

6DFC: Efficiently Planning Soft Non-Planar Area Contact Grasps using 6D Friction Cones

Jingyi Xu^{1,2}, Michael Danielczuk¹, Eckeard Steinbach², Ken Goldberg¹

Abstract—Analytic grasp planning algorithms typically approximate compliant contacts with soft point contact models to compute grasp quality, but these models are overly conservative and do not capture the full range of grasps available. While area contact models can reduce the number of false negatives predicted by point contact models, they have been restricted to a 3D analysis of the wrench applied at the contact and so are still overly conservative. We extend traditional 3D friction cones and present an efficient algorithm for calculating the 6D friction cone (6DFC) for a non-planar area contact between a compliant gripper and a rigid object. We introduce a novel sampling algorithm to find the 6D friction limit surface for a non-planar area contact and a linearization method for these ellipsoids that reduces the computation of 6DFC constraints to a quadratic program. We show that constraining the wrench applied at the contact in this way increases recall, a metric inversely related to the number of false negative predictions, by 17% and precision, a metric inversely related to the number of false positive predictions, by 2% over soft point contact models on results from 1500 physical grasps on 12 3D printed non-planar objects with an ABB YuMi robot. The 6DFC algorithm also achieves 6% higher recall with similar precision and 85x faster runtime than a previously proposed area contact model.

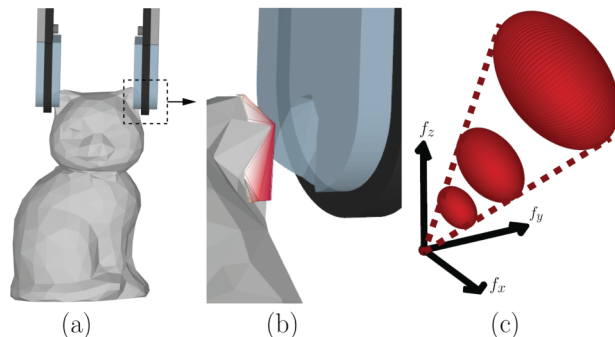


Fig. 1: (a) A non-planar area contact is created when a compliant gripper jaw surface contacts a non-planar object's surface. (b) An enlarged view of the deformed jaw and the contact profile obtained by the REACH model [11]. The non-planar contact area consists of triangles and the redder colors represent higher pressure due to larger deformation of the jaw pad at that point. (c) A projection of the 6D friction cone that constrains the wrenches that can be applied at the contact. Each ellipsoid represents a projection of the friction limit surface for a given gripper closing force and its center corresponds to the wrench created by the contact pressure.

I. INTRODUCTION

Grasping objects of arbitrary geometry with a robotic gripper or hand remains an active area of robotic research and has applications such as warehouse automation and manufacturing as well as household tasks such as decluttering. When exact or approximate 3D models of the objects to be grasped are available, analytic grasp planning models are often used to determine grasp quality given the object and contacts [31]. Among the quality metrics used, wrench-based metrics are popular due to their relative ease of computation and ability to model general and specific tasks [13, 17, 23].

These models rely on accurate estimation of the forces and torques applied at the contacts between the gripper and object to form the Grasp Wrench Space (GWS), or the set of wrenches that can be applied to the object by a set of gripper jaws. By computing the GWS, we can determine which wrenches the grasp can resist. In previous work, area contact models made up of multiple points or regions have been considered, but these either assume a planar contact area [9, 10, 16] or can be inefficient for fine surface geometries [11]. In this paper, we consider non-planar soft area contacts from a compliant gripper and formulate constraints for the wrenches that can be applied at each contact, as shown in Figure 1. We extend the 6D ellipsoidal model proposed by

Xu et al. [43] by more efficiently sampling wrenches on the 6D *friction limit surface* (FLS) without Finite Element Analysis (FEA) and combining the 6D FLS with the normal wrench imposed at the contact. We present a novel 6D friction cone (6DFC) that fully constrains the normal and frictional wrenches that can be applied at a contact by varying the grasp force. We then show how a linearized GWS can be formed from the 6D friction cones at each contact and can be directly used to evaluate grasp quality as a quadratic program. We believe that the 6DFC is particularly relevant to the family of compliant grippers both for grasp planning and for grasp robustness analysis when the exact grasp force is unknown. 6DFC subsumes both the friction cone created from planar contacts from rigid grippers and grasp analysis with known closing forces.

This paper provides three contributions:

- 1) A generalization of the 3D friction cone to a 6D friction cone for grasp reliability computation.
- 2) The 6DFC sampling algorithm for efficiently constructing the 6D friction limit surface and 6D friction cone for a non-planar area contact.
- 3) Results comparing 1500 physical grasps of 12 3D printed non-planar objects on an ABB YuMi robot with predictions from 4 algorithms that suggest the 6DFC algorithm can decrease false negatives by 17% over soft point contacts and 6% over a previously proposed area contact model.

¹The Autolab at University of California, Berkeley. ²Technical University of Munich, Chair of Media Technology. {jingyi_xu, mdanielczuk, goldberg}@berkeley.edu, eckeard.steinbach@tum.de

II. RELATED WORK

A. Contact Models

The contact between a robot gripper jaw and an object can be described as the jaw exerting a 6D wrench on the object with 3D force and 3D torque components, expressed in an object reference frame.

