

Yi Herng Ong¹, J

Abstract—Grasping a simple object — unless the object is almost perfectly aligned with the gripper — requires precise positioning of the gripper with respect to the object. We show that this challenge — adopting a simple approach which enable them to succeed in these awkward poses. We adopt a straightforward method to implement our human-studies approach to grasping task by either “puppeteering” that is identical (geometrically) to underactuated robotic manipulators, this enables us to directly capture and compare human manipulation strategies with robotic ones. Our observation is that, while humans employ underactuation, how they use it is fundamentally different (and more effective) than that found in existing hardware.

I. INTRODUCTION

It is difficult to grasp objects with uncertainties, such as hardware and sensor noise, combined with environmental constraints. For example, often robot end-effectors end up in undesirable positions relative to objects due to errors in sensing and control. A solution to this problem requires effective control strategies for handling undesirable scenarios immediately prior to grasping object.

Research in robotic manipulation tends to focus on grasp detection [1], [2], ideal pre-grasp pose planning [3], [4], and pre-grasp manipulation before actual grasping if noise and constraints exist in the environment [5], [6]. All of these grasp planners **fail** to focus on the interaction between robot fingers and the object, as these grasp planners simply close the robot fingers at the same rate. In this paper we demonstrate how a few simple finger movement strategies that are learned from humans reduce the risk of the object being “knocked out” of the hand during the final stage of grasping.

In this paper we study how humans control robot fingers to grasp a handful of simple objects that span a range of sizes and contact surfaces (see Figure 2, top-left) at increasingly challenging starting hand poses. We impose constraints on the system by systematically moving the wrist away from the “ideal” pose — where the principal axes of the object is aligned with the palm [7].

*This work was supported in part by NSF grants CNS 1730126 and CNS 1659746

¹Oregon State University, Collaborative Robotics and Intelligent Systems, Corvallis, Oregon OngYi, MorrowJo, QiuY, GuptaK, Ravi.Balasubramanian, Cindy.Grimm @ OregonState.edu




Fig. 1. From left to right: Hand for puppeteering, original Barrett hand, participant grasping object with the puppeteered hand, and slider box used to control fingers.

Contributions: 1) We develop a human-studies methodology that supports analysis of the grasping *strategies* employed. This methodology allows us to largely separate the strategies from the “hardware” (i.e., human hand) used to implement them. 2) The ability to collect human demonstration data on a robotic hand in a form that is directly applicable to design a simple robotic grasping controller, without the need for an intermediate mapping. 3) Enumeration of several strategies for extending the range of grasping in noisy or constrained environments.

II. RELATED WORKS

Human grasping has been studied extensively in an effort to improve robot systems. Approaches include detailed case studies [8] and crowd-sourced surveys [9]. These findings were then organized into taxonomies based on object size and shape, such as the GRASP taxonomy [10]. This taxonomy represents grasps in optimal conditions — giving us an intuition on hand shaping and positioning for various ‘primitive’ objects.

Grasping has been studied with motion capture systems [11]; however, these studies only used the data for finger placement on the object, not finger closing. These studies let the human position the wrist in the ideal position.

In non-optimal, or cluttered scenarios Dogar et al. implemented a push-grasping strategy to successfully grasp objects in clutter [12], [13]. This technique was also robust to uncertainty to object position within a graspable region [12]. Others developed frameworks to reposition objects into a graspable region [5], [14]. Other approaches utilize hand pre-shape adjustments and environmental constraints to constrain objects during grasping when in non-optimal scenarios [15], [16]. A caging grasp can also work in these scenarios [4]. Chang et al show that humans tend to apply manipulation strategies to adjust the object pose if necessary before they pick up an object [6]. All of these efforts focus on positioning the wrist or pre-manipulating the object to improve grasping;




Fig. 2. All objects (upper left) and the base pose (lower left). Remaining figures demonstrate min and max ranges for each perturbation (small cube).

our focus is on improving finger movement once the wrist position is given.

III. STUDY METHODOLOGY

In this section we describe the data collection process from the robot and the human subjects study protocol.

