

# Robust Toppling for Vacuum Suction Grasping

Christopher Correa<sup>1</sup>, Jeffrey Mahler<sup>1</sup>, Michael Danielczuk<sup>1</sup>, Ken Goldberg<sup>1,2</sup>

**Abstract**—When robust vacuum suction grasps are not accessible, toppling can change an object’s 3D pose to provide access to suction grasps. We extend prior toppling models by characterizing the robustness of toppling a 3D object specified by a triangular mesh, using Monte Carlo sampling to account for uncertainty in pose, friction coefficients, and push direction. The model estimates the resulting distribution of object poses following a topple action. We run 700 physical toppling experiments using an ABB YuMi and find that the model outperforms a baseline model by an absolute 26.9% when comparing the total variation distance between each model’s predicted probability distribution and the empirical distribution. We use the robust model as the state transition function in a Markov Decision Process (MDP) to plan optimal sequences of toppling actions to expose access to robust suction grasps. Data from 20,000 simulated rollouts suggest the policy can increase suction grasp reliability by 33.6%, computed using grasp analysis from the Dexterity Network (Dex-Net) 3.0. We generate a dataset of the predicted reliability of toppling at 1,257,000 candidate points on 189 3D meshes. Code and datasets can be found at <https://sites.google.com/view/toppling>

## I. INTRODUCTION

Grasping a wide variety of objects is essential in home de-cluttering, e-commerce warehouse manipulation, manufacturing, and service robotics. However, the ability to perceive and execute reliable grasps may be limited due to visibility of graspable points on the object, kinematic feasibility or environment collisions. Non-prehensile motions can be attempted before grasping to reorient objects into poses with increased visibility and access to robust grasps.

Previous research has shown that planar pushing can be used when parallel jaw grasps are not accessible [1], [2], [3], but planar pushing is less effective in increasing vacuum suction grasp accessibility [1]. In contrast, toppling, the act of pushing an object onto a new static resting pose with a robotic point contact, can be used to expose new surfaces for robust vacuum suction grasps.

Models of toppling for extruded 2D shapes have been developed in prior work by Lynch [4], [5] and Zhang *et al.* [6]. In this paper, we extend these models to 3D objects by predicting planar rotations and topple actions that are not perpendicular to the topple edge, and estimate topple reliability under uncertainty in contact point position, friction coefficients, and push direction. We also explore toppling policies to increase access to robust vacuum suction grasps, using suction grasp analysis from Dex-Net 3.0 [7].

<sup>1</sup>Department of Electrical Engineering and Computer Science

<sup>2</sup>Department of Industrial Engineering and Operations Research

<sup>1–2</sup>The AUTOLAB at UC Berkeley; Berkeley, CA 94720, USA  
{chris.correa, jmahler, mdanielczuk, goldberg}@berkeley.edu

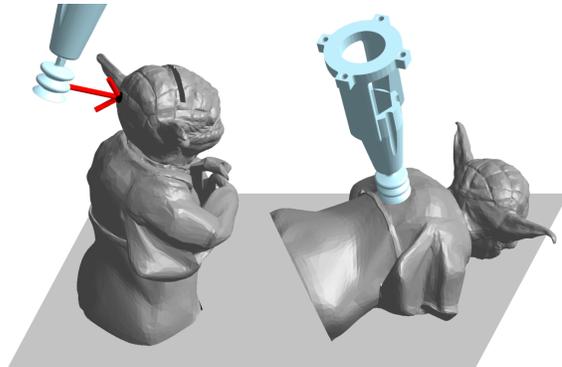


Fig. 1: Toppling the “Yoda Piggybank” object exposes access to a robust vacuum suction grasp.

This paper contributes:

- 1) A quasi-static model that estimates topple reliability for a 3D polyhedral mesh when pushed at a given point under uncertainty in contact point position, friction coefficients, and push direction.
- 2) A formulation of sequential toppling as an Markov Decision Process and a Value Iteration Policy to optimize predicted suction grasp reliability.
- 3) A dataset of 1,257,000 candidate toppling points on 189 3D CAD models in a total of 1257 stable poses, each labelled with associated toppling reliabilities.
- 4) Empirical data from 700 physical and 20,000 simulated experiments evaluating the proposed model and policy, respectively.

We compare performance of the model against baselines without the additions of planar rotations and topple reliability in 700 toppling attempts on an ABB YuMi robot.

Results suggest the model outperforms the toppling baseline by an absolute 26.9% when comparing the total variation of each predicted distribution with the empirical distribution. Results from 20,000 simulated experiments suggest the Value Iteration Toppling Policy can increase suction grasp reliability by an absolute 28.4% over a Max-Height Baseline.

## II. RELATED WORK

There is a wide range of work on feeding and orienting parts, either through planar rotations (i.e. the object stays in the same stable pose) or by reorienting objects into new stable resting poses. We focus on toppling as a means to increase access to suction grasps.

### A. Reorienting in the Horizontal Plane

Approaches to reorienting in the horizontal plane include pushing, squeezing with robotic grippers, and horizontal part

feeders. Mason and Lynch explored the use of a series of stable pushes to reorient parts [8], [9], [10]. Other work explores using grippers to reorient parts with squeezing [11], [12] and dexterous manipulation [13]. Part feeding can be used to orient parts on a conveyor belt [14], [15], [16]. While these approaches are successful in reorienting in the horizontal plane, we consider toppling to reorient parts into new stable resting poses, to increase suction grasp reliability.

### B. Reorienting to a New Stable Pose

One approach to reorient objects to new stable poses designs robotic grippers with multiple contact points designed specifically for the objects of interest, in order to execute several pinching motions while closing the gripper around the object [17], [18]. These grippers perform well with a single object, but do not generalize to the arbitrary objects on which robots are expected to operate. Pivoting, another reorienting motion primitive, is the act of loosely picking up objects so that they can rotate around the parallel jaw gripper’s axis due to gravity [19], [20]. Holladay *et al.* [21] introduce a pivot dynamics model and Hou *et al.* [22] combine pivoting and “rolling”, rotating the arm while stiffly gripping the object, to more efficiently reorient objects. However, these papers assume two-finger grasps are available on the object, in which case toppling to find robust grasps is unnecessary. We extend these papers by reorienting when grasps are not available.

