Grasping objects big and small: Human heuristics relating grasp-type and object size

Ammar Kothari¹, John Morrow¹, Victoria Thrasher², Kadon Engle³, Ravi Balasubramanian¹, and Cindy Grimm¹

Abstract— This paper presents an online data collection method that captures human intuition about what grasp types are preferred for different fundamental object shapes and sizes. Survey questions are based on an adopted taxonomy that combines grasp pre-shape, approach, wrist orientation, object shape, orientation and size which covers a large swathe of common grasps. For example, the survey identifies at what object height or width dimension (normalized by robot hand size) the human prefers to use a two finger precision grasp versus a three-finger power grasp. This information is represented as a confidence-interval based polytope in the object shape space. The result is a database that can be used to quickly find potential pre-grasps that are likely to work, given an estimate of the object shape and size.

I. INTRODUCTION

Enabling robots to grasp and manipulate objects robustly is critical to expanding the use of robots in everyday living. Significant progress has been made in the robotic grasping and manipulation domain over the last few decades [17] [13]. Some newer methods to synthesize grasps include learning grasping strategies from experience based on rich sensor data [13], and using human demonstrations to learn good grasps [11]. However, there is space for improvement. Specifically, even in perfect laboratory settings where object shape and location are exactly known, automatically generated robotic grasps fail one in four times [1]. Furthermore, the search space for grasping is likely too vast for a brute force computational approach to be successful. Specifically, the grasp planning algorithm has to find a near-optimal grasp across the near-infinite continuous space of object shapes, textures, locations, orientations, topologies, tasks, and robot hand morphologies [6] [16].

Learning from humans can reduce the impact of these challenges. Prior work has explored learning "instancebased" grasping, where wrist pose and finger angles specified by humans are used [14]. Other work has also considered human-specified grasp ranges, rather than individual grasp instances, in order to view human grasp choices as continuous clusters in the search space [9]. In this paper, we learn how humans evaluate grasps as a function of object shape

@OregonState.edu

²University of California, Irvine thrashv@uci.edu

³Springfield High School, Springfield

and size. This information can be utilized in automatic grasp planning algorithms to reduce the search space.

Only a few studies have focused on human preference of robotic grasps [18]. Instead, most studies have focused on human preference for *human* grasps. They also focus on only a handful of possible object shapes and sizes, making the results difficult to generalize. Our approach, in contrast, treats object dimensions as a continuously varying space, and learns preferences for *ranges* of sizes instead of single instances.

Unfortunately, obtaining information on grasping heuristics from humans through physical experiments is difficult, time-consuming, and expensive. We and others have shown previously that collecting such human heuristic information through online crowd-sourcing is effective [20] [19] [7]. In this paper, we also use online crowd-sourcing through Amazon's Mechanical Turk service in order to conduct surveys. One challenge with online surveys is participants lack a clear understanding of the physical properties of the robotic hand; we address this problem through a series of training videos.

For simplicity and generalizability, our initial work focuses on a small number of fundamental shapes. We restrict hand morphology to a three fingered hand. We also use survey cross validation instead of physical trials to verify our results. This paper makes two key contributions: 1) it identifies that humans have clear preferences and transition points in the type of grasp based on object shape and size; 2) it presents a scientific methodology in the human-robot interaction domain for collecting grasp pre-shape preference from humans through carefully created online surveys.

With respect to the first contribution, human preferences for grasp type based on object size and shape was previously unavailable to robotic grasp planning algorithms. Current grasp planning software do not know when to transition grasp types on object shape or size [4] [8]. While only the first step, this works brings attention to how grasp planners can use human heuristics to transition from one grasp type to another. With respect to the second contribution, the surveys were created by combining previously developed human grasping classifications, which provide an understanding of different human grasp types [2] [5], and an understanding of the human ability for perceiving the shapes and size of objects from two-dimensional images [12]. Additional paper details can be found at https://goo.gl/s8vwdk.

