# The Busboy Problem: Efficient Tableware Decluttering Using Consolidation and Multi-Object Grasps

Kishore Srinivas<sup>1</sup>, Shreya Ganti<sup>1</sup>, Rishi Parikh<sup>1</sup>, Ayah Ahmad<sup>1</sup>, Wisdom Agboh<sup>1,2</sup>, Mehmet Dogar<sup>2</sup>, Ken Goldberg<sup>1</sup>

*Abstract*—We present the "Busboy Problem": automating an efficient decluttering of cups, bowls, and silverware from a planar surface. As grasping and transporting individual items is highly inefficient, we propose policies to generate grasps for multiple items. We introduce the metric of Objects per Trip (OpT) carried by the robot to the collection bin to analyze the improvement seen as a result of our policies. In physical experiments with singulated items, we find that consolidation and multi-object grasps resulted in an 1.8x improvement in OpT, compared to methods without multi-object grasps. See https://sites.google.com/berkeley.edu/busboyproblem for code and supplemental materials.

#### I. INTRODUCTION

The post-meal task of clearing a dining table, commonly referred to as "bussing," requires moving cups, bowls, and utensils that are dispersed across the surface into a bin or tray to be cleaned in the kitchen. This is a common task that occurs after any event involving food service and dish collection, from daily household meals to casual picnics to formal cocktail parties and dinners. Automating this tedious and repetitive task could reduce fatigue and busy work for the skilled waiters who typically perform it.

We define the "Busboy Problem" as the efficient transfer of cups, bowls, and utensils (collectively called tableware) from the table into a designated collection bin while minimizing the time required for completion. This is an interesting problem for automation because the tableware are of varying shape, requiring low-level planning to execute grasps and high-level planning to consolidate tableware for efficient transport. Even small inaccuracies can lead to toppling or dropping delicate and expensive tableware, so the system must be extremely reliable.

Previous work in multi-object grasping, object manipulation, and grasp candidate generation highlight the efficiency of grasping pre-stacked objects as well as objects manually oriented for multi-object grasps [1], [2]. Whereas these works explore situations with objects are already positioned for said grasps, our work investigates methods of stacking and clustering objects into these favorable positions for multiobject grasps.

In this paper, we present a framework and algorithms for the Busboy Problem. We consider a scenario where multiple items are placed on a work surface (see Fig. 2), under an RGBD camera. We use the concept of multi-object grasping, which enables the robot to move multiple items



Fig. 1: The Busboy Problem. We present a combination of the robust action primitives to declutter a workspace of cups, bowls, and utensils.

simultaneously, thus reducing the number of pick-and-place actions needed.

This paper makes the following contributions:

- 1) Formulation of the Busboy Problem.
- 2) Action primitives for rearranging and grasping cups, bowls, and utensils.
- Two algorithms that leverage consolidation and multiobject grasps.
- 4) Experimental results indicating a 1.8x improvement in OpT.

#### II. RELATED WORK

#### A. Multi Object Grasping

Prior work on multi-object grasping includes different grasping techniques to facilitate multi-object grasps [3], detecting the number of objects in a grasp [4], decluttering surfaces [5], and multi-object grasping to place objects in virtual reality [6]. Yamada et al. considered the simplified multi-object grasping problem, where the objects are already in a configuration where they can be grasped at once [2]. Agboh et. al. [7] showed that friction can increase picks per hour for convex polygonal objects.

Some prior work has focused on the design of grippers for multi-object grasping. Jiang et. al. [8] proposed a vaccum gripper with multiple suction cups, while Nguyen et. al. [9] proposed a soft gripper based on elastic wires for multiobject grasping.

<sup>&</sup>lt;sup>1</sup>The AUTOLab at UC Berkeley (automation.berkeley.edu).

<sup>&</sup>lt;sup>2</sup>University of Leeds, UK.



Fig. 2: We employ the following robust action primitives to declutter a workspace through consolidation and multli-object grasps: single-object grasps, multi-object grasps, push-grasps, and stack-grasps.

Object stacking [1], [10], [11] has the potential to improve the number of objects per trip. We take inspiration from these works to include a stacking primitive.

## B. Pulling

Prior work by Berretty et al. has examined the use of inside-out pulling to orient convex polygonal parts [12]. We utilize a similar technique for circular cups and bowls. Furthermore, a planner for ensuring convergence to the final pose of pulling trajectories is proposed by Huang et al. [13], where they examine the motion of planar objects undergoing quasi-static movement.