### C. Toppling

Previous work has developed physics-based quasi-static models of toppling behavior on conveyor belts [4], [5], [6] and with robotic hands to move heavy objects [23]. Yamashita *et al.* [24] explore the use of multiple cooperative robots to topple objects. However, these models only predict motions of 2D extruded shapes, do not model object resistance to planar rotations, and only model pushes perpendicular to the toppling edge. Lee *et al.* [25] discretize all possible poses of the object, search for the contact points necessary to hold the object in each pose, and determine a trajectory to topple the object over. Learning-based approaches have been developed to predict object motions from just RGB images [26], [27], [28], and to learn a toppling policy from demonstrations [29]. These approaches do not require knowledge of the object geometry, but also only operate on 2D extruded shapes. In this paper, we present a model which predicts the toppling behavior of pushes in arbitrary directions for 3D objects beyond 2D extrusions.

## III. PROBLEM STATEMENT

Given a rigid polyhedral object, our goal is to find a sequence of topple actions for a point pusher to topple the object to a new pose with maximal suction grasp reliability. We formulate the toppling problem as a Markov Decision Process (MDP) and present a model for the reliability of a topple action under uncertainty to act as the state transition function in the MDP.

### A. Assumptions

We make the following assumptions:

- 1) Known object geometry (defined by a two-manifold triangular mesh) and center of mass.
- 2) The object rests in a stable pose [30] on a planar surface.
- 3) The object has a constant friction coefficient across the toppling edge and at each contact point. The choice of friction coefficient is described in Section VI-B.
- 4) The object’s pressure distribution is modelled as two point contacts on the endpoints of the toppling edge.
- 5) Quasi-static physics (i.e., inertial terms are negligible).
- 6) The forces and torques applied by the point pusher are consistent with the *point contact with friction* model [31].

### B. Definitions

1) *State*: Let  $\mathbf{x}_t = (\mathcal{O}, T_o) \in \mathcal{X}$  represent the state of the object at time  $t$ .  $\mathcal{O}$  represents the object’s geometrical properties, material properties, and center of mass, and  $T_o$  represents the stable resting pose of the object.

2) *Topple Action*: Let  $\mathbf{u}_t = (\mathbf{p}, \mathbf{q}) \in \mathcal{U}$  be the linear trajectory of the point pusher manipulator in 3D space between the start point  $\mathbf{p} = (x, y, z)$  and the end point  $\mathbf{q} = (x', y', z')$ .

3) *Topple Reliability*: Let  $\mathbb{P}[\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{u}_t]$  be the probability that the object topples into state  $\mathbf{x}_{t+1}$  when the robot executes action  $\mathbf{u}_t$ . The probability that the object topples to a different pose is  $\mathbb{P}[\mathbf{x}_{t+1} \neq \mathbf{x}_t | \mathbf{x}_t, \mathbf{u}_t]$ .

4) *Rewards*: The reward for executing each topple action is 0. After executing all topple actions, the robot gets a reward  $V_s(\mathbf{x}_T)$ , the predicted reliability of the best accessible suction grasp.

### C. Objective

Our ultimate goal is to develop a policy which plans a sequence of actions  $\{\mathbf{u}_i\}_{i=1}^T$  that maximizes the reliability of the best available grasp at time  $T$ . For the purpose of this analysis, the policy picks a fixed time horizon trajectory, and executes “no-op” actions  $\emptyset$  if it predicts no action will increase access to suction grasps.

$$\begin{aligned} \max_{\mathbf{u}_{1..T}} \quad & \mathbb{E}[V_s(\mathbf{x}_T)] \\ \text{s.t.} \quad & \mathbf{u}_t \in \mathcal{U} \cup \{\emptyset\} \end{aligned}$$

To achieve this goal, we develop a model to robustly estimate the transition distribution  $\mathbb{P}[\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{u}_t]$  to compute the expected reward for the toppling MDP under uncertainty in contact point position, friction coefficients, and push direction. This formulation differs from the standard MDP formulation in that the reward is only applicable for the final action.

## IV. TOPPLING MODEL

We develop a quasi-static models to estimate the topple reliability for a given contact point and push direction. We define the *base* of the object to be the convex hull of the object’s vertices touching the workspace. For each edge

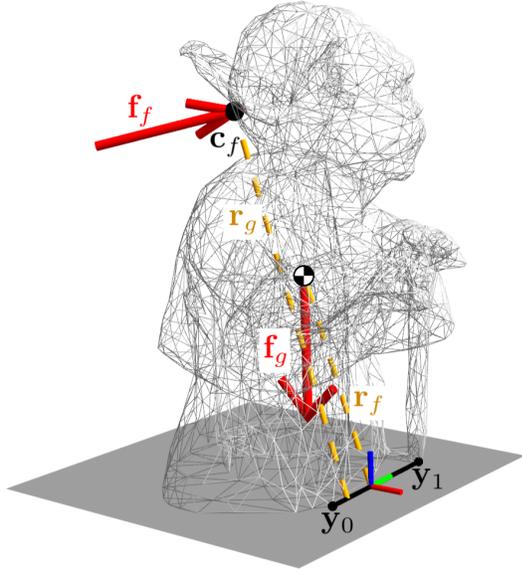


Fig. 2: Illustration of the forces acting on the object. The point pusher produces a wrench according to the force  $\mathbf{f}_f$  and the moment arm  $\mathbf{r}_f$ , defined between the contact point and its projection on the edge. Gravity produces a wrench according to the gravitational force  $\mathbf{f}_g$  and the moment arm  $\mathbf{r}_g$ , defined between the center of mass and its projection onto the edge.

on the base, we check the following three conditions, to determine if  $\mathbf{u}_t$  will topple the object.