^{*}Supported in part by NSF Grant REU Site: Robots in the Real World, CNS 1359480,

¹Oregon State University KothariA, MorrowJ, Cindy.Grimm, Ravi.Balasubramanian

II. RELATED WORK

We discuss prior related work in three related domains: 1) Human grasp classification and robotic grasp planning; 2) Learning from humans through crowd-sourcing. 3) Human perception.

A. Human grasp classification and robotic grasping planning

One of the earliest attempts to classify human grasping focused on machine shop tasks and the use of machine tools [2]. This taxonomy found sixteen unique grasps. More recently, the GRASP taxonomy aggregated previous taxonomies across multiple domains like medical rehabilitation, robotic grasping, and others to partition non-prehensile grasps into thirty three grasps [5].

Complete grasp taxonomies have not yet been fully exploited in grasp planning software like GraspIt! [15] and OpenRAVE [3]. This paper seeks to combine the grasp taxonomies with grasp planning using human heuristics. We utilize OpenRAVE [3] to create the grasp image sets for the surveys.

B. Learning from humans through crowd-sourcing

Alongside the rise of simulation environments, crowd source data collection has also gained traction [7]. Data collection through services like Amazon's Mechanical Turk is much faster and cheaper than in-person studies. The service has been used to evaluate the validity of physics based simulation for grasping [10], and provide supervision for learning object detection and grasp planning [7]. Data collected through crowd sourcing yields similar quality and accuracy as data collected through physical verification [20].

C. Human Perception

One concern when gathering performance data in a teleoperation environment is the ability of the operator to understand physical properties of the system like affect of interaction forces and surface finish when presented with only a two-dimensional image. Lau et al. [12] has shown that people can consistently identify features of an object from an image. Based on previous experience, people are able to determine properties of an object solely based on the object's 2D representation. We leverage this fact by providing training videos and images that inform participants to the physical characteristics of the robot hand and object such as weight, surface friction, and gripping strength.

III. METHODS: TAXONOMY AND DATA COLLECTION

The purpose of our study is to capture human preference of robotic manipulator grasps. We focused on human evaluation informed by training videos of how the hand operates to extract human preference of pre-grasps.

We adapted the GRASP taxonomy [5] in order to enumerate a representative set of grasp pre-shapes and objects. The goal was to cover most common manipulator interactions in the real world. The taxonomy was the basis of our online surveys, which asks people "could you pick up this object from this direction with this pre-grasp shape?".

In the following sections we outline how we developed the taxonomy, how we set up the human subjects data collection, and how we represent and use the resulting data.

A. Developing the parameterized taxonomy

Traditionally, grasp taxonomies have focused on the final grasp shape, largely ignoring the role of the object shape and size or presenting participants with only a small number of objects. Since we are interested in enumerating a large portion of the grasp space, we make two deviations from traditional taxonomy approaches. First, we parameterize by object shape and size in addition to grasp. Second, we enumerate by grasp pre-shape. This approach enables us to ask "how likely is it that this grasp will work?" without specifying details, making the results more broadly applicable as a filtering mechanism. One disadvantage is that we no longer are explicitly differentiating between how the fingers close around the object. Since this data could be collected later (either from human studies or calculated from simulation), we feel this is an acceptable trade-off.

In order to keep our taxonomy at a feasible size, we focused on a three-fingered manipulator, the BarrettHand, because we have one available for physical validation tests. However, one of the grasps we tested (Figure1-Precision F) is equivalent to a binary gripper more commonly seen on industrial applications. Since we only consider grasp pre-shape, our data is broadly applicable to pincher and 3-fingered hands, even if the gripper morphology and actuation is different.

