



UNIVERSITY OF LEEDS

This is a repository copy of *Learning to Efficiently Plan Robust Frictional Multi-Object Grasps*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/202110/>

Version: Accepted Version

Proceedings Paper:

Agboh, W.C., Sharma, S., Srinivas, K. et al. (7 more authors) (2023) Learning to Efficiently Plan Robust Frictional Multi-Object Grasps. In: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems. 2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 01-05 Oct 2023, Detroit, USA. IEEE , pp. 10660-10667. ISBN 978-1-6654-9191-4

<https://doi.org/10.1109/IROS55552.2023.10341895>

This is an author produced version of a conference paper accepted for publication in Proceedings of the 2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), made available under the terms of the Creative Commons Attribution License (CC-BY), which permits unrestricted use, distribution and reproduction in any medium, provided the original work is properly cited.

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Learning to Efficiently Plan Robust Frictional Multi-Object Grasps

Wisdom C. Agboh^{*1,2}, Satvik Sharma^{*1}, Kishore Srinivas¹, Mallika Parulekar¹
Gaurav Datta¹, Tianshuang Qiu¹, Jeffrey Ichnowski³, Eugen Solowjow⁴, Mehmet Dogar², Ken Goldberg¹

Abstract—We consider a decluttering problem where multiple rigid convex polygonal objects rest in randomly placed positions and orientations on a planar surface and must be efficiently transported to a packing box using both single and multi-object grasps. Prior work considered frictionless multi-object grasping. In this paper, we introduce friction to increase the number of potential grasps for a given group of objects, and thus increase picks per hour. We train a neural network using real examples to plan robust multi-object grasps. In physical experiments, we find a 13.7% increase in success rate, a 1.6x increase in picks per hour, and a 6.3x decrease in grasp planning time compared to prior work on multi-object grasping. Compared to single-object grasping, we find a 3.1x increase in picks per hour.

I. INTRODUCTION

When skilled waiters clear tables, they grasp multiple utensils and dishes in a single motion. Similarly, it is inefficient for robotic picking systems in warehouses and fulfillment centers to only handle a single object at a time. Picking multiple objects at once can significantly increase picks per hour (PPH), the total number of objects picked from a scene in an hour. In prior work on multi-object grasping [1], PPH was increased compared to single-object picking. This improvement was limited due to a frictionless grasping assumption and no considerations of robustness. In this work, we find that considering friction and quickly generating robust grasps can lead to significant improvements in PPH. For example, grasps like those shown in Fig. 1 cannot exist without appropriate friction between objects. An important question that then arises is how to generate such robust frictional grasps.

Robust grasps have been generated in prior work [2, 3, 4, 5], but only for single objects. Inspired by these works, we develop a robust multi-object grasping system for planar convex polygonal objects. Instead of using a physics simulator, we propose to collect data entirely on a physical robot and use it to train a multi-object grasping function, *MOG-Net*, which is robust to state and control uncertainty and predicts the number of objects that will be grasped out of a target object group. We train in real to avoid the sim-to-real gap [6, 7, 8]. We also propose a necessary condition for frictional multi-object grasping to filter out inadmissible grasps and show that this




Figure 1. The decluttering problem (top) where objects must be transported to a packing box. We find robust frictional multi-object grasps (bottom) to efficiently declutter the scene.

filtering leads to a high quality dataset and saves valuable physical robot time during data collection.

We use *MOG-Net* in a novel grasp planner to generate robust multi-object grasps. The planner maximizes the predicted number of objects grasped per pick attempt in cluttered scenes. To improve robustness to state and control uncertainty, we weight *MOG-Net*'s predictions with the probability of satisfying multi-object grasping necessary conditions, obtained via Monte-Carlo sampling.

We find a 13.7% increase in success rate, a 1.6x increase in picks per hour, and an 6.3x decrease in grasp planning time compared to prior work [1] on multi-object grasping. Compared to single-object grasping we find a 3.1x increase in picks per hour. This work makes 4 contributions:

- 1) The derivation of a frictional multi-object grasping necessary condition to filter inadmissible grasps.
- 2) *MOG-Net*: a robust multi-object grasp neural network, self-supervised in real to predict the number of objects grasped out of a target group, given a grasp candidate.
- 3) A grasp planning algorithm, μ -MOG, that generates grasps that are robust to state and control uncertainty, by considering the probability of necessary conditions being satisfied.
- 4) Physical experiments evaluating the methods with ran-

¹University of California, Berkeley, USA.

²University of Leeds, UK.

³Carnegie Mellon University, Pittsburgh, USA.

⁴Siemens Research Lab, Berkeley, USA.

*equal contribution

domized scenes of 58 polyhedral objects totaling 2532 grasp attempts.

II. RELATED WORK

Decluttering or picking multiple objects from a table is a common robotics problem [9] which has mainly been addressed with single object grasps [10, 11]. In this section, we discuss prior work on multi-object grasping, frictional single object grasping, and robust grasp synthesis.

A. Multi-object grasps

Harada and Kaneko [12] proposed some of the earliest works on multi-object grasping. In [12, 13, 14] they develop conditions for enveloping grasps of multiple objects using a multi-fingered robot hand, and under a rolling contact assumption. Not long after, Yamada et al. [15] proposed a series of methods [15, 16, 17, 18, 19] to evaluate the grasp stability of multiple planar objects grasped by a multi-fingered robot hand. While these were the pioneering works on multi-object grasping, their results focused on numerical simulations, without physical robot multi-object grasps. This paper derives conditions for equilibrium multi-object grasps, under the frictional point-contact model and shows physical robot multi-object grasps.

Recently, Chen et al. [20] investigated the problem of dipping a robot hand inside a pile of identical spherical objects, closing the hand and estimating the number of objects remaining in the hand after lifting. Shenoy et al. [21] focused on the same problem but with a goal of transferring the picked spherical objects to another bin. Our work is focused on frictional multi-object grasps of arbitrary convex polygonal objects spread over a plane, so physics-based planning is required [22, 23, 24, 25].