- 1) **Contact Slip:** The point pusher does not slip on the object.
- 2) **Workspace Slip:** The contact between the toppling edge and the workspace can resist the wrench applied by the point pusher.
- 3) **Minimum Push Force:** The wrench applied by the point pusher has a magnitude large enough to rotate the object over the given edge.

If multiple edges satisfy these conditions, the model predicts the object will topple over the edge which requires the least force applied by the point pusher. We then use quasi-static analysis to predict the pose of the object after it topples over the predicted edge [30].

The model presented in this section expands on the model by Lynch *et al.* [5] by using the Friction Limit Surface [32] to predict planar rotations and Monte Carlo sampling to predict topple reliability under uncertainty.

#### A. Model Definitions

We formally describe the conditions for toppling with the following quantities, as shown in Figure 2 and 3.

1) *Toppling Edge:*  $\mathbf{s}_n = n\mathbf{y}_0 + (1-n)\mathbf{y}_1, \forall n \in [0, 1]$ , the parametrized line representing the toppling edge.

2) *Instantaneous Velocity Vectors:*  $\hat{\mathbf{v}}(\mathbf{s}_n)$ , the unit instantaneous velocity vector at each point along the toppling edge of the object.

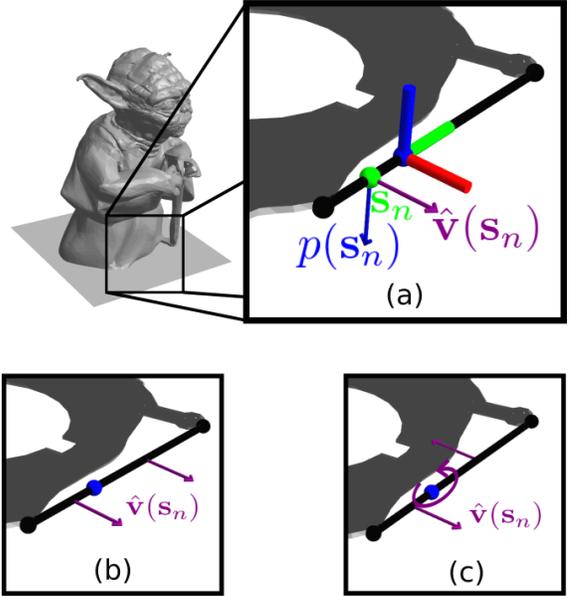


Fig. 3: (a) The topple action produces unit velocity vectors  $\hat{\mathbf{v}}(\mathbf{s}_n)$  across the toppling edge, which the object-workspace interaction resists via its pressure distribution  $p(\mathbf{s}_n)$ . The coordinate frame origin is at the projection of the center of mass onto the toppling edge, and the y-axis is parallel to the toppling edge. (b) and (c) demonstrate the dichotomy between the maximal  $\mathbf{f}_t$  (b) and  $\tau_z$  (c) the object can resist, depending on the direction of  $\hat{\mathbf{v}}(\mathbf{s}_n)$ .

3) *Pressure Distribution:*  $p(\mathbf{s}_n)$ , the pressure distribution of the object along the toppling edge, determined by the distribution of the object's mass and geometry along the toppling edge.

4) *Contact Point:*  $\mathbf{c}_f$ , the point where the point pusher contacts the object.

5) *Forces:*  $\mathbf{f}_f$  and  $\mathbf{f}_g$ , the forces applied to the object by the point pusher and gravity, respectively.

6) *Friction Coefficients:*  $\mu_f$  and  $\mu_T$ , the coefficients of friction of the object-point pusher and the object-workspace contacts, respectively.

7) *Friction Wrench:*  $\tau_z$  and  $\mathbf{f}_t$ , the torques around the z-axis and frictional force that the object-workspace contact exerts on the object, illustrated in Figure 3.

8) *Moment Arms:*  $\mathbf{r}_f$  and  $\mathbf{r}_g$ , the moment arms between the push contact point and the toppling edge, and the center of mass and the toppling edge, respectively.

#### B. Condition 1: Contact Slip

If the point pusher slips on the contact point, the proposed quasi-static model predicts it will not be able to apply the required force to topple the object. The point pusher will not slip on the object as long as it presses on the object in a direction within the object's friction cone at the point of contact.

$$\begin{aligned} [(\mathbf{f}_f)_{x'} \quad (\mathbf{f}_f)_{y'} \quad (\mathbf{f}_f)_{z'}]^T &= T_c \mathbf{f}_f \\ \sqrt{(\mathbf{f}_f)_{x'}^2 + (\mathbf{f}_f)_{y'}^2} &\leq \mu_f (\mathbf{f}_f)_{z'} \end{aligned}$$

$T_c \in SE(3)$  is the rigid transformation from the world coordinate system to the coordinate system defined at the contact point  $\mathbf{c}_f$ .  $(\mathbf{f}_f)_{x'}$ ,  $(\mathbf{f}_f)_{y'}$ ,  $(\mathbf{f}_f)_{z'}$  are the x, y, and z components of  $\mathbf{f}_f$  relative to a coordinate frame defined at the contact point.

### C. Condition 2: Workspace Slip

If the object slips either tangentially on the workspace or rotationally along an axis perpendicular to the workspace, the force applied at the contact point will result in the object moving in the plane instead of toppling. To determine whether a predicted push will topple the object or cause the object to slip on the workspace, we apply the friction limit surface model described in [32]. The friction limit surface describes the set of wrenches the object-workspace contact can resist before the object slips on the workspace. If the wrench applied by the point pusher is within this set of resistable wrenches, the object will topple.