1) Grasp pre-shape classification: To select the possible grasp pre-shapes, we started with the GRASP taxonomy [5]. Figure 1 shows those grasps converted to a 3-fingered hand. Because the robotic hand has significantly fewer degrees of freedom than a human hand, a variety of the human grasps are represented by the same gripper configuration. For example, grasps 1,2, and 3 from the GRASP taxonomy are all wrap grasps around different-sized objects. All of these can be achieved by starting from the same grasp pre-shape (Figure 1-Power F).

In addition to the grasp pre-shape, we considered the approach direction and orientation for the grasp. We limited the approach directions to one of the three primary axes of the shape. This excludes potential grasps such as picking up a cube from the corner. For orientation, we oriented the thumb along the primary axes (four possible orientations). The one exception to this is, for the cube from the top with spread fingers, we also oriented the thumb/fingers along the diagonal.

Many potential approaches and orientations are identical to each other because of hand and object symmetries. The total number of possible pre-shapes is $4 \times 5 \times 4 = 80$ assuming the object is resting on the table. However, taking symmetries into account this number is reduced to between 8 and 64.

2) Object classification and parameterization: For our object shape, we have identified four simple shapes (cube,



Fig. 1. Grasp taxonomy adapted for Barrett hand, based on GRASP (Human) Taxonomy [5]. Grasps with red boundaries, such as those utilizing the side of the finger as the main opposing force, were not achievable due to limitations with the Barrett hand's kinematics. Four grasp preshapes were identified for the Barrett hand and were applied to all achievable grasp types. These grasp types are named at the bottom of the figure based on the main opposing forces and human inferred intent of each preshape grasp.

Manipulator	Dimension (cm)	Parameterized Object Dimensions (percentages of manipulator features)					
reature		Min	MinMid	Mid	MidMax	Max	
Maximum Finger Span	32.0	0.03	0.28	0.53	0.78	1.03	
Minimum Palm Dimension	6.2	0.16	1.45	2.74	4.03	5.32	
Maximum Palm Dimension	8.9	0.11	1.01	1.91	2.81	3.71	

TABLE I. PARAMETRIZED DIMENSIONS FOR OBJECTS USED IN THIS STUDY. THE RANGE OF DIMENSIONS WERE CHOSEN BASED ON MAXIMUM FINGER SPAN.

ellipsoid, cone, and cylinder), and two complex ones (hourglass and handle). The simple shapes can be parameterized by width, height, and extent, *relative to gravity and hand in a base configuration*. Extent is how far the object extends from the palm, width is relative to the direction of the fingers, and height is the third direction which is the direction gravity acts in the base configuration. All object dimensions are relative to a feature of the hand. The parametrized dimensions can be found in Table I.

For the cone, hourglass, and handle, additional parameters are used to allow for full variation of the shape.

In addition to object shape and size, we also consider whether the object is on an end or on side. As before, we exclude symmetries resulting in two orientations for the cylinder, cone, and hourglass, but only one for the cube and ellipsoid. The handle has three orientations.

Again, taking symmetries into account, we have 154 unique hand-object configurations, of which 11 are for the cube, 21 for the cylinder, 8 for the ellipsoid, 18 for the cone, 32 for the hourglass, and 64 for the handle, with three parameters for the cube, cylinder, ellipsoid, and cone, and five each for the hourglass and handle. A subset of the final questions for the cube is given in Figure 2. Other object shapes produce similar trees.

B. Data Collection, Human Studies

Our goal with the human-subject studies is to define for each unique combination in the taxonomy the range of object



Fig. 2. Taxonomy for Cube shape for all possible orientations and approach directions. Taxonomy levels are shown on the right. Grasp preshape names correspond to labels given in Figure 1. Wrist orientations are angle relative to each pre-shape's base configuration.

sizes that are graspable (if any). The most interesting parts of this space are at the boundaries where small changes in object dimensions causes grasps to go from successful to unsuccessful. Our human study surveys are designed to quickly find the boundary.

We first describe our *framing* of the problem for participants, then our different methods for collecting data, designed to both reduce the overall number of questions needed and to provide validation through data triangulation.

