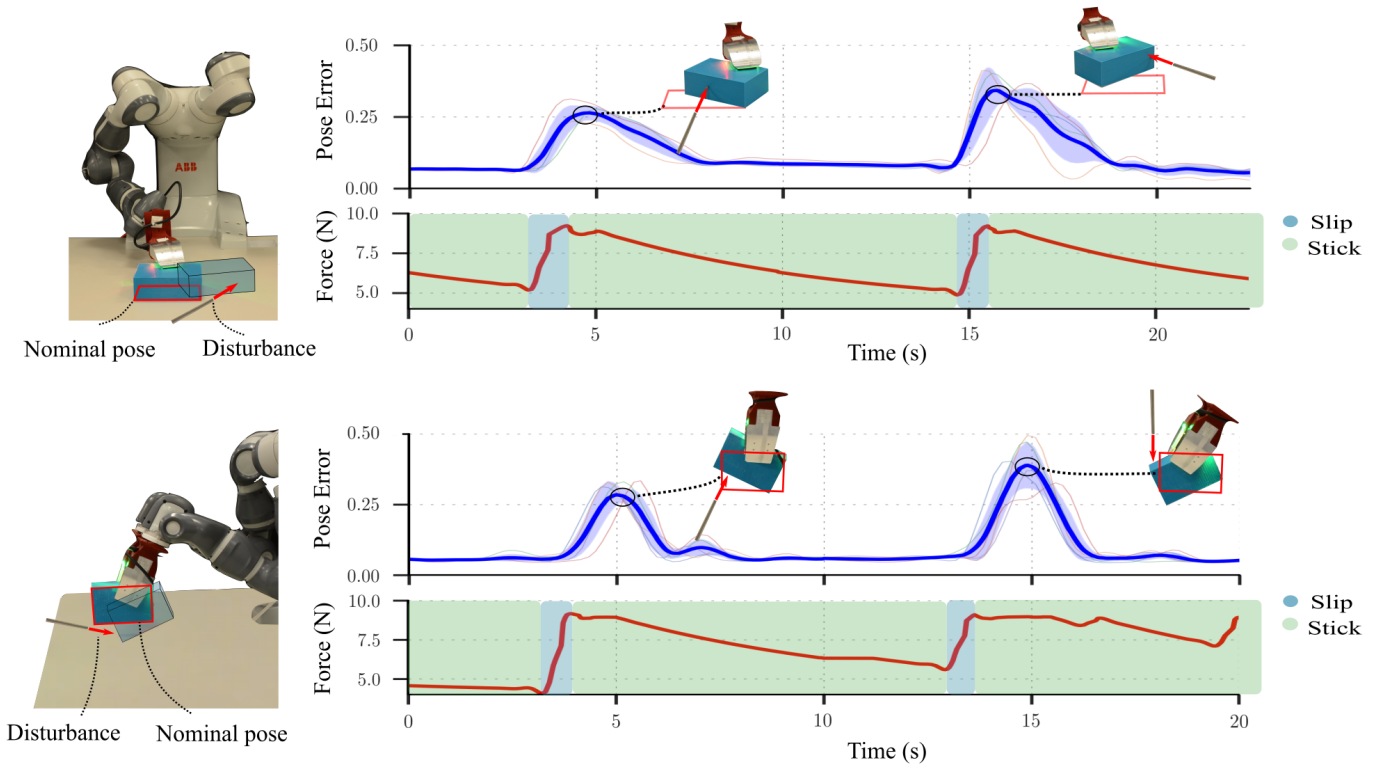


Fig. 6. **Closed-loop evaluation of the tactile controller.** We consider the task of maintaining the object in a stationary pose under external perturbations. When perturbed, the contact state controller increases the normal force on the object to prevent slippage while the contact state controller replans robot motions to bring the object back to its nominal pose.

that have simple rigid body mechanics models and efficient planners. We show the ability of the tactile controllers to modulate contact forces, track the pose of the object, and handle external perturbations by replanning in real-time.

This research relies on a number of important assumptions about the structure of the problem:

- **Known environment.** The manipulation primitives and planning algorithms are tailored to a flat surface environment.
- **Manipulation primitives.** We use human intuition to develop a library of manipulation primitives that are amenable for planning and control (for controllability).
- **Tactile Pose Estimation.** We assume that the pose of the object can be reconstructed from tactile observations (for observability).

These assumptions facilitate the development of state estimation, planning, and control frameworks capable of exploiting tactile sensing for dexterous manipulation. They constitute both the limitations, as well as strengths of the approach.

We believe that there are natural extensions to generalize this approach to dexterous manipulation that puts tactile feedback at the center:

**General Manipulation Primitives.** Developing manipulation primitives with richer contact interactions (contact switches, stick/slip interactions, etc.) will allow more expressive behavior. Due to increased complexity of the mechanics

of these primitives, they will likely benefit even more from real-time tactile observations of contact. This line of research will require controller architectures that can reason across hybrid contact switches, e.g. developed in [37], [38].