Sakamoto et al. [26] proposed a picking system that first uses robot pushing [27, 28, 29], to move one cuboid to the other and thereafter grasps both cuboids in separate actions. In this paper, we take a single-step push-grasp action and derive conditions for multi-object grasping under the frictional point contact model. Given the uncertainty in grasping systems [30, 31], we plan robust multi-object grasps to improve PPH.

B. Frictional grasps

Some prior works on single object grasping have studied the effect of friction. Golan et al. [32] develop a gripper that can switch between being frictionless and frictional, and show that having a frictional gripper provides more secure and robust grasps in real. Hang et al. [33] propose a measure of friction sensitivity to assess how grasp quality varies with the coefficient of friction. Inspired by nature, Roberge et al. [34] use a gecko-inspired adhesive on a gripper to apply large shear forces with low normal forces for grasping a single object. Agboh et al. [1] did not consider friction in multi-object grasping. However, for single objects, friction can increase the number of stable grasps. Using a point contact model, the friction cone is larger for a higher coefficient of friction, meaning that more stable single object grasps can be found.

We show that the same is true for multi-object grasping. We focus on frictional multi-object grasps where objects are pushed together before they are grasped.

C. Robust grasping

Uncertainty in state and model disparity result in grasp failure. An important line of grasping work focuses on generating grasps that are robust to uncertainty. To the best of our knowledge there has been no prior work on robust multi-object grasping. Prior work on robust single object grasping typically falls into one of two categories — analytic or data-driven.

One analytic approach to robust grasp synthesis is the use of funnels [31, 35]. These are primitives that can analytically reduce or ‘funnel’ uncertainty from some initial set to a smaller target set. Bhatt et al. [36] use a sequence of funnels to perform open-loop and robust in-hand manipulation. Another analytic method is caging [37, 38, 39] where a grasped object’s mobility is bounded such that it ‘follows’ when the gripper moves. These analytic methods exploit the object’s environment and shape, using carefully chosen primitives to generate robust grasps. Our work focuses on data-driven methods to generate robust grasps.

Popular data-driven single-object robust grasp synthesis approaches are Dex-Net 2.0 and Dex-Net 4.0 [2] [40]. They train a grasp quality convolutional neural network (GQ-CNN) with synthetic data to predict grasp success probability. Similarly, Dex-Net 3.0 trains a GQ-CNN for robust grasps on single objects but for suction cups [41].

There is a sim-to-real gap [6, 7, 8] when grasps are trained in simulation. Robustness in these simulation settings has been achieved through domain randomization or general Monte-Carlo sampling [42, 43, 44].

In this work, we propose a frictional multi-object grasp necessary condition and use it to filter inadmissible multi-object grasp candidates. To be robust to state and control uncertainty, we estimate the probability of satisfying these necessary conditions through Monte-Carlo sampling. We avoid the sim-to-real problem by training *MOG-Net* entirely in real.

III. PROBLEM STATEMENT

We consider a decluttering problem where multiple rigid convex polygonal objects rest in randomly placed positions and orientations on a planar surface, visible from an overhead camera, and must be transported to a packing box. The objective of this work is to develop a decluttering algorithm that maximizes picks per hour (PPH) for this problem, using robust frictional grasps. Note that we do not consider rearrangement actions (e.g. pushing actions) for arranging the groups before multi-object grasping in this work. Finding the optimal rearrangement plan in such scenes is a challenging long-horizon problem that deserves a separate thorough treatment.

A. Assumptions

We assume that the gripper is a parallel-jaw gripper. We acknowledge that multi-fingered grippers provide more opportunities for multi-object grasps but we focus on parallel-jaw grippers as they are common. We assume that objects




Figure 2. An overview of the decluttering system proposed in this paper. It finds the maximum group of objects that can fit in the gripper and generates a robust grasp for that group. First, it generates candidate grasps, and for each grasp \mathbf{u}^k , estimates a probability of satisfying multi-object grasp necessary conditions (γ^k), under state and control uncertainty. Thereafter, it uses *MOG-Net* which was trained in real to predict the number of objects (N_g^k) that will be grasped using \mathbf{u}^k . The chosen robust grasp maximizes the product $\gamma^k \cdot N_g^k$. We execute the robust grasp and continue to the next object group until the table is cleared of all objects.

are extruded convex polygons laying on a flat, uniform color surface, and that we have a set value for the lower bound for μ , the coefficient of (Coulomb) friction, for all contact interactions. We also assume antipodal multi-object grasps where each object is kept in equilibrium by two neighbouring objects, or one object and a gripper jaw. We further assume that a group of objects in force closure will be securely grasped during motion, and neither the grasping force nor the speed of the motion will dislodge the objects.

B. State and action

The state \mathbf{x} is a list of all convex polygonal objects, where each object, \mathbf{o}_i , in the list is represented by its vertices:

$$\mathbf{x} = [\mathbf{o}_i] \quad \forall i \in [0 \dots N_o - 1]$$

$$\text{where } \mathbf{o}_i = [\{x_i^0, y_i^0\}, \dots, \{x_i^v, y_i^v\}, \dots, \{x_i^{N_v-1}, y_i^{N_v-1}\}]$$

Here, N_o is the number of objects on the table, $\{x_i^v, y_i^v\}$ represents the 2D position of vertex v of object i , provided by an overhead camera, and N_v is the maximum number of vertices for each object i .

We represent single and multi-object grasp actions in the same way: $\mathbf{u} = [x_g, y_g, \theta_g]$, where x_g, y_g , and θ_g represent the desired grasp pose of the gripper, after which the jaws close with a maximum force f_g .

