

Manipulation of unknown objects via contact configuration regulation

Neel Doshi*,¹ Orion Taylor* Alberto Rodriguez

Massachusetts Institute of Technology

*Authors contributed equally, names listed in alphabetical order

Abstract—We present an approach to robotic manipulation of unknown objects through regulation of the object’s *contact configuration*: the location, geometry, and mode of all contacts between the object, robot, and environment. A contact configuration constrains the forces and motions that can be applied to the object; however, synthesizing these constraints generally requires knowledge of the object’s pose and geometry. We develop an object-agnostic approach for estimation and control that circumvents this need. Our framework directly estimates a set of wrench and motion constraints which it uses to regulate the contact configuration. We use this to reactively manipulate unknown planar polygonal objects in the gravity plane.

I. INTRODUCTION

Regulation of an object’s *contact configuration* – the location, mode and geometry of all contacts between the object, robot, and environment – is a fundamental abstraction of object manipulation (Fig. 1). Imagine tumbling a heavy box or tightening a screw. Both tasks can be better described/executed by prescribing/regulating the location, geometry, and mode of all contacts. In these cases, the object can be sufficiently controlled without using pose, inertial, or shape information. Even when this information is available, contact configuration regulation simplifies control, for example during non-prehensile [1], [2] or deformable-object [3] manipulation.

As such, contact configuration regulation can be used to manipulate unknown objects (Fig. 1b). This is a *joint estimation and control* problem. The robot must estimate the kinematic and frictional constraints imposed by the contact configuration and regulate the contact forces and object motion accordingly. This is challenging, as not all contacts are directly observable, and the robot’s control authority is limited by the underactuated mechanics of friction. We focus on manipulating unknown planar polygonal objects on a flat surface using robot proprioception and force/torque sensing at the wrist for feedback (Section III). This minimal system has a diverse set of contact configurations that highlight the challenges discussed above.

Contributions We manipulate the object by regulating it through a (currently predesignated) sequence of contact configurations. For each contact configuration, we

¹ Neel Doshi is supported by an appointment to the Intelligence Community Postdoctoral Research Fellowship Program at the Massachusetts Institute of Technology, administered by Oak Ridge Institute for Science and Education through an interagency agreement between the U.S. Department of Energy and the Office of the Director of National Intelligence.

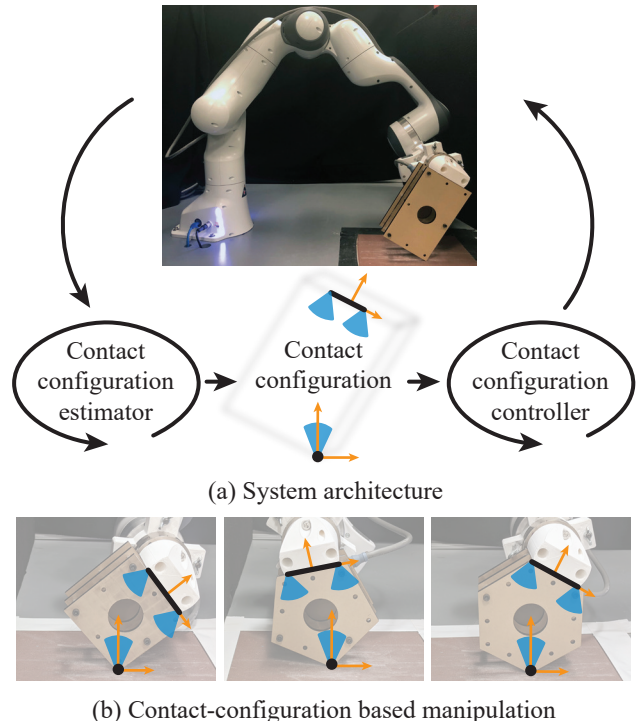


Fig. 1. a) The robot interacts with an object by estimating and controlling its contact configuration. (b) This enables manipulation of unknown objects without explicit knowledge of their geometric and inertial properties.

- **Estimate** contact locations, geometries, and modes, as well as a conservative set of wrenches that can be transmitted to the object (Section IV).
- **Control** the applied wrenches to maintain the contact geometries and modes, as well as the applied motions to regulate the contact locations (Section V).

Both the estimator and controller are object-agnostic. They work together to manipulate the object: the controller maintains the contact configuration to facilitate consistent estimation, and the estimated parameters improve the controller’s performance (Section VI). We experimentally demonstrate that our framework can reactively manipulate planar polygonal objects along a horizontal surface (Section VII).

II. RELATED WORK

Prior work on contact configuration estimation and control has focused on a robot interacting with a static [4], [5] or single degree-of-freedom [6], [7] unknown environment. The ability to directly observe all contacts facilitates contact configuration regulation. Less attention has been given to




Fig. 2. (a) System diagram with parameters, unit vectors, and state variables labelled. (b) Pivoting while maintaining sticking at both hand and ground contacts. (c) Hand sliding while maintaining sticking contact with the ground. (d) Object sliding while maintaining sticking contact with the hand.

cases of robot/object/environment interactions, where not all contacts are directly observable (e.g., during non-prehensile manipulation). In this area, most prior research either focuses on contact configuration control assuming a known model of the world [2], [8], [9], or contact configuration estimation assuming stable interactions [10]. Work on joint estimation and control either uses simplified (e.g., frictionless) models of contact [11], [12], or learns task-specific policies from data (e.g. for cable manipulation [3] or part insertion [13]). Our contribution is an object-agnostic joint estimation and control framework that reasons about all frictional interactions between the robot, object, and environment.

