

AO-Grasp: Articulated Object Grasp Generation

Carlota Parés Morlans^{*1}, Claire Chen^{*1}, Yijia Weng¹, Michelle Yi¹, Yuying Huang¹, Nick Heppert²,
 Linqi Zhou¹, Leonidas Guibas¹, and Jeannette Bohg¹

Abstract— We introduce AO-Grasp, a grasp proposal method that generates 6 DoF grasps that enable robots to interact with articulated objects, such as opening and closing cabinets and appliances. AO-Grasp consists of two main contributions: the AO-Grasp Model and the AO-Grasp Dataset. Given a segmented partial point cloud of a single articulated object, the AO-Grasp Model predicts the best grasp points on the object with an Actionable Grasp Point Predictor. Then, it finds corresponding grasp orientations for each of these points, resulting in stable and actionable grasp proposals. We train the AO-Grasp Model on our new AO-Grasp Dataset, which contains 78K actionable parallel-jaw grasps on synthetic articulated objects. In simulation, AO-Grasp achieves a 45.0% grasp success rate, whereas the highest performing baseline achieves a 35.0% success rate. Additionally, we evaluate AO-Grasp on 120 real-world scenes of objects with varied geometries, articulation axes, and joint states, where AO-Grasp produces successful grasps on 67.5% of scenes, while the baseline only produces successful grasps on 33.3% of scenes. To the best of our knowledge, AO-Grasp is the first method for generating 6 DoF grasps on articulated objects directly from partial point clouds without requiring part detection or hand-designed grasp heuristics. The AO-Grasp Dataset and a pre-trained AO-Grasp model are available at our project website: <https://stanford-iprilab.github.io/ao-grasp/>.

I. INTRODUCTION

Human environments are filled with articulated objects, i.e., objects that have movable parts essential to their function, such as storage furniture and appliances. For example, a typical household contains objects like cabinets, dishwashers, and boxes. For robots to autonomously perform tasks in such spaces, they must be able to interact with articulated objects. In this work, we consider the first crucial but challenging step of interacting with any articulated object: determining how a robot can grasp it to enable downstream tasks.

Interacting with articulated objects via grasps allows robots to move the end-effector in any direction without losing contact with an object. In contrast to grasping, non-prehensile manipulation requires the end-effector to move in a limited set of directions in order to maintain contact with the object. Consequently, non-prehensile manipulation constrains the ways in which a robot can interact with articulated objects. Moreover, [1], which only considers non-prehensile manipulation, concedes that grasp prediction for articulated objects is important future work because “not all tasks can be solved through non-prehensile manipulation”. Finally, interacting with articulated objects via grasps simplifies actuating the object part after contact is made, because




Fig. 1: AO-Grasp consists of (a) the AO-Grasp Dataset, which contains 78K actionable grasps on synthetic articulated objects, and (b) the AO-Grasp Model, which takes a partial point cloud of an articulated object and generates stable and actionable 6 DoF grasps that facilitate downstream manipulation. AO-Grasp not only outperforms baselines in simulation, but also achieves zero-shot sim-to-real transfer (c), enabling interactions with real-world objects with different articulation axes and geometries.

it enables the use of compliant control in a local task frame [2]–[6], an approach we describe in Section V-B, instead of of planning complex, object-specific trajectories.

Grasping articulated objects presents two unique challenges compared to grasping non-articulated objects. Firstly, grasps not only need to be stable, but also need to be actionable. Grasping any arbitrary point on an articulated object may not be sufficient; the grasp must be on the actuated part of the object to facilitate downstream tasks. For example, to open a microwave, a robot must achieve a stable grasp on the door; grasping the microwave body would be useless. Secondly, a single articulated object can exist in an infinite number of joint configurations, which may have different graspable regions. For example, to open a microwave with a closed door, a robot needs to grasp its handle, while opening the same microwave with an open door also allows the robot to grasp the edge of the door. While there are numerous works on grasping non-articulated, i.e. rigid, objects [7, 8], these are not directly applicable to grasping articulated objects because of the two properties described above. Instead, we need to develop grasp prediction methods specific to articulated objects.

Guided by this insight, we introduce AO-Grasp, which is, to the best of our knowledge, the first method for generating 6 degree-of-freedom (DoF) grasps on articulated objects directly from partial point clouds without requiring part detection or hand-designed grasp heuristics. AO-Grasp consists of two main contributions. First, we present the AO-Grasp Dataset (Fig. 1a), which contains 78K parallel-jaw 6 DoF grasps on 84 Box, Dishwasher, Microwave, Safe, TrashCan, Oven, and StorageFurniture instances from the PartNet-Mobility dataset [9, 10]. The AO-Grasp Dataset only contains grasps on actionable parts of artic-

^{*}Equal contribution.

¹Stanford University, CA, USA. ²University of Freiburg, Germany.
 Toyota Research Institute provided funds to support this work.

ulated objects, as well as grasps on objects in diverse joint configurations. Next, we introduce the AO-Grasp Model (Fig. 1b), which generates stable and actionable 6 DoF parallel-jaw grasps for articulated objects given a partial point cloud of an object. The AO-Grasp Model predicts where on objects a robot should grasp with an Actionable Grasp Point Predictor, trained on the AO-Grasp Dataset. Unlike other works [1, 11, 12], the Actionable Grasp Point Predictor does not require additional semantic part segmentation to predict grasp points on actionable parts of articulated objects. Finally, we generate grasp orientations for the top predicted grasp points.