IV. DECLUTTERING WITH MULTI-OBJECT GRASPS

We present an overview of the decluttering system in Fig. 2. Given an initial cluttered scene, we take a greedy approach and find an object group with the maximum number of objects that can fit in the gripper. The next step is to plan a robust

multi-object grasp for this object group. We sample candidate grasps within the convex hull of the objects (see Sec. IV-B). Thereafter, we estimate the probability γ^k that the k th grasp \mathbf{u}^k will satisfy the multi-object grasp necessary conditions (see Sec. V), under state and control uncertainty. We also query *MOG-Net* (see Sec. VI) to predict the number of objects, N_g^k that the given grasp will successfully pick. Finally, we choose the grasp that maximizes the robust prediction $\gamma^k \cdot N_g^k$. In the following subsections, we provide details of the decluttering algorithm and the robust multi-object grasp planner.

A. Decluttering

Alg. 1 details the decluttering algorithm, which is similar to the picking algorithm in prior work [1]. We estimate the current state \mathbf{x} , containing N_o objects, using color segmentation to isolate objects and using their convex hulls to find vertices (line 2). The algorithm then uses the subroutine *CreateObjGroups(.)* (line 3) to create a set of distinct object groups. It loops through center points of objects and creates groups of all objects that are half a gripper width radius away. This also includes all single object groups. Then, the *RankObjGroups(.)* subroutine (line 4) ranks the list of object groups by their size. The *RobustGraspPlanner(.)* subroutine (line 6) in section IV-B finds a grasp for the largest object group. The grasp is executed and the whole process repeats until the table is cleared or a time limit is reached.

B. Robust multi-object grasp planning

One main distinction from prior work is the robust multi-object grasp planner μ -MOG, which is detailed in Alg. 2.

Algorithm 1: Decluttering Algorithm

```

1 do
2    $\mathbf{x}, N_o \leftarrow \text{GetCurrentState}(\cdot)$ 
3    $\text{obj\_groups} \leftarrow \text{CreateObjGroups}(\mathbf{x})$ 
4    $\text{ranked\_obj\_groups} \leftarrow \text{RankObjGroups}(\text{obj\_groups})$ 
5   for  $\text{obj\_group}$  in  $\text{ranked\_obj\_groups}$  do
6      $\mathbf{u} \leftarrow \text{RobustGraspPlanner}(\mathbf{x}, \text{obj\_group})$ 
7     if  $\mathbf{u} \neq \{\}$  then
8       Execute  $\mathbf{u}$ 
9       break
10  while  $N_o > 0$  and time remaining;
```

Algorithm 2: μ -MOG

```

Input      :  $\mathbf{x}$ : Current state
              $\text{obj\_group}$ : Objects in the potential grasp
Output     :  $\mathbf{u}^r$ : A robust grasp action
              $N_g^r$ : Predicted number of objects with  $\mathbf{u}^r$ 
1   $\text{grasp\_cands} \leftarrow \text{GenGraspCands}(\mathbf{x}, \text{obj\_group})$ 
2  for  $\mathbf{u}^k$  in  $\text{grasp\_cands}$  do
3   |  $\gamma^k \leftarrow \text{NecessaryCondsProba}(\mathbf{x}, \mathbf{u}^k)$ 
4   |  $N_g^k \leftarrow \text{MOG-Net}(\mathbf{x}, \mathbf{u}^k)$ 
5   $\mathbf{u}^r \leftarrow \arg \max_{\mathbf{u}^k} (\gamma^k \cdot N_g^k)$ 
6  return  $\mathbf{u}^r, N_g^r$ 
```

The algorithm generates multiple grasp candidates (line 1) using `GenGraspCands(.)`. It finds the convex hull of a given group of objects and generates N_p points that uniformly cover the convex hull. At each point, it generates N_θ orientation samples. It rejects grasp samples that result in collisions between the gripper jaws and any object. Next, it loops through grasp candidates (lines 2-4) and estimates (i) the probability γ^k of satisfying necessary conditions using `NecessaryCondsProba(x, u^k)`, and (ii) the predicted number of objects that u^k will successfully grasp N_g^k , using `MOG-Net`.

To calculate γ^k , `NecessaryCondsProba(.)` performs Monte-Carlo sampling so that the relative position of the grasp candidate varies with respect to the position of the objects in the group. Specifically, we consider samples $u' = u + \delta u$, and $x' = x + \delta x$ where $\delta u \sim \mathcal{N}(0, \sigma_u^2)$, $\delta x \sim \mathcal{N}(0, \sigma_x^2)$, and σ_u , σ_x are standard deviations for control and state respectively. Then, it returns the ratio of grasp samples (γ^k) that satisfy the necessary conditions, under state and control uncertainty. Finally, we choose the robust grasp u^r such that:

$$u^r = \arg \max_{u^k} (\gamma^k \cdot N_g^k) \quad (1)$$

C. Robustness to Frictional Uncertainty

We assumed a lower bound on the coefficient of friction for all objects in Sec. III-A. This is a conservative assumption to allow for frictional uncertainty. Lower values of μ mean fewer admissible grasp candidates. Thus, any grasp candidate that satisfies the necessary conditions for the lower bound will also satisfy the conditions for higher values of μ .

V. NECESSARY CONDITIONS FOR MULTI-OBJECT GRASPING

Prior work [1] studied frictionless multi-object grasping. In this work, we extend the analysis to include friction and derive the frictional necessary conditions to achieve frictional force closure.

A. Frictional equilibrium multi-object grasps

Under the frictionless point contact model, the number of possible antipodal grasp configurations for a polygon is limited. For example, a triangle has only one possible equilibrium grasp configuration: a vertex and an opposing edge. With friction, it is possible to achieve more equilibrium grasps, which is dependent on the coefficient of friction. We first analyze equilibrium grasps for a single object under the frictional point contact model and extend those results to multiple objects.