Estimation Several researchers have proposed methods for localizing robot/object contacts where the known geometry of the manipulator can be leveraged [14], [15], [16], [17]. We extend these approaches to localize object/environment contacts, formalized as *extrinsic* contact sensing [10]. Prior work on wrench-constraint estimation has focused on the problem of planar pushing [18], [19], [20]. Of particular relevance is the work of Zhou et. al [20], who estimate the set of wrenches that can be transmitted through planar frictional contact (i.e., the limit-surface [21]). We focus instead on estimating the intersection of both frictional and kinematic wrench constraints (i.e., the generalized friction cone [22]).

Control We highlight two approaches that are used to simultaneously regulate force and motion: indirect force control (IFC) [23], [24] and hybrid force-velocity control (HFVC) [25], [26]. IFC regulates wrenches and motions by prescribing the interaction dynamics, while HFVC directly regulates wrenches and motions in orthogonal subspaces. Our controller can be viewed through both lenses. From the IFC perspective, it's similar to hybrid impedance control [27], in which wrenches and velocities are controlled in orthogonal subspaces via impedance modulation. From the HFVC perspective, it's similar to parallel force/velocity control [28], in which separate wrench and velocity controllers are summed, with priority given to the wrench controller. However, our approach is unique in that, instead of defining orthogonal subspaces or a priority list, we use a quadratic program to determine wrench and velocity control directions and priorities online.

III. SYSTEM OVERVIEW

As described in Section I, we consider quasi-static manipulation of unknown planar polygonal objects on a horizontal surface. The three main system components are: (a) the robot **hand**, which is a line of length $2l_h$, (b) the **object**, which is a planar convex polygon that moves in the gravity plane, and (c) the **ground**, which is a fixed horizontal line. The system's state consists of the planar poses of the hand (x_h, y_h, θ_h) and object (x_o, y_o, θ_o) as shown in Fig. 2a. We focus on states in which the object contacts both the hand and the ground.

Each contact configuration imposes kinematic and frictional constraints on the system. For example, in Fig. 2b, sticking point contact with the ground and sticking line contact with the hand constrains the system to only admit rigid rotations of the hand and object about the ground contact (say by $\Delta\theta_o$). Other contact configurations (Fig. 2c,d) can admit relative sliding between the object and the hand (Δs_h) or ground (Δs_g). We combine these admissible motions (i.e., $\Delta\theta_o$, Δs_h , and Δs_g) to explore the system's state-space.

Under the quasi-static assumption, the net force on the object is zero:

$$\sum \mathbf{w}_{\text{net}} = \mathbf{w}_h + \mathbf{w}_g + \mathbf{w}_{\text{grav}} = \mathbf{0}, \quad (1)$$

where \mathbf{w}_h , \mathbf{w}_g , and \mathbf{w}_{grav} are the wrenches exerted on the object by the hand, ground, and gravity, respectively. We also assume that all contact interactions follow a simple contact model with constant friction coefficients. Given a line segment of length $2l$ contacting a longer surface with friction coefficient μ , the wrench (f_n, f_t, τ) transmitted through this contact obeys:

$$-\mu f_n \leq f_t \leq \mu f_n \quad (\text{Coulomb friction}) \quad (2)$$

$$-l f_n \leq \tau \leq l f_n \quad (\text{Line contact}) \quad (3)$$

Here f_n and f_t , are the forces normal and tangential to the contact surface, and τ is net torque with respect to the center of the line. Satisfaction of the strict inequality corresponds to sticking for (2) and flush contact for (3), while an equality admits sliding left/right for (2) and pivoting clockwise/counter-clockwise about one of the boundaries of the line segment for (3).




Fig. 3. An illustration of estimated wrench cone constraints (i.e., equations (5)-(8) and (11)-(12)). Measured wrenches at the hand/ground are shown as blue/purple dots. A conservative estimate of the hand/ground friction constraints are shown in dashed blue/purple, and the prescribed torque constraint is shown in solid red.