Among the many models introduced [4, 21, 36, 38], the most common models used in practice are point contact models with friction or soft point contact models [2, 28]. Under the Coulomb friction model, these contacts can exert forces in the plane tangent to the contact surface and a torsional moment (for soft point contacts) about the contact normal [9, 10, 19, 20, 22]. The tangential force and torsional moment that a soft point contact can exert can also be jointly constrained by the FLS [15, 25, 33], which Howe *et al.* [19] approximated with an ellipsoid for computational efficiency.

Planar area contact models are used to construct a 3D ellipsoidal FLS [10, 18, 40]; these area contact models jointly constrain forces in the contact plane and torque about the normal, but do not consider non-planar area contacts. Xu *et al.* [42] analyzed a 3D subspace of 6D friction constraints for curved contact areas and generalized the 3D FLS ellipsoid to 6D to model friction for non-planar surfaces. The 6D FLS is computed by densely sampling body twists and fitting the downsampled wrenches with an ellipsoid via convex optimization [43]. We also use a sampling method to form the 6D FLS, but without densely sampling over the entire space for increased efficiency. Danielczuk *et al.* [11] considered non-planar soft contacts by discretizing the area contact as a triangular mesh, but the REACH model did not consider the coupling between contact triangles, instead formulating constraints for each triangle independently. This formulation also resulted in slow runtime due to the number of contact wrench constraints in the grasp optimization problem scaling with the number of triangles in the contact area.

B. Grasp Analysis

Evaluating grasp quality requires determining the grasp's ability to constrain the motion of the object by applying forces and torques at the contacts to resist external disturbances without violating the frictional constraints at each contact [3]. To evaluate this quality, many metrics have been developed based on the grasp wrench space (GWS), or the space of wrenches that the contacts can apply to the object [37]. For example, force-closure ensures the contacts can resist any external wrenches with arbitrarily high grasp forces [26, 32, 34]. While metrics analyzing the entire GWS can be useful for unknown tasks [8, 37, 41], force-closure grasps are conservative for many tasks, such as lifting an object, as they require the grasp to be able to resist wrenches that will not be applied to the object during the task.

For unknown or complex tasks, formulating the 6D Task Wrench Space (TWS) ellipsoid can be complicated [26], but for simple tasks such as lifting an object, the task wrench space can be formulated as the wrenches applied to the surface of the object (object wrench space) or its center of

mass (mass wrench space) that must be resisted [5, 17, 27, 39]. Mahler *et al.* used a mass wrench space of the gravity wrench, corresponding to a grasp's ability to lift and hold the object [28, 29]. We use the same metric as Mahler *et al.* in that we characterize a grasp as successful if the grasp can resist the gravity wrench.

C. Grasp Wrench Space Formulation

A common approximation of the GWS is to find the convex hull of the union or Minkowski sum of the discretized friction cones at each contact [5, 13, 32]. Krug *et al.* noted that the independent contact bounds via the Minkowski sum more accurately represent fully-actuated grippers than that of the sum-bounded union [24]. However, for soft point contact models or area contact models that consider a 3D friction limit surface, only the union or Minkowski sum of the ellipsoids created by the maximal normal force is considered (e.g., the ellipsoids generated by normal forces between zero and the maximal normal force are ignored) [8, 9, 24]. In contrast, we formulate the GWS using independent contact bounds and the full 6D friction limit surface ellipsoid for each normal force that can be applied at the contact. The 6D friction cone captures both normal wrench and frictional wrench constraints.

D. Contact Wrench Cones

In addition to grasping applications, contact models are also commonly used to form contact wrench cones in legged robotics. The contact wrench cone (CWC) describes wrenches acting on the center of mass of a humanoid robot, determined by the friction cones at discrete contact points where the robot contacts the ground [1, 35]. Caron *et al.* [6] built the 6D CWC for non-coplanar contacts using surface contact wrenches. Carpentier and Mansard [7] approximated the CWC with a 6D cone. While the algorithms in [6, 7] consider 3D wrench constraints at each point or planar area contact, the proposed 6DFC models 6D wrench constraints for each non-planar area contact from compliant grippers.

III. PROBLEM STATEMENT

We consider the problem of predicting grasp reliability, or probability of grasp success, by building a 6D friction cone of each non-planar area contact from compliant jaws.

A. Assumptions

We make the following assumptions:

- 1) Quasi-static physics (inertial terms are negligible) and Coulomb friction with constant friction coefficient μ over the contact area.
- 2) Objects to be grasped are rigid with known geometry.
- 3) The gripper has known geometry and two parallel jaws, each with a linear-elastic material at the tips.
- 4) Both gripper jaws make contact simultaneously.
- 5) Force is applied normally to the object surface at each point within the contact area.

B. Definitions

We define a state \mathbf{x} that contains a single object \mathcal{O} (including its geometric, material, and frictional properties) and its pose $T_{\mathcal{O}}$. We also define a parallel-jaw grasp action \mathbf{u} parametrized by a nominal grasp center $\mathbf{p} \in \mathbb{R}^3$ and an angle $\varphi \in \mathcal{S}^3$. The jaws close with force of magnitude f_C around the grasp center and are oriented according to the grasp angle. A binary reward function R describes the grasp success, where $R = 1$ if the grasp lifts the object (meaning the contacts resist the wrench applied to the object by gravity) and $R = 0$ otherwise. To account for uncertainty in the state as well as imprecision in control of the robot, we consider a grasp reliability distribution $Q(\mathbf{x}, \mathbf{u}) = \mathbb{P}(R|\mathbf{x}, \mathbf{u})$ that describes the probability of grasp success for a state \mathbf{x} and action \mathbf{u} [29]. We evaluate reliability in simulated environments by perturbing object pose, mass, and frictional properties and in physical experiments by repeating the same nominal grasp multiple times under uncertainty in the robot grasp pose accuracy and object registration. We approximate $Q(\mathbf{x}, \mathbf{u})$ with the sample mean of N Monte Carlo samples: $Q(\mathbf{x}, \mathbf{u}) = \frac{1}{N} \sum_{i=1}^N R_i(\mathbf{x}, \mathbf{u})$ [30].

