

Design of Parallel-Jaw Gripper Tip Surfaces for Robust Grasping

Menglong Guo¹, David V. Gealy¹, Jacky Liang², Jeffrey Mahler²,
Aimee Goncalves¹, Stephen McKinley¹, Juan Aparicio Ojea³, Ken Goldberg⁴

Abstract—Parallel-jaw robot grippers can grasp almost any object and are ubiquitous in industry. Although the shape, texture, and compliance of gripper jaw surfaces affect grasp robustness, almost all commercially available grippers provide a pair of rectangular, planar, rigid jaw surfaces. Practitioners often modify these surfaces with a variety of ad-hoc methods such as adding rubber caps and/or wrapping with textured tape. This paper explores data-driven optimization of gripper jaw surfaces over a design space based on shape, texture, and compliance using rapid prototyping. In total, 37 jaw surface design variations were created using 3D printed casting molds and silicon rubber. The designs were evaluated with 1377 physical grasp experiments using a 4-axis robot (with automated reset). These tests evaluate grasp robustness as the probability that the jaws will acquire, lift, and hold a training set of objects at nominal grasp configurations computed by Dex-Net 1.0. Hill-climbing in parameter space yielded a grid pattern of 0.03 inch void depth and 0.0375 inch void width on a silicone polymer with durometer of A30. We then evaluated performance of this design using an ABB YuMi robot grasping a set of eight difficult-to-grasp 3D printed objects in 80 grasps with four gripper surfaces. The factory-provided gripper tips succeeded in 28.7% of the 80 trials, increasing to 68.7% when the tips were wrapped with tape. Gripper tips with gecko-inspired surfaces succeeded in 80.0% of trials, and gripper tips with the designed silicone surfaces succeeded in 93.7% of trials.

I. INTRODUCTION

”Building more general hands for robots that require very little customization, that can dynamically grasp millions of different sized and shaped objects, that can do so quickly, that have a long lifetime over millions of cycles, and that just work would have significant impact on deployment of robots in factories, in fulfillment centers, and in homes.” - Rod Brooks, February 2017 [1]

Parallel-jaw grippers are widely used in the current generation of human-compliant robots, such as Sawyer from Rethink Robotics [2] or the YuMi from ABB [3], due to their low complexity, long lifetime, and ability to precisely manipulate objects [4]. Despite the intention that these robots be used in unstructured environments, such as homes and warehouse order fulfillment centers, most robots come equipped

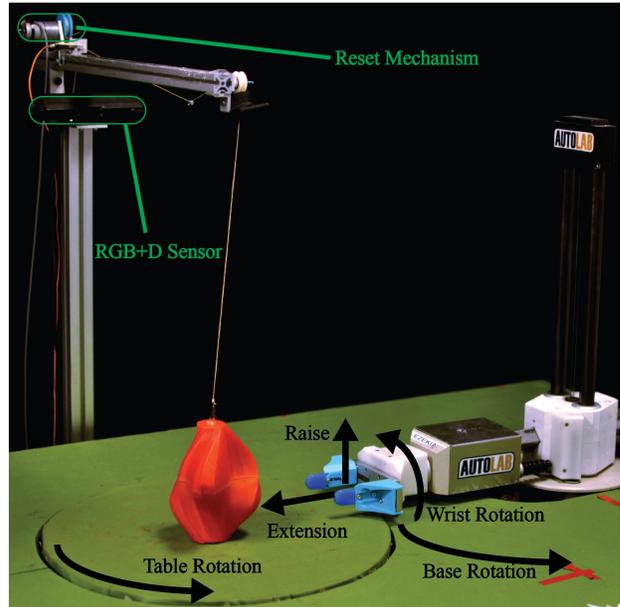


Fig. 1: **Top:** Grasp robustness for 37 gripper surfaces was evaluated using a 4-dof robot with a custom “reset” mechanism that includes an activated rewind mechanism and cable that lifts and replaces the object onto the worksurface in an upright position after each grasp. **Bottom:** Each gripper surface was cast in silicone polymer using custom 3D-printed molds.

with parallel-jaw grippers that conform to the industrial paradigm of planar, rigid jaws [5]. Robots operating in these unstructured environments may benefit from compliant end-effectors that are designed to successfully manipulate a wide variety of shapes and textures while resisting torques due to contact and gravity [6].

We explore options for adding compliant and high-friction jaw surfaces to standard parallel-jaw grippers. A variety of designs have been proposed such as rubber coverings, polymer pancakes [7], and human-inspired skin, bone, nail structures [8, 9], and gecko-inspired surfaces that resist tangential forces [10]. The design process for these surfaces has been largely guided by human intuition and optimization in simulation [11], and often only one or a small number of designs have been physically realized due to the time and

¹University of California, Berkeley, Mechanical Engineering;
{m.guo,dgealy,asgoncalves,mckinley}@berkeley.edu

²University of California, Berkeley, EECS;
{jackyliang,jmahler}@berkeley.edu

³Siemens Research;
Juan.aparicio@siemens.com

⁴University of California, Berkeley, IEOR & EECS;
goldberg@berkeley.edu

The AUTOLAB at UC Berkeley (automation.berkeley.edu)

material cost of manufacturing. However, the gripper/object surface interaction is difficult to predict in simulation and requires modeling assumptions that may not be met in practice [12].