This model allows us to describe these force-motion relationships using a series of complementarity constraints. We begin by decomposing the contact velocities, \dot{s}_h , \dot{s}_g , and the relative angle $\Delta\theta = \theta_o - \theta_h$, into their positive and negative components:

$$\dot{s}_h = \dot{s}_h^+ - \dot{s}_h^-, \quad \Delta\theta = \Delta\theta^+ - \Delta\theta^-, \quad \dot{s}_g = \dot{s}_g^+ - \dot{s}_g^- \quad (4)$$

We can now relate the hand contact wrench to the relative velocity/orientation of the hand with respect to the object:

$$0 \leq \mu_h f_{h,n} + f_{h,t} \perp \dot{s}_h^+ \geq 0 \quad (5)$$

$$0 \leq \mu_h f_{h,n} - f_{h,t} \perp \dot{s}_h^- \geq 0 \quad (6)$$

$$0 \leq l_h f_{h,n} - \tau_h \perp \Delta\theta^+ \geq 0 \quad (7)$$

$$0 \leq l_h f_{h,n} + \tau_h \perp \Delta\theta^- \geq 0 \quad (8)$$

The \perp symbol is a shorthand, where $0 \leq a \perp b \geq 0$ implies $a \geq 0$, $b \geq 0$, and $ab = 0$. Moreover, \mathbf{f}_h and τ_h are the force and torque exerted by the hand, μ_h is the friction coefficient at the hand contact, and the subscripts n and t indicate the normal and tangential directions to the hand contact surface. Equations (5) and (6) relate the forces applied by the hand to relative sliding between the hand and object. Similarly, equations (7) and (8) relate the torque applied by the hand to the relative rotation between the hand and object.

We impose similar constraints at the ground contact:

$$0 \leq \mu_g f_{g,J} - f_{g,I} \perp \dot{s}_g^+ \geq 0 \quad (9)$$

$$0 \leq \mu_g f_{g,J} + f_{g,I} \perp \dot{s}_g^- \geq 0, \quad (10)$$

where \mathbf{f}_g is the force exerted by the ground, μ_g is the friction coefficient at the ground contact, and the subscripts J and I indicate the normal and tangential directions to the horizontal ground. We can use static equilibrium (1) to express (9) and (10) in terms of the force exerted by the hand:

$$0 \leq -\mu_g (f_{h,J} + f_{grav,J}) + (f_{h,I} + f_{grav,I}) \perp \dot{s}_g^+ \geq 0 \quad (11)$$

$$0 \leq -\mu_g (f_{h,J} + f_{grav,J}) - (f_{h,I} + f_{grav,I}) \perp \dot{s}_g^- \geq 0 \quad (12)$$

IV. OBJECT-AGNOSTIC ESTIMATION

The estimator uses a time history of the measured hand wrench ($\mathbf{w}_{h,meas}$) and pose ($[x_h, y_h, \theta_h]^T$) to estimate:

- The **generalized friction cone**: the set of all wrenches that can be applied to the object (Section IV-A).

- The **contact mode**: whether each contact is sticking, sliding left, or sliding right (Section IV-A).
- The **contact geometry**: whether the hand contact is flush or pivoting (Section IV-A).
- The **contact locations**: the relative sliding positions of the hand (s_h) and object (s_g), Section IV-B.

These items constitute what is required to regulate the contact configuration. In addition, we also use this information to estimate and regulate θ_o , the object's orientation.

A. Wrench constraint estimation

The complementarity conditions in equations (5)-(8) and (11)-(12) allow us to estimate/regulate the contact mode and geometry by measuring/regulating the hand contact wrench. The first step is to estimate the wrench space boundaries on the left-hand side of the complementarity (LHS).

We want to construct a conservative online estimate of the wrench boundaries that is robust to sensor noise and possible outliers. A conservative estimate is desirable, as misclassifying a sliding/pivoting contact as a sticking/flush contact compromises both kinematic estimation and the controller's ability to enforce sticking/flush contact. Our estimator relies on the following: (a) that every measured wrench satisfies the wrench constraints, (b) the wrench boundaries are convex, (c) a conservative estimate is sufficient, and (d) the wrench constraints are constant in either the contact (LHS of (5)-(8)) or world (LHS of (11)-(12)) frames.

Hand contact friction constraints We rewrite the friction constraints at the hand contact (5)-(6) as:

$$\mu_{meas} = |(f_{h,t}/(f_{h,n} + \epsilon))| < |f_{h,t}/f_{h,n}| \leq \mu_h, \quad (13)$$

where μ_{meas} is a lower-bound on μ_h and $\epsilon > 0$ ensures that μ_{meas} is both conservative and well-defined, even when the hand wrench is close to zero. To estimate μ_h , we compute μ_{meas} for each measured wrench point. We can use all measured data to estimate a single μ_{meas} because the hand contact friction cone is symmetric. We then use the LiveStats [29] online quantile estimator to find an approximation for the 99% quantile of μ_{meas} , which is our estimate for μ_h . The estimated friction constraints are shown in Fig. 3 (left). In practice, this estimation scheme has been robust to outliers.

Robot contact torque constraints We assume that the object surface always contains the length of the hand. Consequently, the torque boundaries at the hand contact (LHS of (7)-(8)) are known apriori (Fig. 3, center). In the future, we intend to estimate the left and right torque boundaries during overhang or point contact between the hand and object.