Our overall goal is to improve finger movement strategies by studying how humans adapt to structured perturbation from the location where the object is centered relative to the palm. We chose a side grasp as our base case and systematically moved the wrist away from this ideal location. We also used both flat and curved objects at the smaller and larger limits of graspability (see Figure 2) to capture variation in object shape and contact surfaces.

A. Objects

The objects are shown in Figure 2; note that the cylinders were placed both upright and on their sides, for a total of $3 + 3 + 3 = 9$ unique objects (3 sizes for each object). The object's sizes were based on a previous study [9] that established the maximum and minimum graspable sizes for the BarrettHand. For each object we selected three sizes (small, average, and large) based on the graspable width/length for a side grasp in [9]. The objects were printed from PLA plastic, except for the largest objects, which are made from cardboard and covered with duct tape. Each object weighs less than 1 lb, with the large objects' weight roughly twice the smaller ones.

B. Pose perturbations

As described in [9], hand pose perturbations are expressed in terms of the geometric relationship between the object and the hand, and on graspable ranges. Sampling in each direction was chosen to minimize the number of perturbations while ensuring a gradual transition in grasping difficulties, as measured by success rates for the human and robot. Refer to Figure 2, right, which shows the extremes of the perturbation ranges.

All objects start at the position aligned with the robot hand, with the distance between each object and the palm varying with the object size: $d = 1.5, 2.5, \text{ and } 3.5\text{cm}$ for the small, medium and large objects respectively. Our three position variations are as follows. Moving left to right,

$x \in (-6\text{cm}, 6\text{cm})$ with step-size $\delta_x = 2\text{cm}$. Moving up to down, $y \in (-2\text{cm}, 8\text{cm})$ with increment $\delta_y = 2\text{cm}$. Note that 0cm in y puts the middle of the hand roughly in the middle of the object, and -2cm puts the hand resting on the table. Moving in-out with $z \in (2\text{cm}, 8\text{cm})$ by step-size $\delta_z = 2\text{cm}$.

The three orientation variations were: Tilting left-right (roll), $\psi \in (-45\text{deg}, 45\text{deg})$, step-size $\delta_\psi = 15\text{deg}$, rotating left-right (yaw) $\theta \in (-30\text{deg}, 30\text{deg})$, step-size $\delta_\theta = 15\text{deg}$, and rotating up-down (pitch) $\gamma \in (-15\text{deg}, 45\text{deg})$, step-size $\delta_\gamma = 15\text{deg}$. Note that pitch is not symmetrical because the hand bumps into the table.

These variations resulted in a total of $7 + 4 + 4 + 7 + 4 + 4 = 29$ pose perturbations per object, for a total of $9 * 29 = 261$ grasp trials for each condition. To ensure consistent initial placement of each object we used a marked piece of paper.

C. Robot protocol

We used a commercial Barrett WAM arm with the patented, underactuated Barrett hand (BH-280). The Barrett hand was moved to each sampled pose with the fingers fully opened. The robot grasped the object by closing all of its fingers simultaneously over approximately 1 second. When the joint angles quit changing (no change for 4 samples, recorded at 9 Hz) we declared the grasp “done” and applied a “shake test” [17]. This test consisted of lifting the robot hand, then shaking and rotating it slowly, over approximately 5 seconds (see video) [17]. A grasp was marked as successful only if the object was still in the grasp after the test.

D. Human subject protocol

Two robot hands were used in the human studies: 1) The commercial Barrett with custom slider control box to control each finger (right of Figure 1), and 2) A 3D printed replica of the Barrett hand without joint limitations which study participants could puppet (as described in previous work [18]) (middle of Figure 1). The purpose of using these two hands is to compare human strategies between hands with limited and completely free finger movement.