In single object grasping with a parallel-jaw gripper, frictional equilibrium grasps occur at pairs of contacts where the friction cones contain opposing forces that lie on the line passing through them. Recall that μ is the lower bound on the coefficient of friction at the left and right contacts. Then, the friction cones (C_l and C_r) are characterized by $\alpha_l = \alpha_r = \tan^{-1}(\mu)$, and are centered on the contact normals (\hat{n}_l , \hat{n}_r). If the line L_g that passes through both contact points is contained in both friction cones, the parallel-jaw grasp is in equilibrium.




Figure 3. A top-down view of a 3-object frictional equilibrium grasp. We check that each object i is in an equilibrium grasp by inspecting their left (C_{l_i}) and right (C_{r_i}) friction cones and the line (L_{g_i}) passing through both contact locations. Also, we ensure that all connecting lines lie on the same line ($\hat{L}_{g_i} = \hat{L}_{g_{i+1}} = \dots = \hat{L}_{g_{n_o}}$). The friction cones are centered on the contact normal and are defined by the coefficient of friction μ . $\alpha = \tan^{-1}(\mu)$.

For a convex polygonal object, it is then possible to consider discrete points along the object's surface, and enumerate equilibrium grasps by considering opposing contact pairs using the friction cones. The number of equilibrium grasps can be infinite depending on the size of the friction cone (α).

We require each individual object in a frictional multi-object grasp to be in an equilibrium grasp. Consider Fig. 3. It shows a sample frictional multi-object equilibrium grasp for 3 objects. The left (C_{l_i}) and right (C_{r_i}) friction cones for object i are characterized by $\alpha_{l_i} = \alpha_{r_i} = \tan^{-1}(\mu)$. To achieve an equilibrium multi-object grasp for a group of n_o objects, both friction cones for each object $i \in \{0, 1, \dots, n_o - 1\}$, must contain opposing forces that lie on line L_{g_i} , connecting their contact locations. Since the forces at object-object contacts are reactionary, all connecting lines (L_{g_i} in Fig. 3) must lie on the same line:

$$\hat{L}_{g_i} = \hat{L}_{g_{i+1}} = \dots = \hat{L}_{g_{n_o}} \quad (2)$$

where \hat{L}_{g_i} is the unit vector of L_{g_i} , for object i .

B. Frictional multi-object grasp diameter

Every object has a *final* diameter d_f at which a stable frictional single-object grasp will occur. This is the distance between gripper jaws when they become stationary in a stable grasp. By sampling stable frictional contact pairs on an object, we can enumerate multiple final single object grasp diameters to find the minimum, d_f^* . Similarly, every object group has a *final* multi-object diameter h_f at which a stable frictional multi-object grasp will occur. We compute the minimum *final* multi-object grasp diameter h_f^* given n_o objects as:

$$h_f^* = \sum_{i=0}^{n_o-1} d_{f_i}^* \quad (3)$$

This is true because we showed in Eq. 2 that all connecting lines must lie on the same line. We use h_f^* in the multi-object grasping necessary condition detailed in Sec. V-C.

C. Necessary conditions for multi-object grasping

Prior work on frictionless multi-object grasping [1] developed two necessary conditions — intersection area and multi-object grasp diameter. These conditions are used to filter

inadmissible grasps in a multi-object grasp planner. In this section we summarize these conditions for completeness.

Given a grasp, let the internal rectangular region between the gripper jaws be S . Let $o_i^s = S \cap o_i$, be the intersection polygon between S and object o_i .

1) *Intersection area*: Let $A_i(t) = \text{Area}(o_i^s)$, be the area of the intersection polygon for object i , during a multi-object grasp at time t for a grasp. The intersection area condition from prior work [1] can be written as:

$$A_i(0) > 0, i \in \{0, 1, \dots, n_o - 1\}. \quad (4)$$

We directly use this intersection area condition in this work.

2) *Multi-object grasp diameter*: Prior work [1] defined the multi-object grasp diameter necessary condition. We restate it here for completeness. Let $w_g(t)$ be the gripper width at time t . Let $b_l(t)$ be the shortest distance between o_0^s and the left jaw (where o_0^s is the closest object to the left jaw), and $b_r(t)$ be the shortest distance between $o_{n_o-1}^s$ and the right jaw (where $o_{n_o-1}^s$ is the closest object to the right jaw). Then, the multi-object grasp diameter as a function of time is: $h(t) = w_g(t) - (b_l(t) + b_r(t))$. Let h_0 be the *initial* multi-object grasp diameter at time t_0 , and h_f be the corresponding *final* multi-object grasp diameter, at time t_f when the grippers become stationary after closing. Given a group of n_o objects, one can compute the minimum possible diameter h_f^* , such that any multi-object grasp must satisfy:

$$h_0 \geq h_f^*. \quad (5)$$

Prior work [1] computed h_f^* for the frictionless case. Here we provide a method to compute it for the frictional case. Specifically, we compute $d_{f_i}^*$ in Eq. 3 by sampling N_s contact points along an object’s edge. Then, we generate all contact pairs resulting from these points. Next, we check if a contact pair is stable by ensuring that the left and right friction cones contain the line connecting the contact points. We pick $d_{f_i}^*$ as the stable contact pair with the minimum diameter.

Note that given a group of n_o objects, h_f^* is smaller in the frictional case compared to the frictionless case. This allows for more multi-object grasps to satisfy the necessary conditions when friction is considered. We further note that these necessary conditions are independent of the contact area, given the Coulomb friction assumption.

VI. LEARNING A MULTI-OBJECT GRASP NEURAL NETWORK

We train *MOG-Net* with self-supervised learning in real to predict the number of objects (N_g) that can be successfully grasped from a target object group. It takes the state of all objects in a target group, and a grasp action \mathbf{u} as inputs. In Sec. VI-A we detail our data collection process, and in Sec. VI-B we explain details of the neural network model.

