

REACH: Reducing False Negatives in Robot Grasp Planning with a Robust Efficient Area Contact Hypothesis Model

Michael Danielczuk¹, Jingyi Xu^{1,2}, Jeffrey Mahler¹,
Matthew Matl¹, Nuttapong Chentanez³, Ken Goldberg¹

¹UC Berkeley ²Technical University of Munich ³NVIDIA
{mdanielczuk, jingyi.xu, jmahler, mmatl, goldberg}@berkeley.edu
nchentanez@nvidia.com

Abstract. Although point contact models are ubiquitous for robot grasp planning, they do not model the range of wrenches that finite-area soft contacts provide. This approximation leads to many false negatives. To reduce these, we propose REACH, a Robust Efficient Area Contact Hypothesis model. We consider its potential benefits and investigate two potential drawbacks: increased computational complexity and increased false positives. The REACH model computes the contact profile using constructive solid geometry intersection and barycentric integration and estimates the contact’s ability to resist external wrenches (e.g., gravity) under perturbations in object pose and material properties. We evaluate the performance of REACH with 2,625 physical grasps of 21 diverse objects with an ABB YuMi robot. We compare performance of a soft point contact model, an elliptical area contact model, and a rigid-body dynamic simulation model using NVIDIA Flex. The REACH model reduces false negatives by 17% compared to the point contact model, achieving 72% average recall. The REACH model also compares favorably to full dynamic simulation in Flex and is two orders of magnitude faster, with 50 ms average computation time. Experimental data and supplementary material are available at <https://sites.google.com/berkeley.edu/reach>.

1 Introduction

In nearly all robotic grasping applications, gripper contacts are covered with compliant material to better resist disturbing wrenches [8, 12]. Compliant gripper fingertips also create contact areas that can deform around local surface geometry [22]. However, the majority of research on grasp planning uses a point contact model that does not take contact area into account [25]. Point contact models can produce a high false-negative rate, predicting failure for grasps that are successful, especially when grasping objects at edges and corners, as shown in Fig. 1. A high false negative rate may lead to a grasp planner being unable to find grasps on an object, even when they exist. Additionally, when training a deep network for grasp planning, a large dataset of true positives and true negatives are needed to avoid overly confident or conservative predictions [38].




Fig. 1. Top: An example of a grasp on the corner of a cube. With a point contact model (left) the grasp cannot resist the gravity wrench since the contacts slip; however, with the REACH model (right) the grasp can resist the gravity wrench under perturbations in pose and material properties as the gripper deforms around the object corners to form a sticking contact. Bottom: Predicted grasp reliability (red, yellow, green $\rightarrow$ low, medium, high) under the point contact model (left) and under the REACH model (right) for a bracket, illustrating that the REACH model is far less conservative. The REACH model is significantly more accurate in predicting the true outcome of physical grasp experiments on this bracket.

Reducing the contact area to idealized surfaces [5, 14, 15, 46] or using the Finite Element Method (FEM) [6, 44, 45] are other approaches for modelling the interaction between a compliant gripper and objects, but the former does not model non-parametric surface geometries, while the latter is computationally expensive.

To address these issues, we propose the REACH model, a novel Robust Efficient Area Contact Hypothesis model for robot grasping that estimates contact area through the constructive solid geometry intersection between the object mesh and the compliant volumetric model of each gripper jaw. REACH represents the pressure distribution using a linear model based on the deformation of the jaw and decomposes the contact surface into a triangular mesh to compute the friction constraints for each triangle based on the ellipsoidal limit surface [14]. “Hypothesis” refers to the inherent uncertainty in these constraints for multi-point frictional contacts. The wrench contributions from all triangles form the grasp wrench space, which is used to predict grasp reliability (probability of suc-

cess). As in Dex-Net 1.0 [24], REACH uses Monte-Carlo sampling to evaluate grasp robustness under perturbations in gripper pose and material properties.

This paper makes two contributions:

1. REACH, a novel and efficient area contact model that considers full six-dimensional wrenches applied at the contact by modeling frictional constraints at each triangle.
2. A labeled dataset from physical experiments with 2,625 grasps of 21 objects by an ABB YuMi robot with a compliant parallel-jaw gripper, with cross-validation experiments comparing REACH to point and dynamic contact models across results from physical experiments.

2 Related Work

There is a substantial literature on contact modeling for robot grasping, with notable surveys by Kao, Lynch, and Burdick [17], Rimon and Burdick [34], and Bicchi and Kumar [2].

2.1 Point Contact Models

Point contact models provide constraints on forces and torques applied at the contact between a point and a surface. A point contact can be modeled as frictionless point contact, frictional, or soft [2, 37]. A frictionless point contact can only exert a force along the inward contact normal. A frictional point contact can additionally exert a tangential force, and a soft point contact also applies a pure torsional moment about the inward contact normal. The tangential force is commonly constrained by Coulomb’s Law, and the torsional moment is similarly constrained, scaling with a *torsional friction coefficient* [29]. The tangential force and torsional moment that a soft point contact can exert can also be jointly constrained by the so-called *friction limit surface (FLS)* [11, 22, 31], which Howe et al. approximated with an ellipsoid for computational efficiency [14].