Another interesting possibility is to learn manipulation primitives from experience, and defining a primitive, broadly speaking, as a behavior for which the robot can learn a stable controller. This would potentially reduce the need for human intuition.

**Object tracking.** A formal approach to characterizing object observability directly from tactile feedback could assist with identifying discriminative features and grasps of the object. Furthermore, combining touch with vision is a natural extension of this research that would allow the robot to recover from larger object disturbances.

**Task-and-Motion Planning.** Dealing with more complex environments and tasks (e.g., pick up the box and put it on a shelf, or inside a closed drawer), would require to integrate the system with a high level task planner that can sequence manipulation primitives, as is commonly done in task and motion planning, while still being mindful of the mechanics of contact [25], [39]. Ideally, these planners would be reactive and allow the robot to switch primitives in response to unexpected deviations.

## REFERENCES

- [1] Johansson, R.S., Westling, G.: Roles of glabrous skin receptors and sensorimotor memory in automatic control of precision grip when lifting rougher or more slippery objects. *Experimental brain research* **56**(3) (1984) 550–564 [1](#)
- [2] Correll, N., Bekris, K.E., Berenson, D., Brock, O., Causo, A., Hauser, K., Okada, K., Rodriguez, A., Romano, J.M., Wurman, P.R.: Analysis and observations from the first amazon picking challenge. *IEEE Transactions on Automation Science and Engineering* **15**(1) (2016) 172–188 [1](#)
- [3] Koval, M.C., Pollard, N.S., Srinivasa, S.S.: Pose estimation for planar contact manipulation with manifold particle filters. *The International Journal of Robotics Research* **34**(7) (2015) 922–945 [1](#)
- [4] Woodruff, J.Z., Lynch, K.M.: Planning and control for dynamic, non-prehensile, and hybrid manipulation tasks. In: 2017 IEEE International Conference on Robotics and Automation (ICRA), IEEE (2017) 4066–4073 [1](#)
- [5] Hogan, F.R., Grau, E.R., Rodriguez, A.: Reactive planar manipulation with convex hybrid MPC. In: 2018 IEEE International Conference on Robotics and Automation (ICRA), IEEE (2018) 247–253 [1, 3](#)
- [6] Donlon, E., Dong, S., Liu, M., Li, J., Adelson, E., Rodriguez, A.: Gelslim: A high-resolution, compact, robust, and calibrated tactile-sensing finger. In: 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), IEEE (2018) 1927–1934 [1, 2, 3](#)
- [7] Veiga, F., Van Hoof, H., Peters, J., Hermans, T.: Stabilizing novel objects by learning to predict tactile slip. In: 2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), IEEE (2015) 5065–5072 [2](#)
- [8] Veiga, F., Edin, B.B., Peters, J.: In-hand object stabilization by independent finger control. arXiv preprint arXiv:1806.05031 (2018) [2](#)
- [9] Dong, S., Ma, D., Donlon, E., Rodriguez, A.: Maintaining grasps within slipping bounds by monitoring incipient slip. In: 2019 International Conference on Robotics and Automation (ICRA). (May 2019) 3818–3824 [2, 3](#)
- [10] Li, M., Bekiroglu, Y., Kragic, D., Billard, A.: Learning of grasp adaptation through experience and tactile sensing. In: 2014 IEEE/RSJ International Conference on Intelligent Robots and Systems, Ieee (2014) 3339–3346 [2](#)
- [11] Tian, S., Ebert, F., Jayaraman, D., Mudigonda, M., Finn, C., Calandra, R., Levine, S.: Manipulation by feel: Touch-based control with deep predictive models. arXiv preprint arXiv:1903.04128 (2019) [2](#)
- [12] Li, R., Platt, R., Yuan, W., ten Pas, A., Roscup, N., Srinivasan, M.A., Adelson, E.: Localization and manipulation of small parts using gelsight tactile sensing. In: 2014 IEEE/RSJ International Conference on Intelligent Robots and Systems, IEEE (2014) 3988–3993 [2](#)
- [13] Izatt, G., Mirano, G., Adelson, E., Tedrake, R.: Tracking objects with point clouds from vision and touch. In: 2017 IEEE International Conference on Robotics and Automation (ICRA), IEEE (2017) 4000–4007 [2](#)
- [14] Bauza, M., Canal, O., Rodriguez, A.: Tactile mapping and localization from high-resolution tactile imprints. arXiv preprint arXiv:1904.10944 (2019) [2](#)
- [15] Erdmann, M.: An exploration of nonprehensile two-palm manipulation. *The International Journal of Robotics Research* **17**(5) (1998) 485–503 [2](#)