We can integrate the infinitesimal forces applied by the object along the toppling edge to find the tangential force  $\mathbf{f}_t$  and torque around the z-axis  $\tau_z$  which the object-workspace contact can resist due to friction:

$$\mathbf{f}_t = - \int_{n=0}^1 \mu_T \hat{\mathbf{v}}(\mathbf{s}_n) p(\mathbf{s}_n) dn$$

$$\tau_z = \int_{n=0}^1 \mu_T \|\mathbf{s}_n \times \hat{\mathbf{v}}(\mathbf{s}_n)\|_2 p(\mathbf{s}_n) dn$$

The tangential force that the object-workspace contact can resist is maximized when  $\hat{\mathbf{v}}(\mathbf{s}_n)$  is constant across the toppling edge (Figure 3b), and the torque around the z-axis is maximized when the instantaneous velocities all produce torques in the same direction around the z-axis (Figure 3c), where  $F_n = mg$ .

$$\|\mathbf{f}_t\|_2 \leq \mu_T F_n$$

$$\tau_z \leq \int_{n=0}^1 \mu_T \|\mathbf{s}_n\|_2 p(\mathbf{s}_n) dn$$

$$\leq (\tau_z)_{max}$$

For our experiments, we assume the object's pressure distribution is defined as two point masses on the endpoints of the toppling edge, though other distributions can be substituted. In this case,  $(\tau_z)_{max} = \frac{\mu_T F_n}{2} [\mathbf{y}_1 + \mathbf{y}_0]$ .

We approximate the set of wrenches the friction contact can exert (the Friction Limit Surface) with an ellipse in wrench space [32], defined by the maximum tangential force and rotational torque:

$$\frac{\|\mathbf{f}_t\|_2^2}{(\mu_T F_n)^2} + \frac{\tau_z^2}{(\tau_z)_{max}^2} \leq 1$$

### D. Condition 3: Minimum Push Force

The applied force produces a torque around the toppling edge. If the magnitude of this torque is less than the magnitude of the torque caused by gravity, then the object will not topple. An object such as the one in Figure 2 will rotate at the toppling edge if  $\|\mathbf{r}_f \times \mathbf{f}'_f\|_2 \geq \|\mathbf{r}_g \times \mathbf{f}_g\|_2$ , where  $\mathbf{f}'_f$  is the component of  $\mathbf{f}_f$  orthogonal to the toppling edge.

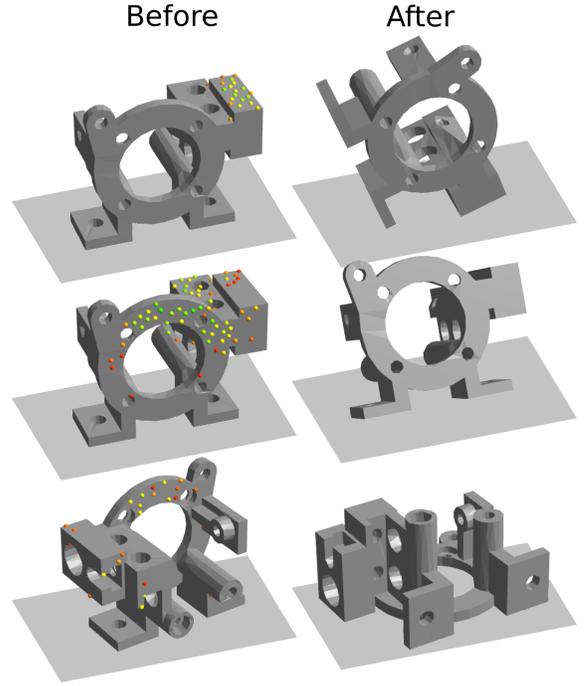


Fig. 4: Predicted reliability of each toppling action. The color of each contact point in the left column corresponds to the model's predicted probability that the object will reach the pose in the right column when pushed normal to the surface at each contact point. Contact points with 0% predicted reliability are omitted.

### E. Predicting Toppling Final Resting Pose

If the conditions in Sections IV-B, IV-C, and IV-D are satisfied, we predict the object's pose after toppling by rotating the object around the toppling edge until either:

- 1) The center of mass passes over the toppling edge (i.e. the apex of the topple). After this point, the object is free to fall unaided. We use the analysis from Goldberg *et al.* [30] to predict which face the object settles on after it falls.
- 2) The topple is blocked before it reaches the apex of the topple (i.e. some point other than the toppling edge touches the workspace). In this case, we return the object to its original pose.

All trials which fail to topple the object over any edge are mapped to the starting pose. Figure 4 shows the predicted poses of an object when toppled over three edges from various candidate points.

### F. Toppling Reliability

To account for uncertainty in gripper pose and object geometric and material properties, we introduce a robust metric for computing toppling probabilities. We consider 100 topple trials for each candidate contact point, with Gaussian noise added to the contact position, push direction, and friction coefficients:  $\mu_T \sim \mathcal{N}(0.43, 0.1)$ ,  $\mu_f \sim \mathcal{N}(0.98, 0.2)$ ,  $\mathbf{c}_f \sim \mathcal{N}(\mathbf{c}_f, 5.05 \cdot 10^{-4} \cdot I)$ , and  $\mathbf{f}_f \sim \mathcal{N}(\mathbf{f}_f, 0.055 \cdot I)$ . Section VI-B describes how we choose these noise distributions. These

trials serve as Monte-Carlo estimates for the true underlying toppling distribution. Figure 4 illustrates these Monte-Carlo estimates for each sampled contact point on the object, as well as the predicted final pose of the object if toppled at each contact point.

We sum the probabilities of toppling over edges which result in the same final pose. We label these probabilities the Toppling Distribution  $\mathbb{P}[\mathbf{x}_{t+1}|\mathbf{x}_t, \mathbf{u}_t]$ , defined in Section III-B, where  $\mathbf{x}_t$  is the current state,  $\mathbf{u}_t$  is the topple action, and  $\mathbf{x}_{t+1}$  is the predicted final pose of the object.

### G. Toppling Model Baselines

The presented model expands on prior work by predicting planar rotations, and predicting the toppling reliability under uncertainty. To evaluate the utility of these additions, we compare the model against three baselines:

- 1) **Baseline:** The baseline predicts whether a 3D object will topple over any edge or slide tangentially on the workspace, similar to prior work.
- 2) **Baseline + Rotations:** This model applies the friction limit surface to the baseline to additionally predict planar rotations.
- 3) **Baseline + Robustness:** This model adds Monte Carlo sampling to the baseline to predict the reliability of a topple action under uncertainty in pose, friction coefficients, and push direction.