A. Data collection

Physical robot time is expensive and we would like to quickly generate a high quality dataset for *MOG-Net*. An important question is what grasps do we execute during

data collection. Given an object group, instead of randomly sampling grasps, we propose to use the frictional multi-object grasp necessary conditions to filter out inadmissible grasps.

Our data collection algorithm is similar to the decluttering algorithm with two key differences: i) unlike in Alg. 1, during data collection, *obj_group* consists of only multi-object groups (i.e. $\text{len}(\text{obj_group}) > 2$), and is chosen at random, ii) in Alg. 2, we use a heuristic instead of *MOG-Net* — the total intersection area, $A_T = \sum_{i=0}^{n_o-1} A_i$ (see Sec. V-C1). We pick the grasp with $\arg \max_{\mathbf{u}^k} (\gamma^k \cdot A_T^k)$. After grasp execution, the data collection system uses an overhead image of the scene and the gripper jaw position to count the number of objects grasped. In this way, data collection is self-supervised.

B. Multi-object grasp neural network

MOG-Net predicts $N_g \in \{0, \dots, N_g^{max}\}$, where N_g^{max} is the maximum number of objects that can be grasped. We train a separate classifier for each N_g prediction class, using the same dataset collected in real. Specifically, data for a specific class is created by setting only occurrences of the desired N_g label to true while others are false. Thereafter, we train a feedforward neural network model for each class to perform binary classification. At test time, given a target object group of size $n_o \leq N_g^{max}$, we query neural network models for classes between 0 and n_o . Then, we pick the prediction with the maximum probability as N_g . We provide further details on *MOG-Net* in VII-C.

VII. PHYSICAL EXPERIMENTS

We conduct physical experiments to evaluate the data collection and decluttering algorithms. Our goal is to investigate the effect of friction on multi-object grasping for different methods. In the following subsections, we explain the general setup, experimental details, baselines, and results.

A. Experimental setup

The setup is as shown in Fig. 1 where we use a UR5 robot with a Robotiq 2F-85 gripper. In experiments, we have two sets of objects – low friction and frictional. Each contains a total of 58 objects from 3-sided to 8-sided convex polygons. To get frictional objects we wrap low friction objects with transparent, *non-stick*, high friction tape.

During data collection and decluttering experiments, we generate initial scenes with randomized object poses. We begin by repeatedly creating random object clusters. Each scene contains 17 non-overlapping object clusters that have a random

Table I. In this table, we show the number of grasped objects in the dataset of 1545 grasp samples collected on the physical robot for *MOG-Net* and *Rand-Net*. We see that *MOG-Net* produces a more balanced and higher quality dataset compared to *Rand-Net*. This saves valuable robot time.

	Number of grasped objects in the dataset				
	0	1	2	3	4
Rand-Net	543	665	250	74	13
MOG-Net	278	322	474	341	130

Table II. Physical decluttering experimental results for 10 scenes, each with 58 objects randomized as described in VII-A. We reset each scene precisely by hand to compare the methods. Errors here are within 95% confidence interval of the mean. Compared to prior work (Frictionless MOG[1]), *MOG-Net* achieved **13.7%** higher grasp success, **1.6x** PPH, and plans grasps **6.3x** faster. We also record a **3.1x** improvement in PPH compared to Frictional SOG.

Methods	Success rate (%)	Picks per hour	Grasped Objs.	Planning time (s)	Cleared (%)	Pick attempts
Frictional SOG	51.5 ± 2.2	188.9 ± 9.2	0.59 ± 0.0	0.11 ± 0.0	98.5 ± 0.7	80.7 ± 3.6
Rand-Net	67.0 ± 3.0	321.9 ± 19.9	0.93 ± 0.1	0.53 ± 0.1	99.5 ± 0.7	54.9 ± 3.3
Frictionless MOG [1]	68.9 ± 5.6	370.5 ± 48.3	1.30 ± 0.2	2.21 ± 0.2	99.0 ± 0.8	43.0 ± 5.6
Frictionless MOG-Net	71.5 ± 5.7	392.1 ± 24.8	1.26 ± 0.1	0.59 ± 0.1	97.8 ± 1.4	44.0 ± 2.4
MOG-Net	82.6 ± 3.8	580.7 ± 29.1	1.83 ± 0.1	0.35 ± 0.1	100.0 ± 0.0	30.6 ± 1.6

center point. Within each cluster, we randomly sample ¹ the number of objects, their types, positions, and orientations.

We use an RGBD camera (Intel Realsense Camera D435) to get a top-down image of the cluttered scene and then extract vertices of all objects to get the state \mathbf{x} . The grasp action \mathbf{u} involves four steps. (1) Moving the open gripper above the desired grasp pose and lowering until just above the table. (2) Closing the gripper jaws. (3) Moving the gripper upwards and above the packing box. (4) Opening the jaws so the objects fall into the packing box. All parameters used in this work are detailed here (i) grasp sampling parameters: $N_p = 25$, and $N_\theta = 12$. (ii) Monte-Carlo sampling parameters to estimate γ : $\sigma_{\mathbf{u}} = [2mm, 2mm, 2^\circ]$, and $\sigma_{\mathbf{x}} = 2 \cdot \{1\}mm$. (iii) friction: $\mu = 0.5$ for frictional, $\mu = 0.01$ for frictionless, and $N_s = 5$ for contact point sampling. (iv) N_v , the maximum number of vertices per object is set to 8 in our object set.

B. Baseline methods

1) *Frictional SOG*: We use Alg. 1 but restrict object groups to contain only single objects. It plans a frictional single-object grasp with the frictional point-contact model. This is similar to state of the art single-object grasping methods such as Dex-Net.

2) *Rand-Net*: This trains the same neural network model as ours but with a different dataset. The dataset is comprised of random grasps to execute during training from *grasp_cands* in Alg. 2 but without filtering with the necessary conditions. We trained and tested this baseline with frictional objects.