We show, in simulation and through zero-shot transfer to the real-world (Fig. 1c), that AO-Grasp generates actionable grasps on articulated objects with diverse handle geometries and articulation axes, even on categories not seen in training. AO-Grasp also generates grasps for objects in both closed and open joint configurations and viewed from varied camera angles. In simulation, we show that AO-Grasp achieves an average 45.0% success rate across unseen instances from both train and test categories, while a rigid object grasping baseline CGN [13] and articulated object interaction baselines VAT-Mart [12] and Where2Act [11] achieve average success rates of 35.0%, 5.21%, and 3.9% respectively. In the real-world, we compare AO-Grasp and CGN on a custom-built re-configurable cabinet with 4 handles and 4 articulation axes, 2 microwaves, a toaster oven, prismatic-jointed drawer, and cardboard box. On 120 real-world scenes, AO-Grasp produces successful grasps on 67.5% of scenes, while CGN only produces successful grasps on 33.3% of scenes.

II. RELATED WORK

This work focuses on interacting with articulated objects via grasps. As such, we situate this work amongst both the literature on interacting with articulated objects and the body of work on grasping rigid objects. We also differentiate the AO-Grasp Dataset from existing grasp datasets.

Interacting with articulated objects: While there are many existing works on interacting with articulated objects, none specifically addresses the task of generating stable grasps for objects in diverse configurations without requiring part detection or hand-designed grasp heuristics.

Some works consider interacting with articulated objects via grasps [4, 14]–[20], but use various strategies to simplify grasp generation. Unlike AO-Grasp, V-MAO [14] only predicts contact points, as opposed to 6 DoF poses, and requires part segmentation of the object, whereas AO-Grasp does not. V-MAO also does not show real-world results. Other works use hand-designed grasp heuristics [4, 15]–[20] that would not scale well to a larger variety of objects, motivating a more general grasp generation method like AO-Grasp.

Another category of work focuses on learning interaction policies that do not require grasps [1, 11, 12, 15, 21]–[26]. Where2Act (W2A) [11] and VAT-Mart [12] are most similar to AO-Grasp, as they also predict per-point interaction poses given partial point clouds. Unlike AO-Grasp, however, W2A and VAT-Mart do not specifically check that the robot

achieves stable grasps on objects, resulting in interactions that are predominantly non-prehensile. In our evaluation, we confirm that these methods rarely produce successful grasps, which limits how a robot can interact with an articulated object. Some of these works [16, 24] even further simplify making contact with the object by using a suction cup end-effector, which constrains the types of geometries the robot can interact with to smooth and flat surfaces. Furthermore, although these works all showcase interactions on an impressive variety of articulated objects in simulation, most only show limited or no real-world results. Of the works that do include real-world results [1, 12, 15, 21, 24]–[26], only one [25], which uses a suction cup, conducts quantitative real-world evaluations. In contrast, we show that AO-Grasp can be used in the real-world by quantitatively evaluating AO-Grasp on 16 variations of a reconfigurable cabinet and 5 additional common household articulated objects.

Datasets: A line of work [27, 28] has been devoted to building large-scale grasp datasets, as they are essential for data-driven grasping approaches. Most related to us are datasets for parallel-jaw grippers [29]–[34], which label grasps either analytically [29, 31, 32] or by running simulation trials [30, 33, 34]. However, these datasets only target rigid objects and focus on grasp stability, not actionability. Both W2A [11] and VAT-Mart [12] release data generation methods for 6 DoF interactions with articulated objects, but these interactions do not guarantee stable grasps, limiting how a robot can interact with objects. To the best of our knowledge, we are the first to build a grasp dataset for articulated objects that contains both stable and actionable grasps.

Grasping non-articulated objects: Despite the unique challenges that come with grasping articulated objects, an important part of the grasping problem remains the same as in grasping rigid objects: understanding the local geometries of an object that make those areas suitable for grasping. There is a deep body of work studying this grasping problem [7, 8]. Contact-GraspNet (CGN) [13] and AnyGrasp [35] are two recent works that use large datasets to learn grasp prediction models. We find that once we have predicated good locations at which to grasp articulated objects, we can leverage works like these to predict grasp orientations that match local geometry. AO-Grasp uses CGN [13], as it is a representative rigid object grasp generation method.

III. AO-GRASP DATASET

We introduce the AO-Grasp Dataset, a dataset of simulated, actionable grasps on articulated objects. It contains 78K 6 DoF grasps for 84 instances from 7 common household furniture/appliance categories (Box, Dishwasher, Microwave, Safe, TrashCan, Oven, StorageFurniture) from the PartNet-Mobility dataset [9, 10]. Table I summarizes the per-category statistics of our dataset. Across its 7 categories, the AO-Grasp Dataset exhibits considerable variation in geometries, articulation axes, joint configurations, and number of actionable parts, as shown in Fig. 1a.