2.2 Area Contact Models

The first area contact models used Hertz’s contact theory [13], which described the contact profile for two elastic bodies. Xydas et al. extended the model to include materials that are not linear-elastic [18, 46]. Other compliant contact models have also been introduced [15, 28]. Barbagli et al. provide an excellent summary of several of these models, and how well each one approximates the human finger [1]. However, these models assume a spherical or hemispherical object contacting a planar surface, and require finite-element analysis for nontrivial surfaces [33, 45]. Some work has focused on adapting these models to non-planar geometries. Ciocarlie et al. applied constraints from the ellipsoidal FLS model to area contacts by summing over elements in a FEM simulation [6] and by approximating the contact area with an ellipse and applying Hertzian and Winkler

pressure distributions [5]. The 2D ellipse is later generalized to quadric surfaces in [40] to approximate the contact area. Xu et al. [44] analyze a 3D subspace of 6D friction constraints for curved contact area results from soft-finger grasping, where the contact profile is obtained from a FEM simulation. Ghafoor et al. applied a different model consisting of circular elastic point contact contours to find the 6D grasp stiffness matrix, which relates the force applied by the finger to the 6D wrench applied at the contact [9].

Of these, Ciocarlie et al.’s work is most similar to the work presented here, but our work precisely determines the area of contact without approximating the surface. Charusta et al. also find a patch contact by intersecting the gripper geometry with the object, but only consider spherical gripper pads and do not consider the pressure distribution formed by the patch contact [4]. In contrast, we consider arbitrary planar gripper geometries and explicitly model the pressure distribution created by the patch contact. Sinha and Abel analyze non-planar contact surfaces through discretization into finite area patches [39], but they applied their results only to objects with simple geometries (such as cubes and cylinders) and required a variational approach to solve for a quasi-static equilibrium. In contrast, we estimate the contact area without a smooth approximation of the object surface or FEM simulations. The proposed model also considers the full 6D wrench applied at the contact, whereas Ciocarlie et al. [5] and Charusta et al. [4] consider a 4D wrench and Xu et al. [44] study a 3D subspace.

2.3 Grasp Wrench Space Analysis

Our wrench-space analysis extends the analysis presented by Mahler et al. [26, 27] in that we characterize a grasp as successful if the grasp resists the gravity wrench. Previous work has considered many metrics including probability of force closure and ability to resist any applied disturbing wrenches [7, 35, 43]. Krug et al. perform a complete analysis of wrench-based grasp quality metrics [19]. In this work, we develop constraints for each patch contact applied to the object and solve a quadratic program to determine if the forces applied at the contacts can resist the gravity wrench applied to the object.

2.4 Grasp Datasets

Bohg et al. provided a survey of the most commonly-used grasp datasets [3]. Large-scale datasets for grasping are commonly created via hand-labeling, such as Team MIT-Princeton’s dataset for the 2017 Amazon Picking Challenge [47] and Lenz et al.’s Cornell Grasping Dataset [20]. Other training datasets label successful robot grasps either in simulation [16] or on a physical system [21, 32]. Mahler et al. proposed a hybrid approach, using soft point contact models to efficiently generate a labeled dataset of millions of grasps for 3D models in simulation, and training a network on a labeled dataset of simulated grasps [25, 26, 27]. In contrast, to create our physical dataset, we sample grasps in simulation and gather ground-truth labels on the physical system.

3 Problem Statement

We consider the problem of predicting grasp reliability, or probability of grasp success, based on the ability of the grasp contacts applied by a robot gripper to resist the gravity wrench. In this paper, we consider a parallel-jaw gripper, but the model can be applied to grippers with any number of jaws.

3.1 Assumptions

We make the following assumptions:

1. Quasi-static physics (inertial terms are negligible) and Coulomb friction.
2. Objects to be grasped are rigid with known geometry.
3. The gripper has known geometry and two parallel jaws, each covered with a linear-elastic material whose geometry can be approximated by an extruded planar polygon.
4. The friction coefficient μ is constant over the contact area.
5. Both gripper jaws make contact simultaneously.
6. The jaw applies a force normal to each triangle's face in the contact patch.