Inspired by recent advances in 3D printing and rapid prototyping, we explore the possibility of guiding the design process empirically by evaluating success on a physical system for a large number of prototype fingertip designs. Our primary contribution is an extensive evaluation of gripper surface texture and stiffness for compliant robotic fingers (as shown in dark blue in Fig. 1) across 37 iterations of surface features (Fig. 4). Each design was parametrized and prototyped using 3D printing and molded silicone. We iteratively evaluated the probability of grasp success for our design set, inspired by resampling-based optimization methods such as zooming [13] and the cross entropy method [14], and expanded our design set around the most promising design from the previous evaluations. We collected between 21 and 35 grasp trials for each design across three 3D printed objects on a 4-dof Zymark Zymate robot for a total of 1377 total evaluations.

Initial Assumptions: The Series Type A Pro 3D printers extrude Poly(lactic acid) (PLA). Platinum-cure silicon rubber was used for a compliant material because it is robust, easy to manufacture and can be washed. Grasps were tested with a Zymark Zymate 2 laboratory robot (equipped with parallel jaws) with positional uncertainty of up to 5mm and a pointcloud-based vision system with positional error of up to 1cm . Known objects from the Dex-Net 1.0 dataset [15] were used to test grasps.

II. RELATED WORK

A. Related work in gripper design

An important tool for design inspiration is the idea of utilizing nature as a model. This technique, known as biomimicry, is highly useful for robotic gripper design. Two specific areas this design concept has been applied to are structural and surface features.

Structurally, a typical industrial robotic finger consists of a single, uniform, rigid material throughout, regardless of its shape. However, anthropomorphic observations reveal that human fingers consist of several layers: bone, soft tissue, skin, and nails, where each of these characteristics provide a unique function for human grasping. To increase the ability of a robotic hand, Murakami et al. [8] explored adding a hard nail with a strain gauge to a fingertip covered by soft elastic. Hosoda et al. evaluated these characteristics as well by using a metal bar and two types of silicone rubber to replicate the bone-body-skin structure [9]. To study the benefits of compliance on surface contact gripping, Berselli et al. designed and tested soft fingertip covers with four varied internal geometric structures, relying on rapid prototyping for inexpensive fingertip production [16]. Additionally, in both [9] and [17], embedded strain gauges were added (randomly and strategically-placed, respectively) into silicone gripper pads to replicate sense of touch and provide useful tactile feedback during manipulation. All of these

anthropomorphic-inspired designs revealed the effectiveness of utilizing multiple structural materials, geometries, and features when designing grippers for manipulation purposes.

Nature-inspired surface features have been investigated for their various attractive properties. Initially motivated by the fingerprint surfaces on human hands, Cutkosky et al. found that textured and compliant gripper surfaces improved object handling [18]. Research stemming from fingerprint analysis has led to the development of gecko-inspired adhesives, which replicate the strong adhesion gecko feet possess [19]. Pairing this material with grippers is a successful technique for manipulation. For example, Hawkes et al. developed grippers that use shear adhesion of gecko-inspired fibrillar film micro-structures to grasp curved objects [20].

From the previous work on bio-inspired adhesion, it has been suggested that interface geometry is important for adhesion. To try and provide insight on adhesion mechanisms, uniform surface patterns were explored by Crosby et al. [7]. Described as polymer 'pancakes', tests were run for surfaces with a range of cylindrical posts with varying heights, diameters, and grid-spacing. These patterned dimensions were linked to material properties which increase grip force for sliding contact; from this, adhesion effect relationships were extracted between surface dimensions and material properties. This conceptual surface feature inspired initial design concepts in Section III.



Fig. 2: The 37 gripper surface designs, with 3D printed molds (in white) used for fabrication.

B. Designing from Soft Contact Model Results

One strategy for compliant robotic finger design is to use analytic soft contact models [21]. These models take into account contact area and tangential friction forces in addition to normal friction forces, where the Soft Contact Model utilizes the power law to relate contact radius to the normal force exerted by compliant fingertips.

Design research in this area has focused on matching empirical results to the analytic contact models. For example, Han et al. [22] developed a model for maximum static friction on human fingers by first measuring friction of a human finger and fitting the data to the Hertz contact model. The friction model's results were then compared to resulting friction properties of a silicone finger. Similarly,

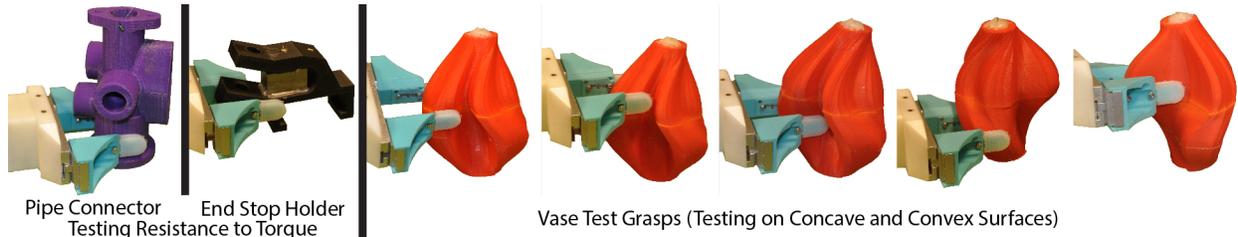


Fig. 3: Difficult-to-grasp objects selected from the online part database ‘Thingiverse’. The ‘Pipe Connector’ (206g) and ‘End Stop Holder’ (287g) were specifically chosen to investigate resistance to torque over curved and flat surfaces respectively. The ‘Vase’ (155g) object was chosen because of its convex and concave surfaces. For reference, the vase is 15.8 cm tall.