The human trials used the same poses and variations as the robotic one. For the shake test with the 3D printed robot hand (posable hand) participants were asked to maintain the grasp during the lift and shake, with gravity compensation

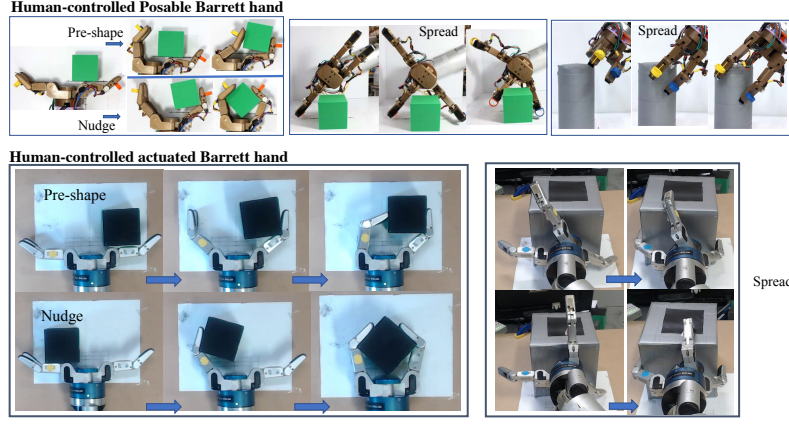


Fig. 3. Observed human strategies: From left to right: Spreading the fingers to increase surface contact, nudging the object in with one finger before grasping, pre-shaping the fingers to the object and closing them so they come into contact at the same time.

Hardware	Total successful poses	Success rate
Robot	173	66.28%
Human Barrett Slider	207	79.31%
Human Posable	220	84.29%
Barrett PID	188	72.03%

TABLE I
SUCCESS RATE OUT OF 261 PERTURBATION POSES.

on to assist in the lift. Note that the study overseer placed the fingers back to the starting position and ensured participants did not apply excessive force during the shake test. Participants were allowed two tries and a maximum of two minutes to complete each grasp trial. Note that a grasp trial is successful if the majority of the people succeeded.

We captured the joint angles of the posable hand at a sampling rate of 2 Hz (versus 9 Hz for the Barrett Hand). This lower sampling rate still suffices to capture what happened because the humans tended to move the fingers fairly slowly.

1) *Human subjects*: We had 12 human participant and they showed no qualitative nor quantitative differences in ability. We split the starting poses into two groups; those the robot succeeded at, and those it failed. 5 subjects used the posable hand on the successful trials, 5 subjects used the posable hand on the failed trials, and the last 2 subjects did the failed trials with the Barrett hand using the sliders. The small sample sizes are generally not an issue because humans tend to perform very similarly [18], [19].

Each grasp trial took 30 seconds on average. Total session time for each human subject was 1-2 hours. In total we captured 1305 trials with 3D printed posable hand, and 176 trials with Barrett hand from human subjects.

IV. DATA ANALYSIS

We use the joint angles of the proximal link, the distal link, and finger spread to analyze the movement of the fingers, and the moments when they curl in, move out or stop. Note that Figure 5 shows the joint angles in Figure 4. Each trial was labeled with the object size, shape and pose perturbation. We additionally aggregated data based on variance from the “base” grasp (value - base value / max). The fingers are labeled in Figure 1.




Fig. 4. A grasp at one extreme (positive 6 cm in x) where the robot failed but human participants succeeded with the Barrett (top) and posable (bottom) hands.

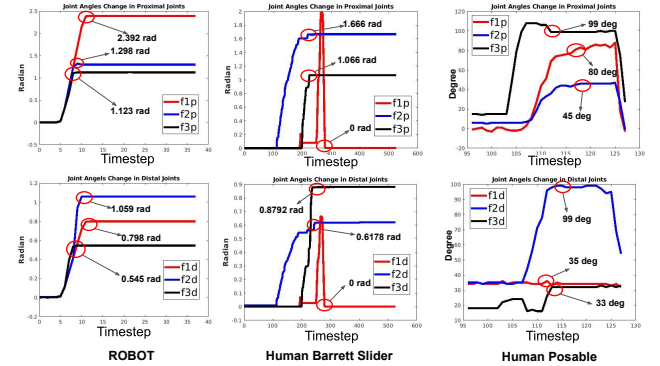


Fig. 5. Finger movement (joint angles) for controlled robot, human using slider, and human using posable hand during a grasp. Red circles indicate joint angles of final grasp poses

A. Barrett hand data analysis

For each finger of the Barrett hand we identified the first contact (when the distal joints start to change drastically and the proximal joints stop increasing) and the final grasp (both joints stop moving). We fit an actuation curve to the sampled data to better estimate the actual time of contact and of movement stoppage.