1) Participant Guidance and Framing: We used online surveys, which necessitated textual and contextual clarity. Additionally, we need to familiarize participants with the robot hand capabilities (finger strength, surface finish, kinematics) so that we get consistent answers. We pilot tested different phrasing and presentation, looking for both clarity and balancing conciseness of the instructions. Our final format (shown in the accompanying video) consists of: an example question, a video showing the real robot hand demonstrating the acceptable motions, a video of a simulated hand demonstrating the acceptable motions, video and text describing the gripping ability of robot, and surface texture of object, and five questions to verify that the participants understood how the hand is allowed to move.

2) The surveys: We developed three types of survey. All surveys use stimulus of the form shown in Figure 3, with the hand offset from the object and a floor plane. The inset shows a top-down view of the same scene. Images were cropped and sized so that the hand shape remains constant across all images. The image size was set so that all of the images were visible at the same time, given a reasonable sized laptop screen. Potential survey takers were warned that the survey required a decent Internet connection and screen size. Questions were presented in randomized order to prevent learning and fatigue effects, and survey time limited to 25 minutes. The human-subject study was approved by OSU's Institutional Review Board.

3) One-dimension bracketing: Each question consists of a unique configuration (object and grasp). The two fixed dimensions were set to the middle of their respective ranges (Table I). Five images were generated by evenly sampling



Fig. 3. Example survey question where height dimension of cylinder changes for a single manipulator pre-shape, approach, and object type.

that dimension from the smallest feasible to the largest. Sliders are more sample efficient than yes/no questions for bracketing boundariess of a range. Participants were asked to pick the *smallest* and *largest* graspable size by setting two sliders (see Figure 3).

We asked bracketing questions for width, height, and extent per unique configuration, yielding a total of 6 data points per configuration (averaged across the participants). Note that this survey technique is not limited to axis-aligned sampling, but could be applied to any line in parameter space.

4) Two-dimensions sampling: In this survey type, we simultaneously vary two dimensions around a point of interest creating a grid of 9 object sizes. Participants selected the images showing object sizes that they think are graspable. This sampling method essentially captured the transition from graspable to not graspable.

We generated the points of interest and sampling directions from our one-dimensional bracketing surveys as follows. The six (averaged) data points form a polytope in shape parameter space (see Figure 5). Take the center of each triangle in the polytope as the point of interest, and the *normal* of the triangle as one of the sampling directions. Take as the other sampling direction a vector lying in the triangle; by using two orthogonal vectors (two questions) we generate samples that better refine the boundary around the triangle.

5) Grasp validation: Our third survey type serves two purposes: Data validation and an initial ranking of preferred grasps given a fixed object size and shape. Participants are shown all of the potential grasp pre-shapes and orientations for a given object, and are asked to select the ones they think will work. For each selected pre-grasp, they are further asked to provide a confidence value (0-100%) percent as well as sort the valid pre-grasps by preference. For each unique, parameterized configuration this provides three confidence values: How many people ranked it as valid, the confidence value, and where it appeared in the ranking.

We enumerated every object shape, orientation, and grasp approach direction, but with object size fixed at the mid point for all dimensions (Table I). We used this to independently verify our one-parameter data by comparing to the bracketing results. To verify the boundary results, we generated three groups of grasps (definitely valid, on the boundary, and definitely not valid) to compare against.

C. Data representation

In this section, we focus on how we represent (and use) the continuous parameter data from the human subjects study. For illustrative purposes, assume that we are working with an object defined by three parameters; the representation generalizes to higher dimensions.

Every question in the surveys results in a data point with a valid/not valid label. Our goal is to use this data to define a region where the grasp is valid. However, the boundary of this region is fuzzy because people have varying preferences. Our solution to this is to define a nested set of polytopes that represent confidence intervals (see Figure 5). Given an object's size (a point in the space), we can return a valid (inside the inner polytope), not valid (outside the outer polytope) or a confidence value (relative distance to the bracketing polytopes).