Kao and Yang [23] derived an expression for nonlinear stiffness of soft finger contact from the power law theory (also presented as a generalized Hertzian contact theory), which was compared to experimental results of force tests. Controzzi et al. [24] developed a 3D finite element model to simulate internal behavior of a proposed finger prototype. In this work, they compared the results of their modeled finger to the experimental results of the physically prototyped finger and a real finger. In comparison, we explore the use of a data-driven design approach to potentially reduce modeling errors.

C. Optimization Methods

This work is also closely related to work on optimization methods for gripper design. Convex optimization of gripper parameters may be desirable when such a parametrization is available [11, 25]. However, many design problems have more than one locally optimal solution. Methods for non-convex optimization can be difficult to apply in design because of the limited opportunity to iterate on physical prototypes. Techniques for gripper design include simulated annealing [26], gradient descent [27], nonlinear programming [28], evolutionary algorithms [29], and sequential convex programming [30]. Recently Ruiz and Mayol-Cuevas [31] compared four predefined compliant robot hand designs using performance across 3600 physical gripper-object interactions to avoid accumulating errors due to unrealistic modeling assumptions. In comparison, we iteratively use performance on physical trials to update the set of designs to test, optimizing over gripper tip texture and stiffness.

On a physical system, multiple trials may be needed to evaluate success due to imprecision in sensing and control. Past research in robotics has focused on optimizing a sampled estimate of an objective [32, 33]. To minimize the number of samples, Multi-Armed Bandits [34] and Bayesian Optimization [35] can adaptively allocate samples to more promising alternatives over a predefined design space. Our method was inspired by resampling-based optimization methods such as zooming [13] and the cross entropy method [14], which iteratively resample the design space in more promising regions and allow us to adaptively allocate manufacturing effort to more promising designs. In particular, we iteratively evaluated the success for a design set on a physical robot and resampled the design set around the design with the highest sample mean as discussed in

Section IV.

III. PROBLEM FORMULATION

We consider the problem of finding a parallel-jaw fingertip design that maximizes the likelihood that a grasp is successful on a physical system. We assume that the probability of success for a particular design and grasp is stationary; e.g. the robot’s control and perception calibration is constant over time.

Our design problem had the following attributes: probability of success could be determined quickly through experimental trials, fabrication could be iterated quickly using rapid prototyping, and the space of parameters was relatively small. We note that our empirical design approach can only be expected to work in design problems with similar properties to that of the fingertip design discussed in this paper - a design space with a small number of parameters, a short cycle time for fabrication and testing, and a clearly-defined success criterion.

Our objective was to maximize the mean likelihood of success for grasps in our training set over a space of possible designs, which we formalize below for concreteness.

A. Design Space

Let \mathcal{D} be the *design space*, a set of parameters specifying all possible designs. For example, $\mathcal{D} \subset \mathbb{Z} \times \mathbb{R}$ could represent the width of dots and depth of the fingertip surface structures. We assumed the design space is bounded and fixed. We called an element $d \in \mathcal{D}$ a design.

B. Object and Parallel-Jaw Grasping Model

Let \mathcal{O} be an object to grasp with center of mass $\mathbf{z} \in \mathbb{R}^3$. For clarity, we assumed the vertices of the mesh are specified with respect to a reference frame $T_{\mathcal{O}} = (R_{\mathcal{O}}, \mathbf{t}_{\mathcal{O}}) \in SE(3)$ centered on \mathbf{z} and oriented along the principal axes of the object [15].

We parametrized parallel-jaw grasps as $\mathbf{g} = (\mathbf{x}, \mathbf{v}, \theta)$ where $\mathbf{x} \in \mathbb{R}^3$ is the grasp center, $\mathbf{v} \in \mathcal{S}^2$ is the grasp axis, and $\theta \in \mathcal{S}$ is the approach angle. We assumed a Bernoulli distribution \mathbb{P} modeling the probability of achieving a successful grasp on a given object due to imprecision in sensing and actuation.

Let $\Gamma = \{(\mathbf{g}_1, \mathcal{O}_1), \dots, (\mathbf{g}_n, \mathcal{O}_n)\}$ be a given set of training grasps and objects sampled from a larger set of test objects, for which the design should be expected to perform well. For example Γ could contain grasps sampled on a set of industrial parts.

C. Design Objective

Let $S_i(d)$ be a binary random variable measuring the success of using design d to execute grasp \mathbf{g}_i on \mathcal{O}_i . Our goal was to find the design that maximizes the mean likelihood of success for grasps in our test set (shown in Figure 3):

$$d^* = \operatorname{argmax}_{d \in \mathcal{D}} \frac{1}{n} \sum_{i=1}^n \mathbb{P}(S_i(d) = 1) \quad (1)$$

which is the expected number of successes for a uniform distribution over the dataset Γ .

D. Methodology

Solving Equation 1 may be very difficult in practice due to the large number of possible designs, grasps, and objects. Our approach to this problem was inspired by resampling optimization methods [13, 14], which iteratively evaluate a set of points and resample near the best points. We first formed an initial discrete set of designs sampled from a number of concepts, such as different surface features and fingertip geometries. We then evaluated the probability of success for all designs on all grasps and objects in our training set by sampling. Specifically, we estimated the probability of success by taking the percentage of successes over m total trials

$$P_S(d, \mathbf{g}_i, \mathcal{O}_i) = \frac{1}{m} \sum_{j=1}^m \mathbf{1}(\hat{S}_{i,j}(d) = 1)$$

where $\hat{S}_{i,j}(d)$ is the j -th sample of $S_i(d)$ and $\mathbf{1}(\cdot)$ is the indicator function. We then expanded our design set by sampling a grid of design parameters around the best performing design from the initial set. Finally, we repeatedly evaluated and resampled for k rounds, choosing the design with the highest sample mean as d^* at termination. We used $k = 3$ and $m = 3$ in our experiments based on the amount of time to run each trial and round of evaluations.