3.2 Definitions

The state $\mathbf{x}$ of the grasping scene consists of the geometric, mass, and frictional properties describing an object $\mathcal{O}$, as well as the pose of the object $T_{\mathcal{O}}$. The set of actions are a set of grasps $\mathcal{U}$ that are available to a robot with a parallel-jaw gripper. Each grasp $\mathbf{u} \in \mathcal{U}$ is parametrized by a nominal grasp center $\mathbf{p} \in \mathbb{R}^3$ and an angle $\varphi \in \mathcal{S}^3$. The jaws close with closing force f_C around the grasp center and are oriented according to the grasp angle. Success is measured through a binary reward function R , where $R = 1$ if the grasp successfully lifts the object and $R = 0$ otherwise. To model uncertainty in state, imprecision in robot control and imperfect knowledge of external wrenches, we define grasp reliability as a distribution $Q(\mathbf{x}, \mathbf{u}) = \mathbb{P}(R | \mathbf{x}, \mathbf{u})$ that describes the probability of success for a state $\mathbf{x}$ and grasp $\mathbf{u}$ [26]. We evaluate reliability in simulated environments by varying object pose, mass, and frictional properties, and by taking $Q(\mathbf{x}, \mathbf{u})$ to be the sample mean of N trials: $Q(\mathbf{x}, \mathbf{u}) = \frac{1}{N} \sum_{i=1}^N R_i(\mathbf{x}, \mathbf{u})$.

3.3 Objective

Given a nominal state $\mathbf{x}$ and grasp $\mathbf{u}$, we seek an accurate, efficiently-computable estimation of reliability $Q(\mathbf{x}, \mathbf{u})$, while optimizing average precision (AP). AP is defined as the weighted mean of precision values at each recall threshold, with the weights being the increase in recall from the previous threshold, and measures area under the precision-recall curve. This metric optimizes binary classification performance for an imbalanced dataset (e.g., more successes than failures) [36].




Fig. 2. (a) Under the REACH model, the contact patch is formed via the constructive solid geometry intersection between the extruded pad polygon and the object mesh. $\hat{\mathbf{a}}$ indicates the approach vector of the jaw. Redder colors represent more deformation of the pad at that point and hence higher pressure. (b) The deformation of the gripper and (c) the triangle frame for a single triangle in the contact patch. The wrenches that can be applied at each triangle sum to form the total wrench applied to the object in its frame. (d) The pressure distribution of the patch under the REACH model and (e) as measured by a Weiss Robotics WTS tactile sensor [41].

4 REACH Model

To efficiently approximate grasp reliability, the REACH model estimates the contact wrenches a gripper can apply to the object. We first estimate the contact area, which we decompose into multiple triangles, and the pressure distribution. The wrenches that can be applied at each triangle are computed individually, and the sum of each triangle’s wrenches forms the total wrench that the grasp can apply. Each triangle contributes frictional and normal wrenches, where the frictional wrenches are constrained with an elliptical FLS model [14, 31].

The wrench set that the contacts can exert in the object frame is given by $\Lambda = \{\mathbf{w} \in \mathbb{R}^6 \mid \mathbf{w} = G\boldsymbol{\alpha}, \boldsymbol{\alpha} \in \mathcal{F}\}$, where $\boldsymbol{\alpha} \in \mathbb{R}^{6n}$ for n triangles, and $G \in \mathbb{R}^{6 \times 6n}$ is a set of $6n$ basis wrenches in the object frame. If each triangle has m constraints, $\mathcal{F} \subseteq \mathbb{R}^{mn}$ is the set of constraints for the grasp. The contacts then can resist an external wrench $\mathbf{w}$ if $-\mathbf{w} \in \Lambda$ [19, 26, 31]. In practice, we solve this system using a quadratic program, which modern solvers can efficiently find an exact solution for given that the set $\mathcal{F}$ is defined by linear equality and inequality constraints. In this section, we describe the REACH model for efficient computation of the contact profile and the constraints $\mathcal{F}$.

4.1 Contact Area Computation

We model the object’s geometry as a triangular mesh and define the contact area as the constructive solid geometry intersection of the extruded polygon of the jaw with the object. To model deformability of the jaw material, we find the gripper deformation depth $d_{def} = \kappa f_C$, where κ is a material constant specific to the jaw material and f_C is the closing force the jaw can apply. If the material has a physical limit for how much it can deform, we introduce this constraint by

saturation of d_{def} at this limit. We define the contact distance d_{con} as the minimum positive distance before contacting the object and slice the intersection mesh with planes parallel to the gripper at distances d_{con} and $d_{con} + d_{def}$ along the approach vector, as shown in Fig. 2. If the resulting surface patch(es) on the object have a finite non-zero area, then the jaw contacts the object.