C. Objective

We evaluate the accuracy on a dataset of physical grasp experiments. Specifically, we seek to maximize average precision (AP), defined as the mean of precision values at each recall threshold weighted by the difference in recall thresholds, which measures the area under the precision-recall curve. Additionally, we seek to maximize the average recall (AR) on the dataset for a given AP, as this metric indicates the ability to correctly predict positive grasps. Both metrics measure binary classification performance and are commonly used in computer vision for unbalanced datasets [12]. A combination of high AP and AR indicates that the algorithm predicts few false positives while also predicting few false negatives. We select AP and AR as metrics to evaluate the algorithm in reducing false positives and negatives.

IV. NON-PLANAR AREA CONTACT CONSTRAINTS

A. Background

We define the 6D wrench that can be applied at the contact as $\mathbf{w} = [f_x \ f_y \ f_z \ \tau_x \ \tau_y \ \tau_z]^T$, which corresponds to a force and torque that can be applied about each axis defined in the contact frame. The set of these contact wrenches (e.g., all wrenches that can be applied at the contact) forms the *contact wrench space (CWS)*. Additionally, we define a maximum closing force $f_{C,max} \in \mathbb{R}_+$ with which the jaw can contact the object such that $0 \leq f_C \leq f_{C,max}$.

For frictional point contacts using Coulomb friction, we can define the axis for the contact such that the z -axis is aligned with the negative surface normal and the x and y axes are in the plane. Then, $f_z = f_C$ and the CWS is defined by [33]:

$$\mathcal{W} = \{\mathbf{w} \in \mathbb{R}^6 \mid 0 \leq f_z \leq f_{C,max}, \sqrt{f_x^2 + f_y^2} \leq \mu f_z, \tau_x = \tau_y = \tau_z = 0\} \quad (\text{IV.1})$$

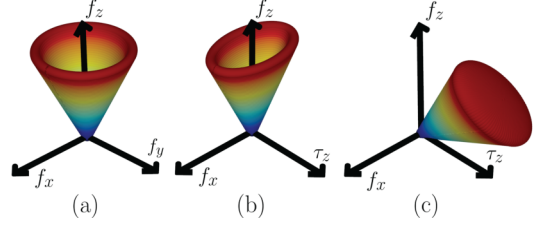


Fig. 2: In each image, warmer colors indicate increasing gripper closing force magnitude f_C . (a) For frictional point contacts, the friction cone is circular and extends along the f_z axis. (b) For soft point contacts, the cone is 4D and elliptical; f_x , f_y and τ_z are jointly constrained at each value of f_z by an ellipsoid approximating the friction limit surface. Here we show a projection of the 4D cone. (c) In the non-planar area contact case, the cone is 6D and may not align with any axis, resulting in non-axis-aligned 6D ellipsoids approximating the friction limit surface at each value of the closing force. Here we show a projection of the 6D cone.

Equation IV.1 limits the forces that can be applied at the contact with a 3D friction cone, as shown in Figure 2(a), where the tangential forces f_x and f_y that can be applied increase with f_z .

This idea can be extended to the case of soft point contacts, where the jaw can also apply a torque τ_z around the z -axis. In this case, the torque τ_z and the tangential forces f_x and f_y are jointly constrained by the so-called friction limit surface (FLS) [15, 25, 33], which can be approximated as a 3D ellipsoid [19]:

$$\mathcal{W} = \left\{ \mathbf{w} \in \mathbb{R}^6 \mid \frac{f_x^2 + f_y^2}{\mu^2} + \frac{\tau_z^2}{\gamma^2} \leq f_{C,max}, \tau_x = \tau_y = 0 \right\}$$

In previous work, the largest ellipsoid is used to constrain the wrench applied at the soft point contact [9, 24]. Thus, the CWS is a 3D ellipsoid. We note that this formulation can easily be expressed in the same way as the previous case by replacing $f_{C,max}$ with f_z and independently constraining $0 \leq f_z \leq f_{C,max}$, resulting in a 4D cone that extends along the f_z axis and 3D ellipsoidal cross sections at each value of f_C as shown in Figure 2(b).

In both cases, the CWS can be discretized into k samples:

$$\mathbf{W} = [\mathbf{w}_1, \dots, \mathbf{w}_k]$$

and the GWS for n contacts is formed either through the convex hull of the union or Minkowski sum of the n discretized contact wrench spaces [13].