IV. EXPERIMENTAL DETAILS

A. Fabrication of Gripper Surfaces

Molds were created with *Series 1 Pro Type A* 3D printers from Polylactic acid (PLA) filament. This manufacturing step limited the resolution of surface textures to 0.1 mm. 3D printed mounts on the parallel jaws of the robot allowed for quick swapping of grippers for testing and minimized the development cost per iteration. Each gripper had an identical base printed from PLA filament to index into the mount. The base served as a hard plastic structure for the soft silicone rubber tips to be cast around. The unique texture surface was created by casting the base in 3D printed molds (Fig. 5).

B. Dataset

Our design test set Γ consisted of seven total grasps across three objects (Fig. 3). The dataset was selected to test (a) resistance to torques about the principal grasp axis, which are difficult to resist with two fingers [5], and (b) adaptivity to varying geometric features such as concavities, convexities, and ridges.

TABLE I: Silicone Rubber Properties (Manufactured by Smooth-On)

| Silicone Rubber | Softness (Durometer) | Stiffness (Elongation at Break %) |
|-----------------|----------------------|-----------------------------------|
| Mold Star 30 | 30A | 339% |
| Dragon Skin 30 | 30A | 364% |
| Dragon Skin 10 | 10A | 1000% |
| Eco Flex 00-20 | 00-20 | 845% |

The two grasps on the end stop holder and pipe connector were chosen to test (a) torque resistance and the five grasps on the vase were chosen to test (b) geometric adaptivity. Each object was also labeled with a single stable pose on the table chosen for reachability with our 4 degree of freedom arm. All grasps were hand-selected from a set of contact points generated using the antipodal grasp sampling of Dex-Net 1.0 [15], and the approach axis was constrained to be parallel to the table for the given stable pose.

C. Experimental Platform

Grasping trials were run on a Zymark Zymate 2 robot with 4 degrees of freedom plus gripper control and a rotating turntable for 5 total controllable degrees of freedom (Fig. 1). To begin an experiment, a test object was placed onto the workspace table in a pre-defined stable pose and attached to a reset mechanism. For each grasp trial, a PrimeSense Carmine 1.09 depth sensor was used to register the pose of the object. After registration, the robot proceeded to perform a chosen grasp by planning a straight line trajectory to the desired grasp pose, moving to the pose, and closing its jaws. The robot then attempted to raise the object by 17.5mm, at which point the PrimeSense camera took a color picture for labeling. This concluded a single trial. To begin the next trial, the reset mechanism then raised and lowered the test object back to the known stable pose on the work table.

Registration was performed using convolutional neural networks for a coarse pose estimate [36] and Iterated Closest Point matching with a weighted point-to-plane objective [37] for fine pose estimation in the plane of the table. The registration system had a mean X translational error of 4.2mm, a mean Y translational error of 1.0 mm, and a mean angular error of 5.1° in the plane. The standard deviations were 3.1mm, 3.3mm, and 8.6° for X translation, Y translation, and rotation in the plane, respectively.

D. Grasp Success Criteria

Each grasp was considered a success or failure based on the criteria illustrated in Fig. 6. We considered a grasp attempt a failure if it fell into one of three modes:

- 1) **Drop:** The gripper failed to lift/hold the object.
- 2) **Slip:** The gripper lifted the object but the object rotated by more than 10 degrees about the principal grasp axis.
- 3) **Case:** The gripper lifted the object upright but leveraged a part of the gripper other than the fingertip surface.

Therefore a grasp was considered successful if it lifted the object in an upright position using only the fingertips.