We define a frame of reference for the object, where the origin lies at the center of mass of the object and the orientation depends on the object, and a frame of reference for each triangle in the contact patch, where the origin lies at the barycenter of the triangle, the z -axis is along the face normal, and the x -axis points to the first vertex. Fig. 2 shows these frames for one triangle in a contact patch. Then, for a deformation function $\delta : \mathbb{R}^3 \rightarrow \mathbb{R}$ that describes the distance the jaw presses into the object at each point on the contact surface, we consider a pressure distribution $p(\mathbf{r}) = k\delta(\mathbf{r})(\hat{\mathbf{n}} \cdot \hat{\mathbf{a}})$ that scales with deformation and with the angle between the object surface normal $\hat{\mathbf{n}}$ and approach vector $\hat{\mathbf{a}}$. The normalizing constant k is determined by integrating the pressure distribution over the contact area. The pressure distribution $p_i : \mathbb{R}^2 \rightarrow \mathbb{R}$ in the i -th triangle frame is given by transforming from object frame into the i -th triangle frame. Note that the pressure distribution in triangle frame varies only over two coordinates, as the triangle is planar. Fig. 2 shows that this pressure distribution reasonably models that measured by a Weiss Robotics WTS tactile sensor [41] for the given contact.

4.2 Wrench Space Constraints

Once we find the contact patch and its pressure distribution, we construct the ellipsoidal FLS for each triangle. The 6D wrench $[f_{x_i} \ f_{y_i} \ f_{z_i} \ \tau_{x_i} \ \tau_{y_i} \ \tau_{z_i}]^T$ for the i -th triangle in its local triangle frame is given in terms of the friction coefficient μ , the unit instantaneous velocities $\hat{\mathbf{v}}_i = [v_{x_i} \ v_{y_i} \ v_{z_i}]^T$, and the pressure distribution $p_i(\mathbf{r}_i)$ over the triangle area A_i parametrized by coordinates $\mathbf{r}_i$ [17]:

$$\begin{aligned} f_{t_i} &= \begin{bmatrix} f_{x_i} \\ f_{y_i} \end{bmatrix} = - \int_{A_i} \mu \begin{bmatrix} v_{x_i} \\ v_{y_i} \end{bmatrix} p_i(\mathbf{r}_i) dA_i, \\ f_{z_i} &= \int_{A_i} p_i(\mathbf{r}_i) dA_i, \\ \tau_{x_i} &= \tau_{y_i} = 0, \\ \tau_{z_i} &= - \int_{A_i} \mu \|\mathbf{r}_i \times \hat{\mathbf{v}}_i\| p_i(\mathbf{r}_i) dA_i. \end{aligned}$$

Note that τ_{x_i} and τ_{y_i} are constrained to be zero in the local triangle frame, but are typically non-zero after transformed into the object frame for the whole contact patch. The torsional moment and tangent friction force are jointly limited by a 3D ellipsoidal FLS:

$$\frac{\|f_{t_i}\|^2}{(f_{t_i,\max})^2} + \frac{|\tau_{z_i}|^2}{(\tau_{z_i,\max})^2} \leq 1, \quad (4.1)$$

where $f_{t_i, \max}$ and $\tau_{z_i, \max}$ are the maximal frictional force and torque, which are determined by maximizing the integral over the set of possible velocities $\hat{\mathbf{v}}_i$. Then the tangential force has maximum magnitude with $\hat{\mathbf{v}}_i = [v_{x_i} \ v_{y_i} \ 0]^T$ completely in the tangent plane:

$$\begin{aligned} f_{t_i} &= \begin{bmatrix} f_{x_i} \\ f_{y_i} \end{bmatrix} \leq \mu \begin{bmatrix} f_{z_i} v_{x_i} \\ f_{z_i} v_{y_i} \end{bmatrix}, \\ \|f_{t_i}\|_2^2 &\leq (\mu f_{z_i})^2 = (f_{t_i, \max})^2, \text{ with } |f_{z_i}| \leq f_{z_i, \max}, \end{aligned}$$

where $f_{z_i, \max}$ is the fraction of the closing force f_C that is applied normal to the triangle. The torsional moment is maximized by $\hat{\mathbf{v}}_i = [-y_i \ x_i \ 0]^T / \sqrt{x_i^2 + y_i^2}$:

$$\tau_{z_i, \max} = - \int_{A_i} \mu \|\mathbf{r}_i \times \hat{\mathbf{v}}_i\| p_i(\mathbf{r}_i) dA_i \leq \mu \int_{A_i} \sqrt{x_i^2 + y_i^2} p_i(\mathbf{r}_i) dA_i.$$

Next, we compute the symbolic integrals for $f_{x_i, \max}$, $f_{y_i, \max}$, $f_{z_i, \max}$, and $\tau_{z_i, \max}$, so that they can be efficiently computed as a function evaluation at runtime. We determine the lower bound for $\tau_{z_i, \max}$, as its constraint integral can only be numerically integrated. This conservative approximation is used to limit false positives in grasp planning, and by avoiding numerical integration, we greatly increase efficiency in grasp prediction.