B. Friction Cones in 6D

In this section, we generalize the friction cone to the 6D space. When considering a non-planar area contact, the friction wrenches that can be applied at the contact and the wrench impressed by the normal pressure, defined as the normal wrench, are 6D, as the force and torque are both 3D. Xu et al. suggest that the friction wrenches are bounded with a 6D ellipsoid centered at the origin [43]. We propose that the total wrench applied at a contact with a grasp force f_C can be modeled with a 6D ellipsoid that is centered at the 6D normal wrench, so that varying the value of f_C results in a

6D friction cone, whose center lies along the vector $\mathbf{f}_N \in \mathbb{R}^6$ and has 6D frictional ellipsoids as contours for each value of f_C , similar to the one shown in Figure 1. We can express this cone as:

$$\begin{aligned} (\mathbf{w} - f_C \mathbf{f}_N)^T A (\mathbf{w} - f_C \mathbf{f}_N) &\leq f_C^2, \\ 0 &\leq f_C \leq f_{C,max} \end{aligned} \quad (\text{IV.2})$$

Figure 2(c) shows a 3D projection of the 6D friction cone that is produced from a non-planar area contact.

C. Finding 6D Friction Limit Surface Cone Constraints

For the 3D frictional point contact or 4D soft point contact case discussed above, the ellipsoid that approximates the 2D or 3D friction limit surface can be easily found, since it is axis-aligned. Then, by finding the maximum values for f_x , f_y , and τ_z (in the soft point contact case), we can construct the ellipsoid. However, in the 6D case, the ellipsoid that approximates the friction limit surface may be rotated, since when one friction wrench component reaches its maximum value, the other dimensions might not be zero. This scenario can occur from an asymmetric pressure distribution over the contact area or from the geometry of the contact area itself [18]. Thus, we find an equation that describes the ellipsoid by explicitly sampling its surface, then fitting an ellipsoid to the sampled wrenches. We describe this process in detail in the following sections.

1) *Computation of Friction Wrench Samples:* First, we extract the contact area patch on the object using the same method as in [11]: the constructive solid geometry intersection between the gripper pad at its maximum depth of deformation and the object is the contact area patch. This resulting contact patch can be discretized into m triangles. However, unlike [11], which treats each triangle in the discretized patch as a separate planar contact, we find the 6D friction limit surface for the entire patch.

To find the 6D friction limit surface for the patch contact, we sample points that lie on the surface, similar to Xu et al. [43]. However, unlike [43], we do not evenly sample axes of rotation at various distances from the contact, since finding extreme points on the surface requires sampling axes at both small distances from the contact (to maximize torques) and at infinite distances (to maximize forces) [18]. Thus, we sample wrenches with large torques or large forces separately.

First, we sample k_1 tuples consisting of a unit axis of rotation ω_i and a 3D center point \mathbf{c}_i , where ω_i is sampled uniformly from the unit sphere and \mathbf{c}_i is sampled randomly within a radius r of the pressure center of the contact patch. For each tuple, we find the instantaneous unit velocity vector in the plane normal to ω_i at each triangle. We project the unit velocity vector onto the surface of each triangle to find the projected velocity vector $\hat{\mathbf{v}}_{i,t}$ at the t -th triangle in the contact area patch. For wrenches with maximum forces, we uniformly sample k_2 unit velocities \mathbf{v}_i , similar to sampling ω_i , and project them onto each triangle plane to find $\hat{\mathbf{v}}_{i,t}$.

Then, we can calculate the magnitude of force that can be resisted in the triangle plane, depending on the force applied normal to that triangle, denoted as f_{N_i} . The force that can

be applied in the frame of the triangle is $-\mu f_{N_i} \hat{\mathbf{v}}_{i,t}$. By transforming each of these forces into the contact patch frame using the triangle's adjoint matrix Ad_t , we can then form the full 6D wrench that can be applied at the contact:

$$\mathbf{w}_i = - \sum_{t=1}^m \mu Ad_t f_{N_t} \hat{\mathbf{v}}_{i,t}$$

2) *Fitting the 6D Ellipsoid:* Given the k samples $\{\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_k\}$, we then fit a 6D ellipsoid to the data using linear least squares. To fit the ellipsoid, we find the positive semidefinite matrix A^* :

$$A^* = \arg \min_A \sum_{i=1}^k \|\mathbf{w}_i^T A \mathbf{w}_i - 1\|_2^2$$

We can solve this equation exactly using least squares since it is linear in A . Note that this method of solving for A^* is not guaranteed to result in an ellipse if the matrix is not positive semidefinite (in fact, a hyperboloid could also be returned), so in practice, we first determine the dimensionality of the data by using principal component analysis (e.g., if the contact area is planar, then the ellipsoid will only be 3D), then fit an ellipsoid of the determined dimensionality to the data rotated to the PCA frame. We fill out the remaining dimensions with axes lengths of a small value $\epsilon > 0$, then rotate the fitted 6D ellipsoid back to the original non-PCA frame. Thus, this case subsumes the soft point contact case above; if the contact is planar, it reduces to the case where the normal wrench is along f_z and the FLS is 3D.

3) *Linearizing the Ellipsoidal Constraints:* Although the ellipsoidal constraint is itself a convex constraint, quadratic programs with linear constraints can be solved more efficiently. We formulate a quadratic program to determine if a grasp resists gravity. We approximate the ellipsoid with a set of linear constraints for computational efficiency.