3) *Frictionless MOG*: This is state-of-the-art in multi-object grasping from prior work [1]. It filters grasp candidates with frictionless necessary conditions and uses a physics simulator (Mujoco) to find grasps.

4) *Frictionless MOG-Net*: This baseline uses *MOG-Net* but computes necessary conditions using a low friction value.

Note that we use low-friction objects for the two frictionless baselines above since they rely on frictionless necessary conditions from prior work [1] to generate grasps.

We left one set unmodified to use as low-friction objects for the frictionless baselines. We added clear grip tape to the objects of the other set to use as the high-friction counterparts

¹For each cluster, we first randomly select an orientation for the diameter of that cluster and an ordered subset of all 58 objects without replacement to place in that cluster. We then randomly sample points along the diameter of that cluster, with uniform noise $[-0.9, 0.9]$ cm perpendicular to that diameter, to be the center of the longest edge for that object. For each object in the cluster, we sample a random orientation in $[-\pi/2, \pi/2]$.

for the frictional methods. The low-friction objects do not generate frictional grasps, but instead get filtered out as inadmissible grasps using the frictionless necessary conditions from prior work [1].

C. Experimental details

1) *Data collection*: We collected 1545 grasp samples in real for *MOG-Net* and *Rand-Net*, using the random initialization process described in VII-A.

2) *Decluttering*: We create 10 decluttering scenes with the same process as data collection. In each scene, we use the four different baselines and *MOG-Net* to generate grasps. We replicated the randomly generated scene manually in each case, leading to a total of 50 physical robot experiment scenes and a total of 2532 robot grasp samples. A failed grasp attempt is where the robot misses a grasp (all objects escape), or where all objects fall out of the gripper before they reach the packing box.

3) *Neural network details*: We use a feedforward neural network with 4 hidden layers. Specifically, we use the `MLPClassifier(.)` from scikit-learn [45] with default parameters and *hidden_layer_sizes* = (500, 300, 150, 50). We limit the maximum number of objects for the input vector to $N_g^{max} = 4$ given the gripper’s size. Each object can have at most ($N_v = 8$) vertices. The input vector contains x and y points for each object’s vertex taken with respect to the grasp center. It also contains the grasp orientation. Therefore, the input vector size is fixed at $16 \times 4 + 1 = 65$. If an object has vertices less than N_v , we pad the input vector with the last vertex in the list. Similarly if the number of objects in a group (n_o) is less than N_g^{max} , we pad the input vector with the last object’s vertices in the list. Recall that we train 5 classifiers with the same model architecture to predict 5 different *number of objects grasped* (N_g) classes (0 to 4). We then pick the N_g prediction with the highest probability as the output of *MOG-Net*.

D. Metrics

We compare methods with (i) *Success rate*: percentage of grasp attempts that moved at least one object into the box, (ii) *Picks per hour*: total number of objects picked per hour, through single or multi-object grasps, (iii) *Grasped Objs*: the average number of objects grasped per pick attempt, (iv) *Planning time*: time to plan a grasp, (v) *Cleared*: fraction of objects that were moved to the box from the cluttered scene, (vi) *Pick attempts*: the average number of pick attempts.

E. Results

1) *Data collection and model*: Please see Table I for a distribution of the data collected for *MOG-Net* and *Rand-Net*. We see that *MOG-Net* collects a well distributed dataset with more balanced samples per class as opposed to *Rand-Net*, thereby saving valuable physical robot time. Classification accuracy was 71.8% for *MOG-Net* and 45.1% for *Rand-Net* on a combined test set containing 20% of the data samples.

2) *Decluttering*: The results can be found in Table II. We see that introducing friction and learning a grasp function in real significantly increases PPH. Compared to prior work [1], *MOG-Net* achieved 13.7% higher grasp success, 1.6x picks per hour, and plans grasps 6.3x faster. We also record a 3.1x improvement in PPH compared to single object grasping (Frictional SOG, which is very similar to other state of the art grasping methods such as Dex-Net [46]) as opposed to a 1.6x improvement in prior work [1]. A sample rollout on the same scene for 3 of the methods can be seen in Fig. 4.

MOG-Net outperforms *Rand-Net* on all metrics, suggesting the importance of our data collection system that generates a more balanced dataset for training by using the necessary conditions. Frictionless MOG uses a physics simulator and that led to a lower success rate, compared to *MOG-Net* which was trained in real. Frictionless *MOG-Net* outperforms Frictionless MOG on PPH. This demonstrates the importance of a reduced planning time by using a learned model, instead of a physics simulator. One mode of failure for all systems is where a grasp attempt topples objects, resulting in difficult-to-grasp poses.

VIII. LIMITATIONS AND FUTURE WORK

This work has the following limitations: i) *Pushing to rearrange objects*: The experiments have randomly generated object groups, but the robot could consider pushes to rearrange objects before planning multi-object grasps. This is explored in [47, 48]. ii) *Contact models*: To derive the necessary conditions to speed up *MOG-Net*'s training and grasp planning, we assumed a frictional point-contact model. However, for more general 3D objects we will explore soft contact models that account for torsional frictional forces. iii) *Non-polygonal objects*: We assume scenes with extruded convex polygons but household objects can be curved, non-convex, and non-polygonal. We will explore multi-object grasps for household objects and in more diverse backgrounds.

In this work, we consider the decluttering problem where multiple convex polygonal objects are grasped and moved to a packing box. We leverage a novel frictional multi-object grasping necessary condition to train *MOG-Net*, a neural network model using real examples. It predicts the number of objects grasped out of a target object group. We use *MOG-Net* in a novel grasp planner to generate robust multi-object grasps. Experiments suggest that introducing friction and considering robustness in multi-object grasping leads to improvements in success rate and picks per hour, compared to prior work.