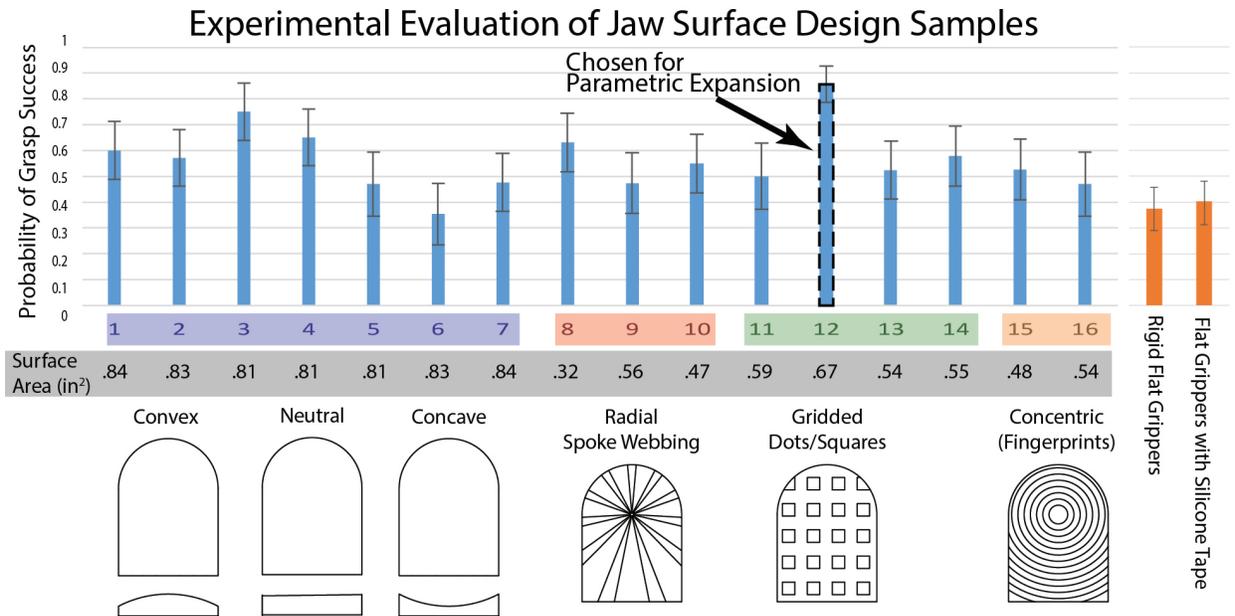


Fig. 4: Initial explorations for gripper surfaces were made for conceptual designs based on related work and the manufacturing limits of the 3D printers. Within each class of patterns were different surfaces with differences in parameters. Initial designs were compared against default rigid grippers and grippers covered in adhesive tape (as shown in orange at right).

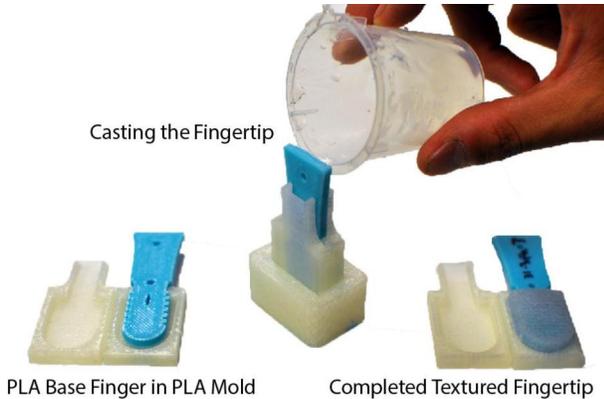


Fig. 5: Fingertips were cast around a PLA 'bone' structure at the center of each finger. Mold components were 3D printed from PLA.



Fig. 6: Examples of the three failure cases.

To provide labels for each grasp, we collected a single image of the grasp after the arm had attempted to lift the object for each grasp trial. Then, a single human labeler was shown the image of each grasp and asked to label the grasp as a success or failure based on the above criteria or reject the datapoint. Datapoints were rejected if the robot pushed the object out of the way and thus failure could not be attributed to the fingertips themselves.

V. DESIGN EVALUATION

The study included $k = 3$ rounds of design evaluation. Our design space \mathcal{D} consisted of the following parameters: curvature of the fingertip, angular resolution of spoke webbing patterns, radial resolution of concentric patterns, fingertip softness, and depth, shape, and width of gridded fingertip indentations. The cross-sectional dimensions of the fingertip were modeled after the dimensions of the human thumb (width = 0.8in, height = 1.1in, and depth = 0.35in).

A. Round 1: Initial Design Concepts

Initial design concepts were chosen to reflect a study of related works (Section II). Designs were intended to maximally resist torque around the fingertip surface. Designs 1 through 7 (Fig. 4) were parametrized by radius of surface curvature (radius = 0.93, 1.36, 2.68, flat, -2.68, -1.36, -0.93 in). Designs 8 through 16 (Fig. 4) were parametrized by width of surface features, distance between features, and depth of the surface features. For each design, we evaluated each of the 7 grasps in Fig. 3 for $m = 3$ for a total of 21 binary success trials per design. As a baseline, standard rigid flat grippers and flat grippers with silicone tape underwent the same evaluation (Fig. 4). Design 12, which had a grid of square surface indentations reminiscent of a waffle, had an 86% probability of success as the best in moment test and second best in geometric test and was chosen for expansion.

B. Round 2: First Parametrized Grid Expansion

In the second round, a 3x3x3 grid-search was employed to explore the parametric design space near initial design 12. These 27 design permutations investigated the relationship between gripper stiffness (elongation at 100% strain), void depth, and void width.

First Parametric Gripper Design Expansion

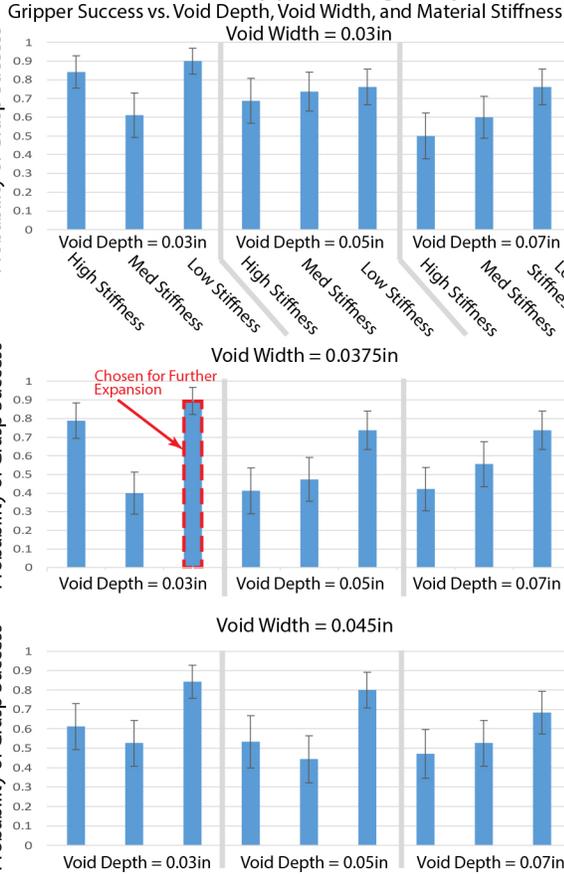


Fig. 7: The first round of parametric expansion investigates the effect of material stiffness, void depth, and void width on the most successful design from the initial set (Fig. 4) through a $3 \times 3 \times 3$ cube of possibilities. Success was found to be linked to lower stiffness and diminished void width.