Ground contact friction constraints Unlike the hand contact friction constraints, the ground friction boundaries (LHS of (11)-(12)) cannot be rewritten in a highly structured form, as they depend on the unknown mass of the object. Instead, we directly estimate a convex hull of the force measurements ($f_{h,J}, f_{h,I}$) as a conservative approximation of the ground contact friction constraints (Fig. 3, right). This estimate is also robust to different ground angles. We have developed a heuristic for generating an online estimate of the convex

hull of a set of data points that is robust to sensor noise and outliers. It uses the same LiveStats [29] online quantile estimator to infer a set of supporting hyperplanes of the convex hull. The intersections of these hyperplanes are used to identify candidate corners, which are then refined into the final approximation of the ground contact friction cone (Fig. 3, right).

Contact mode and geometry estimation Based on the complementarity conditions, we use our estimates of the wrench constraints to infer the contact mode of the system. We predict that the hand is sliding/pivoting if the measured wrench is in (near) violation of our estimates of the LHS of (5)-(6)/(7)-(8), and in sticking/line contact otherwise. Similarly, we predict that object is sliding along the ground if the measured wrench is in (near) violation of our estimates of the LHS of (11)-(12) and is sticking otherwise. As discussed above, this estimation scheme is conservative: sliding/pivoting contact is rarely misclassified as sticking/flush contact, but the reverse misclassification can be frequent. This is intentional as undetected sliding/pivoting introduces error into the kinematic estimator and can result in the system moving into an unrecoverable state.

B. Object pose and contact location estimation

The kinematic estimator synthesizes measurements from robot proprioception and the contact mode estimate into estimates of the object's orientation, as well as the hand and ground contact locations. We exploit the kinematic constraints that sticking/line contacts place on the system's motion. These are the right-hand side of the complementarity in (5)-(8) and (11)-(12) (RHS).

When the hand and object are in line contact, their contact faces are parallel (i.e., $\Delta\theta = 0$, meaning $\theta_h = \theta_o$). In this case, we relate the hand pose to the ground contact as follows:

$$-s_h \mathbf{e}_{h,t} - d \mathbf{e}_{h,n} + \mathbf{r}_o = \mathbf{r}_h, \quad (14)$$

where $\mathbf{r}_h = [x_h, y_h]^T$, $\mathbf{r}_o = [x_o, y_o]^T$, and d is the constant distance between the ground contact and the object face in contact with the hand (see Fig. 2a). Due to line-contact, $\mathbf{e}_{h,n}$ and $\mathbf{e}_{h,t}$ are unit vectors normal and tangent to both hand and object contact surfaces. The location of the ground contact, \mathbf{r}_o , is constant across periods of sticking point contact with the ground. During these periods, we regress d , x_o , and y_o from the dot product of (14) with $\mathbf{e}_{h,n}$:

$$d(-1) + x_o(\mathbf{I} \cdot \mathbf{e}_{h,n}) + y_o(\mathbf{J} \cdot \mathbf{e}_{h,n}) = \mathbf{r}_h \cdot \mathbf{e}_{h,n}, \quad (15)$$

where $\mathbf{e}_{h,n}$ and \mathbf{r}_h are measured. Note that $s_g = x_o$.

Once the ground contact location is inferred, we can estimate the sliding position of the hand s_h from the hand's position during line-contact via the relation:

$$s_h = (\mathbf{r}_o - \mathbf{r}_h) \cdot \mathbf{e}_{h,t} \quad (16)$$

During periods of sticking line contact at the hand (i.e., $\Delta\theta = \dot{s}_h = 0$), the hand and object move as a single rigid body. We can use this to update x_o, y_o during periods of sliding ground contact using the hand pose and our previous estimates of

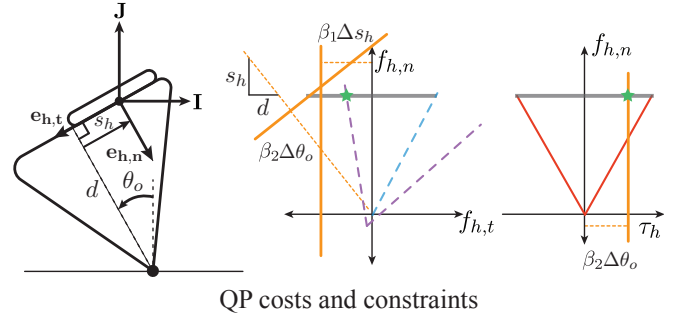


Fig. 4. An illustration of QP cost and constraints in the hand frame for maintaining sticking point contact with ground and object orientation while encouraging sliding right at the hand. The ground friction constraints (dashed purple) are appropriately transformed into the hand frame and only the hand friction constraint preventing sliding left is enforced (dashed blue). Solid orange lines show minima for individual cost terms, dashed orange lines correspond to projected error, and solid gray lines are prescribed normal force limits. A potential optimum for $\Delta f_{h,n}$, $\Delta f_{g,n}$, and $\Delta \tau_h$ is shown using green stars.

d and s_h . When the object is in point contact with the hand and sticking with the ground, we can still use the system's kinematics to estimate θ_o and s_h . In all other contact mode combinations, the motion of the system is indeterminate, and we reset the estimator.