### C. Grasp Candidates

Satish et al. discuss using a synthetic data sampling distribution that combines grasps sampled from the policy action set with guiding samples from a robust grasping supervisor to construct grasp candidates [14]. Additionally, Mahler et al. [15] discuss the use of energy-bounded caging to evaluate grasp candidates. They efficiently compute candidate rigid configurations of obstacles that form energy-bounded cages of an object, where the generated push-grasps are robust to perturbations. Mousavian et al. describe the process of using a variational autoencoder to generate grasps by mapping the partial point cloud of an observed object to a diverse set of grasps for the object [16]. Because of the relative simplicity of our setup, we found that an analytical approach to constructing grasp candidates is sufficient. In the case of bowls and cups, we sample a random point uniformly on the

rim and then orient the gripper perpendicular to the tangent of the circle at that point. In the case of utensils, we identify the axis of the utensil, and pick the highest depth point along that line, with the gripper perpendicular to the axis.

#### D. Object Manipulation in Cluttered Environments

Efficiently finding object manipulation plans in highdimensional environments with a large number of objects is a challenging problem. Hasan et al. [17] addressed this problem by identifying high-level manipulation plans in humans, and transferring these skills to robot planners. Other work by Tirumala et al. [18] used tactile sensing to singulate layers of cloth from a stack. Different from these works, our goal in the cluttered environment is to bring objects together, or stack them, to enable multi-object grasps.

## III. THE BUSBOY PROBLEM

The Busboy Problem involves the task of decluttering a workspace containing cups, bowls, and utensils, with the objective of minimizing both the time and number of trips required for completion.

### A. Assumptions

In the initial configuration, a planar workspace is defined in a cartesian grid (x, y) and has  $n_c$  cups,  $n_b$  bowls, and  $n_u$ utensils scattered across its surface. All items are assumed to be face up, visible by camera, and within a workspace defined by the constraints of the robot arm. These items may be initially stacked on top of one another or resting

## INTERNAL PULLING



Fig. 3: Bowls and cups are pulled with an internal pull, and utensils are grasped with a cage grip.

individually on the surface, and we assume that the initial state meets the following criteria:

- All items are of known dimensions, and cups and bowls are circular when viewed from top-down. Cups have radius 4.5cm, bowls have radius 8.5cm, and utensils are at most 17cm × 1.8cm.
- Cups and bowls are upright, and utensils are laid flat on the surface.
- Any stacks that exist are stable, such that  $r_0 \ge r_1 \ge \dots \ge r_s$ , where  $r_0$  represents the radius of the vertically lowest item, and  $r_s$  the highest one.
- Initially, no two items are touching (items are singulated).

#### B. State

We use cups, bowls, and utensils (forks and spoons) as the tableware set - collectively called "tableware" - in this work. Each cup and bowl has a position [x, y], and each utensil has a position [x, y] and orientation  $\theta$ .

#### IV. DECLUTTERING TABLEWARE

#### A. Action primitives

We propose to use a combination of manipulation primitives to solve the Busboy Problem. We specifically propose to use single object grasps, multi-object grasps, pull-grasps, and stack-grasps to efficiently clear a work surface of items (Figure 2).

1) Grasp: We use both single and multi-object grasps in this work. Let  $\mathbf{u}_G$  be the grasp to pickup objects — single or multiple. We represent this action as:

$$\mathbf{u}_G = \left[\mathbf{p}_G, \theta_G\right] \tag{1}$$

where  $\mathbf{p}_G = [x_G, y_G, z_G]$  is the center point of the grasp, and  $\theta_G$  is the grasp orientation.

2) *Pull-Grasp:* A pull-grasp action involves two steps: a pull of one object to another, then a multi-object grasp of both objects. We represent a pull action as:

$$\mathbf{u}_P = \left[\mathbf{p}_S, \theta_S, \mathbf{p}_E, \theta_E\right] \tag{2}$$

where  $\mathbf{p}_S = [x_S, y_S, z_S]$  is the pull start point,  $\theta_s$  is the gripper orientation at the state point,  $\mathbf{p}_E = [x_E, y_E, z_E]$  is the pull end point, and  $\theta_E$  is the gripper orientation at  $\mathbf{p}_E$ . For circular objects such as bowls and cups, the gripper pulls outwards from the center of the dish using an internal pull, and for utensils, the gripper cages the utensil around its center point while moving it (Figure 3). Then, we denote a pull-grasp action as:

$$\mathbf{u}_{\mathbf{PG}} = \left[\mathbf{u}_P, \mathbf{u}_G\right] \tag{3}$$

*3) Stack-Grasp:* A stack-grasp action involves two steps: a stack of one object onto another, then a multi-object grasp of both objects. We represent a stack action as:

$$\mathbf{u}_{\mathbf{S}} = \left[\mathbf{u}_{G_i}, \mathbf{p}_L, \theta_L\right] \tag{4}$$

where  $\mathbf{u}_{G_i}$  is a grasp on the lifted object, and  $\mathbf{p}_L = [x_L, y_L, z_L]$  is the placement point on the stationary object, and  $\theta_L$  is the gripper orientation at  $\mathbf{p}_L$ . Then, we denote a stack-grasp action as:

$$\mathbf{u_{SG}} = \begin{bmatrix} \mathbf{u}_S, \mathbf{u}_G \end{bmatrix} \tag{5}$$

## B. Determining allowable actions

1) Grasp: A single-object grasp is always allowable. We can safely assume this since any dish or stack of items is already top-down graspable. When no other actions are allowed, the single-object grasp action is used as a default to clear the workspace.

A multi-object grasp is allowable when the grasp heights of both items are similar (within an adjustable threshold value) and if the lateral distance between the grasp points of both items is less than the width of the gripper. If the grasp heights of the items are significantly different, the gripper will have to either collide with the taller dish while attempting to grasp the shorter dish or grasp only the taller dish to avoid the collision, and either case results in a failure of grasping multiple items at once. Similarly, if the items are separated by more than the maximum inside width of the grippers, an attempt to grasp both at the same time will fail.

2) *Pull:* A pull of two items is allowable if a multi-object grasp can be executed on those items and if no other objects lie between the two items on the workspace. We disallow pull actions of items for which a multi-object grasp cannot be executed, since the pull becomes a wasted action. We also disallow pull actions of items with other objects between them to ensure that the intermediate objects are not displaced in a non-deterministic manner.

3) Stack: A stack of dish  $d_a$  with radius  $r_a$  onto dish  $d_b$  with radius  $r_b$  is allowable if  $r_a \leq r_b$ . This means that a cup can be stacked onto a bowl, but not vice versa, and that a utensil can be stacked onto any other dish, including another utensil. This is to ensure that the stack stability assumption present at the initial state remains valid after each action.

## ROBUSTNESS



Fig. 4: Uncertainty in a cup or bowl's location decreases once the arm has interacted with the dish. The position of the dish relative to the gripper's location is known with greater certainty after the action.

## C. Robustness of action primitives

We present three primitives to robustly execute the above actions. This design makes the primitives more robust.

1) Grasp: When executing a grasp at location x, y, z, the robot will open its grippers centered around x, y, and then move down to the appropriate height, as measured by the depth sensor, before closing the gripper to grasp the object. The affordances granted by max gripper opening, gripper height, and gripper width mean that an off-center grasp point x, y, z will still successfully complete the single-object or multi-object grasp of the object (Figure 4).

2) Pull: For cups and bowls, the gripper pulls outwards from the center of the dish, contacting the inner surface of the dish (Figure 3). This action is successful as both  $r_b$  and  $r_c$  are larger than the width of the gripper when closed. If the gripper is anywhere within the opening of the object, it will be able to move the target object to a specified location. For utensils, the gripper cages the utensil around its center point while moving it, preventing unwanted rotation and moving the utensil to its specified location.

3) Stack: For bowls and cups, the top lip radius is larger than the radius of the base, giving the sides a taper. Because a dish  $d_a$  is only stacked onto another dish  $d_b$  of equal or larger size, the base radius of  $d_a$  is guaranteed to be smaller than the top radius of  $d_b$ , allowing the tapered sides of the items to funnel  $d_a$  into place even if there is slight error in the placement of the dish. Placing a utensil onto a bowl is extremely robust to error because of the relative radii of the items, and placing a utensil onto another utensil is robust due to the curvature of the utensils themselves which slide a misplaced utensil into place, making them naturally conducive to stacking.