## V. TOPPLING POLICIES

The goal of the policies in this section is to pick a toppling action to transition 3D objects into poses with improved grasp accessibility. We use the probabilities from the Toppling Model in Section IV to predict the pose of objects after each topple action  $\mathbf{u}_t$  and use the grasp analysis from Mahler *et al.* [7] to predict suction grasp accessibility. We compare the proposed policy against Max-Height and Greedy Baselines.

### A. Max-Height Baseline

To evaluate the utility of modelling the toppling behavior of objects in attempting to increase suction grasp access, we develop a baseline with no knowledge of the toppling transition probabilities. However, the policy does have knowledge of the predicted grasp accessibility of the object in each pose. This baseline picks the highest possible point on the object’s surface with a surface normal within  $15^\circ$  of the workspace and pushes at this point. The Max-Height policy executes “no-ops” if the object is in the pose with the highest reliability of accessible grasps.

### B. Greedy Baseline

The Greedy Baseline uses the proposed toppling model in Section IV to compute the topple probabilities for only the current time-step and picks the action which maximizes the expected suction grasp reliability  $V_s(\mathbf{x}_{t+1})$  immediately after the topple action:

$$\pi(\mathbf{x}_t) = \underset{\mathbf{u}_t}{\operatorname{argmax}} \left( \mathbb{E}_{\mathbf{x}_{t+1} \sim \mathbb{P}[\mathbf{x}_{t+1}|\mathbf{x}_t, \mathbf{u}_t]} [V_s(\mathbf{x}_{t+1})] \right)$$

The policy executes “no-ops” if  $V_s(\mathbf{x}_{t+1}) \leq V_s(\mathbf{x}_t) \forall \mathbf{x}_{t+1} \in \mathbb{P}[\mathbf{x}_{t+1}|\mathbf{x}_t, \mathbf{u}_t]$ . The Greedy Baseline benefits from fast computation times because it only computes the actions for the current time-step, but does not account for sub-optimal topples that could allow for better future topple actions.

### C. Value Iteration Policy

The Value Iteration Toppling Policy considers sub-optimal poses by assigning a value to each pose defined as the maximum suction grasp reliability for any pose reachable with linear topple actions, according to the transition probabilities generated by the model in Section IV.

The policy generates a graph of the toppling MDP, such as the one in Figure 5, where the nodes represent the object’s stable poses, and the edges represent topple actions. The policy assigns each action a value  $Q_T(\mathbf{x}_t, \mathbf{u}_t)$ , based on the discounted suction grasp reliability of future object states following action  $\mathbf{u}_t$ , using Value Iteration [33]:

$$Q_T(\mathbf{x}_t, \mathbf{u}_t) = \mathbb{E}_{\mathbf{x}_{t+1} \sim \mathbb{P}[\mathbf{x}_{t+1}|\mathbf{x}_t, \mathbf{u}_t]} [V_T(\mathbf{x}_{t+1})]$$

$$V_T(\mathbf{x}_t) = \max \left( V_s(\mathbf{x}_t), \gamma \max_{\mathbf{u}_t} Q_T(\mathbf{x}_t, \mathbf{u}_t) \right)$$

We choose  $\gamma = 0.95$  as our discount factor. After computing Value Iteration, the policy executes the action with the highest q-value  $Q_T(\mathbf{x}_t, \mathbf{u}_t)$ , or “no-ops” if  $V_T(\mathbf{x}_{t+1}) \leq V_T(\mathbf{x}_t) \forall \mathbf{x}_{t+1} \in \mathbb{P}[\mathbf{x}_{t+1}|\mathbf{x}_t, \mathbf{u}_t]$ .

## VI. RESULTS

### A. Toppling Dataset Generation

Our goal was to predict the toppling behavior of objects with dense samples of contact points across the object’s surface and with a high number of Monte-Carlo trials per sampled contact point. Since this can be computationally expensive, we pre-computed the topple probabilities and final poses for every stable pose of the desired objects, and store these values in a database. The pre-computation allowed us to quickly retrieve the toppling probabilities for a wide variety of objects.

We generated a dataset of 1000 candidate topple points on 189 objects in 1257 stable poses, for a total of 1,257,000 candidate topple points. We chose objects which satisfy the following conditions: 1) The object has a stable pose with high grasp reliability ( $> 50\%$ ) and a stable pose with low grasp reliability ( $< 25\%$ ) and 2) the object’s center of mass is higher in the high grasp reliability pose than in the low grasp reliability pose. We chose a complex set of objects whose convex hulls had an average of 1437 faces. The average time to compute the toppling probabilities and final poses per object stable pose is 40.47 seconds on a desktop computer running Ubuntu 16.04 with a 3.6 GHz Intel Core i7-6850k CPU.

### B. Physical Model Experiments

To compare the predictions of the model and the baselines in Section IV-G with empirical outcomes, we ran physical experiments of topple actions on the seven 3D printed objects

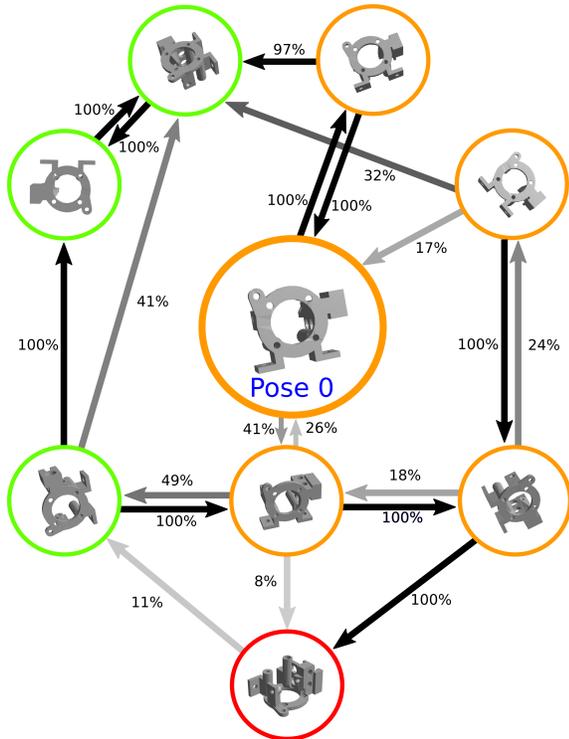


Fig. 5: Toppling Graph: The object in Figure 4 starts at Pose 0. All poses that are attainable via linear topple actions are shown. The border color corresponds to the reliability of the best available suction grasp of the object in the pose, and the edges are labelled with the probability of success for the best linear topple action.

shown in Figure 6. We executed 10 topple actions per object and repeated each topple action 10 times.