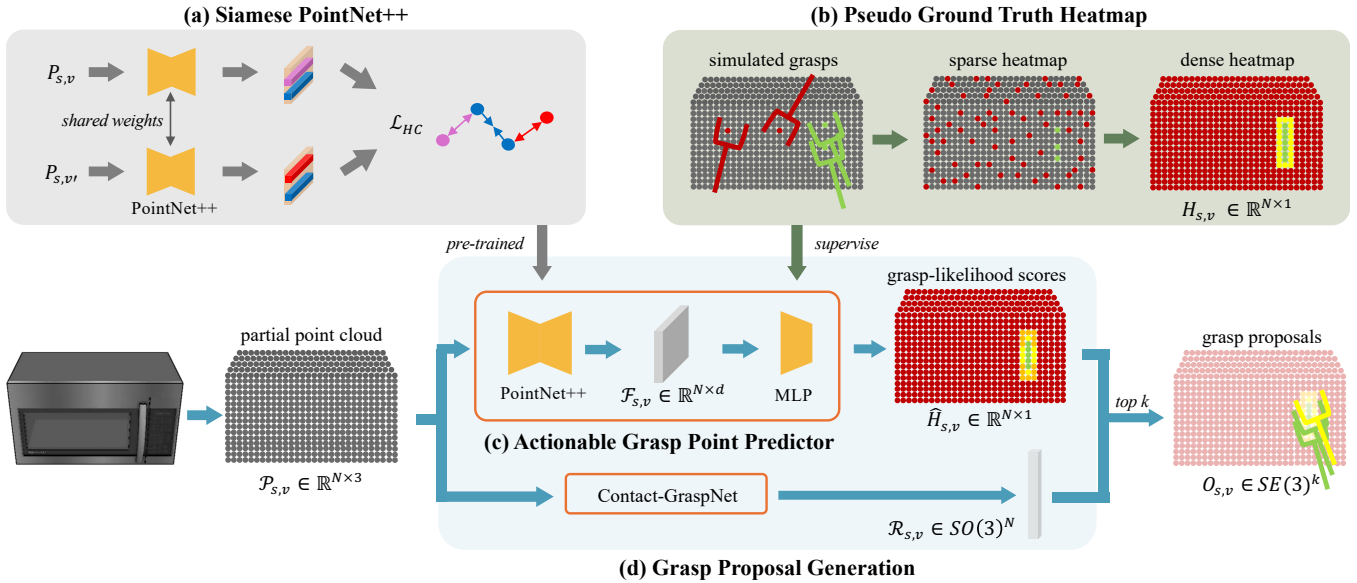


Fig. 2: An overview of the AO-Grasp Model. (a) Siamese PointNet++: We find positive and negative correspondences between partial point clouds of an object in joint state s captured from two different object views v and v' to train the network with a hardest contrastive loss L_{HC} (1). (b) Pseudo ground truth heatmap $H_{s,v}$ (2): We supervise training the Actionable Grasp Point Predictor on dense heatmaps computed from grasps in the AO-Grasp Dataset. (c) Actionable Grasp Point Predictor: Given a partial point cloud of an articulated object $P_{s,v}$, the AO-Grasp Point Predictor outputs a heatmap of grasp-likelihood scores $\hat{H}_{s,v}$. It first processes the partial point cloud $P_{s,v}$ with the pre-trained PointNet++ module, resulting in point features $F_{s,v}$. Then, these features are passed through an MLP which returns the grasp-likelihood scores $\hat{H}_{s,v}$. (d) Grasp Proposal Generation: We leverage Contact-GraspNet [13] to get per-point grasp orientations. We take the top k scores from $\hat{H}_{s,v}$ to get a set of stable and actionable 6 DoF grasp poses $O_{s,v} \in SE(3)^k$.

Category	All							
# Instances	84	9	17	11	11	13	6	17
Closed state grasps	15632	516	1396	1546	372	2493	2715	6594
Open state grasps	62641	8091	8020	8022	6152	11669	5400	15287
Total grasps	78273	8607	9416	9568	6524	14162	8115	21881

TABLE I: The number of instances and grasps for Box, Dishwasher, Microwave, Safe, TrashCan, Oven, and StorageFurniture categories in the AO-Grasp Dataset. In total, the AO-Grasp Dataset spans 7 categories and contains 78K grasps for 84 instances, each in 1 closed state and 9 randomly-sampled states.

A. Grasp parametrization and labeling criteria

The AO-Grasp Dataset uses a two-fingered Robotiq 2F-85 gripper. A grasp is parameterized by a 6 DoF pose $g = (t, R) \in SE(3)$, with a position t and rotation R .

In contrast to rigid object grasping, where the stability of grasps is usually verified by shaking objects or applying disturbance forces [30, 33, 34], we require semantically meaningful interactions with articulated objects, such as opening the door of a microwave. Consequently, we design our grasp evaluation procedure to test not only for a grasp’s stability, but also its actionability.

We label each grasp by executing a grasp episode with a floating gripper in PyBullet [36]. We first spawn the fully-open gripper at g , close the gripper to complete the grasp if no collision is detected, then start action execution by moving the gripper in the optimal direction to actuate the object part, which we obtain via the object’s ground-truth joint type and axis. After a fixed number of steps, we terminate the action

and label the grasp as successful if 1) the gripper is still in contact with the object, indicating stability and 2) the grasped part has moved a certain distance, indicating actionability.

B. Grasp sampling

For each object instance, we randomly sample grasp poses and label each one using the procedure described in the previous section. For each instance, we generate grasps for the object in its canonically-closed state and 9 randomly-sampled open states, and capture each state from 20 randomly-sampled camera viewpoints. In total, we obtain 78,273 positively-labeled grasps on 84 instances.

IV. AO-GRASP MODEL

The AO-Grasp Model takes a partial point cloud of an articulated object as input and outputs a set of 6 DoF grasp poses (see Fig. 2). First, it determines where on the object a robot should grasp with the novel Actionable Grasp Point Predictor, trained on the AO-Grasp Dataset to predict per-point grasp-likelihood scores. We design the Actionable Grasp Point Predictor loss and training strategies to facilitate generalization to new viewpoints, joint states, and objects. Then, to generate grasp orientations for grasp points predicted by the Actionable Grasp Point Predictor, the AO-Grasp Model leverages rotations generated by Contact-GraspNet [13], a state-of-the-art rigid object grasping method. Finally, from the per-point grasp-likelihood scores and grasp orientations, we compose the final set of grasp proposals by selecting the points with the highest grasp-likelihood scores.