We start by transforming each integral to barycentric coordinates [42]. Given a function $f(\mathbf{r})$, with $\mathbf{r} = (x, y, z)$ in Cartesian coordinates, to integrate over a triangle T with vertices $\mathbf{t}_1 = (x_1, y_1, z_1)$, $\mathbf{t}_2 = (x_2, y_2, z_2)$, $\mathbf{t}_3 = (x_3, y_3, z_3)$, and area A ,

$$\iint_T f(\mathbf{r}) dA = 2A \int_0^1 \int_0^{1-\lambda_2} f(\lambda_1 \mathbf{t}_1 + \lambda_2 \mathbf{t}_2 + (1 - \lambda_1 - \lambda_2) \mathbf{t}_3) d\lambda_1 d\lambda_2.$$

Then, we symbolically compute the maximal normal force of the i -th triangle $f_{z_i, \max}$ based on its pressure distribution $p_i = a_i x_i + b_i y_i + d_i$, where a_i, b_i are based on the transform to triangle frame, and d_i is an offset from the first point of contact to the barycenter of the i -th triangle:

$$\begin{aligned} f_{z_i, \max} &= \int_{A_i} k(a_i x_i + b_i y_i + d_i) dA_i \\ &= \frac{A_i k}{3} (a_i(x_{i,1} + x_{i,2} + x_{i,3}) + b_i(y_{i,1} + y_{i,2} + y_{i,3}) + 3d_i). \end{aligned}$$

The lower bound of the maximal torsional moment $\tau_{z_i, \max}$ for the i -th triangle is computed with:

$$\begin{aligned} \tau_{z_i, \max} &= \int_{A_i} k(a_i x_i + b_i y_i + d) \sqrt{x_i^2 + y_i^2} dA_i \\ &\geq \frac{k}{\sqrt{2}} \left(\left| \int_{T_i} x_i (a_i x_i + b_i y_i + d) dA_i \right| + \left| \int_{T_i} y_i (a_i x_i + b_i y_i + d) dA_i \right| \right). \end{aligned}$$

The proof of the lower bound $\tau_{z_{i,\max}}$ and complete derivations and expressions for $f_{z_{i,\max}}, \tau_{z_{i,\max}}$ can be found in the supplementary material.

Once these per-triangle integrals are complete, we form the final set of constraints for each triangle. First, the normalizing constant k is found using the condition that $f_C = \sum_{i=1}^n f_{z_{i,\max}}$. This equality follows from the condition that the force applied must be equivalent to the integral of the pressure distribution over the contact area [17]. We additionally introduce equality constraints on k so that the pressure distribution is consistent across triangles. Next, we linearize the constraint in Equation 4.1 by approximating the ellipsoid with a series of linear constraints found by sampling j points on the ellipsoid [5]. For our implementation, $j = 18$, so there are $m = 23$ constraints for each triangle. In the experiments below, we use the ten triangles with the largest area for our constraints (for a total of 230 constraints per contact), as we empirically found these accurately describe the contact patch for nearly all objects without sacrificing prediction accuracy.

5 Experiments

5.1 Physical Experiments Dataset

To evaluate the precision and recall of the REACH model, we collected a set of 2,625 grasps across a set of 21 3D-printed objects (seen in Fig. 3) on a physical ABB YuMi robot with a compliant parallel-jaw gripper. These objects test generalization of our method to a diverse set of object geometries. To generate grasps for each object, we first computed all stable poses with a probability greater than 1% [10]. Then, in simulation, we placed each object in each stable pose on a flat workspace and sample parallel-jaw grasps on and around the object (within a radius equal to the width of the gripper). Grasps are sampled by choosing a point on the surface of the object and shooting rays along vectors within an angular threshold of the negative surface normal to find a second point on the surface. Grasps that only contact at one point or are wider than the maximum gripper width are discarded. Additionally, grasps that fall within a distance threshold of another grasp are discarded, along with stable poses that do not have any grasps that can be sampled (for example, the object is too wide for the gripper in the stable pose). Grasps within the gripper pad width of the object are sampled by adding uniform noise samples to the generated grasp centers, resulting in a set of sampled grasps such as those in Fig. 1.

When evaluating grasps on the physical system, a stable pose for the object is chosen from its stable pose distribution. A random grasp for the chosen stable pose is planned in simulation and is executed on the physical system through a point-cloud registration system built around Super4PCS [30]. The transformation between the simulated point-cloud and the point-cloud captured on the physical system is applied to the simulated grasp pose, and grasp success is recorded by a human operator. A grasp is successful if the object is stably lifted (e.g., does not fall from the jaws during lifting). We sample 25 grasps for each object and each sampled grasp is repeated 5 times to measure




Fig. 3. The 21 objects (top left) used to create the labeled dataset of real-world grasps, along with the percentage of sampled grasps that succeeded on the physical system for each object (bottom left). The objects are all 3D printed so that grasps could be indexed from simulation and were chosen to cover a wide range of geometries and grasp success rates. An example real-world grasp (right) with our compliant gripper, showing its ability to deform around complex geometries.

robustness to small errors in registration and actuation. The 21 objects used and percentage of sampled grasps that succeeded are shown in Fig. 3. The complete dataset of 2625 object and grasp poses with success labels can be found at <https://sites.google.com/berkeley.edu/reach>, along with videos of each grasp trial.