Parameters Explored:

Materials: Dragon Skin 30 (DS30), Mold Star 30 (MS30), and a 50:50 mixture between DS30 and MS30.

Void Width: 0.03 0.0375 0.045 in.

Void Depth: 0.03, 0.05, 0.07 in.

Again, for each design we evaluated each of the 7 grasps in Fig. 3 for $m = 3$ samples. The results of the second round are illustrated in Fig. 7. The results suggest that lower stiffness materials and shallower indentation depth performed better. The best performing design had the same indentation depth and width as design 12 with lower stiffness, which had a 90% probability of success.

C. Round 3: Second Parametrized Grid Expansion

In this grid expansion, we investigated the square indented gripper with further exploration of void depth. We also explored material softness.

Parameters Explored:

Materials: Dragon Skin 30, Dragon Skin 10, Eco-Flex 00-20.

Void Depth: 0.02, 0.03 in.

We were forced to use a backup Zymark robot in our design evaluations because the robot we tested on in the

first two rounds malfunctioned. Thus, we ran the same 7 grasps but with $m = 5$ samples each instead of 3 because we found that the backup robot was noisier and we needed to reject more samples for each design. The results are illustrated in Fig. 8. We found that further softening the material and reducing the depth of the indentations decreased the probability of success. Furthermore, all designs had a lower probability of success on the backup robot.

The most successful design \hat{d}^* was the winner of round two, which had a void depth of 0.03in, a void width of 0.0375in, and a durometer of A30.

Second Parametric Gripper Design Expansion

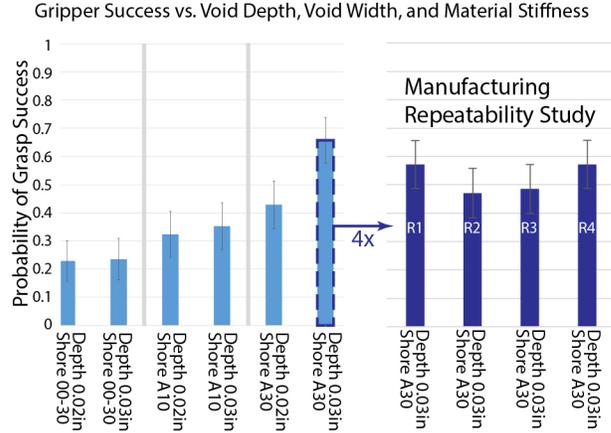


Fig. 8: A second round of parametric expansion further explored the effect of smaller void depth and material softness on grasping success. The most successful design was fabricated four more times and tested on the same conditions to verify manufacturing repeatability (shown at right in dark blue).

D. Manufacturing Repeatability

We made 4 additional copies of \hat{d}^* and evaluated the probability of success for each independently in order to measure the repeatability of our manufacturing process. The results are illustrated in the right panel of Fig. 8. We found that the designs have different success probabilities ranging from 47% to 65% with a mean of 55% and a standard deviation of approximately 7%. This suggests that variability in our manufacturing process affects the success of the design. However, even in the worst case the design had a higher estimated success probability than the other designs from the second grid expansion.

E. Manufacturing and Evaluation Time

Considerable time was saved throughout this process by using rapid prototyping to make many unique molds and using a reset mechanism to autonomously test the grippers with a real robot. The manufacturing process to create 16 unique grippers in the initial design phase took approximately 11.5 hours, including a 10 hour time period where the operator had to wait for the machine to print and the silicone to set. During grasp evaluation, each grasp took approximately 1.5 minutes to execute. With 1377 total grasps, the total grasp experiment time was 31 hours. However, the reset mechanism allowed the test to be run autonomously, so the

operator only spent about 2 hours changing grippers between tests.

F. Comparing four gripper tip surfaces on 320 grasp trials.

The highest-performing gripper surface was evaluated using eight 3D printed objects using the ABB Yumi industrial robot (Fig. 9). The factory-provided gripper tips succeeded in 28.7% of the 80 trials, increasing to 68.7% when the tips were wrapped with tape. Gripper tips with gecko-inspired surfaces succeeded in 80.0% of trials, and gripper tips with the designed silicone surfaces succeeded in 93.7% of trials. (Fig. 10).



Fig. 9: (above) After the gripper tip surface design process, an ABB Yumi industrial robot was used to compare four gripper tip surfaces using a test set of eight objects: 80 trials for each gripper tip surface. (below) The four gripper tip surfaces from left to right: ABB default grippers, ABB grippers wrapped with electrical tape, the silicone design described above, and the gecko-inspired tip surfaces.