V. OBJECT-AGNOSTIC CONTROL

The controller drives the system along admissible motion directions (a subset of $\Delta\theta_o$, Δs_h , Δs_g) and regulates the hand contact wrench to maintain the desired contact mode/geometry. The regulation of both forces and motions is achieved through a low-level impedance control layer. We prefer an IFC over explicitly defining force and motion control subspaces as it is more tolerant of uncertainty in the kinematic estimates.

For a quasi-static system, the low-level impedance controller can be approximated by the following compliance law:

$$\mathbf{w}_h = \mathbf{K}(\mathbf{x}_{\text{tar}} - \mathbf{x}_h) \quad (17)$$

Here \mathbf{x}_{tar} is the planar impedance target for the hand, \mathbf{x}_h is the planar pose of the hand, \mathbf{K} is a stiffness matrix, and \mathbf{w}_h is the wrench exerted by the robot. Our controller regulates both \mathbf{x}_h and \mathbf{w}_h by updating \mathbf{x}_{tar} while keeping \mathbf{K} constant.

The controller relies on the intuition that, for a fixed contact mode/geometry, moving the impedance target will only change the equilibrium state of the system along the admissible motion directions. For the contact mode/geometry shown in Fig. 4 (left), these are object rotation ($\Delta\theta_o$) and hand sliding (Δs_h). The controller can reduce object orientation error by moving the impedance target along the direction $\mathbf{v}_\theta = [-d, s_h, 1]^T$ in the hand contact frame. The product $\Delta\theta_o \mathbf{v}_\theta$ corresponds to the predicted change in the hand pose for an incremental rotation by $\Delta\theta_o$ of the object about the estimated ground contact. Similarly, the controller can reduce error in s_h by moving the impedance target in the hand-frame direction $\mathbf{v}_h = [-1, 0, 0]^T$, where $\Delta s_h \mathbf{v}_h$ corresponds to an incremental change in the hand pose due to sliding by Δs_h . Finally, the error in s_g can be regulated using similar reasoning with $\mathbf{v}_g = [1, 0, 0]^T$ in the world frame.




Fig. 5. Images from a contact configuration sequence during manipulation of an equilateral triangle with an addition of 500 g mass.

Regulation of the contact wrenches ensures that the contact mode/geometry is maintained and the equilibrium state of the system changes as expected. Maintaining the contact mode/geometry shown in Fig. 4 (left), requires enforcing a strict inequality in the LHS of (7)-(8) (i.e., flush hand contact) and (11)-(12) (i.e., sticking at the ground contact). In this example, the controller also enforces the constraint preventing sliding in the direction opposite to the one commanded: (5) when sliding left or (6) when sliding right. As such, the specific wrenches acting on the object do not matter to the controller, so long as they satisfy the wrench constraints defined by the contact mode/geometry.

For this example, the controller computes the incremental change of the hand impedance target ($\Delta \mathbf{x}_{\text{tar}}$) as the solution to the following quadratic program (QP), visualized in Fig. 4:

$$\min_{\Delta \mathbf{w}_h, \Delta \mathbf{x}_{\text{tar}}} (\Delta \mathbf{x}_{\text{tar}}^T \mathbf{v}_h - \Delta s_h)^2 + (\Delta \mathbf{x}_{\text{tar}}^T \mathbf{v}_\theta - \Delta \theta)^2 \quad (18)$$

$$s.t. \quad \mathbf{n}_j^T (\mathbf{w}_{h, \text{meas}} + \gamma_j \Delta \mathbf{w}_h) \leq b_j \quad \forall j \quad (19)$$

$$\Delta \mathbf{w}_h = \mathbf{K} \Delta \mathbf{x}_{\text{tar}} \quad (20)$$

The sum $\mathbf{w}_{\text{meas}} + \Delta \mathbf{w}_h$ is the predicted wrench that the hand will exert after the control is applied. The set of constraints (19) correspond to the estimated value of the wrench space constraints discussed above plus a maximum on the normal force applied by the hand. The scaling factors γ_j are used to amplify or attenuate wrench corrections. The equality constraint (20) is an incremental approximation of (17) that relies on the assumption that $\Delta \mathbf{x}_h$ is small along the directions where the hand applies a wrench $\Delta \mathbf{w}_h$ to the object (i.e., the constrained motion directions). Finally, as discussed, the cost terms (18) reduce control error along the admissible motion directions by moving the impedance target along corresponding hand-frame incremental motion directions.

The cost function for the QP shown above can be applied to any of the contact modes/geometries discussed as:

$$\alpha_0 \|\Delta \mathbf{x}_{\text{tar}}\|^2 + \sum \alpha_i (\Delta \mathbf{x}_{\text{tar}}^T \mathbf{v}_i - \beta_i \Delta \varepsilon_i)^2 \quad (21)$$

Here $\Delta \varepsilon_i$ are the errors along the admissible motion directions, the $\|\Delta \mathbf{x}_{\text{tar}}\|^2$ term regularizes the cost, and α_i and β_i are controller gains. The incremental change of the hand impedance target is computed by minimizing this cost subject to the equality constraint (20) and the appropriate subset of wrench constraints (19).