## D. Policies

1) Pull Policy: The pull policy combines Pull-Grasp and Grasp actions. From the initial scene, it checks if any multiobject grasps can be executed right away, and executes those first. Then, it runs the Pull-Grasp action for all remaining items, pulling together items that don't cause collisions and executing multi-object grasps to clear them from the workspace. If any items remain after all possible multi-object grasps are executed, those items are cleared with single-object Grasp actions. After each action, a new image of the workspace is taken and the state representation is updated to reflect the new state of the workspace, including any tableware that has been moved or left behind by the previous action. This policy is formalized in Algorithm 1.

2) Stack Policy: The stack policy combines Stack-Grasp and Grasp actions. It repeatedly executes the Stack-Grasp action to clear the workspace, and if there are any remaining items they are cleared with single-object Grasp actions. It prioritizes stacking utensils onto bowls and transporting them to the bin, and then tries to stack the remaining dishes. Stacking utensils first is an efficient way to improve the number of OpT for this policy. The policy is formalized in Algorithm 2.

After utensils are cleared, the stacks created by this policy are limited to be a combination of at most 2 existing stacks (i.e. once a Stack action is executed, the next action is necessarily a Grasp on the resulting stack, not another Stack action onto that stack). This is because when 4 or more bowls or cups are stacked, the height difference between the lip of the top dish and the lip of the bottom dish exceeds the height of the gripper jaws, causing many attempted grasps to fail. By limiting stacks to at most 2 existing stacks, we significantly reduce the chances of creating a stack with more than 3 dishes.

If I unorasphossible $(a_i, a_j) \forall i, j \in len(D)$ then
$Pull(d_i, d_j);$
$Grasp(d_i, d_j);$
$D \leftarrow UpdateStateFromImage();$
continue;
$  Grasp(d_i) \forall i \in len(D):$

 $D \leftarrow UpdateStateFromImage();$ 

#### V. EXPERIMENTS AND RESULTS

We evaluate through physical experiments the robustness of the pulling action primitive and then evaluate the pull and Algorithm 2 Stack Policy

**Input:**  $D \leftarrow [d_1, d_2, ..., d_n]$  - initial distribution of *n* items to be cleared.

0	utput:	Executes	an	action	sequence	at	each	iteration	ı.
---	--------	----------	----	--------	----------	----	------	-----------	----

stack policies on a real-world table clearing task.

## A. Experimental Setup

We use a UR5 robot arm with a Robotiq 2F-85 gripper and Intel RealSense 455D RGBD camera mounted 83cm above the workspace. The workspace is a flat 78cm x 61cm surface with 4 cups, 4 bowls, and 4 utensils,  $n_b = n_c = n_u = 4$ . In our experimental setup, we calculated a max gripper opening of w = 8.5cm, gripper height of h = 4.5cm, bowl radius  $r_b = 8.5cm$ , cup radius  $r_c = 4.5cm$  and utensil width  $r_u = 1.8cm$ .

We identify and locate tableware on the workspace with a vision pipeline. Since the surface of the workspace is white, we use darker colored tableware to be easily visible. To locate cups and bowls, we first use edge detection, contour forming, and HoughCircles to identify circular shapes on the workspace, then filter these circles based on the known image radius of cups and bowls. We cluster these circles by their centers and remove circles that overlap beyond a specified threshold, allowing an unambiguous detection of cups and bowls. To locate utensils, we use edge detection and contour forming, and then filter out the contours that are too "square", as determined by the aspect ratio of the identified contour. We draw an imaginary line through the lengthwise center of bounding rectangle of the contour, and sample depth values along that line; we use the highest depth point as the grasp point of the utensil to allow the gripper maximum clearance with the surface.

We define three tiers to evaluate the performance of our algorithm on scenes of increasing complexity.

- Tier 0: scenes contain 6 items, either all cups, all bowls, or all utensils, with no stacks in the initial state.
- Tier 1: scenes contain 4 items each of cups, bowls, and utensils, and have no stacks in the initial state.
- Tier 2: scenes contain 4 items each of cups, bowls, and utensils, but we allow stacks of at most 3 objects in the initial state.

For Tier 2, we limit initial stacks to at most 3 objects because of the dimensions of the gripper, as mentioned in Section IV-D.2. The number of objects in a stack, and not the actual dimensions of individual dishes, is the main limiting factor for the grasp, because we grasp dishes from the rim. The dishes could actually be much larger and still be graspable as long as the walls are thin enough to allow the gripper to slide over them, and the weight of the dish does not exceed the payload limitations of the gripper itself. We limit ourselves to a small set of known kitchenware objects for consistency in our experiments.