B. Posable hand data analysis

The posable data is substantially more stochastic than the underactuated Barrett one because participants were allowed to move each finger joint freely. We first filter the data with




Fig. 6. Example perturbations where the participants utilized spreading: Pitch (cylinders), y (side-ways cylinders), and roll (big cube).

a moving average of 5 samples to mitigate noise, and then label the moments where finger movements start and stop.

We again identify a first-contact and final-grasp time point. The final-grasp time point is easy to identify because the joint angles remain stable during the 5 second shake test. We define the first-contact point as when the proximal joint has been moved. Although it does not always seem to be the case, this is typically the first contact from observation.

In addition we further analyzed the data to determine the ratio of time each finger moved and spread as opposed to move simultaneously. We also measured the total radians moved by each finger to quantify how participants interacted with fingers and object during a grasp trial. Specifically, for the posable hand, we further classify grasping motion by the change of distal joint angles in each grasp trial. This captured a commonly-seen behavior where participants straightened fingers before or after bending them. To analyze how the finger pre-shapes used in the posable hand differed from the Barrett we note that the distal joint moves with the proximal joint at a reduced rate (1:3 radians). Any shapes that differ from this ratio (pre-contact) are not achievable with the Barrett Hand (this was true most of the time).

V. RESULTS

We present three sets of results: 1) Overall statistics, patterns and differences for the robotic and human data sets. 2) Three observed human strategies (shown in Figure 3) that can be mapped to either hand. 3) The increase in capabilities of the 3D-printed posable hand over the Barrett hand.

A. Overall performance in all grasp trials

Our main observation is that humans were able to improve performance at perturbation by 13% and 18% with Barrett hand and posable hand respectively (see Table I). The robot success rate was 66%. The participants succeeded with the same hand on an additional 13% of all trials when able to control each finger individually. Full control of both distal and proximal joints resulted in an additional 5% success rate. Overall, both humans and robots had similar success rates at hand positions in x, y, z for the small and medium objects, but not for the large objects. Across the board, humans were more successful with orientation variations.

B. Control strategies

We identify three control strategies (see Figure 3) in both conditions (human using posable and underactuated hands). 1) **[pre-shape]** Move the robot fingers to align with the object's principle axes; 2) **[nudge]** Use the robot fingers

to manipulate the object closer to the center of the hand; 3) **[spread]** Spread the robot fingers to improve the contact surface. Note that these control strategies were implemented before the final grasp.

The choice of strategy depended on object size, weight and position. Nudging was used to shift the object toward the palm and was used mostly with the smaller, lighter objects. Pre-shaping was used on the larger, heavier objects. These strategies appear more frequently as the perturbations become higher.

1) *Pre-shape and nudge:* In Figure 7, nudging smaller and light-weight objects was achieved with either one of the fingers, with the final grasp being relatively symmetric (object in the middle of the hand). In contrast, as the object's size and weight increase, participants began to execute pre-shaping by moving the robot fingers to align with the object. This appears as a larger difference in the joint angles of the final grasp pose with increasing perturbation in x , and is also more pronounced with the bigger cube, as shown in Figure 3). Note that using the simple finger control algorithm (closing with the same rate) did not show significant joint angle differences in the final pose.

Note, too, that the distance travelled (expressed as ratio of the movement) by the two fingers versus the thumb shifts. In other words, note that at -6 cm, the two fingers have to shift more to reach the object, whereas at 6 cm it is the thumb that has to travel more. This pattern is very clear in the big cube; while the medium cube is a hybrid strategy of nudge and pre-shape.