- [16] Trinkle, J.C., Abel, J.M., Paul, R.P.: An investigation of frictionless enveloping grasping in the plane. *The International journal of robotics research* **7**(3) (1988) 33–51 [2](#)
- [17] Trinkle, J.C.: On the stability and instantaneous velocity of grasped frictionless objects. *IEEE Transactions on Robotics and Automation* **8**(5) (1992) 560–572 [2](#)
- [18] Erdmann, M.A., Mason, M.T.: An exploration of sensorless manipulation. *IEEE Journal on Robotics and Automation* **4**(4) (1988) 369–379 [2](#)
- [19] Paljug, E., Yun, X.: Experimental results of two robot arms manipulating large objects. In: Robotics and Automation, 1993. Proceedings., 1993 IEEE International Conference on, IEEE (1993) 517–522 [2](#)
- [20] Yun, X.: Object handling using two arms without grasping. *The International journal of robotics research* **12**(1) (1993) 99–106 [2](#)
- [21] Bicchi, A., Melchiorri, C., Balluchi, D.: On the mobility and manipulability of general multiple limb robots. *IEEE Transactions on Robotics and Automation* **11**(2) (1995) 215–228 [2](#)
- [22] Bicchi, A., Sorrentino, R., et al.: Dexterous manipulation through rolling. In: ICRA. Volume 95. (1995) 452–457 [2](#)
- [23] Kaelbling, L.P., Lozano-Pérez, T.: Hierarchical planning in the now. In: Workshops at the Twenty-Fourth AAAI Conference on Artificial Intelligence. (2010) [2](#)
- [24] Woodruff, J., Lynch, K.: Planning and control for dynamic, non-prehensile, and hybrid manipulation tasks. In: Robotics and Automation (ICRA), 2017 IEEE International Conference on. (2017) [2](#)
- [25] Toussaint, M., Allen, K., Smith, K.A., Tenenbaum, J.B.: Differentiable physics and stable modes for tool-use and manipulation planning. In: Robotics: Science and Systems. (2018) [2, 6](#)
- [26] Siciliano, B., Khatib, O.: Springer handbook of robotics. Springer (2016) [3](#)
- [27] Goyal, S., Ruina, A., Papadopoulos, J.: Planar sliding with dry friction part 1. limit surface and moment function. *Wear* **143**(2) (1991) 307–330 [3](#)
- [28] Howe, R.D., Cutkosky, M.R.: Practical force-motion models for sliding manipulation. *The International Journal of Robotics Research* **15**(6) (1996) 557–572 [3](#)
- [29] Chavan-Daffe, N., Holladay, R., Rodriguez, A.: In-hand manipulation via motion cones. arXiv preprint arXiv:1810.00219 (2018) [3](#)
- [30] Hou, Y., Jia, Z., Mason, M.T.: Fast planning for 3d any-pose-reorienting using pivoting. In: 2018 IEEE International Conference on Robotics and Automation (ICRA), IEEE (2018) 1631–1638 [3, 5](#)
- [31] Aceituno-Cabezas, B., Mastalli, C., Dai, H., Focchi, M., Radulescu, A., Caldwell, D.G., Cappelletto, J., Grieco, J.C., Fernández-López, G., Semini, C.: Simultaneous contact, gait, and motion planning for robust multilegged locomotion via mixed-integer convex optimization. *IEEE Robotics and Automation Letters* **3**(3) (2017) 2531–2538 [4](#)
- [32] Stewart, D.E., Trinkle, J.C.: An implicit time-stepping scheme for rigid body dynamics with inelastic collisions and coulomb friction. *International Journal for Numerical Methods in Engineering* **39**(15) (1996) 2673–2691 [4](#)
- [33] Holladay, R.M., Srinivasa, S.S.: Distance metrics and algorithms for task space path optimization. In: 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), IEEE (2016) 5533–5540 [4](#)
- [34] Wan, W., Mason, M.T., Fukui, R., Kuniyoshi, Y.: Improving regrasp algorithms to analyze the utility of work surfaces in a workcell. In: Robotics and Automation (ICRA), 2015 IEEE International Conference on, IEEE (2015) 4326–4333 [5](#)
- [35] Zhou, J., Hou, Y., Mason, M.T.: Pushing revisited: Differential flatness, trajectory planning, and stabilization. *The International Journal of Robotics Research* (2017) 0278364919872532 [5](#)
- [36] Olson, E.: Apriltag: A robust and flexible visual fiducial system. In: 2011 IEEE International Conference on Robotics and Automation, IEEE (2011) 3400–3407 [5](#)
- [37] Hogan, F.R., Rodriguez, A.: Feedback control of the pusher-slider system: A story of hybrid and underactuated contact dynamics. arXiv preprint arXiv:1611.08268 (2016) [6](#)
- [38] Doshi, N., Hogan, F.R., Rodriguez, A.: Hybrid differential dynamic programming for planar manipulation primitive. arXiv preprint arXiv:1911.00175 (2019) [6](#)
- [39] Holladay, R., Lozano-Pérez, T., Rodriguez, A.: Force-and-motion constrained planning for tool use. In: International Conference on Intelligent Robots and Systems (IROS). (2019) [6](#)