A. Actionable Grasp Point Predictor

For each point in the input point cloud, the Actionable Grasp Point Predictor outputs a grasp-likelihood score that signifies how likely that point will afford a stable and actionable grasp. It consists of a PointNet++ (PN++) [37] backbone, which extracts per-point features, and an MLP output head that learns to predict the grasp-likelihood scores of points given their features. We use the following 2 strategies to achieve generalization across viewpoints, joint states, and object instances and categories:

1) Learning viewpoint-independent point correspondences:

To achieve generalization across different camera viewpoints, the Actionable Grasp Point Predictor must understand that grasp-likelihood scores are viewpoint-independent. To facilitate this understanding, we pre-train PN++ using a Siamese Network in a self-supervised manner using the hardest-contrastive loss proposed in [38]. We finetune this pre-trained PN++ backbone when training the grasp-likelihood output head.

Inspired by [39], we use a Siamese Network architecture to learn viewpoint-independent point-wise features (see Fig. 2a). One training pass of the Siamese Network takes, as input, two partial view point clouds, $\mathcal{P}_{s,v}$ and $\mathcal{P}_{s,v'}$, of a single instance in the same joint state s , but captured from different camera viewpoints v and v' . These are passed through PN++ to compute point-wise features $F_{s,v}$ and $F_{s,v'}$. Next, we uniformly sample a set \mathcal{Z} of pairs of matched points (i, j) where $i \in \mathcal{P}_{s,v}$ and $j \in \mathcal{P}_{s,v'}$, with corresponding features $f_{s,v}^i \in F_{s,v}$ and $f_{s,v'}^j \in F_{s,v'}$. Then, we compute the following hardest contrastive loss across all point pairs,

$$\mathcal{L}_{HC} = \sum_{(i,j) \in \mathcal{Z}} \left[\frac{\left[\left\| f_{s,v}^i - f_{s,v'}^j \right\|_2 - m_p \right]_+^2}{|\mathcal{Z}|} + \frac{\left[m_n - \min_{k \in \mathcal{N}_{s,v'}} \left\| f_{s,v}^i - f_{s,v'}^k \right\|_2 \right]_+^2}{2|\mathcal{N}_{s,v'}|} + \frac{\left[m_n - \min_{k \in \mathcal{N}_{s,v}} \left\| f_{s,v'}^j - f_{s,v}^k \right\|_2 \right]_+^2}{2|\mathcal{N}_{s,v}|} \right], \quad (1)$$

where $[\cdot]_+$ denotes $\max(0, \cdot)$, $\mathcal{N}_{s,v}$ and $\mathcal{N}_{s,v'}$ are sets of randomly sampled negative points from $\mathcal{P}_{s,v}$ and $\mathcal{P}_{s,v'}$, respectively, and m_p and m_n are margins for positive and negative pairs, respectively. We set $|\mathcal{Z}|=64$ and $|\mathcal{N}|=10$. Following the values used by Xie et al. in [40], we set the margins m_p and m_n to 0.1 and 1.4, respectively.

2) Computing dense “pseudo ground truth” heatmaps: We found that training the Actionable Grasp Point Predictor directly on the binary grasp labels in the AO-Grasp Dataset results in poor generalization to test categories, as we show in our ablation studies in Section V-A. We hypothesize that this is because the model is susceptible to overfitting, due to the number of labeled points being a relatively small percentage of the total number of points in the input point cloud. To mitigate this overfitting, we augment our data by assigning pseudo ground truth labels h_{pgt}^i to each point i in a point

cloud with

$$h_{pgt}^i = \min \left(1, \left(\sum_{j \in G^i} \max(0, w^{(i,j)}) \right) * \frac{1}{k} \right) \quad (2)$$

$$w^{(i,j)} = \begin{cases} 1 - \frac{1}{r} * \lambda_+ d(i, j) & \text{if } j \text{ is positive} \\ 1 - \frac{1}{r} * \lambda_- d(i, j) & \text{if } j \text{ is negative,} \end{cases} \quad (3)$$

where h_{pgt}^i is a weighted average of the k closest labeled points to point i , denoted by the set G^i . Each labeled point $j \in G^i$, if within distance r to point i , contributes a weight $w^{(i,j)}$ to the pseudo ground truth label of i inversely proportional to its distance to i , where $\lambda_{\{+,-\}}$ is the amount to weight positive and negative points, and d denotes the Euclidean distance between two points. Point i is labeled as negative if the closest ground truth positive point is farther than r . We use $k=15$, $\lambda_+=2$, $\lambda_-=0$, and $r=4\text{cm}$. Fig. 2b illustrates the difference between raw binary labels and dense heatmaps.

Total loss: To train the Actionable Grasp Point Predictor, we use the loss function

$$\mathcal{L}_{\text{total}} = \lambda_{HC} \mathcal{L}_{HC} + \lambda_{MSE} \mathcal{L}_{MSE}, \quad (4)$$

which combines the hardest contrastive loss \mathcal{L}_{HC} and the mean squared error \mathcal{L}_{MSE} between per-point predicted scores and pseudo ground truth heatmap labels, weighted with λ_{HC} and λ_{MSE} respectively, to learn generalizable feature encodings. We set $\lambda_{HC}=3$ and $\lambda_{MSE}=1$.