We placed the 3D printed objects in front of an ABB YuMi robot on a planar workspace. The robot acquired a 3D point cloud of the mesh using a Photoneo PhoXi depth sensor, and used the Super4PCS algorithm [34] to match the pose of the object in simulation to the pose of the object in front of the robot.

We then chose a topple action and recorded the pose of the object before and after the action. Since the predicted topple reliability is 0 at most contact points, we sampled actions with probabilities proportional to the topple probabilities from Section IV-F, where  $\epsilon = 10^{-4}$  to get topple actions with a broader range of reliabilities.

$$\mathbb{P}[\mathbf{u}_t] = \frac{\mathbb{P}[\mathbf{x}_{t+1} \neq \mathbf{x}_t | \mathbf{x}_t, \mathbf{u}_t] + \epsilon}{\sum_{\mathbf{u}'_t} \mathbb{P}[\mathbf{u}'_t]}$$

1) *Predicting Whether the Object Will Topple:* For each model and baseline, we considered 100 sets of noise distributions to perturb the contact point position, friction coefficients, and push direction. We performed 6-fold Cross Validation, each with a different object held out, so that the model comparison is not dependent on the choice of noise distributions. We averaged the performance on the held out object for each fold.

Topple Predictions	
Model	mAP
Baseline	0.741
Baseline + Rotations	0.752
Baseline + Robustness	0.852
<b>Robust Model</b>	<b>0.848</b>

TABLE I: Mean Average Precision (mAP) for each model’s prediction of toppling into any new pose, based on the empirical data collected. We perform k-fold cross validation, and average the mAP of each of the held-out folds.

Final Pose Predictions		
Model	TV	mAP
Baseline	0.424	0.412
Baseline + Rotations	0.494	0.381
Baseline + Robustness	0.247	0.568
<b>Robust Model</b>	<b>0.211</b>	<b>0.589</b>

TABLE II: Total variation distance (TV) and mean Average Precision (mAP) between each model and empirical toppling distributions. We perform 6-fold cross validation, and average the total variation and mAP of each of the held out folds. Model distributions with low total variation distance are consistent with the empirical distribution.

Table I shows how well each model and baseline was able to predict whether a given push will topple the object into a pose other than the start pose, when compared against empirical data.

2) *Predicting the Object Pose Distribution:* We also computed an empirical distribution of final poses and estimated distributions of final poses from the proposed model and the baseline models. We computed the total variation distance between each model’s predictions and the empirical distribution to quantify the accuracy of the proposed model’s predicted toppling behavior. We averaged the total variation distance across every state and action:

$$TV = \frac{1}{|\mathcal{X}||\mathcal{U}|} \sum_{\substack{\mathbf{x}_t \in \mathcal{X} \\ \mathbf{u}_t \in \mathcal{U}}} \sup_{\mathbf{x}_{t+1} \in \bar{\mathcal{X}}} \left| \hat{\mathbb{P}}[\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{u}_t] - \mathbb{P}[\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{u}_t] \right|$$

In the physical experiments,  $\mathcal{X}$  represents every initial state,  $\mathcal{U}$  represents the actions executed, and  $\bar{\mathcal{X}}$  represents the set of 10 final poses the object reaches.

Again, we consider 100 models, each with different noise distribution parameters and performed 6-fold Cross Validation. Table II shows both the mean Average Precision and Total Variation distance averaged over the held out sets for the best model in each of the six folds.

### C. Simulated Policy Experiments

To test the policies’ abilities to effectively use toppling to increase suction grasp accessibility, we ran all policies in a simulated environment and recorded the predicted reliability

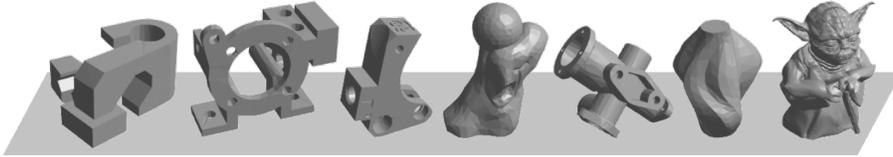


Fig. 6: The objects used in physical experiments.

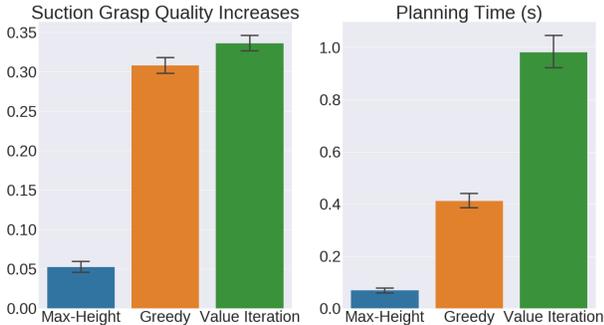


Fig. 7: Average difference in suction grasp reliability after each policy in Section V executes topple actions on a randomly-placed mesh in 3654 simulations. The planning times are the average planning time per action of each policy.

of accessible suction grasps before and after each policy finished execution. In each of the 20,000 trials, we dropped a random mesh into a simulated environment into a random pose, weighted according to its stable pose probability [30], used each policy to plan an action on the object, and used the Toppling Model to forward simulate the action. We drew a next state according to  $\mathbb{P}[\mathbf{x}_{t+1}|\mathbf{x}_t, \mathbf{u}_t]$ .