To find these constraints, we first resample the ellipsoid. For many contacts, the initial random sampling process results in data that is not evenly sampled due to the geometry of the contact surfaces; thus, we resample points evenly over the surface of the ellipsoid to better approximate it. Then, for each resampled point \mathbf{x}_i on the ellipsoid, its outward-facing normal vector is given by $A^* \mathbf{x}_i$. This normal vector defines the hyperplane tangent to the ellipsoid at \mathbf{x}_i ; by finding many of these hyperplanes, we can construct a set of linear constraints that approximate the ellipsoid. The constraint set is of the form:

$$\mathcal{F} = \{\mathbf{z} \in \mathbb{R}^6 : \mathbf{z}^T A^* \mathbf{x}_i \leq \mathbf{x}_i^T A^* \mathbf{x}_i = 1, \forall i\}$$

4) *Formulating Cone Constraints:* Once the constraints are found for the unit closing force, we can shift the planar constraints along the normal force vector and scale them by an arbitrary closing force:

$$\mathcal{F} = \left\{ \begin{array}{l} \mathbf{z} \in \mathbb{R}^6 : \mathbf{z}^T A^* \mathbf{x}_i - f_C (1 + \mathbf{f}_N^T A^* \mathbf{x}_i) \leq 0, \forall i \\ f_C \in \mathbb{R} : 0 \leq f_C \leq f_{C,max} \end{array} \right\}$$

Note that although the ellipsoidal constraint in Equation IV.2 is quadratic, by first approximating the unit closing force

ellipse with a set of planes, we can take advantage of the fact that the ellipse grows linearly with increasing closing force to express the planar constraints linearly with f_C . This set of linear constraints can be used directly in a quadratic program to determine the ability of contacts that apply wrenches $\tilde{\mathbf{z}} = [\mathbf{z}_1 \ f_{C,1} \ \dots \ \mathbf{z}_n \ f_{C,n}]^T$ to resist the gravity wrench \mathbf{t} :

$$\min_{\tilde{\mathbf{z}}} \|G\tilde{\mathbf{z}} + \mathbf{t}\|_2^2 \quad \text{s.t.} \quad F\tilde{\mathbf{z}} \leq \mathbf{h} \quad (\text{IV.3})$$

Here, F and \mathbf{h} are generated by concatenating the constraints \mathcal{F} for each contact. The matrix $G \in \mathbb{R}^{6 \times 7n}$ transforms $\tilde{\mathbf{z}}$ from the contact frames to the object frame.

V. EXPERIMENTS

We evaluate the 6DFC algorithm proposed in Section IV-C against three baseline algorithms that generate frictional and normal wrench constraints for each contact to determine which most accurately can predict grasp success.

A. Baseline Algorithms

1) *Soft Point Contact*: As described in Section IV-A, this algorithm constrains the wrench applied at the contact through a 3D ellipsoid representing the friction limit surface corresponding to the maximum normal force that can be applied at the contact. We linearize these constraints to solve the quadratic program in Equation IV.3.

2) *REACH*: The area contact model proposed by Danielczuk *et al.* [11] discretizes the contact area into a triangular mesh and develops constraints of a similar form to the soft point contact model for each triangle. It can be modified to consider only a maximum number of triangles for each contact (e.g., the 10 largest triangles), which increases computational efficiency without significantly reducing accuracy.

3) *Maximum Force Ellipsoid (MFE)*: This algorithm constructs the 6D FLS ellipsoid as described in Section IV-C, but constrains the wrench applied at the contact to the ellipsoid generated by the maximum closing force. This method is similar to that proposed by Xu *et al.* [43].

B. Soft Non-Planar Area-Contact Physical Robot Grasps

To evaluate the precision and recall of 6DFC and the baseline algorithms, we use a subset of 1,500 grasps on 12 3D-printed objects from the Soft Area-Contact Physical Robot Grasp Dataset, collected on a physical ABB YuMi robot with a compliant parallel-jaw gripper [11]. The exerted forces are not actively controlled, as the gripper does not have force sensors. The subset of objects chosen from the dataset have non-planar contacts for all grasps. Note that the assumptions in Section III do not necessarily hold for the physical grasps that the algorithms are tested on (e.g., the jaws do not always make contact simultaneously, quasi-static assumptions do not hold).

We measure both average precision (AP) and average recall (AR) for each object using the dataset’s ground-truth physical grasp labels and each algorithm’s predictions. We consider a physical grasp to be successful if it lifts and transports an object from a bin to a receptacle. We measure performance with mean average precision (mAP)

Algorithm	Precision	Recall	Runtime (ms/grasp)
Point	0.80 ± 0.01	0.50 ± 0.01	13.0 ± 2.7
REACH	0.83 ± 0.01	0.61 ± 0.01	207.1 ± 35.3
MFE	0.83 ± 0.02	0.60 ± 0.01	247.9 ± 9.6
6DFC	0.82 ± 0.01	0.67 ± 0.01	251.6 ± 14.2

TABLE I: Mean average precision (mAP) and mean average recall (mAR) and their standard deviations for each algorithm’s predictions of grasp quality, for 5 runs of the 1,500 grasps collected on the physical robot. 6DFC outperforms the point, REACH, and MFE algorithms by 17%, 6%, and 7% in mAR, respectively, suggesting that cone constraints can reduce false negatives on objects with non-planar surfaces.

and mean average recall (mAR), which are the AP and AR of the algorithm averaged over all objects to account for discrepancies in the number of successful grasps for each object. We also measure runtime per grasp computation for each model on the an Ubuntu 16.04 machine with a 12-core 3.7 GHz i7-8700k processor. For each algorithm, parameters such as the friction coefficient, elasticity coefficient, and robustness sample standard deviation were chosen using leave-one-out cross validation.