B. Predicting grasp orientations

While considering an articulated object’s actionability and joint configuration is critical for predicting good grasp points, these properties matter much less for predicting grasp orientations. Instead, if given a good grasp point, an object’s local geometry is the most important factor in determining a suitable grasp orientation, regardless of whether the object is rigid or articulated. As such, we leverage orientation predictions from Contact-GraspNet (CGN) [13], a state-of-the-art grasp generation method for rigid objects. As CGN is trained on point clouds with 2048 points, whereas our point clouds have 4096 points, we assign each point in our point cloud the orientation corresponding to the closest point in the down-sampled CGN point cloud.

Although the AO-Grasp Dataset contains 6 DoF grasps, the performance of the models we trained on our data to predict orientations could not match that of CGN’s. We believe this is because of the difference in training data quantity and density between the AO-Grasp Dataset and the ACRONYM dataset [34], which CGN is trained on. The ACRONYM dataset is not only substantially larger than AO-Grasp Dataset, containing over 17 million grasps, but its grasps are also on rigid objects much smaller in scale than the articulated objects in the AO-Grasp Dataset, resulting in much denser coverage of ground truth grasps on objects. It remains exciting future work to develop less data-hungry methods for learning grasp orientations, as we have already done for learning grasp points.

V. EXPERIMENTAL RESULTS

A. Simulation evaluation

In simulation, we compare AO-Grasp to baseline methods and ablations. In the comparison against baseline methods, we investigate how state-of-the-art models for both rigid object grasping and interacting with articulated objects compare to AO-Grasp on the task of grasping articulated objects. In our ablation study, we explore the effect that pre-training PointNet++ (PN++) and training on dense heatmap labels has on the Actionable Grasp Point Predictor performance.

Evaluation setup: We train and test AO-Grasp on partial point clouds captured from camera viewpoints sampled randomly from a range of 120° about the object’s yaw axis and 10° about the object’s pitch axis. We train on 47 instances (6 Box, 14 Dishwasher, 9 Microwave, 9 Safe, 9 TrashCan), with 8 states per instance (1 closed, 7 randomly-sampled open) and 16 randomly-sampled viewpoints per state. We test each model on 480 partial point clouds of 14 held-out training category instances (3 Box, 3 Dishwasher, 2 Microwave, 2 Safe, 4 TrashCan), and 920 partial point clouds of 23 held-out test category instances (6 Oven and 17 StorageFurniture). AO-Grasp is only trained on objects with one movable part, but several test category instances have more than one movable part. For each partial point cloud, we evaluate the top-1 highest scoring grasps by executing them in simulation with the procedure described in Section III-A. We report the success rates of the top-1 grasps for each test scene.

Baselines: To highlight why existing grasp generation methods for rigid objects are not suitable for grasping articulated objects, we compare AO-Grasp to Contact-GraspNet (CGN) out-of-the-box, using both its proposed grasp points and orientations. We also compare AO-Grasp to Where2Act (W2A) [11] and VAT-Mart [12], two existing methods that also consider interactions with articulated objects, but do not focus specifically on prehensile manipulation. We evaluate pre-trained VAT-Mart¹ and W2A models using their ‘pull-door’ and ‘pull’ action policies, respectively. We evaluate each method both with and without ground truth segmentation masks. First, to match how W2A and VAT-Mart conduct their evaluations², we use ground truth part segmentation masks to filter out grasps on unmovable object parts. Secondly, as our method does not need segmentation masks, we evaluate the proposals from W2A and VAT-Mart without masks.

Results against baselines: We show in Table II that AO-Grasp consistently achieves higher success rates than CGN, W2A, and VAT-Mart on both closed and open states. This trend holds true even on categories not seen by AO-Grasp during training, demonstrating that AO-Grasp achieves gen-

¹VAT-Mart does not provide model checkpoints and the released code contains errors that prevent the user from training policies. After fixing these errors, training VAT-Mart, and running their evaluation, we verified that we were able to match the scores reported in their paper [12]. This is the policy we use in our evaluation.

²Although the use of segmentation masks may not be stated clearly in their papers, we found that the released W2A and VAT-Mart code use ground truth part segmentation masks during both training and evaluation.

eralization across categories, even to objects that contain multiple joints. Compared to CGN, which often proposes grasps on unmovable object parts, AO-Grasp predicts grasp scores that capture actionability and thus proposes grasps that are on the movable parts of objects, as shown in Fig. 3. While CGN suffers from this failure mode in all states, closed states, particularly of the test categories, emphasize this more as the number of graspable points in closed states is much smaller compared to open states (i.e. the points on the handle are much fewer than those on the edges of a door).

Without using ground truth masks, W2A, VAT-Mart exhibit similar failure cases to CGN, where grasps are often on unmovable object parts (Fig. 3). However, even when we use ground truth segmentation masks to ensure that grasps are on movable object parts, both methods still achieve very low success rates, as the proposed grasp orientations do not lead to stable grasps. In general, unlike CGN, W2A and VAT-Mart hardly find any successful grasps because they are not trained specifically for grasping and thus struggle to perform well under our grasp evaluation criteria. Compared to W2A, VAT-Mart achieves a slightly higher success rate, likely because it considers task-aware pulling trajectories which improves the quality of the gripper orientation predictions.