We use two methods to build the polytopes, the first is used to "boot strap" subsequent sampling. Method one uses statistical values from the one-dimensional survey results to produce six points. The points can correspond to different confidence intervals to define an upper and lower bound polytope.

Method two works directly with sampled data, and creates an iso-surface that separates the valid from the invalid space. An alternative approach is to use Support Vector Machines to define partitions in the data.

IV. RESULTS

Due to space constraints, only data for a single scenario, namely a cube approached from the side with a Figure 1-Claw pre-shape with one finger pointing directly up, is shown in this paper.

A. Data Analysis

The surveys were conducted through Amazon's Mechanical Turk with survey generation done in Qualtrics. All counted responses passed verification checks (data completeness and consistency).

1) One-dimensional survey: The total number of participants was 48, with each participant answering 50 questions, with an average of 13 answers per question (total 166 configurations). Results for cube configurations are in Table II. Non-parametric statistics were used to calculate the boundary values for the polytope; we used the 75% confidence values for the "valid" polytope.

Grasp Type	Approach	Viable	Height		Width		Extent	
		Grasp	Min	Max	Min	Max	Min	Max
3 Finger Pinch	Тор	Y	0.39	0.62	0.03	0.59	0.03	0.48
3 Finger Pinch	Side	Y	0.09	0.49	0.03	0.63	0.03	0.66
Equidistant	Тор	Y	0.13	1.03	0.04	0.53	0.03	0.78
Equidistant	Side	Y	0.16	0.99	0.06	0.53	0.03	0.66
2 Finger Pinch	Тор	Y	0.16	1.03	0.13	0.78	0.03	0.58
2 Finger Pinch	Side	Y	0.16	0.66	0.03	0.59	0.03	0.77

TABLE II. PARAMETRIZED EXTREME DIMENSIONS FOR HUMAN PREFERENCE FOR DIFFERENT TYPES OF GRASPS FOR PICKING UP A CUBE FOR THE 75% CONFIDENCE AROUND THE MEDIAN OF RESPONSES



Fig. 4. Grasp validation. Left, histogram of number of valid votes per configuration. Right, per person consistency, comparing confidence score to ranking (normalized to confidence ranges).

2) Verification survey: Mid-sized objects: The total number of participants was 11, with one excluded. The distribution of number of valid votes for each configuration is shown in Figure 4. The majority of the not valid grasps are side hooks; we hypothesize that this is due to the size range used. Side Palm F are most applicable for skinny, long shapes.

We compared three measures: Ratio of valid votes to total, confidence values, and ranking. For ranks, we used a linear mapping to convert rankings to a scale between 25

3) Two-parameter survey: We collected responses from 16 participants for a subset of the cube polytopes. We use majority vote to label each sample point as valid or invalid.

B. Grasp Validation

We collected 5 responses to the survey with none excluded. Respondents choose all points inside the 75% confidence polytope and more than 75% of boundary points as likely to succeed. This signifies that the boundary was reflective of human preference.

C. Data Representation and Usage

The data enables the construction of polytopes in shape parameter space which designate the graspable region. Figure 5 shows the resulting polytope for different confidence intervals after the one dimension bracketing survey. Vertices are a specified confidence interval from the median survey response. As confidence decreases, the size of the region grows. Figure 6 shows objects at vertices of the region.

The 2 dimensional sampling survey further resolves the boundaries. Figure 5 shows the point of interest for the 2 dimensional survey with red markers.

Figure 8 shows the refined polytope with data from the 2 dimension survey. A region with different confidence intervals can be built from these responses.