2) *Spread finger to adjust contact surface:* Our human participants spread the fingers to increase surface area/contact, mainly for the purpose of stabilizing the grasp. Plots of the spread angle versus perturbation for these cases are shown in Figure 6, with example grasps shown in Figure 10. In the example of large and medium cube (leftmost plot in Figure 10), increasing ψ increased finger spread in order to create a wider grip around the object. Finger spread was also used when the hand was pitched around the vertical cylinder or raised up vertically on the horizontal cylinder. Note that the use of spread for pitch on vertical cylinder, as well as move vertically on horizontal cylinder, was dependent upon the size of the object. It is less effective on the large vertical cylinder (*pitch*), but exhibits stabilization on the large horizontal cylinder (y).

C. Full actuation versus underactuation

Prior work showed that humans do not use underactuated finger shapes for difficult grasping tasks if given a choice

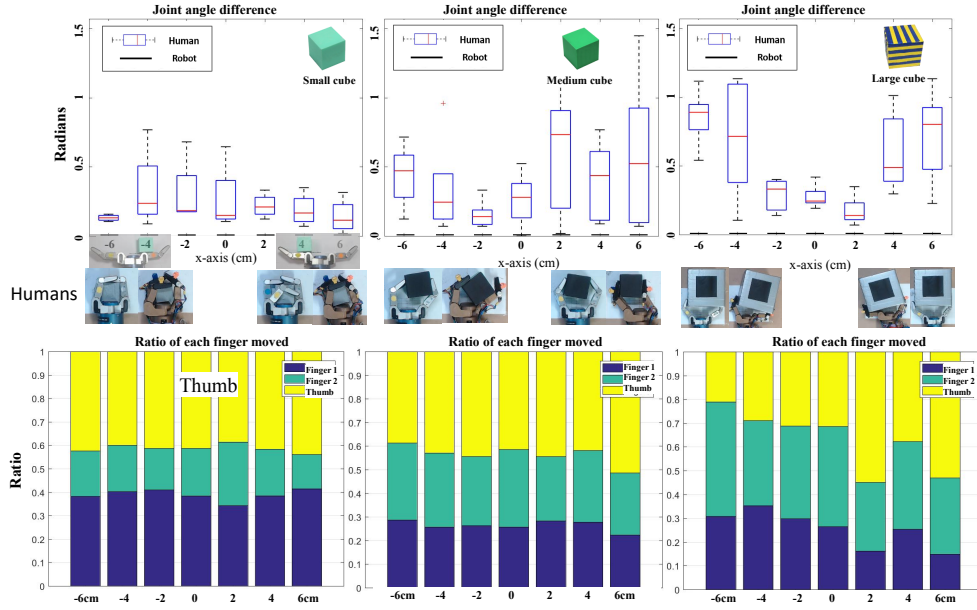


Fig. 7. Example plots showing how the joint angles between the fingers differ as the cube is shifted left and right (x). Top row: Differences between the fingers in their final joint angles (fingers that make contact with the object in grasp position). Bottom row: Ratio of the time each finger was moved. From left to right: Small to big cube. Note that the robot (black lines on bottom) only exhibited final joint angle differences for the small (at 6cm) and medium (at 2cm) cubes. Images in middle row shows final grasp pose at extreme x poses.

Average number of fingers that participants disobey underactuation in grasp trials




Fig. 8. Measuring finger-shape deviation from under actuated finger shapes. Left: Several example perturbation/object combinations that demonstrated substantial deviation, particularly for the middle-sized shapes. Right: Deviation




Fig. 9. An example where human participants failed using the Barrett hand but succeeded with the posable Barrett hand. Note that the posable hand formed larger contact surface areas than the Barrett hand (red circles indicates contacts made by hand)

[18]. Overall performance (Table I) shows that even in a simple side grasp full actuation increases success rates by 5%.

Control Scheme: Our setup allowed participants to move each joint angle independently, or simultaneously through placing the hand along the length of the finger. Using the finger loop to shape the distal joint also tended to move the

proximal joint at the same time. That said, the percentage of time participants moved both the proximal and distal joints at the same time was just 28%, $SD = 35$. They moved *just* the distal joint first approximately 20% of the time (F1: 23%, F2: 22.5%, Th: 20.0%). Therefore, from a control strategy standpoint, participants often focused on moving one finger joint at a time.