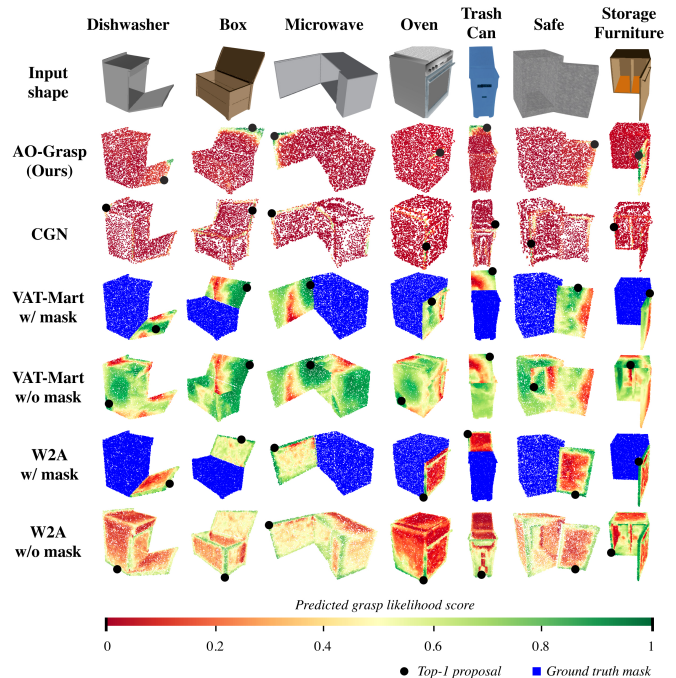


Fig. 3: A comparison of predicted grasp-likelihood scores from AO-Grasp and baselines CGN, VAT-Mart, and W2A on synthetic point clouds, with top-1 proposals highlighted with black dots. Given that VAT-Mart and W2A use part-segmentation masks² in their training and evaluation, while AO-Grasp does not, we include their predictions with and without ground truth masks. Compared to AO-Grasp, all baselines tend to propose non-actionable points more often. Point cloud sizes are 4K, 2K, 10K, and 10K for AO-Grasp, CGN, VAT-Mart, and W2A, respectively.

Ablations: In our ablation studies, we explore how pre-training PN++ on viewpoint-independent point correspondences and training the Actionable Grasp Point Predictor on dense heatmap labels influence model performance. Table IV shows the grasp success rates for the top-1 proposals















Model	Simulation success rates (%) \uparrow																		
	All	Train categories										Test categories							
	All states	Closed state					Open states					Closed state		Open states					
	All n = 1400	All 200						All 280						All 460			All 460		
AO-Grasp (Ours)	45.0	40.0	65.0	48.3	35.0	10.0	36.3	56.1	55.0	55.0	85.0	42.5	50.0	37.4	45.0	34.7	48.0	29.2	54.7
CGN [13]	35.0	29.0	30.0	0.0	30.0	0.0	57.5	46.8	58.3	56.7	57.5	50.0	23.8	21.7	7.50	26.8	43.7	27.5	49.4
VAT-Mart w/ mask [12]	5.21	3.0	0.0	10	0.0	0.0	0.0	6.8	6.7	1.7	7.5	2.5	12.5	0.6	2.5	0.0	9.8	15.8	7.6
VAT-Mart w/o mask [12]	2.0	1.0	0.0	3.3	0.0	0.0	0.0	5.4	3.3	6.7	2.5	0.0	10.0	0.0	0.0	0.0	2.4	3.3	2.1
W2A w/ mask [11]	3.9	2.0	5.0	1.7	0.0	0.0	2.5	8.9	6.7	5.0	22.5	10.0	6.3	0.0	0.0	0.0	5.4	3.3	6.2
W2A w/o mask [11]	2.4	0.5	0.0	1.7	0.0	0.0	0.0	4.3	0	0.0	10	10	5	0.0	0.0	0.0	4.4	0.0	5.9

TABLE II: Simulation grasp success rates (%) for the top-1 grasps generated by AO-Grasp, CGN, VAT-Mart [12], and Where2Act (W2A) [11]. Results are broken down by Actionable Grasp Point Predictor train/test categories (all instances are unseen during training) as well as open and closed states. Note that pre-trained W2A and VAT-Mart have different train/test categories: W2A trains on *Box*, *Microwave*, *TrashCan*, and *StorageFurniture* and VAT-Mart trains on *StorageFurniture* and *Microwave*. Given that VAT-Mart and W2A use part segmentation masks² during training and evaluation, we include results for both with and without part segmentation. AO-Grasp achieves a higher average success rate than all baselines. We note that the success rates for closed-state *Safe* instances are lower than success rates of other categories because the handles on them are very small and shallow, making them difficult to grasp. Even still, AO-Grasp achieves a 10% success rate, while all other baselines fail to find any successful grasps.
















Model	Real-world success rates (%) \uparrow																	
	All states	Closed states								Random states								
	All n = 120	All 56									All 64							
AO-Grasp (Ours)	67.5	57.1	87.5	62.5	75.0	37.5	37.5	62.5	37.5	76.6	100	100	37.5	87.5	62.5	87.5	50.0	87.5
CGN [13]	33.3	10.7	0.00	12.5	0.00	0.00	0.00	25.0	37.5	53.1	62.5	50.0	37.5	87.5	62.5	62.5	0.00	62.5

TABLE III: Real-world grasp success rates (%) for AO-Grasp compared to baseline CGN. The first column of results lists the average success rates over all 120 scenes, where AO-Grasp achieves a success rate double that of CGN. The following columns break down results by joint state, reconfigurable cabinet variation, and object type. For closed-state reconfigurable cabinet variations, we break down results by handle type (cylindrical, L-shaped, smooth climbing hold, textured climbing hold), as local handle geometry is the most critical factor for grasping objects in closed states. For open-state reconfigurable cabinet variations, we break down results by articulation axes (hinge-right, hinge-left, hinge-up, hinge-down), as an object’s articulation axis is more relevant than handle geometry when grasping objects in open states. Notably, AO-Grasp dramatically outperforms CGN on closed states, where CGN fails to generate any successful grasps for 4 out of the 8 closed state object variations. Please see the supplementary video for more details on real-world experiments.