5.2 Benchmark Estimators

In addition to the model proposed here, we also benchmark three other models; two analytical models used in the literature, and a full rigid-body dynamic simulation implemented in NVIDIA Flex. Robust versions of each of the analytical models are also considered, and are calculated as described in Sect. 3.2 by taking the sample mean of N grasp trials over perturbations in object pose and frictional properties. Each object mesh contains between 12 and 3400 triangles, with an average of 1200 triangles.

Soft Point Contact Model The soft point contact model applies a 4D wrench at the contact point, $[f_x \ f_y \ f_z \ \tau_z]^T$ [37]. The constraints on f_x and f_y are the linearized friction cone, determined by friction coefficient μ , and f_z is constrained




Fig. 4. An example grasp on object 1 in the Flex dynamic simulator, with perturbations applied to each trial. We use a prototype GPU rigid-body global solver based on [23] to move to the grasp pose (left), close the gripper jaws, and lift (right). Success is recorded for each trial if the object remains in the gripper’s jaws after lifting.

by the closing force of the gripper. The torque τ_z applied at the contact is determined by a separate *torsional coefficient of friction* γ and the constraint is given by γf_z .

Elliptical Area Contact Model The elliptical contact model is adapted from Ciocarlie et al. [5]. It is implemented here by first fitting a 2D ellipse to the extracted surface patch. The Hertzian pressure distribution is assumed to compute the maximal frictional force and torque; these constraints are solved for exactly by Ciocarlie et al. and depend only on the closing force of the gripper and the ellipse axis lengths. In this case, the linearized 3D ellipsoidal limit surface model for each contact forms the grasp wrench space.

Flex Simulation We use a prototype version of a GPU-based rigid-body global solver [23] in NVIDIA Flex to run a full dynamic simulation of each grasp in the physical dataset. The solver uses a preconditioned conjugate residual method (PCR) to solve an approximated linear system for each outer iteration. We use a time step of 1/60 sec with 2 sub-steps, 6 outer iterations and 40 PCR iterations per each outer iteration, with a relaxation factor of 0.75. Each gripper consists of 720 triangles, 300 for each finger and 120 for the base. Each trial consists of the gripper moving to the grasp pose, closing the gripper, and attempting to lift the object. Since the simulator allows for many simultaneous trials, we run 100 trials of each grasp in parallel, with perturbations of the same magnitude as we use for our robustness score applied to each trial. Fig. 4 shows an example grasping scene.

5.3 Metrics

Average precision (AP) and average recall (AR) are calculated using the ground-truth grasp labels and each model’s predictions for the dataset of physical grasps. Average Precision is defined as the mean of the precision at each recall threshold, weighted by the difference between thresholds. Thus, AP is a measure of the area under the precision-recall curve and indicates overall performance of a binary classifier. AR is inversely related to the false negative rate (grasps predicted to fail but succeed) and is calculated as the recall averaged across 11 equally-spaced prediction thresholds for success from 0.0 to 1.0. Since each model has several

parameters that can be tuned for each object, a cross-validation approach was used to avoid overfitting the model to the choice of objects. We sampled model parameters from uniform distributions: $U(0.3, 1.5)$ for the tangential friction coefficient, and $U(0.01, 1)$ for the point contact model’s torsional friction coefficient or $U(0.0001, 0.01)$ for the area contact models’ material constant. We performed leave-one-out cross-validation with 100 sets of sampled parameters, choosing the best performing model on the training objects and applying it to the held-out object. We also measure run time for each model on an Ubuntu 16.04 machine with a 12-core 3.7 GHz i7-3700k processor and a NVIDIA V100 GPU.

Table 1. Average Precision (AP) and Average Recall (AR) for each model’s predictions of grasp quality, for the 525 grasps collected on the physical robot. We perform leave-one-out cross validation, and average the AP and AR on each of the held-out folds. The non-robust REACH and REACH models achieves an absolute 17% and 11% higher AR than the non-robust and robust point contact models, respectively, with similar AP. Adding robustness increases AP and still allows the REACH model to recall grasps that the other models cannot. Flex outperforms all models in AP, but requires 20x more computation time.

Model	AP	AR	Run Time (s/grasp)
Point	0.76	0.55	0.006
Point Robust	0.82	0.55	0.040
Elliptical Area	0.73	0.47	0.024
Elliptical Area Robust	0.73	0.47	0.220
Flex Simulation	0.86	0.54	10.300
REACH (non-robust)	0.77	0.72	0.053
REACH	0.82	0.66	0.520

5.4 Discussion

The results in Tab. 1 suggest that the REACH model is able to correctly predict grasps that are classified as false negatives by the other methods, at the cost of several more false positive predictions. The non-robust REACH model achieves an absolute 17% higher AR than the soft point contact model while maintaining a similar AP. Adding robustness increases AP for both models, but the REACH model still increases AR by an absolute 11% over the point contact model. Flex outperforms both the robust soft point contact model and the REACH model in the AP metric by 4%, but at the cost of a run time that is 20x that of the REACH model and 250x that of the robust point contact model. Thus, despite similar AP performance to the robust soft point contact model, the REACH model shows promise for use in a grasp planner, as it can recall grasps that the other models are unable to recall.