VI. JOINT ESTIMATION AND CONTROL

The estimator and controller work together to manipulate an unknown object. The hand starts in line-contact with

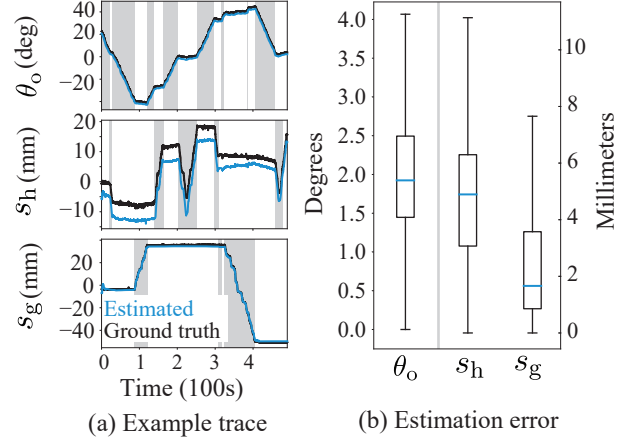


Fig. 6. (a) Ground truth (black) versus estimated (blue) time-traces for object orientation (top), hand sliding position (middle), and object sliding position (bottom) during manipulation of a triangle with no additional mass. Highlighted gray regions indicate when the corresponding variable was commanded to change. (b) Box and whisker plots show absolute estimation error for all ten trials. The median is in blue, top and bottom edges of the box are the 75% and 25% quantiles, and each whisker is $1.5 \times$ the box length. Outliers constitute 5% of the data and are not shown.

the object, and consequently, its orientation is known. We initialize the estimator with conservative guesses of the friction constraints (i.e., small friction coefficients). The object geometry, relative pose of the hand, and pivot location are unknown. We manipulate the object in a reactive fashion through a hand-scripted sequence of contact configurations.

Initially, we focus on exploration to improve our estimates of the system's kinematics and wrench constraints. We use the controller described above with a naive guess of $\mathbf{v}_\theta = [-d_{\text{guess}}, 0, 1]^T$, and noting that \mathbf{v}_h and \mathbf{v}_g are already geometry-agnostic. While initial commanded motions are limited by our conservative estimate of the friction constraints, we more accurately estimate each constraint by commanding sliding in that direction. For example, commanding right sliding at the hand (i.e., $s_h^+ > 0$) removes the LHS of (5) from the QP and drives the measured wrench towards the true constraint boundary. This allows the estimator to more accurately estimate that constraint. We also command object rotations to estimate the pivot location.

As the robot manipulates the object, its estimates of the wrench constraints and system kinematics become more accurate. This improves the controller's performance by refining \mathbf{v}_θ and expanding the feasible region of the QP.

VII. EXPERIMENTS AND RESULTS

We conduct several experiments to verify our framework using Franka Emika's Panda robot in impedance control

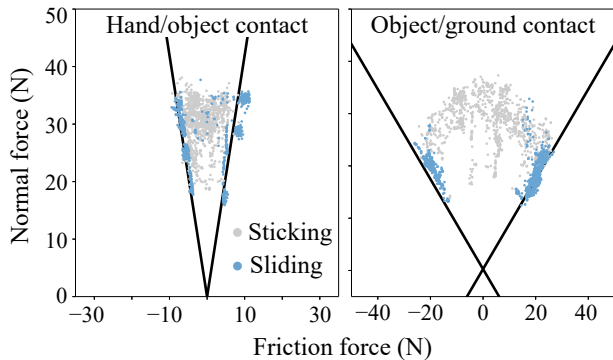


Fig. 7. Estimated friction cone boundaries (black) for the hand (left) and ground (right) contacts during the same trial as in Fig. 6. We superimpose measured wrenches, which are colored gray for sticking and blue for sliding based on the contact mode measured via the ground truth.

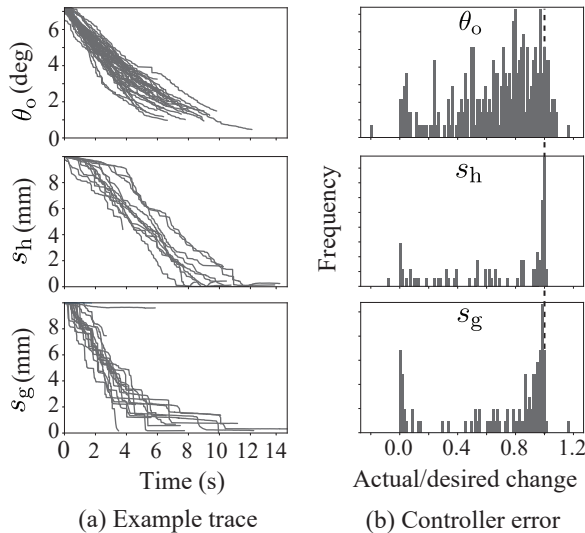


Fig. 8. (a) Time traces of the control error during incremental motions along the admissible motions directions for the same trial as in Fig. 6. (b) A histogram of the actual change in pose normalized by the commanded change in pose along admissible motion directions across all ten trials. A value of one indicates commanded change is fully realized.

mode. The estimators and controller run at 100 Hz and 30 Hz.