The results for each algorithm are shown in Table I. These results suggest that formulating the 6DFC algorithm can increase the number of successful grasps on the physical system that can be recalled by as much as 17% over existing algorithms while maintaining a similar number of false positives predicted. We find that grasps found only by 6DFC are often of the kind shown in Figure 1, where the contact area is small with high surface curvature, and the grasp may include slight dynamic effects such as a bowing of the jaws. We hypothesize that allowing for reduced closing forces in this scenario could more accurately model the changing contact profile that occurs.

C. Grasp Planning Results

We also evaluate the reliability of each algorithm as part of a grasp planning policy via a second experiment. We hypothesize that our algorithm can find grasps in scenarios where a point contact algorithm would not return any high-quality grasps due to predicted collisions with other objects, environmental constraints, or motion-planning constraints. To test this hypothesis, we place each of the 12 objects in their 3 most probable stable poses [14] and attempt the top 3 grasps that each algorithm labels as the highest reliability grasps for that stable pose. We remove nearby grasps so that the algorithm cannot choose 3 similar grasps, simulating other objects blocking the grasp from being executed. If multiple grasps are labeled with the same probability, one of them is chosen at random.

We evaluated 875 unique grasps in total for 32 stable poses of the 12 objects. Both algorithms found successful grasps in at least one stable pose that the other could not (Figure 4 shows two examples), but overall the point grasps surprisingly succeeded more often (339 successes compared to 272 for the 6DFC algorithm). This result may be due to an inflation in grasp quality in the 6DFC algorithm; we

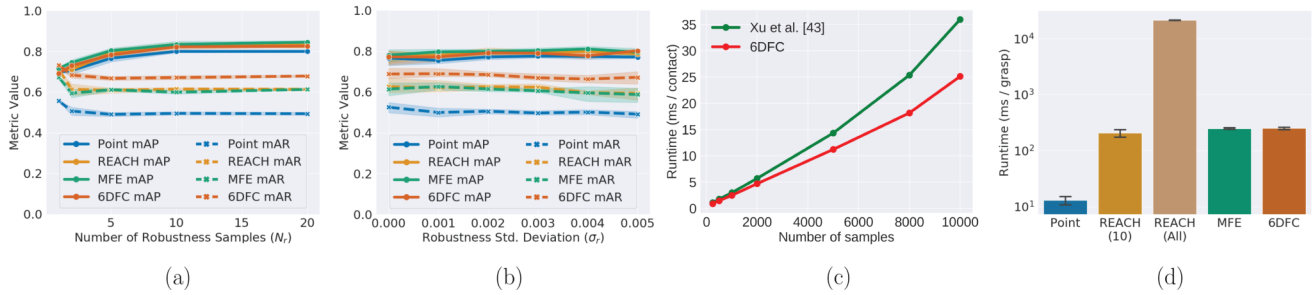


Fig. 3: Mean average precision (mAP) and mean average recall (mAR) for each algorithm as a function of (a) the number of robustness samples N_r , (b) the robustness standard deviation σ_r with error bars showing the standard deviation of 5 runs of each algorithm with the given parameter for the dataset of 1500 physical grasps. Adding robustness samples and spreading the samples increases mAP up to 10 samples and $\sigma_r = 0.003$ m. After this point, mAR continues to decline, but mAP remains constant or decreases (in the case of σ_r). (c) Runtime analysis of the sampling algorithms with different numbers of wrenches, averaged over 1000 runs. 6DFC is up to 30% faster than [43]. Timing analysis for each algorithm is shown in (d), averaged over 6000 grasps. The numbers below the REACH algorithm indicate the number of triangles considered in computation. The 6DFC algorithm is of similar time to the other algorithms that analyze area contacts, but is 85x faster when analyzing the same number of triangles as REACH.

found that many grasps that had large contact areas were rated highly but did not succeed, as shown in the bottom row of Figure 4. We will investigate this discrepancy further in future work.

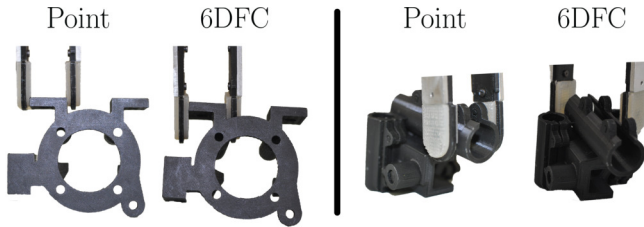


Fig. 4: Left: An example where the 6DFC algorithm finds a robust grasp when the Point algorithm does not. Right: An example where the Point algorithm finds a robust grasp when the 6DFC algorithm does not. In the first case, the thin part of the object results in a low-quality prediction for the point contact algorithm, whereas in the second case, the large contact area produces a false positive prediction from the 6DFC algorithm.

D. Sensitivity Analysis

As each of the algorithms contains several parameters such as friction and elasticity coefficients, robustness parameters, we include an analysis of each algorithm’s sensitivity to a subset of these parameters.

1) *Effect of Robustness Parameters:* All of the algorithms benefit from “robustness”, or sampling grasp poses around the nominal pose and averaging the predicted grasp qualities. We sample a random 3D translation from a zero-mean Gaussian with variance σ_r^2 and a uniformly random 3D rotation with angle θ proportional to σ_r^2 and apply them to the nominal grasp. Figure 3(a-b) shows sensitivity of the algorithms to the number of samples N_r and the standard deviation σ_r . At $N_r = 10$ and $\sigma_r = 0.003$, mAP and mAR are maximized. Adding more samples after this point does not increase mAR and may moderately increase mAP but does result in the standard deviation (shaded area) decreasing. Increasing σ_r results in lower mAR as sampled grasps no longer resemble the nominal grasp.