These results also suggest that robustness can significantly increase average precision; the AP for both the robust soft point contact model and the REACH model is 6% and 5% higher than their non-robust counterparts. While these results are limited to our dataset of plastic 3D printed objects, adding this

robustness could also help the model generalize to different object masses and frictional properties.

Additionally, we note that the sampled parameters that best fit our physical dataset for the REACH model often are unrealistic for the physical gripper that we use. In particular, the sampled material constant often leads to deformation depths much larger than the material can actually deform. We hypothesize that this choice of parameters leads to better performance because more deformation accounts for more dynamic grasps where the object may rotate slightly into alignment or where one jaw contacts the object first.

6 Conclusions and Future Work

We introduce REACH, a Robust Efficient Area Contact Hypothesis model that estimates contact profile and contact wrench constraints by leveraging triangular mesh models and barycentric integration for computational efficiency to increase recall over traditional point contact models by 17%. To extend the REACH model further, we will study remaining false positives and false negative predictions from the REACH model, which mainly stem from dynamic effects, such as the object initially slipping before being grasped and the potential change of object pose due to two jaws not in contact simultaneously. Additionally, we will test our model on a wider range of non-3D-printed objects, which will challenge the ability of robustness in frictional. Furthermore, we observe that the plastic support behind each gripper jaw slightly cantilevers outward when grasping. We have begun to model this effect with Bernoulli beam theory.

Acknowledgments

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 the Scalable Collaborative Human-Robot Learning (SCHooL) Project, NSF National Robotics Initiative Award 1734633 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, Weiss Robotics, 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, code, and suggestions, especially Jeff Ichnowski, Ashwin Balakrishna, Daniel Seita, Ajay Tanwani, and Sandra Skaff.

References

- [1] Federico Barbagli et al. “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.

- [2] Antonio Bicchi and Vijay Kumar. “Robotic grasping and contact: A review”. In: *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. 2000.
- [3] Jeannette Bohg et al. “Data-driven grasp synthesis — a survey”. In: *IEEE Trans. Robotics* 30.2 (2014).
- [4] Krzysztof Charusta et al. “Independent contact regions based on a patch contact model”. In: *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. 2012.
- [5] Matei Ciocarlie, Claire Lackner, and Peter 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.
- [6] Matei Ciocarlie, Andrew Miller, and Peter Allen. “Grasp analysis using deformable fingers”. In: *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*. 2005.
- [7] Matei Ciocarlie et al. “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.
- [8] Aaron M Dollar and Robert D Howe. “The highly adaptive SDM hand: Design and performance evaluation”. In: *Int. Journal of Robotics Research (IJRR)* 29.5 (2010).
- [9] Abdul Ghafoor, Jian S Dai, and Joseph Duffy. “Stiffness modeling of the soft-finger contact in robotic grasping”. In: *Journal of mechanical design* 126.4 (2004).
- [10] Ken Goldberg et al. “Part pose statistics: Estimators and experiments”. In: *IEEE Trans. Robotics and Automation* 15.5 (1999).
- [11] Suresh Goyal, Andy Ruina, and Jim Papadopoulos. “Planar sliding with dry friction part 1. limit surface and moment function”. In: *Wear* 143.2 (1991).
- [12] Menglong Guo et al. “Design of parallel-jaw gripper tip surfaces for robust grasping”. In: *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. 2017.
- [13] Heinrich Rudolf Hertz. “Über die Berührung fester elastischer Körper und Über die Harte”. In: *Verhandlung des Vereins zur Beförderung des Gewerbefleißes, Berlin* (1882).
- [14] Robert D Howe, Imin Kao, and Mark R Cutkosky. “The sliding of robot fingers under combined torsion and shear loading”. In: *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. 1988.
- [15] Takahiro Inoue and Shinichi Hirai. “Elastic model of deformable fingertip for soft-fingered manipulation”. In: *IEEE Trans. Robotics* 22.6 (2006).
- [16] Edward Johns, Stefan Leutenegger, and Andrew J Davison. “Deep learning a grasp function for grasping under gripper pose uncertainty”. In: *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*. 2016.
- [17] Imin Kao, Kevin Lynch, and Joel W Burdick. “Contact modeling and manipulation”. In: *Springer Handbook of Robotics*. Springer, 2008.