Height	Width	Extent
0.33	0.48	0.39
0.53	0.48	0.36
0.49	1.49	1.34
0.61	0.40	0.41
0.52	0.30	0.38

TABLE III. POINTS ON THE BOUNDARY OF SHAPE SPACE FOR THE GRASP AS SHOWN IN FIGURE 5. DIMENSIONS ARE PARAMETRIZED TO MAXIMUM FINGER SPAN. VALUES CORRESPOND TO 75% CONFIDENCE INTERVAL.

V. DISCUSSION

A. Grasping Preference and Grasping Success Models

A fundamental challenge in current automated grasp planning algorithm is the inability to choose human-preferred grasp types based on object shape and size. Currently, people struggle to predict which grasp a robot will use. Through daily experience, humans have learned to robustly use different grasp types for different objects and tasks. It will be very useful for robots to learn those preferences.

This work shows that it is feasible to use online surveys to obtain this information from humans. Specifically, the online survey enables us to define a human-preferred "graspable" space for a large number of common grasps. This space includes shape, orientation, approach direction for pre-shape and object. We have validated that the collected data is consistent, both within participant and across participants. We have also used a hierarchical view of the grasp space to make data collection from humans efficient.

We can analyze the resulting polytopes as well. The center point of the region is an object that is about half the max finger span in all dimensions. This is an expected result since very small objects and very large objects are difficult to pick up. Instead, an object that can easily interact with all the features of a manipulator is easiest to pick up.

Another interesting application of the data is to see how the shape space of different grasps correspond. By looking at where overlap occurs of polytopes of different pre-shapes, objects that can be successfully grasped by two different pre-shapes can be determined. The border between regions demonstrates when to change from one grasp to another. Figure 7 shows when a robot should switch grasps to improve likelihood of success.

Furthermore, the approach also provides grasp preferences at continuously varying levels of confidence and a comparison across pre-grasp-object-shape combinations. This flexibility allows planners and controllers to utilize the model more effectively based on demand from context. Finally, data collection using this method is cumulative. Areas of low confidence can be re-sampled as needed or desired.

Future work includes physical testing to further verify survey data. Subsequently, a system will be built based on the collected data to aid grasp planners in reducing the search the space.



Fig. 5. Polytope built from the results of the one parameter dimension survey for a cube with a 3 finger pinch grasp from the side. Red markers are the average points of boundary points for that surface. From left to right: Pre-shape, hand orientation, and approach of grasp, conservative (75% confidence), average, and possible (25% confidence).





Fig. 6. Polytope at 75% confidence interval showing grasps that lie at boundary vertices.

Fig. 7. Transition between when people prefer to change pre-shapes. Approach, pre-shape, and hand orientation shown for each polytope.



Fig. 8. Refined surface of the boundary between successful and unsuccessful grasps in shape space. Unsuccessful grasps are shown with red markers. The configuration is the same as in Figure 5

B. Human Preference Data Collection

This work carefully collects human preference data for how a robot should grasp an object. In addition to logistics issues (cost, tedium, finding motivated subjects), people also struggle to explain their preferences. However, they may be able to demonstrate a task in a certain way. Thus, the surveys in this paper have been created so that researchers can identify critical decision points in the grasp space triggered by object size and shape. Another important step in conducting the surveys is that training materials were used to prime participants about the physical properties of the robot and objects. Future work includes refinement of survey questions to maximize efficiency and efficacy.

REFERENCES

 R. Balasubramanian, L. Xu, P. D. Brook, J. R. Smith, and Y. Matsuoka. Physical human interactive guidance: Identifying grasping principles from human-planned grasps. *IEEE Transactions on Robotics*, 28(4):899–910, 2012.