VI. DISCUSSION AND FUTURE WORK

We evaluated gripper surface texture and stiffness for compliant robotic fingertips across 37 iterations of individual conceptual surface features and 1377 grasping evaluations in a hill-climbing approach to optimizing gripper tip surfaces. We do not claim to have found the optimal gripper in this space but performance was slightly higher than that with gecko-inspired surfaces designed to resist tangential forces. The difference in performance may be attributed to the rough surfaces of the 3D printed objects, as well as the torsional

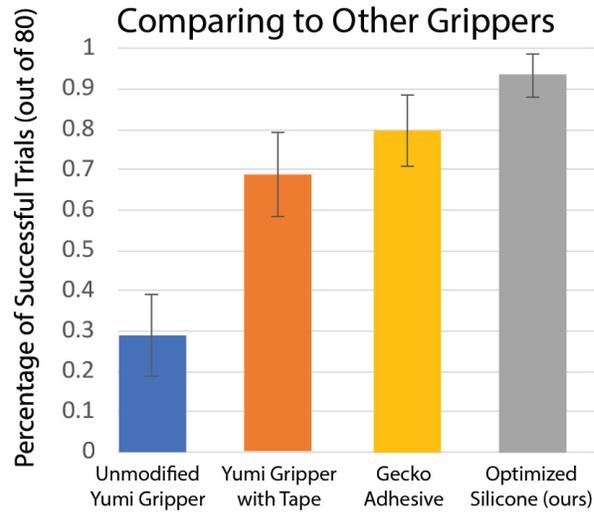


Fig. 10: Results of 320 grasp trials: 80 trials with each gripper surface.

loading of the gripper. Gecko adhesive is optimized for smooth surfaces, and shear activated adhesion is directional which results in a rotational slip failure mode. In future work, we will continue to explore the design space using Multi-Armed Bandit search methods and explore the addition of embedded force sensors.

Acknowledgements

This research was performed at the AUTOLAB at UC Berkeley in affiliation with the AMP Lab, BAIR, and the CITRIS ‘People and Robots’ Initiative (CPAR): (robotics.citris-uc.org).

The authors were supported in part by the U.S. National Science Foundation under NRI Award IIS-1227536: Multilateral Manipulation by Human-Robot Collaborative Systems, by the Department of Defense (DoD) through the National Defense Science & Engineering Graduate Fellowship (NDSEG) Program, and by Siemens Corporation. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the sponsors. We are grateful to our colleagues at Stanford for providing the gecko-inspired gripper surface material, in particular Mark Cutkosky, Elliot Hawkes, and Wilson Ruotolo, with special help from Vincent Duchaine of ETS Montreal.

REFERENCES

- [1] R. Brooks, “Research needed on robot hands,” Jan 2017. [Online]. Available: <https://rodneybrooks.com/research-needed-on-robot-hands/>
- [2] rethink robotics, “Sawyer user guide for intera 3.3 software,” http://mfg.rethinkrobotics.com/mfg-mediawiki-1.22.2/images/1/1a/Sawyer_User_Guide_3.3.pdf, 2016.
- [3] P. Crowther, “Yumi® irb 15000 overview,” <https://library.e.abb.com/public/4bcf2603f76f49088b80f7f1c49045eb/IRB14000ExternalVersionFinal.pdf>, 2015.
- [4] A. Bicchi and A. Marigo, “Dexterous grippers: Putting nonholonomy to work for fine manipulation,” *The International Journal of Robotics Research*, vol. 21, no. 5-6, pp. 427–442, 2002.
- [5] A. Bicchi and V. Kumar, “Robotic grasping and contact: A review,” in *IEEE International Conference on Robotics and Automation*. Citeseer, 2000, pp. 348–353.