We evaluate the performance of the pull and stack policies against a baseline single-item policy, referred to as "Random" in Table I. This policy picks a dish at random, and if the dish is a cup or bowl, it uniformly samples a point on the rim and grasps the dish at that point. If the dish is a utensil, it identifies the grasp point of the utensil as described above and grasps the utensil at that point. This policy is stackagnostic, so even in Tier 2 when there are stacks present in the initial state, it treats each item in the stack as its own object, and clears the stack by transporting one item at a time.

#### B. Scene Generation

In order to evaluate our policies, we generate multiple scenes at each tier, and every policy is run once on each scene. To generate each scene, we use the dimensions of the workspace (78cm  $\times$  61cm), and  $r_b, r_c, r_u$  for the dimensions of the objects. We randomly sample x, y locations within the scene for each object. If an object intersects with another object, we create a stack of the two objects if the maximum number of intersections has not been exceeded, and resample a position for the object if it has. Tiers 0 and 1 allow no such intersections, whereas Tier 2 allows 4 intersections. For each trial we manually reset the scene to maintain consistency.

## C. Evaluation

We evaluated on 9 scenes at Tier 0 (3 scenes per type of dish), 3 scenes at Tier 1, and 3 scenes at Tier 2. A trial is one execution of one policy on one scene, so we have a total of (9+3+3)\*3 = 45 trials. For each trial, we record the time in seconds to clear the table, the OpT, and the number of failures. A failure occurs when the robot is unable to move all items to the collection bin, either because of a perception failure that leaves items behind on the workspace or a policy failure that drops a dish off the workspace. We report our results in Table I.

To evaluate the performance of our policies in more realistic scenario, we present the theoretical improvement in execution time when the bin is placed further away from the workspace, as might be seen in a home or professional kitchen. Given the physical limitation of the UR5 arm length, we simulated the lengthening distance by adding time delays of 3 and 5 seconds in both directions of motion (to and from the collection bin). We find that moving the bin further away causes the stack and pull policies to perform significantly better than the baseline policy because motions to and from the bin are penalized, making policies with fewer total

Tier	Policy	Time (sec)	Objects per Trip	Failures	Time Ratio	<b>OpT Ratio</b>
Tier 0 Cups	Random	78.2	0.8	0	-	-
	Stack	58.5	2.0	0	1.4	2.6
	Pull	48.8	1.6	2	1.6	2.1
Tier 0 Bowls	Random	63.3	1.0	0	-	-
	Stack	60.2	2.0	0	1.1	2.0
	Pull	41.3	1.8	0	1.5	1.8
Tier 0 Utensils	Random	64.1	1.0	0	-	-
	Stack	64.3	1.8	0	1.0	1.8
	Pull	55.3	1.8	1	1.2	1.8
Tier 1	Random	121.3	1.0	1	-	-
	Stack	111.3	2.0	2	1.1	2.0
	Pull	102.1	1.6	3	1.2	1.6
Tier 2	Random	93.5	1.4	0	-	-
	Stack	88.2	2.6	2	1.1	1.8
	Pull	84.2	2.3	3	1.1	1.6

TABLE I: **Physical Experiments** We present the total time and number of trips to clear a table for each policy. We found that compared to the baseline policy, the **stack** policy makes at least a **1.8x** improvement in the number of objects grasped per trip (OpT) and the **pull** policy makes at least a **1.6x** improvement.

actions perform better. We report these results in Table III in the appendix of the project website.

#### VI. DISCUSSION

Results show that using consolidation and multi-object grasps allows clearing the workspace efficiently, with the pull policy transporting at least 1.6x as many objects per trip, and the stack policy at least 1.8x. A discussion of resulting execution time improvement is in the appendix of the project website.

#### VII. LIMITATIONS AND FUTURE WORK

An overhead RGBD camera gives only a clear top view. This affects state estimation and can lead to failures. We assume circular cups and bowls. This makes it easy to compute grasps. For more general dishes, advanced grasp generation methods will be needed. In future work, we will loosen the assumption of starting with singulated objects. We also hope to combine the pull and stack policies into a higher-level policy that can efficiently clear the workspace.

#### VIII. ACKNOWLEDGMENTS

This research was performed at the AUTOLAB at UC Berkeley in affiliation with the Berkeley AI Research (BAIR) Lab, and the CITRIS "People and Robots" (CPAR) Initiative. The authors were supported in part by donations from Toyota Research Institute, Bosch, Google, Siemens, and Autodesk and by equipment grants from PhotoNeo, NVidia, and Intuitive Surgical. Mehmet Dogar was partially supported by an EPSRC Fellowship (EP/V052659).