Of these 20,000 trials, we found toppling to be useful in 3654 trials, i.e., the object was not already dropped in the pose with the highest suction grasp access and there existed a topple action which was predicted to increase suction grasp reliability. The performance of each policy in these useful trials is illustrated in Figure 7. These results suggest that the situations in which toppling is useful is limited, but that toppling can be effective in uncovering robust grasps.

While pushing the object near the top of the object can marginally improve suction grasp access, the Greedy Baseline and Value Iteration Policy are able to increase the reliability of suction grasps by an absolute 25.5% and 28.4%, respectively, over the Max-Height Baseline. While the Value Iteration Policy achieves the highest increase in access, it comes at the cost of planning time, as it has to perform more database retrievals, and compute value iteration on the toppling graph at run-time.

Figure 5 exemplifies a situation in which the Value Iteration Policy outperforms the Greedy Baseline. The Greedy Baseline attempts a suction grasp at Pose 0, since the predicted grasp reliability of the object in all immediately reachable poses is the same as the predicted grasp reliability of the object in Pose 0. However, the Value Iteration Policy recognizes that toppling twice can lead to a state with 100% grasp reliability.

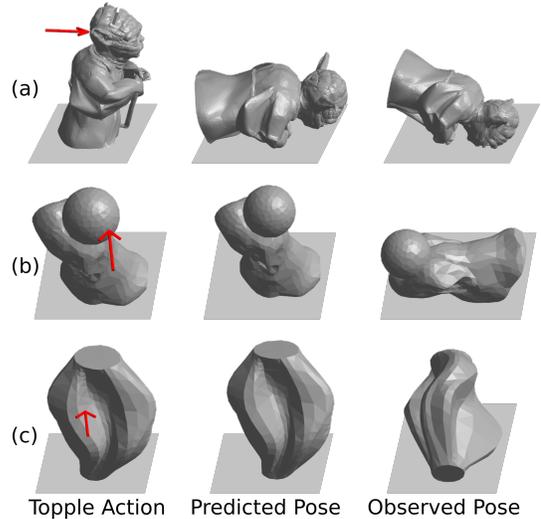


Fig. 8: Failure modes from physical experiments. (a) Momentum causes the object to roll further than predicted (b) the model correctly predicts the point pusher will slip, but the object still topples as the point pusher slips (c) the model correctly predicts the object will slip on the workspace, but the object first rotates then topples.

#### D. Toppling Model Failure Modes

To evaluate our model beyond our analysis, we examined examples in which the model significantly differs from physical outcomes. The three main limitations we identified with the proposed model are illustrated in Figure 8. In all three situations, the model is unable to predict the behavior because of the quasi-static assumption. In Figure 8a, the point pusher is able to topple the object, but momentum causes the object to roll past the predicted pose. The proposed model assumes quasi-statics during the rolling phase of the topple. In Figure 8b the topple action causes the object to initially slip on the workspace and then topples. The proposed model predicts a low topple reliability for this topple action, as it does not model dynamics. In a related scenario, the point pusher slips on the object surface in Figure 8c, but still topples the object. The model assumes the object will not topple if the point pusher slips on the object, so it incorrectly predicts this topple action fails.

## VII. DISCUSSION AND FUTURE WORK

In this paper, we present a model to robustly predict the toppling behavior of rigid 3D polyhedral objects to increase access to vacuum suction grasps. We compare the

proposed model to baselines without rotations and topple reliability predictions and evaluate the utility of the model in a Value Iteration Toppling Policy to increase suction grasp accessibility.

In future work, we would like to relax the model's assumptions, such as the approximation of the object's pressure distribution as two point masses at the endpoints of each edge. We would also like to explore non-linear topple motions. We also hope to explore more complex toppling policies such as multi-arm topple actions, such as the one presented by Yamashita *et al.* [24], and composite toppling and planar pushing policies. Finally, we would like to explore pre-computing a dataset of the topple Q-Values, to decrease planning time for the Value Iteration Policy.

## VIII. ACKNOWLEDGEMENTS

This research was performed at the AUTOLAB at UC Berkeley in affiliation with the Berkeley AI Research (BAIR) Lab, Berkeley Deep Drive (BDD), the Real-Time Intelligent Secure Execution (RISE) Lab, and the CITRIS "People and Robots" (CPAR) Initiative. The authors were supported in part by donations from Siemens, Google, Toyota Research Institute, Autodesk, Knapp, Honda, Intel, Comcast, Hewlett-Packard and by equipment grants from PhotoNeo, NVidia, and Intuitive Surgical. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Sponsors. We thank our colleagues who provided helpful feedback and suggestions, in particular Daniel Seita and Kate Sanders. We also thank Matthew Matl and David Gealy for their help developing the YuMi pushing tool.