Object-agnostic manipulation We test our framework’s ability to manipulate unknown polygons via ten experiments. For each experiment, the robot manipulates either an equilateral triangle, a square, a rectangle, or regular hexagon or pentagon. The objects’ side-length and mass vary between 7 cm to 19 cm and 200 g to 450 g. For two experiments, we add a 500 g mass to the triangle and pentagon. The estimator and controller are initialized as described in Section VI, and the manipulator moves through a long sequence of contact configurations during each trial. A portion of this sequence is shown in Fig. 5. We command changes in θ_o , s_g and s_h in increments of 7.2° , 10 mm, and 10 mm, respectively.

We compare the performance of our kinematic estimator against ground truth provided using the AprilTag vision system [30]. We show an example time-trace in Fig. 6a, and present the estimation error statistics for all ten trials in Fig. 6b. These results indicate that kinematic estimation is sufficiently accurate for control purposes. We also compare the wrench cone constraint estimates to the contact mode

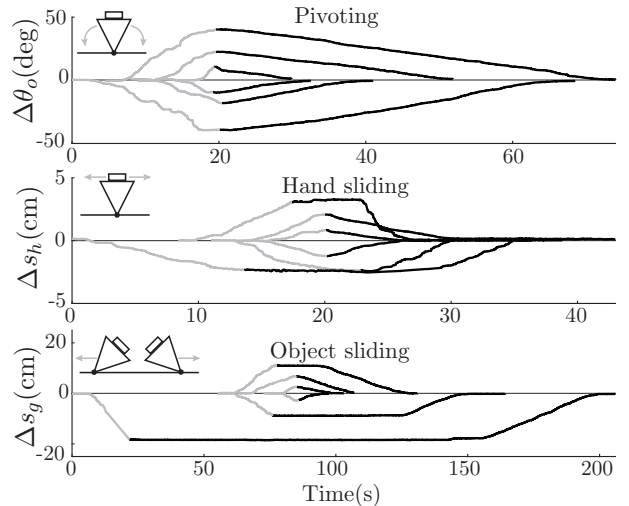


Fig. 9. A demonstration of pose regulation along the admissible motion directions. Time traces of the ground truth error in the state variable being actively regulated are shown. Gray lines correspond to a manually applied perturbation, and black lines show the controller’s response.

as detected via the ground truth in Fig. 7. As expected, the measured wrench is near the edge of the friction cone when sliding occurs. Lastly, we analyze the ability of the controller to execute motions along the admissible directions $\Delta\theta_o$, Δs_h , and Δs_g in Fig. 8. We find that the controller drives the system towards the commanded state.

Perturbation rejection We also perform experiments to test the controller’s ability to reject perturbations along the three admissible motion directions (Fig. 9). We find that the system is able to recover from large perturbations in θ_o , s_h , and s_g as long as the initial state can be reached from the perturbed state under the given contact geometry and mode. For instance, for the given system parameters and initial conditions, ground sliding regulation (bottom, Fig. 9) can only correct perturbations in one direction.

VIII. CONCLUSION

We demonstrate that objects can be manipulated by estimating and controlling their contact configurations. This allows for object manipulation without knowledge of its pose, inertia, or shape. We leverage this to manipulate unknown planar polygonal objects, but posit that the idea of contact configuration regulation is more general.

A limitation of this approach is our inability to estimate the length and location of patch contacts between the object and the hand or ground. The key difficulty is that, unlike the friction parameters, the length and location of these patches will change based on the motion of the hand and object. We plan to address this in the future by incorporating kinematic feedback into the wrench cone estimator.

We also plan to develop a higher level planning framework that automatically sequences contact configurations and intermediate target states to (a) move between states that require sequencing multiple contact configurations, and (b) to more efficiently explore the state space and estimate the kinematic parameters and wrench constraints.