#### REFERENCES

A. X. Lee, C. Devin, Y. Zhou, T. Lampe, K. Bousmalis, J. T. Springenberg, A. Byravan, A. Abdolmaleki, N. Gileadi, D. Khosid, C. Fantacci, J. E. Chen, A. Raju, R. Jeong, M. Neunert, A. Laurens, S. Saliceti, F. Casarini, M. A. Riedmiller, R. Hadsell, and F. Nori, "Beyond pick-and-place: Tackling robotic stacking of diverse shapes," *CoRR*, 2021.

- [2] T. Yamada, S. Yamanaka, M. Yamada, Y. Funahashi, and H. Yamamoto, "Grasp stability analysis of multiple planar objects," in *IEEE International Conference on Robotics and Biomimetics*, 2009.
- [3] Y. Sun, E. Amatova, and T. Chen, "Multi-object grasping types and taxonomy," in *IEEE ICRA*, 2022.
- [4] T. Chen, A. Shenoy, A. Kolinko, S. Shah, and Y. Sun, "Multi-object grasping – estimating the number of objects in a robotic grasp," in *IEEE IROS*, 2021.
- [5] W. C. Agboh, J. Ichnowski, K. Goldberg, and M. R. Dogar, "Multiobject grasping in the plane," in *ISRR*, 2022.
- [6] U. J. Fernández, S. Elizondo, N. Iriarte, R. Morales, A. Ortiz, S. Marichal, O. Ardaiz, and A. Marzo, "A multi-object grasp technique for placement of objects in virtual reality," *Applied Sciences*, 2022.
- [7] W. C. Agboh, S. Sharma, K. Srinivas, M. Parulekar, G. Datta, T. Qiu, J. Ichnowski, E. Solowjow, M. Dogar, and K. Goldberg, "Learning to efficiently plan robust frictional multi-object grasps," *CoRR*, 2022.
- [8] P. Jiang, J. Oaki, Y. Ishihara, and J. Ooga, "Multiple-object grasping using a multiple-suction-cup vacuum gripper in cluttered scenes," 2023.
- [9] V. P. Nguyen and W. T. Chow, "Wiring-claw gripper for soft-stable picking up multiple objects," *IEEE Robotics and Automation Letters*, vol. 8, no. 7, pp. 3972–3979, 2023.
- [10] H. Huang, L. Fu, M. Danielczuk, C. M. Kim, Z. Tam, J. Ichnowski, A. Angelova, B. Ichter, and K. Goldberg, "Mechanical search on shelves with efficient stacking and destacking of objects," *ISRR*, 2022.
- [11] Y. Jiang, M. Lim, C. Zheng, and A. Saxena, "Learning to place new objects in a scene," *IJRR*, vol. 31, no. 9, pp. 1021–1043, 2012.
- [12] R.-P. Berretty, K. Goldberg, M. H. Overmars, and A. van der Stappen, "Orienting parts by inside-out pulling," *IEEE ICRA*, 2001.
- [13] E. Huang, A. Bhatia, B. Boots, and M. T. Mason, "Exact bounds on the contact driven motion of a sliding object, with applications to robotic pulling," in *Robotics: Science and Systems*, 2017.
- [14] V. Satish, J. Mahler, and K. Goldberg, "On-policy dataset synthesis for learning robot grasping policies using fully convolutional deep networks," *IEEE Robotics and Automation Letters*, 2019.
- [15] J. Mahler, F. T. Pokorny, S. Niyaz, and K. Goldberg, "Synthesis of energy-bounded planar caging grasps using persistent homology," *IEEE Transactions on Automation Science and Engineering*, 2018.
- [16] A. Mousavian, C. Eppner, and D. Fox, "6-dof graspnet: Variational grasp generation for object manipulation," *ICCV*, 2019.
- [17] M. Hasan, M. Warburton, W. C. Agboh, M. R. Dogar, M. Leonetti, H. Wang, F. Mushtaq, M. Mon-Williams, and A. G. Cohn, "Humanlike planning for reaching in cluttered environments," *ICRA*, 2020.
- [18] S. Tirumala, T. Weng, D. Seita, O. Kroemer, Z. Temel, and D. Held, "Learning to singulate layers of cloth using tactile feedback," in *IEEE/RSJ IROS*, 2022, pp. 7773–7780.