generated by an Actionable Grasp Point Predictor with a pre-trained and finetuned PN++ and trained on dense heatmaps (our method), as well as 2 ablation models: one without pre-trained PN++ but still trained with dense heatmaps and one with neither pre-trained PN++ nor dense heatmaps. Overall, the model with pre-trained PN++ and trained on dense heatmaps outperforms both ablation models. Further examination of the train and test category results reveals that both pre-training PN++ and supervising on dense heatmap labels improve generalization to test categories. Notably, while the model trained with sparse binary labels is competitive with our full method on train categories, it performs worse on test categories, particularly on closed states. This difference in performance underscores the role that our pseudo ground truth heatmaps play in mitigating model overfitting.

Additionally, in Table V, we report the F-score and accuracy rates for the predicted grasp-likelihood scores. These results show that pre-training PN++ and training on dense heatmap labels improves the model’s ability to capture graspable regions for both train and test categories. In Figure 4, which shows several examples of predicted grasp-likelihood scores from each model on test category instances, we see that both ablation models are more prone to predicting false positives and false negatives.

PT PN++	Dense heatmap	Ablations’ success rate (%) \uparrow				
		All	Train categories		Test categories	
		All states	Closed	Open	Closed	Open
		n = 1400	n = 200	n = 280	n = 460	n = 460
\checkmark	\checkmark	45.0	40.0	56.1	37.4	48.0
\times	\checkmark	41.7	34.5	59.3	28.3	47.4
\times	\times	42.2	41.5	60.4	33.9	39.8

TABLE IV: Simulation top-1 grasp success rates for Actionable Grasp Point Predictor ablations on pre-training PointNet++ (PT PN++) and training on dense pseudo ground truth heatmaps. Without PT PN++ (second row) and without supervising on dense heatmaps (third row), generalization to test categories suffers, particularly on closed states for test categories.

B. Real-world evaluation

To showcase AO-Grasp’s zero-shot sim-to-real transfer, we conduct a quantitative evaluation of AO-Grasp and CGN on 120 scenes of real-world objects with varied local geometries and articulation axes, in different joints states, and captured from different viewpoints. To comprehensively test different local geometries and articulation axes, we design and fabricate a custom reconfigurable cabinet. It features a magnetic handle mounting system that allows for easily exchanging custom door handles and can be flipped on any side to achieve different articulation axes. The top row of Fig. 1c

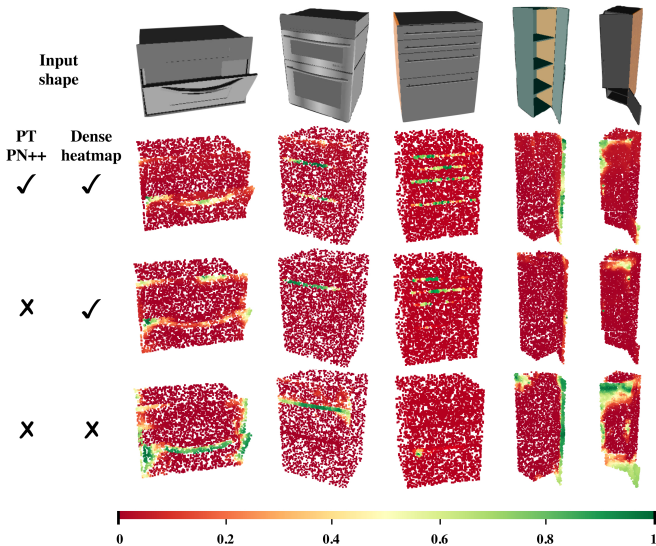


Fig. 4: A comparison of predicted grasp-likelihood scores from AO-Grasp and ablations on synthetic point clouds of test category instances. Despite not being trained on any instances with more than one movable part, AO-Grasp (first row of heatmaps) accurately predicts scores for objects with multiple movable parts. For example, in the second and third columns, it detects points on all the handles, and in the last column it detects points on the edges of both doors. Moreover, it rarely predicts false positives, i.e., it does not predict points on the body of the object. In contrast, the model without pre-training PN++ and dense heatmaps (third row of heatmaps) does not predict grasp points on all handles (second and third columns) and predicts many false positives. Training the model with dense heatmaps (second row of heatmaps) reduces the number of false positives (i.e. top left corner of the cabinet in the fourth column). However, the number of false negatives (missed good points) also increases (ie. bottom door of the right-most cabinet).

shows variations of the reconfigurable cabinet with climbing holds and 3D-printed cylindrical handles, as well as in both hinge-right and hinge-up configurations.

We evaluate both methods on 16 variations of our reconfigurable cabinet (4 handles, with 4 locations of the articulation axes per handle (hinge top, bottom, left, right)), as well as 2 microwaves, a prismatic-jointed drawer, toaster oven, and cardboard box. We test each object in both closed and open states, and viewed from multiple viewpoints. Note that for all of these objects, the models only receive partial point clouds of the object, which we obtain using [41], as input; no additional information such as the type of articulation mechanism is provided. We use a Franka Emika robot arm and a ZED2 camera to capture depth data. We evaluate the top-1 grasp proposed by each model by moving the end-effector to that pose, closing the gripper, and executing an action to actuate grasped part for 3 seconds.