Hand Pre-shape: Qualitatively, we observed participants disobeying underactuation behavior in the following cases: Bending the distal joint backwards so that the entire front face of distal link could touch the object, and straightening distal links to be parallel with proximal link to nudge or grasp the object.

The left plot of Figure 8 shows that the medium-sized shapes show the most deviation from the underactuation scheme, as it offers the most opportunity to match the finger shape to the object's surface. The grasp completed on small shapes by posable hand was similar to the actuated Barrett hand, while fingers were hardly adjusted around large shapes. Figure 9 shows a grasp where the underactuation scheme failed but the full actuation scheme succeeded.

In the right side of Figure 8, we also summarize the overall deviation for each object type (x perturbation). The side-lying cylinder offered the most surfaces that were not aligned with the default underactuation pre-shape.

VI. CONTROLLER IMPLEMENTATION

In this section we describe our three simple PID controllers, one for each observed strategy. Which strategy to use is determined by when the humans used it, which is based on the size of the object (small, medium, large), and perturbation axis ($x, y, z, \text{roll}, \text{pitch}$ and yaw). These controllers are closed-loop controllers using computer vision (Intel realsense ZR300 camera), to track object in real time. See




Fig. 10. Spread on large cube with increasing roll (posable top, Barrett bottom).

video for a demonstration of all controllers.

“Nudge” finger control strategy: This strategy manipulates the object towards the center of the hand by gently pushing it. We designed a PID controller that continuously tracks the object position and moves the finger in until the object is centered. This controller was used in the x and yaw (θ) extremes for the small and medium-sized objects.

“Pre-shape” finger control strategy: The pre-shape strategy first moves the closest finger(s) to just touch the object, which aligns the fingers with it (again using the camera to track the object position). The controller next brings the opposite finger(s) into contact with the object from the opposing side before tightening the grasp. This controller is applied to large objects that are shifted in the x axis or rotated in yaw (θ).

“Spread” finger control strategy: This controller simply spreads the fingers by a certain amount (based on observed human values) before closing its fingers. This is used for the following objects (all sizes): horizontal cylinder (spread in y), vertical cylinder (spread in $pitch$), cube (spread in $roll$).

Validation: We applied the appropriate controller to the perturbations where the simple robot finger closing strategy failed but the human succeeded. This resulted in a 6% improvement (see Table I). Note that, although humans often employed a mix of controllers, for this test we only chose the dominant observed one.

VII. DISCUSSION

In summary, we have demonstrated that a human-subjects study structured around controlled task perturbation can elicit potential control strategies for improving grasping from awkward poses in the real world. We identified three strategies and implemented them as simple PID controllers, improving grasp performance.

Similar to [18], our human studies show that additional degrees of freedom (eg, control of each individual link) both improves grasping performance and changes control strategies. For example, humans tended to straighten distal links to improve the contact surface or provide nudging forces in the correct direction. Participants are also able to provide more nuanced force at the fingertips; future work will investigate what effect this has.

A. Future work

Our PID controllers were very simple and we simply picked one based on the pre-dominant human strategy. We

note that, although we used simple objects and one type of grasp, the controllers are potentially usable for a broader range of grasp directions and objects. In future work we propose 1) Learning when and how to employ a specific controller by analyzing the hand-object geometry and matching it to our examples and 2) Use machine learning to better map the observed behavior to novel contact points.