IX. ACKNOWLEDGEMENT

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 Siemens, Toyota Research Institute, Bosch, Google, and Autodesk and by equipment grants from PhotoNeo, NVidia, and Intuitive Surgical. Mehmet Dogar was partially supported by an EPSRC Fellowship (EP/V052659). For the purpose of open access, the authors have applied a Creative Commons Attribution (CC BY) license to any Accepted Manuscript version arising.

REFERENCES

- [1] W.C. Agboh, J. Ichnowski, K. Goldberg, and M.R. Dogar. Multi-object grasping in the plane. In *ISRR*, 2022.
- [2] J. Mahler, J. Liang, S. Niyaz, M. Laskey, R. Doan, X. Liu, J.A. Ojea, and K. Goldberg. Dex-net 2.0: Deep learning to plan robust grasps with synthetic point clouds and analytic grasp metrics. In *RSS*, 2017.
- [3] B. Shen, Z. Jiang, C. Choy, S. Savarese, L.J. Guibas, A. Anandkumar, and Y. Zhu. ACID: Action-Conditional Implicit Visual Dynamics for Deformable Object Manipulation. In *RSS*, 2022.
- [4] J. Liang, J. Zhang, B. Pan, S. Xu, G. Zhao, G. Yu, and X. Zhang. Visual reconstruction and localization-based robust robotic 6-dof grasping in the wild. *IEEE Access*, 9, 2021.
- [5] P. Jiang, Y. Ishihara, N. Sugiyama, J. Oaki, S. Tokura, A. Sugahara, and A. Ogawa. Depth image-based deep learning of grasp planning for textureless planar-faced objects in vision-guided robotic bin-picking. *Sensors*, 20(3), 2020.
- [6] D. Son, H. Yang, and D. Lee. Sim-to-real transfer of bolting tasks with tight tolerance. In *IEEE IROS*, 2020.
- [7] Z. Ding, Ya-Yen Tsai, W.W. Lee, and B. Huang. Sim-to-real transfer for robotic manipulation with tactile sensory. In *IEEE IROS*, 2021.
- [8] Y. Chebotar, A. Handa, V. Makoviychuk, M. Macklin, J. Issac, N. Ratliff, and D. Fox. Closing the sim-to-real loop: Adapting simulation randomization with real world experience. *IEEE ICRA*, 2019.
- [9] J. Mahler and K. Goldberg. Learning deep policies for robot bin picking by simulating robust grasping sequences. In *CoRL*, 2017.
- [10] D. Morrison, P. Corke, and J. Leitner. Learning robust, real-time, reactive robotic grasping. *IJRR*, 39(2-3), 2020.
- [11] X. Lou, Y. Yang, and C. Choi. Collision-aware target-driven object grasping in constrained environments. In *ICRA*, 2021.
- [12] K. Harada and M. Kaneko. Enveloping grasp for multiple objects. In *IEEE ICRA*, 1998.
- [13] K. Harada and M. Kaneko. Kinematics and internal force in grasping multiple objects. In *IEEE IROS*, 1998.
- [14] K. Harada, M. Kaneko, and T. Tsujii. Rolling-based manipulation for multiple objects. *IEEE Transactions on Robotics and Automation*, 16(5):457–468, 2000.
- [15] T. Yamada, T. Ooba, T. Yamamoto, N. Mimura, and Y. Funahashi. Grasp stability analysis of two objects in two dimensions. In *IEEE ICRA*, 2005.
- [16] T. Yamada, N. Mimura, and Y. Funahashi. Grasp stability analysis of two objects with both friction and frictionless contacts in two dimensions. In *IEEE International Symposium on Micro-Nano Mechatronics and Human Science*, 2005.
- [17] T. Yamada, S. Yamanaka, M. Yamada, Y. Funahashi, and H. Yamamoto. Grasp stability analysis of multiple planar objects. In *ROBIO*, 2009.
- [18] T. Yamada, . Yamada, and H. Yamamoto. Stability analysis of multiple objects grasped by multifingered hands with revolute joints in 2d. In *IEEE ICMA*, 2012.
- [19] T. Yamada and H. Yamamoto. Static grasp stability analysis of multiple spatial objects. *Journal of Control Science and Engineering*, 3, 2015.

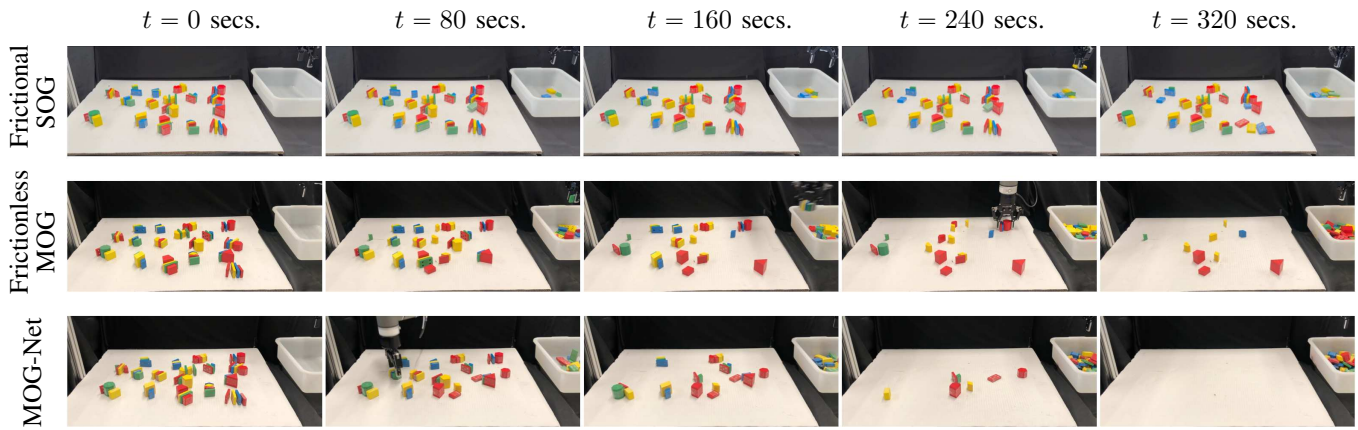


Figure 4. The same decluttering scene, replicated across three methods. *Frictional SOG* is the single object grasping baseline, *Frictionless MOG* is prior work [1], and *MOG-Net* is the method proposed in this paper. *MOG-Net* successfully removed all objects in 320 seconds while the other methods were slower.