- [6] A. M. Dollar and R. D. Howe, "The highly adaptive sdm hand: Design and performance evaluation," *The international journal of robotics research*, vol. 29, no. 5, pp. 585–597, 2010.
- [7] A. J. Crosby, M. Hageman, and A. Duncan, "Controlling polymer adhesion with pancakes," *Langmuir*, vol. 21, no. 25, pp. 11738–11743, 2005.
- [8] K. Murakami and T. Hasegawa, "Novel fingertip equipped with soft skin and hard nail for dexterous multi-fingered robotic manipulation," in *Robotics and Automation, 2003. Proceedings. ICRA'03. IEEE International Conference on*, vol. 1. IEEE, 2003, pp. 708–713.
- [9] K. Hosoda, Y. Tada, and M. Asada, "Anthropomorphic robotic soft fingertip with randomly distributed receptors," *Robotics and Autonomous Systems*, vol. 54, no. 2, pp. 104–109, 2006.
- [10] E. V. Eason, E. W. Hawkes, M. Windheim, D. L. Christensen, T. Libby, and M. R. Cutkosky, "Stress distribution and contact area measurements of a gecko toe using a high-resolution tactile sensor," *Bioinspiration & biomimetics*, vol. 10, no. 1, p. 016013, 2015.
- [11] M. Ciocarlie and P. Allen, "A constrained optimization framework for compliant underactuated grasping," *Mech. Sciences*, vol. 2, no. 1, pp. 17–26, 2011.
- [12] I. Kao, K. M. Lynch, and J. W. Burdick, "Contact modeling and manipulation," in *Springer Handbook of Robotics*. Springer, 2016, pp. 931–954.
- [13] R. Kleinberg, A. Slivkins, and E. Upfal, "Multi-armed bandits in metric spaces," in *Proceedings of the fortieth annual ACM symposium on Theory of computing*. ACM, 2008, pp. 681–690.
- [14] R. Y. Rubinfeld and D. P. Kroese, *The cross-entropy method: a unified approach to combinatorial optimization, Monte-Carlo simulation and machine learning*. Springer Science & Business Media, 2013.
- [15] J. Mahler, F. T. Pokorny, 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," *ICRA*, 2016.
- [16] G. Berselli and G. Vassura, "Differentiated layer design to modify the compliance of soft pads for robotic limbs," in *Robotics and Automation, 2009. ICRA'09. IEEE International Conference on*. IEEE, 2009, pp. 1285–1290.
- [17] D. S. Chaturanga, Z. Wang, Y. Noh, T. Nanayakkara, and S. Hirai, "A soft three axis force sensor useful for robot grippers," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct 2016, pp. 5556–5563.
- [18] M. Cutkosky, J. Jourdain, and P. Wright, "Skin materials for robotic fingers," in *Robotics and Automation. Proceedings. 1987 IEEE International Conference on*, vol. 4. IEEE, 1987, pp. 1649–1654.
- [19] A. Parness, D. Soto, N. Esparza, N. Gravish, M. Wilkinson, K. Autumn, and M. Cutkosky, "A microfabricated wedge-shaped adhesive array displaying gecko-like dynamic adhesion, directionality and long lifetime," *Journal of The Royal Society Interface*, 2009.
- [20] E. W. Hawkes, D. L. Christensen, A. K. Han, H. Jiang, and M. R. Cutkosky, "Grasping without squeezing: Shear adhesion gripper with fibrillar thin film," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 2305–2312.
- [21] I. Kao, K. M. Lynch, and J. W. Burdick, *Contact Modeling and Manipulation*. Cham: Springer International Publishing, 2016, pp. 931–954. [Online]. Available: http://dx.doi.org/10.1007/978-3-319-32552-1_37
- [22] H.-Y. Han, A. Shimada, and S. Kawamura, "Analysis of friction on human fingers and design of artificial fingers," in *Proceedings of IEEE International Conference on Robotics and Automation*, vol. 4, Apr 1996, pp. 3061–3066 vol.4.
- [23] I. Kao and F. Yang, "Stiffness and contact mechanics for soft fingers in grasping and manipulation," *IEEE Transactions on Robotics and Automation*, vol. 20, no. 1, pp. 132–135, Feb 2004.
- [24] M. Controzzi, M. D'Alonzo, C. Peccia, C. M. Oddo, M. C. Carrozza, and C. Cipriani, "Bioinspired fingertip for anthropomorphic robotic hands," *Applied Bionics and Biomechanics*, vol. 11, no. 1-2, pp. 25–38, 2014.
- [25] A. M. Dollar and R. D. Howe, "Joint coupling design of underactuated grippers," in *ASME 2006 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*. American Society of Mechanical Engineers, 2006, pp. 903–911.
- [26] M. M. Shalaby, H. A. Hegazi, A. O. Nassef, and S. M. Metwalli, "Topology optimization of a compliant gripper using hybrid simulated annealing and direct search," in *ASME 2003 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*. American Society of Mechanical Engineers, 2003, pp. 641–648.
- [27] M. Ciocarlie, F. M. Hicks, R. Holmberg, J. Hawke, M. Schlicht, J. Gee, S. Stanford, and R. Bahadur, "The velo gripper: A versatile single-actuator design for enveloping, parallel and fingertip grasps," *The International Journal of Robotics Research*, vol. 33, no. 5, pp. 753–767, 2014.
- [28] T. Zhang and K. Goldberg, "Design of robot gripper jaws based on trapezoidal modules," in *Robotics and Automation, 2001. Proceedings 2001 ICRA. IEEE International Conference on*, vol. 2. IEEE, 2001, pp. 1065–1070.
- [29] R. Saravanan, S. Ramabalan, N. G. R. Ebenezer, and C. Dharmaraja, "Evolutionary multi criteria design optimization of robot grippers," *Applied Soft Computing*, vol. 9, no. 1, pp. 159–172, 2009.
- [30] C. Lanni and M. Ceccarelli, "An optimization problem algorithm for kinematic design of mechanisms for two-finger grippers," *Open Mechanical Engineering Journal*, vol. 3, pp. 49–62, 2009.
- [31] E. Ruiz and W. Mayol-Cuevas, "Towards an objective evaluation of underactuated gripper designs," *arXiv preprint arXiv:1601.04547*, 2016.
- [32] K. Hsiao, M. Ciocarlie, and P. Brook, "Bayesian grasp planning," in *ICRA 2011 Workshop on Mobile Manipulation: Integrating Perception and Manipulation*, 2011.
- [33] J. Schulman, S. Levine, P. Moritz, M. I. Jordan, and P. Abbeel, "Trust region policy optimization," *CoRR*, abs/1502.05477, 2015.
- [34] T. L. Lai and H. Robbins, "Asymptotically efficient adaptive allocation rules," *Advances in applied mathematics*, vol. 6, no. 1, pp. 4–22, 1985.
- [35] G. Hitz, A. Gotovos, M.-É. Garneau, C. Pradalier, A. Krause, R. Y. Siegwart, et al., "Fully autonomous focused exploration for robotic environmental monitoring," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2014, pp. 2658–2664.
- [36] S. Gupta, R. Girshick, P. Arbeláez, and J. Malik, "Learning rich features from rgb-d images for object detection and segmentation," in *European Conference on Computer Vision*. Springer, 2014, pp. 345–360.
- [37] R. A. Newcombe, A. J. Davison, S. Izadi, P. Kohli, O. Hilliges, J. Shotton, D. Molyneaux, S. Hodges, D. Kim, and A. Fitzgibbon, "Kinectfusion: Real-time dense surface mapping and tracking," in *IEEE Int. Symp. on Mixed and augmented reality (ISMAR)*. IEEE, 2011, pp. 127–136.