- [2] M. R. Cutkosky. On Grasp Choice, Grasp Models, and the Design of Hands for Manufacturing Tasks. *IEEE Transactions on Robotics and Automation*, 5(3):269–279, 1989.
- [3] R. Diankov. Automated Construction of Robotic Manipulation Programs. PhD thesis, Carnegie Mellon University, Robotics Institute, August 2010.
- [4] C. Eppner and O. Brock. Planning grasp strategies that exploit environmental constraints. In *Robotics and Automation (ICRA)*, 2015 IEEE International Conference on, pages 4947–4952. IEEE, 2015.
- [5] T. Feix, J. Romero, H. B. Schmiedmayer, A. M. Dollar, and D. Kragic. The grasp taxonomy of human grasp types. *IEEE Transactions on Human-Machine Systems*, 46(1):66–77, Feb 2016.
- [6] C. Ferrari and J. Canny. Planning optimal grasps. In *Robotics and Automation*, 1992. Proceedings., 1992 IEEE International Conference on, pages 2290–2295. IEEE, 1992.
- [7] K. Goldberg and B. Kehoe. Cloud robotics and automation: A survey of related work. *EECS Department, University of California, Berkeley, Tech. Rep. UCB/EECS-2013-5*, 2013.
- [8] K. Hang, M. Li, J. A. Stork, Y. Bekiroglu, F. T. Pokorny, A. Billard, and D. Kragic. Hierarchical fingertip space: A unified framework for grasp planning and in-hand grasp adaptation. *IEEE Transactions on Robotics*, 32(4):960–972, Aug 2016.
- [9] B. John, J. Carter, J. Ruiz, S. K. Allani, S. Dixit, C. M. Grimm, and R. Balasubramanian. Human-planned robotic grasp ranges: Capture and validation. arXiv preprint arXiv:1607.03366, 2016.
- [10] D. Kappler, J. Bohg, and S. Schaal. Leveraging big data for grasp planning. In 2015 IEEE International Conference on Robotics and Automation (ICRA), pages 4304–4311, May 2015.
- [11] M. Kopicki, R. Detry, M. Adjigble, R. Stolkin, A. Leonardis, and J. L. Wyatt. One-shot learning and generation of dexterous grasps for novel objects. *The International Journal of Robotics Research*, 35(8):959– 976, 2016.
- [12] M. Lau, K. Dev, J. Dorsey, and H. Rushmeier. Learning a Human-Perceived Softness Measure of Virtual Objects. pages 2–5, 2016.
- [13] S. Levine, P. Pastor, A. Krizhevsky, and D. Quillen. Learning Hand-Eye Coordination for Robotic Grasping with Large-Scale Data Collection, pages 173–184. Springer International Publishing, Cham, 2017.
- [14] Y. Lin and Y. Sun. Robot grasp planning based on demonstrated grasp strategies. *The International Journal of Robotics Research*, 34(1):26– 42, 2015.
- [15] A. T. Miller and P. K. Allen. Graspit! a versatile simulator for robotic grasping. *IEEE Robotics Automation Magazine*, 11(4):110–122, Dec 2004.
- [16] B. Mirtich and J. Canny. Easily computable optimum grasps in 2-d and 3-d. In *Robotics and Automation*, 1994. Proceedings., 1994 IEEE International Conference on, pages 739–747. IEEE, 1994.
- [17] U. Sahu, S. Upadhyay, N. Singh, and P. Patil. Libo: The grasping robot using object recognition. In 2016 International Conference on Electrical, Electronics, and Optimization Techniques (ICEEOT), pages 3166–3171, March 2016.
- [18] K. B Shimoga. Robot grasp synthesis algorithms: A survey. *The International Journal of Robotics Research*, 15(3):230–266, 1996.
- [19] A. Sorokin, D. Berenson, S. S. Srinivasa, and M. Hebert. People helping robots helping people: Crowdsourcing for grasping novel objects. In *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on*, pages 2117–2122. IEEE, 2010.
- [20] M. Unrath, Z. Zhang, A. Goins, R. Carpenter, W. Wong, and R. Balasubramanian. Using crowdsourcing to generate surrogate training data for robotic grasp prediction. In *Second AAAI Conference on Human Computation and Crowdsourcing*, 2014.