2) *Effect of Number of Samples and Contact Triangles:* Figure 3(c) shows the sampling runtime with different number of wrenches and suggests that the proposed algorithm described in Section IV-C.1 is up to 30% faster than [43]. The runtime of REACH strongly depends on the number of triangles in the area contact analyzed. The 6DFC algorithm runtime scales with the number of contacts as opposed to the number of triangles in each contact, resulting in the 85x faster runtime on the same surface patch mesh, as shown in Figure 3(d).

VI. DISCUSSION AND FUTURE WORK

We present 6DFC, an algorithm that generalizes 3D friction cones to non-planar soft area contacts, constraining both the contact normal and frictional wrenches. We sample the 6D friction cone using projections of instantaneous velocity vectors onto each triangle of the surface patch mesh. We show 6DFC outperforms point contact and area contact model baselines on a dataset of physical grasps.

We note that one reason for predicted false positives is that the 6DFC algorithm allows different contact forces of each jaw, which is not feasible with the current physical setup. In future work, we plan to evaluate 6DFC with a three-jaw gripper that allows fully controllable forces to further investigate this effect. We also plan to relax the simultaneous jaw contact assumption as part of a dynamic contact analysis.

VII. ACKNOWLEDGMENTS

This research was performed at the AUTOLAB at UC Berkeley in affiliation with the Berkeley AI Research (BAIR) Lab. The authors were supported in part by NSF National Robotics Initiative Award 1734633, National Science Foundation Graduate Research Fellowship Program under Grant No. 1752814, and by donations from Google, Siemens, Amazon Robotics, Toyota Research Institute, Autodesk, ABB, Samsung, Knapp, Loccioni, Honda, Intel, Comcast, Cisco, Hewlett-Packard and by equipment grants from PhotoNeo and NVIDIA. 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, especially Jeff Ichnowski and Matt Matl.