## REFERENCES

- [1] M. Danielczuk, J. Mahler, C. Correa, and K. Goldberg, "Linear push policies to increase grasp access for robot bin picking," in *2018 IEEE 14th International Conference on Automation Science and Engineering (CASE)*, pp. 1249–1256, IEEE, 2018.
- [2] L. Y. Chang, S. S. Srinivasa, and N. S. Pollard, "Planning pre-grasp manipulation for transport tasks," in *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, pp. 2697–2704, IEEE, 2010.
- [3] M. Dogar and S. Srinivasa, "A framework for push-grasping in clutter," *Robotics: Science and systems VII*, vol. 1, 2011.
- [4] K. M. Lynch, "Inexpensive conveyor-based parts feeding," *Assembly Automation*, vol. 19, no. 3, pp. 209–215, 1999.
- [5] K. M. Lynch, "Toppling manipulation," in *ICRA*, pp. 2551–2557, 1999.
- [6] T. Zhang, G. Smith, R.-P. Berretty, M. Overmars, and K. Goldberg, "The toppling graph: Designing pin sequences for part feeding," in *Robotics and Automation, 2000. Proceedings. ICRA'00. IEEE International Conference on*, vol. 1, pp. 139–146, IEEE, 2000.
- [7] J. Mahler, M. Matl, X. Liu, A. Li, D. Gealy, and K. Goldberg, "Dexnet 3.0: Computing robust robot vacuum suction grasp targets in point clouds using a new analytic model and deep learning," in *IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2018.
- [8] K. M. Lynch, "The mechanics of fine manipulation by pushing," in *Robotics and Automation, 1992. Proceedings., 1992 IEEE International Conference on*, pp. 2269–2276, IEEE, 1992.
- [9] K. M. Lynch and M. T. Mason, "Controllability of pushing," in *Robotics and Automation, 1995. Proceedings., 1995 IEEE International Conference on*, vol. 1, pp. 112–119, IEEE, 1995.
- [10] K. M. Lynch and M. T. Mason, "Stable pushing: Mechanics, controllability, and planning," *The International Journal of Robotics Research*, vol. 15, no. 6, pp. 533–556, 1996.
- [11] R. C. Brost, *Planning robot grasping motions in the presence of uncertainty*. Carnegie-Mellon University, The Robotics Institute, 1985.
- [12] K. Y. Goldberg, "Orienting polygonal parts without sensors," *Algorithmica*, vol. 10, no. 2-4, pp. 201–225, 1993.
- [13] J. C. Trinkle and J. J. Hunter, "A framework for planning dexterous manipulation," in *Proceedings. 1991 IEEE International Conference on Robotics and Automation*, pp. 1245–1251, IEEE, 1991.
- [14] M. A. Peshkin and A. C. Sanderson, "Planning robotic manipulation strategies for workpieces that slide," *IEEE Journal on Robotics and Automation*, vol. 4, no. 5, pp. 524–531, 1988.
- [15] M. Brokowski, M. Peshkin, and K. Goldberg, "Optimal curved fences for part alignment on a belt," *Journal of Mechanical Design*, vol. 117, no. 1, pp. 27–35, 1995.
- [16] A. F. van der Stappen, R.-P. Berretty, K. Goldberg, and M. H. Overmars, "Geometry and part feeding," in *Sensor Based Intelligent Robots*, pp. 259–281, Springer, 2002.
- [17] M. T. Zhang and K. Goldberg, "Gripper point contacts for part alignment," *IEEE Transactions on Robotics and Automation*, vol. 18, no. 6, pp. 902–910, 2002.
- [18] M. T. Zhang and K. Goldberg, "Designing robot grippers: optimal edge contacts for part alignment," *Robotica*, vol. 25, no. 3, pp. 341–349, 2006.
- [19] B. Carlisle, K. Goldberg, A. Rao, and J. Wiegley, "A pivoting gripper for feeding industrial parts," in *ICRA*, pp. 1650–1655, 1994.
- [20] A. Rao, D. J. Kriegman, and K. Y. Goldberg, "Complete algorithms for feeding polyhedral parts using pivot grasps," *IEEE Transactions on Robotics and Automation*, vol. 12, no. 2, pp. 331–342, 1996.
- [21] A. Holladay, R. Paolini, and M. T. Mason, "A general framework for open-loop pivoting," in *Robotics and Automation (ICRA), 2015 IEEE International Conference on*, pp. 3675–3681, IEEE, 2015.
- [22] Y. Hou, Z. Jia, and M. T. Mason, "Fast planning for 3d any-pose-reorienting using pivoting," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1631–1638, IEEE, 2018.
- [23] Y. Aiyama, M. Inaba, and H. Inoue, "Pivoting: A new method of grasplless manipulation of object by robot fingers," in *Intelligent Robots and Systems '93, IROS'93. Proceedings of the 1993 IEEE/RSJ International Conference on*, vol. 1, pp. 136–143, IEEE, 1993.
- [24] A. Yamashita, T. Arai, J. Ota, and H. Asama, "Motion planning of multiple mobile robots for cooperative manipulation and transportation," *IEEE Transactions on Robotics and Automation*, vol. 19, no. 2, pp. 223–237, 2003.
- [25] G. Lee, T. Lozano-Pérez, and L. P. Kaelbling, "Hierarchical planning for multi-contact non-prehensile manipulation," in *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*, pp. 264–271, IEEE, 2015.
- [26] M. Kopicki, J. Wyatt, and R. Stolkin, "Prediction learning in robotic pushing manipulation," in *2009 International Conference on Advanced Robotics*, pp. 1–6, IEEE, 2009.
- [27] T. Mörwald, M. Kopicki, R. Stolkin, J. Wyatt, S. Zurek, M. Zillich, and M. Vincze, "Predicting the unobservable visual 3d tracking with a probabilistic motion model," in *Robotics and Automation (ICRA), 2011 IEEE International Conference on*, pp. 1849–1855, IEEE, 2011.
- [28] M. Kopicki, S. Zurek, R. Stolkin, T. Mörwald, and J. Wyatt, "Learning to predict how rigid objects behave under simple manipulation," in *2011 IEEE International Conference on Robotics and Automation*, pp. 5722–5729, IEEE, 2011.
- [29] N. S. Pollard and J. K. Hodgins, "Generalizing demonstrated manipulation tasks," in *Algorithmic Foundations of Robotics V*, pp. 523–539, Springer, 2004.
- [30] K. Goldberg, B. V. Mirtich, Y. Zhuang, J. Craig, B. R. Carlisle, and J. Canny, "Part pose statistics: Estimators and experiments," *IEEE Transactions on Robotics and Automation*, vol. 15, no. 5, pp. 849–857, 1999.
- [31] R. M. Murray, *A mathematical introduction to robotic manipulation*. CRC press, 2017.
- [32] K. L. Imin Kao and J. Burdick, "Contact modeling and manipulation," in *Springer Handbook of Robotics*, pp. 946–949, Springer, 2008.
- [33] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.
- [34] N. Mellado, D. Aiger, and N. J. Mitra, "Super 4pcs fast global pointcloud registration via smart indexing," in *Computer Graphics Forum*, vol. 33, pp. 205–215, Wiley Online Library, 2014.