REFERENCES

- [1] N. Chavan-Dafle, R. Holladay, and A. Rodriguez, "Planar in-hand manipulation via motion cones," *The International Journal of Robotics Research*, vol. 39, no. 2-3, pp. 163–182, 2020.
- [2] F. R. Hogan and A. Rodriguez, "Reactive planar non-prehensile manipulation with hybrid model predictive control," *The International Journal of Robotics Research*, vol. 39, no. 7, pp. 755–773, 2020.
- [3] Y. She, S. Wang, S. Dong, N. Sunil, A. Rodriguez, and E. Adelson, "Cable manipulation with a tactile-reactive gripper," *The International Journal of Robotics Research*, 2021.
- [4] M. G. Her and H. Kazerooni, "Automated Robotic Deburring of Parts Using Compliance Control," *Journal of Dynamic Systems, Measurement, and Control*, vol. 113, no. 1, pp. 60–66, 03 1991.
- [5] T. Furukawa, D. Rye, M. Dissanayake, and A. Barratt, "Automated polishing of an unknown three-dimensional surface," *Robotics and Computer-Integrated Manufacturing*, vol. 12, no. 3, pp. 261–270, 1996.
- [6] Y. Karayiannidis, C. Smith, F. E. Vina, P. Ögren, and D. Kragic, "Model-free robot manipulation of doors and drawers by means of fixed-grasps," in *Proceedings of IEEE International Conference on Robotics and Automation (ICRA)*, 2013, pp. 4485–4492.
- [7] G. Niemeyer and J.-J. Slotine, "A simple strategy for opening an unknown door," in *Proceedings of IEEE International Conference on Robotics and Automation (ICRA)*, vol. 2, 1997, pp. 1448–1453.
- [8] Y. Hou, Z. Jia, and M. Mason, "Manipulation with Shared Grasping," in *Proceedings of Robotics: Science and Systems*, Corvallis, Oregon, USA, July 2020.
- [9] F. R. Hogan, J. Ballester., S. Dong, and A. Rodriguez, "Tactile dexterity: Manipulation primitives with tactile feedback," in *Proceedings of International Conference on Robotics and Automation (ICRA)*, 2020.
- [10] D. Ma, S. Dong, and A. Rodriguez, "Extrinsic contact sensing with relative-motion tracking from distributed tactile measurements," in *Proceedings of IEEE International Conference on Robotics and Automation (ICRA)*, 2021.
- [11] T. Lefebvre, H. Bruyninckx, and J. De Schutter, "Polyhedral contact formation modeling and identification for autonomous compliant motion," *IEEE Transactions on Robotics and Automation*, vol. 19, no. 1, pp. 26–41, 2003.
- [12] J. De Schutter, T. De Laet, J. Rutgeerts, W. Decré, R. Smits, E. Aertbeliën, K. Claes, and H. Bruyninckx, "Constraint-based task specification and estimation for sensor-based robot systems in the presence of geometric uncertainty," *The International Journal of Robotics Research*, vol. 26, no. 5, pp. 433–455, 2007.
- [13] S. Dong, D. Jha, D. Romeres, S. Kim, D. Nikovski, and A. Rodriguez, "Tactile-rl for insertion: Generalization to objects of unknown geometry," in *Proceedings of IEEE International Conference on Robotics and Automation (ICRA)*, 2021.
- [14] A. Bicchi, J. K. Salisbury, and D. L. Brock, "Contact sensing from force measurements," *The International Journal of Robotics Research*, vol. 12, no. 3, pp. 249–262, 1993.
- [15] L. Manuelli and R. Tedrake, "Localizing external contact using proprioceptive sensors: The contact particle filter," in *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2016.
- [16] K.-T. Yu and A. Rodriguez, "Realtime state estimation with tactile and visual sensing for inserting a suction-held object," in *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018.
- [17] S. Wang, A. Bhatia, M. T. Mason, and A. M. Johnson, "Contact localization using velocity constraints," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020.
- [18] T. Yoshikawa and M. Kurisu, "Identification of the center of friction from pushing an object by a mobile robot," in *Proceedings of IEEE/RSJ International Workshop on Intelligent Robots and Systems (IROS)*, 1991.
- [19] K. M. Lynch, "Estimating the friction parameters of pushed objects," in *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 1993.
- [20] J. Zhou, M. T. Mason, R. Paolini, and D. Bagnell, "A convex polynomial model for planar sliding mechanics: theory, application, and experimental validation," *The International Journal of Robotics Research*, vol. 37, no. 2-3, pp. 249–265, 2018.
- [21] S. Goyal, A. Ruina, and J. Papadopoulos, "Planar sliding with dry friction part 1. limit surface and moment function," *Wear*, vol. 143, no. 2, pp. 307–330, 1991.
- [22] M. Erdmann, "On a representation of friction in configuration space," *The International Journal of Robotics Research*, vol. 13, no. 3, pp. 240–271, 1994.
- [23] J. K. Salisbury, "Active stiffness control of a manipulator in cartesian coordinates," in *IEEE conference on decision and control*, 1980.
- [24] N. Hogan, "Impedance control: An approach to manipulation: Parts i–iii," *ASME J. Dyn. Syst. Meas. Control*, vol. 107, pp. 1–24, 1985.
- [25] M. T. Mason, "Compliance and force control for computer controlled manipulators," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 11, no. 6, pp. 418–432, 1981.
- [26] M. H. Raibert and J. J. Craig, "Hybrid Position/Force Control of Manipulators," *Journal of Dynamic Systems, Measurement, and Control*, vol. 103, no. 2, pp. 126–133, 06 1981.
- [27] R. J. Anderson and M. W. Spong, "Hybrid impedance control of robotic manipulators," *IEEE Journal on Robotics and Automation*, vol. 4, no. 5, pp. 549–556, 1988.
- [28] S. Chiaverini and L. Sciacivco, "The parallel approach to force/position control of robotic manipulators," *IEEE Transactions on Robotics and Automation*, vol. 9, no. 4, pp. 361–373, 1993.
- [29] S. Cassidy, "Livestats: Online statistical algorithms for python." [Online]. Available: <https://github.com/cxxr/LiveStats>
- [30] E. Olson, "Apriltag: A robust and flexible visual fiducial system," in *2011 IEEE International Conference on Robotics and Automation*, 2011, pp. 3400–3407.