- [20] Tianze Chen, Adheesh Shenoy, Anzhelika Kolinko, Syed Shah, and Yu Sun. Multi-object grasping - estimating the number of objects in a robotic grasp. In *IROS*, 2021.
- [21] A. Shenoy, T. Chen, and Y. Sun. Multi-object grasping - efficient robotic picking and transferring policy for batch picking. *IEEE IROS*, 2022.
- [22] Mehmet Dogar and Siddhartha S. Srinivasa. A framework for push-grasping in clutter. In *RSS*, 2011.
- [23] Michael Danielczuk, Jeffrey Mahler, Chris Correa, and Ken Goldberg. Linear Push Policies to Increase Grasp Access for Robot Bin Picking. In *CASE*, 2018.
- [24] H. Huang, M. Dominguez-Kuhne, V. Satish, M. Danielczuk, K. Sanders, J. Ichnowski, A. Lee, A. Angelova, V.O. Vanhoucke, and K. Goldberg. Mechanical search on shelves using lax-ray: Lateral access x-ray. In *IEEE IROS*, 2021.
- [25] W.C. Agboh and M.R. Dogar. Real-time online re-planning for grasping under clutter and uncertainty. In *Humanoids*, 2018.
- [26] Takumi Sakamoto, Weiwei Wan, Takao Nishi, and Kensuke Harada. Efficient picking by considering simultaneous two-object grasping. In *IEEE IROS*, 2021.
- [27] Eric Huang, Zhenzhong Jia, and Matthew T. Mason. Large-scale multi-object rearrangement. In *IEEE ICRA*, 2019.
- [28] W. Bejjani, W.C. Agboh, M.R. Dogar, and M. Leonetti. Occlusion-aware search for object retrieval in clutter. In *IEEE IROS*, 2021.
- [29] M. Hasan, M. Warburton, W.C. Agboh, M.R. Dogar, M. Leonetti, H. Wang, F. Mushtaq, M. Mon-Williams, and A.G. Cohn. Human-like planning for reaching in cluttered environments. In *IEEE ICRA*, 2020.
- [30] Wisdom C. Agboh and Mehmet R. Dogar. Pushing fast and slow: Task-adaptive planning for non-prehensile manipulation under uncertainty. In *WAFR*, 2018.
- [31] W. C. Agboh and M. R. Dogar. Robust physics-based manipulation by interleaving open and closed-loop execution. *arXiv*, 2021.
- [32] Y. Golan, A. Shapiro, and E. Rimon. Jamming-free immobilizing grasps using dual-friction robotic fingertips. *IEEE Robotics and Automation Letters*, 5(2), 2020.
- [33] Kaiyu Hang, Florian T. Pokorny, and Danica Kragic. Friction coefficients and grasp synthesis. In *IEEE IROS*, 2013.
- [34] J. Roberge, W. Ruotolo, V. Duchaine, and M. Cutkosky. Improving industrial grippers with adhesion-controlled friction. *IEEE Robotics and Automation Letters*, 2018.
- [35] A. Majumdar and R. Tedrake. Funnell libraries for real-time robust feedback motion planning. *IJRR*, 36(8), 2017.
- [36] A. Bhatt, A. Sieler, S. Puhmann, and O. Brock. Surprisingly Robust In-Hand Manipulation: An Empirical Study. In *RSS*, 2021.
- [37] R. Diankov, S. S. Srinivasa, D. Ferguson, and J. Kuffner. Manipulation planning with caging grasps. In *Humanoids*, 2008.
- [38] Satoshi Makita and Weiwei Wan. A survey of robotic caging and its applications. *Advanced Robotics*, 2017.
- [39] A. Rodriguez, M.T. Mason, and S. Ferry. From caging to grasping. *IJRR*, 31(7), 2012.
- [40] J. Mahler, M. Matl, V. Satish, M. Danielczuk, B. DeRose, S. McKinley, and K. Goldberg. Learning ambidextrous robot grasping policies. *Science Robotics*, 4(26), 2019.
- [41] J. Mahler, M. Matl, X. Liu, A. Li, D. V. Gealy, and K. Goldberg. Dex-net 3.0: Computing robust robot suction grasp targets in point clouds using a new analytic model and deep learning. *IEEE ICRA*, 2018.
- [42] H. Dong, J. Zhou, C. Qiu, D. K. Prasad, and I. Chen. Robotic manipulations of cylinders and ellipsoids by ellipse detection with domain randomization. *IEEE/ASME Transactions on Mechatronics*, 2022.
- [43] 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.
- [44] J. Tobin, L. Biewald, R. Duan, M. Andrychowicz, A. Handa, V. Kumar, B. McGrew, A. Ray, J. Schneider, P. Welinder, W. Zaremba, and P. Abbeel. Domain randomization and generative models for robotic grasping. In *IEEE IROS*, 2018.
- [45] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [46] 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. In *IEEE ICRA*, 2016.
- [47] Shrey Aeron, Edith LLontop, Aviv Adler, Wisdom C. Agboh, Mehmet R Dogar, and Ken Goldberg. Push-mog: Efficient pushing to consolidate polygonal objects for multi-object grasping. *IEEE CASE*, 2023.
- [48] K. Srinivas, S. Ganti, R. Parikh, A. Ahmad, W.C. Agboh, M. Dogar, and K. Goldberg. The busboy problem: Efficient tableware decluttering using consolidation and multi-object grasps. *IEEE CASE*, 2023.