- [18] Imin Kao and Fuqian Yang. “Stiffness and contact mechanics for soft fingers in grasping and manipulation”. In: *IEEE Trans. Robotics and Automation* 20.1 (2004).
- [19] Robert Krug, Yasemin Bekiroglu, and Máximo A Roa. “Grasp quality evaluation done right: How assumed contact force bounds affect wrench-based quality metrics”. In: *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. 2017.
- [20] Ian Lenz, Honglak Lee, and Ashutosh Saxena. “Deep learning for detecting robotic grasps”. In: *Int. Journal of Robotics Research (IJRR)* 34.4-5 (2015).
- [21] Sergey Levine et al. “Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection”. In: *Int. Journal of Robotics Research (IJRR)* 37.4-5 (2018).
- [22] Yanmei Li and Imin Kao. “A review of modeling of soft-contact fingers and stiffness control for dextrous manipulation in robotics”. In: *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. 2001.
- [23] Miles Macklin et al. “Non-Smooth Newton Methods for Deformable Multi-Body Dynamics”. In: *ACM Transactions on Graphics (TOG)* 38.5 (2019).
- [24] Jeffrey Mahler et al. “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. Robotics and Automation (ICRA)*. 2016.
- [25] Jeffrey Mahler et al. “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.
- [26] Jeffrey Mahler et al. “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. Robotics and Automation (ICRA)*. 2018.
- [27] Jeffrey Mahler et al. “Learning ambidextrous robot grasping policies”. In: *Science Robotics* 4.26 (2019).
- [28] Duane W Marhefka and David E Orin. “A compliant contact model with nonlinear damping for simulation of robotic systems”. In: *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans* 29.6 (1999).
- [29] Matthew T Mason and J Kenneth Salisbury Jr. *Robot hands and the mechanics of manipulation*. The MIT Press, Cambridge, MA, 1985.
- [30] Nicolas Mellado, Dror Aiger, and Niloy J Mitra. “Super 4pcs fast global pointcloud registration via smart indexing”. In: *Computer Graphics Forum*. Vol. 33. 5. Wiley Online Library. 2014.
- [31] Richard M Murray. *A mathematical introduction to robotic manipulation*. CRC press, 2017.
- [32] Lerrel Pinto and Abhinav Gupta. “Supersizing self-supervision: Learning to grasp from 50k tries and 700 robot hours”. In: *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. 2016.

- [33] Mammy Pouliquen et al. “Real-time finite element finger pinch grasp simulation”. In: *Proc. IEEE Eurohaptics Conf. and S. on Haptic Interfaces for Virtual Environment and Teleoperator Systems*. 2005.
- [34] Elon Rimon and Joel Burdick. *The Mechanics of Robot Grasping*. Cambridge University Press, 2019.
- [35] Máximo A Roa and Raúl Suárez. “Grasp quality measures: review and performance”. In: *Autonomous robots* 38.1 (2015), pp. 65–88.
- [36] Takaya Saito and Marc Rehmsmeier. “The precision-recall plot is more informative than the ROC plot when evaluating binary classifiers on imbalanced datasets”. In: *PloS one* 10.3 (2015).
- [37] J Kenneth Salisbury and B Roth. “Kinematic and force analysis of articulated mechanical hands”. In: *Journal of Mechanisms, Transmissions, and Automation in Design* 105.1 (1983).
- [38] Vishal Satish, Jeffrey Mahler, and Ken Goldberg. “On-Policy Dataset Synthesis for Learning Robot Grasping Policies Using Fully Convolutional Deep Networks”. In: *IEEE Robotics & Automation Letters* (2019).
- [39] Pramath R Sinha and Jacob M Abel. “A contact stress model for multifingered grasps of rough objects”. In: *IEEE Trans. Robotics and Automation* 8.1 (1992).
- [40] Tokuo Tsuji et al. “Grasp planning for constricted parts of objects approximated with quadric surfaces”. In: *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*. 2014.
- [41] *Weiss robotics tactile sensor*. URL: <https://www.weiss-robotics.com/en/produkte/unkategorisiert/wts-en> (visited on 08/26/2019).
- [42] Eric W Weisstein. “Barycentric coordinates”. In: *MathWorld—A Wolfram Web Resource* (2003).
- [43] Jonathan Weisz and Peter K Allen. “Pose error robust grasping from contact wrench space metrics”. In: *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. 2012.
- [44] Jingyi Xu et al. “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. Robotics and Automation (ICRA)*. 2017.
- [45] Nicholas Xydias, Milind Bhagavat, and Imin Kao. “Study of soft-finger contact mechanics using finite elements analysis and experiments”. In: *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. 2000.
- [46] Nicholas Xydias and Imin Kao. “Modeling of contact mechanics and friction limit surfaces for soft fingers in robotics, with experimental results”. In: *Int. Journal of Robotics Research (IJRR)* 18.9 (1999).
- [47] Andy Zeng et al. “Robotic pick-and-place of novel objects in clutter with multi-affordance grasping and cross-domain image matching”. In: *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. 2018.