REFERENCES

- [1] D. J. Balkcom and J. C. Trinkle, "Computing wrench cones for planar rigid body contact tasks," *Int. Journal of Robotics Research (IJRR)*, vol. 21, no. 12, pp. 1053–1066, 2002.
- [2] F. Barbagli, A. Frisoli, K. Salisbury, and M. Bergamasco, "Simulating human fingers: A soft finger proxy model and algorithm," in *Proc. IEEE Int. S. on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, 2004.
- [3] A. Bicchi, "On the closure properties of robotic grasping," *Int. Journal of Robotics Research (IJRR)*, vol. 14, no. 4, pp. 319–334, 1995.
- [4] A. Bicchi and V. Kumar, "Robotic grasping and contact: A review," in *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2000.
- [5] C. Borst, M. Fischer, and G. Hirzinger, "Grasp planning: How to choose a suitable task wrench space," in *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, IEEE, vol. 1, 2004, pp. 319–325.
- [6] S. Caron, Q.-C. Pham, and Y. Nakamura, "Leveraging cone double description for multi-contact stability of humanoid with applications to statics and dynamics," in *Proc. Robotics: Science and Systems (RSS)*, Jul. 2015.
- [7] J. Carpentier and N. Mansard, "Multicontact locomotion of legged robots," *IEEE Trans. Robotics*, vol. 34, no. 6, pp. 1441–1460, 2018.
- [8] M. Ciocarlie, H. Dang, J. Lukos, M. Santello, and P. Allen, "Functional analysis of finger contact locations during grasping," in *Proc. IEEE Eurohaptics Conf. and S. on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, IEEE, 2009.
- [9] M. Ciocarlie, C. Lackner, and P. Allen, "Soft finger model with adaptive contact geometry for grasping and manipulation tasks," in *Proc. IEEE Eurohaptics Conf. and S. on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, 2007.
- [10] M. Ciocarlie, A. Miller, and P. Allen, "Grasp analysis using deformable fingers," in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2005.
- [11] M. Danielczuk, J. Xu, J. Mahler, M. Matl, N. Chentanez, and K. Goldberg, "Reach: Reducing false negatives in robot grasp planning with a robust efficient area contact hypothesis model," in *Int. S. Robotics Research (ISRR)*, 2019.
- [12] M. Everingham, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman, "The pascal visual object classes (voc) challenge," *International journal of computer vision*, vol. 88, no. 2, pp. 303–338, 2010.
- [13] C. Ferrari and J. Canny, "Planning optimal grasps," in *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, 1992.
- [14] K. Goldberg, B. V. Mirtich, Y. Zhuang, J. Craig, B. R. Carlisle, and J. Canny, "Part pose statistics: Estimators and experiments," *IEEE Trans. Robotics and Automation*, vol. 15, no. 5, 1999.
- [15] S. Goyal, A. Ruina, and J. Papadopoulos, "Planar sliding with dry friction part 1. limit surface and moment function," *Wear*, vol. 143, no. 2, 1991.
- [16] K. Harada, T. Tsuji, S. Uto, N. Yamanobe, K. Nagata, and K. Kitagaki, "Stability of soft-finger grasp under gravity," in *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, IEEE, 2014, pp. 883–888.
- [17] R. Haschke, J. J. Steil, I. Steuwer, and H. J. Ritter, "Task-oriented quality measures for dextrous grasping," in *CIRA*, Citeseer, 2005, pp. 689–694.
- [18] R. D. Howe and M. R. Cutkosky, "Practical force-motion models for sliding manipulation," *Int. Journal of Robotics Research (IJRR)*, vol. 15, no. 6, pp. 557–572, 1996.
- [19] R. D. Howe, I. Kao, and M. R. Cutkosky, "The sliding of robot fingers under combined torsion and shear loading," in *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, IEEE, 1988, pp. 103–105.
- [20] I. Kao and M. R. Cutkosky, "Quasistatic manipulation with compliance and sliding," *Int. Journal of Robotics Research (IJRR)*, vol. 11, no. 1, pp. 20–40, 1992.
- [21] I. Kao, K. Lynch, and J. W. Burdick, "Contact modeling and manipulation," in *Springer Handbook of Robotics*, Springer, 2008.
- [22] I. Kao and F. Yang, "Stiffness and contact mechanics for soft fingers in grasping and manipulation," *IEEE Trans. Robotics and Automation*, vol. 20, no. 1, pp. 132–135, 2004.
- [23] D. Kirkpatrick, B. Mishra, and C.-K. Yap, "Quantitative steinitz's theorems with applications to multifingered grasping," *Discrete & Computational Geometry*, vol. 7, no. 3, pp. 295–318, 1992.
- [24] R. Krug, Y. Bekiroglu, and M. A. Roa, "Grasp quality evaluation done right: How assumed contact force bounds affect wrench-based quality metrics," in *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2017.
- [25] Y. Li and I. Kao, "A review of modeling of soft-contact fingers and stiffness control for dextrous manipulation in robotics," in *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, IEEE, vol. 3, 2001, pp. 3055–3060.
- [26] Z. Li and S. S. Sastry, "Task-oriented optimal grasping by multi-fingered robot hands," *IEEE Journal on Robotics and Automation*, vol. 4, no. 1, pp. 32–44, 1988.
- [27] Y. Lin and Y. Sun, "Grasp planning to maximize task coverage," *Int. Journal of Robotics Research (IJRR)*, vol. 34, no. 9, pp. 1195–1210, 2015.
- [28] J. Mahler, J. Liang, S. Niyaz, M. Laskey, R. Doan, X. Liu, J. Aparicio, and K. Goldberg, "Dex-net 2.0: Deep learning to plan robust grasps with synthetic point clouds and analytic grasp metrics," in *Proc. Robotics: Science and Systems (RSS)*, 2017.
- [29] J. Mahler, M. Matl, X. Liu, A. Li, D. Gealy, and K. Goldberg, "Dex-net 3.0: Computing robust vacuum suction grasp targets in point clouds using a new analytic model and deep learning," in *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2018.
- [30] J. Mahler, F. T. Pokorný, B. Hou, M. Roderick, M. Laskey, M. Aubry, K. Kohlhoff, T. Kröger, J. Kuffner, and K. Goldberg, "Dex-net 1.0: A cloud-based network of 3d objects for robust grasp planning using a multi-armed bandit model with correlated rewards," in *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, IEEE, 2016, pp. 1957–1964.
- [31] M. T. Mason, *Mechanics of robotic manipulation*. MIT press, 2001.
- [32] B. Mishra, J. T. Schwartz, and M. Sharir, "On the existence and synthesis of multifinger positive grips," *Algorithmica*, vol. 2, no. 1-4, pp. 541–558, 1987.
- [33] R. M. Murray, *A mathematical introduction to robotic manipulation*. CRC press, 2017.
- [34] V.-D. Nguyen, "Constructing force-closure grasps," *Int. Journal of Robotics Research (IJRR)*, vol. 7, no. 3, pp. 3–16, 1988.
- [35] Z. Qiu, A. Escande, A. Micaelli, and T. Robert, "Human motions analysis and simulation based on a general criterion of stability," in *International Symposium on Digital Human Modeling*, 2011, pp. 1–8.
- [36] E. Rimon and J. Burdick, *The Mechanics of Robot Grasping*. Cambridge University Press, 2019.
- [37] M. A. Roa and R. Suárez, "Grasp quality measures: Review and performance," *Autonomous robots*, vol. 38, no. 1, pp. 65–88, 2015.
- [38] J. K. Salisbury and B. Roth, "Kinematic and force analysis of articulated mechanical hands," *Journal of Mechanisms, Transmissions, and Automation in Design*, vol. 105, no. 1, 1983.
- [39] M. Strandberg, "A grasp evaluation procedure based on disturbance forces," in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, IEEE, vol. 2, 2002, pp. 1699–1704.
- [40] P. Tiezzi and I. Kao, "Modeling of viscoelastic contacts and evolution of limit surface for robotic contact interface," *IEEE Trans. Robotics*, vol. 23, no. 2, pp. 206–217, 2007.
- [41] J. Weisz and P. K. Allen, "Pose error robust grasping from contact wrench space metrics," in *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2012.
- [42] J. Xu, N. Alt, Z. Zhang, and E. Steinbach, "Grasping posture estimation for a two-finger parallel gripper with soft material jaws using a curved contact area friction model," in *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2017, pp. 2253–2260.
- [43] J. Xu, T. Aykut, D. Ma, and E. Steinbach, "6dls: Modeling non-planar frictional surface contacts for grasping using 6d limit surfaces," in <https://arxiv.org/pdf/1909.06885.pdf>, Under Review, 2018.