REFERENCES

- [1] J. Mahler, M. Matl, X. Liu, A. Li, D. Gealy, and K. Goldberg, “Dex-net 3.0: Computing robust robot vacuum suction grasp targets in point clouds using a new analytic model and deep learning,” *arXiv preprint arXiv:1709.06670*, 2017.
- [2] S. Levine, P. Pastor, A. Krizhevsky, J. Ibarz, and D. Quillen, “Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection,” *The International Journal of Robotics Research*, vol. 37, no. 4-5, pp. 421–436, 2018.
- [3] C. Ferrari and J. Canny, “Planning optimal grasps,” pp. 2290–2295, 1992. [Online]. Available: <http://ieeexplore.ieee.org/document/219918/>
- [4] R. Diankov, S. S. Srinivasa, D. Ferguson, and J. Kuffner, “Manipulation planning with caging grasps,” *Humanoid Robotics*, pp. 285–292, 2008.
- [5] L. Y. Chang, S. S. Srinivasa, and N. S. Pollard, “Planning pre-grasp manipulation for transport tasks,” *ICRA*, pp. 2697–2704, 2010.
- [6] L. Y. Chang, G. J. Zeglin, and N. S. Pollard, “Preparatory object rotation as a human-inspired grasping strategy,” *Humanoid Robotics*, pp. 527–534, 2008.
- [7] R. Balasubramanian, L. Xu, P. D. Brook, J. R. Smith, and Y. Matsuoka, “Physical human interactive guidance: Identifying grasping principles from human-planned grasps,” *Springer Tracts in Advanced Robotics*, vol. 95, no. 4, pp. 477–500, 2014.
- [8] T. Feix, I. M. Bullock, and A. M. Dollar, “Analysis of human grasping behavior: Object characteristics and grasp type,” *IEEE transactions on Haptics*, vol. 7, no. 3, pp. 311–323, 2014.
- [9] A. Kothari, J. Morrow, V. Thrasher, K. Engle, R. Balasubramanian, and C. M. Grimm, “Grasping objects big and small: Human heuristics relating grasp-type and object size,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2018.
- [10] T. Feix, J. Romero, H. Schmiedmayer, A. M. Dollar, and D. Kragic, “The grasp taxonomy of human grasp types,” *IEEE Transactions on Human-Machine Systems*, vol. 46, no. 1, pp. 66–77, Feb 2016.
- [11] E. Oztog, L.-H. Lin, M. Kawato, and G. Cheng, “Dexterous skills transfer by extending human body schema to a robotic hand,” in *Humanoid Robots*. IEEE, 2006, pp. 82–87.
- [12] M. R. Dogar and S. S. Srinivasa, “A Framework for Push-Grasping in Clutter,” *Transit*, pp. 1–6, 2011. [Online]. Available: <http://vault.willowgarage.com/wgdata1/vol1/muu11/WS-F-12-13.pdf>
- [13] —, “A planning framework for non-prehensile manipulation under clutter and uncertainty,” *Autonomous Robots*, vol. 33, no. 3, pp. 217–236, 2012.
- [14] D. Kappler, L. Chang, M. Przybylski, N. Pollard, T. Asfour, and R. Dillmann, “Representation of pre-grasp strategies for object manipulation,” *Humanoid Robots*, pp. 617–624, 2010.
- [15] C. Eppner and O. Brock, “Planning grasp strategies That Exploit Environmental Constraints,” *ICRA*, pp. 4947–4952, June 2015.
- [16] R. Deimel, C. Eppner, J. Álvarez-Ruiz, M. Maertens, and O. Brock, “Exploitation of environmental constraints in human and robotic grasping,” *Springer Tracts in Advanced Robotics*, vol. 114, pp. 393–409, 2016.
- [17] A. K. Goins, R. Carpenter, W.-K. Wong, and R. Balasubramanian, “Evaluating the efficacy of grasp metrics for utilization in a gaussian process-based grasp predictor,” in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2014, pp. 3353–3360.
- [18] J. Morrow, A. Kothari, Y. H. Ong, N. Harlan, R. Balasubramanian, and C. Grimm, “Using human studies to analyze capabilities in underactuated and compliant hands in manipulation tasks,” in *IROS 2018*, Madrid, Spain, 2018.
- [19] S. K. Allani, B. John, J. Ruiz, S. Dixit, J. Carter, C. Grimm, and R. Balasubramanian, “Evaluating human gaze patterns during grasping tasks,” *ACM Symposium on Applied Perception*, pp. 45–52, 2016.