To actuate the part, we use a compliant controller to move the end-effector along a direction of motion, parameterized as a vector defined in the local object part frame. Because AO-Grasp’s prehensile grasps result in the end-effector being rigidly attached to the object part, we can define the object part frame to be the end-effector frame and directly specify the direction of motion in the end-effector frame, removing the need to track the object part pose at every timestep. The compliant controller moves the end-effector along the local direction of motion while being compliant in all other directions. This approach allows us to use same vector to

		<i>F</i> -score (%) \uparrow / Accuracy (%) \uparrow				
		All	Train categories		Test categories	
		<i>All states</i>	<i>Closed</i>	<i>Open</i>	<i>Closed</i>	<i>Open</i>
PT	Dense	n=	n=	n=	n=	n=
PN++	heatmap	1400	200	280	460	460
\checkmark	\checkmark	36.4/73.1	45.4/71.6	40.3/85.7	54.1/72.8	15.0/62.0
\times	\checkmark	33.5/67.4	36.8/68.3	41.0/80.0	43.5/66.5	11.3/56.1
\times	\times	2.80/50.5	2.42/50.5	3.24/50.4	3.32/50.4	2.00/50.2

TABLE V: *F*-score and accuracy rates of predicted grasp-likelihood scores on synthetic point clouds for Actionable Grasp Point Predictor ablations on pre-training PointNet++ (PT PN++) and training on our dense pseudo ground truth heatmaps. Without PT PN++ (second row), and without supervising on dense heatmaps (third row), models achieve lower *F*-scores, indicating a higher rate of false positives and false negatives. For all models, the *F*-score for open-state test categories is lower than *F*-scores for other splits. This is because the test categories have objects with multiple movable parts and therefore more graspable areas, resulting in the sampled grasps in the AO-Grasp Dataset that are more spread out. Upon visualizing the predicted grasp-likelihood scores (Fig. 4), we find that the lower *F*-scores are because the models are able to capture good graspable regions that are not sampled in the AO-Grasp Dataset.

parameterize a controller for opening a revolute door and a controller for opening a prismatic drawer, instead of needing to plan different trajectories for each joint type. In our experiments, we define the direction of motion to be normal to the object part surface, computed from the object’s point cloud, but we note that with the prehensile grasps proposed by AO-Grasp, a robot could also move its end-effector in any direction without losing contact with the object.

We use the same success criteria that we used in our data generation and simulation evaluation, where a grasp is labeled a success if the target joint of the object is actuated and the end-effector maintains contact with the object for the entire interaction. The results, presented in Table III, show that AO-Grasp achieves an overall success rate of 67.5% while CGN achieves a success rate of 33.3%. Our real-world success rate is higher than our simulation success rate because the grasp execution procedure used in simulation is less forgiving than the procedure used for real-world experiments. In simulation, we initialize the end-effector directly at grasp poses and label grasps as failures if the end-effector intersects the object by any amount. In the real-world, we use a compliant controller to move the end-effector to grasp poses. This is more forgiving than spawning the end-effector at grasp poses and checking for collision, resulting in fewer failed grasps in the real-world. Please see the project website to view some real-world grasps.

VI. CONCLUSION

We introduce AO-Grasp, which generates stable and actionable 6 DoF grasps on articulated objects, and the AO-Grasp Dataset, which contains 78K simulated grasps. Although AO-Grasp achieves higher grasp success rates than baselines and shows promising sim-to-real transfer on a variety of objects, we acknowledge that there is still much room for improvement in both object diversity and model performance, underscoring the difficulty of grasp generation for articulated objects and leaving room for future work.

REFERENCES

- [1] G. Schiavi, P. Wulkop, G. Rizzi, L. Ott, R. Siegwart, and J. J. Chung, "Learning agent-aware affordances for closed-loop interaction with articulated objects," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 5916–5922.
- [2] O. Khatib, "A unified approach for motion and force control of robot manipulators: The operational space formulation," *IEEE Journal on Robotics and Automation*, vol. 3, no. 1, pp. 43–53, 1987.
- [3] H. Bruyninckx and J. De Schutter, "Specification of force-controlled actions in the "task frame formalism"—a synthesis," *IEEE Transactions on Robotics and Automation*, vol. 12, no. 4, pp. 581–589, 1996.
- [4] M. Stuede, K. Nuelle, S. Tappe, and T. Ortmaier, "Door opening and traversal with an industrial cartesian impedance controlled mobile robot," in *2019 International Conference on Robotics and Automation (ICRA)*, 2019, pp. 966–972.
- [5] W. Meeussen, M. Wise, S. Glaser, S. Chitta, C. McGann, P. Mihelich, E. Marder-Eppstein, M. Muja, V. Eruhimov, T. Foote, J. Hsu, R. B. Rusu, B. Marthi, G. Bradski, K. Konolige, B. Gerkey, and E. Berger, "Autonomous door opening and plugging in with a personal robot," in *2010 IEEE International Conference on Robotics and Automation*, 2010, pp. 729–736.
- [6] M. Prats, S. Wieland, T. Asfour, A. P. del Pobil, and R. Dillmann, "Compliant interaction in household environments by the armar-iii humanoid robot," in *Humanoids 2008 - 8th IEEE-RAS International Conference on Humanoid Robots*, 2008, pp. 475–480.
- [7] R. Newbury, M. Gu, L. Chumbley, A. Mousavian, C. Eppner, J. Leitner, J. Bohg, A. Morales, T. Asfour, D. Kragic, D. Fox, and A. Cosgun, "Deep learning approaches to grasp synthesis: A review," *IEEE Transactions on Robotics*, pp. 1–22, 2023.
- [8] J. Bohg, A. Morales, T. Asfour, and D. Kragic, "Data-driven grasp synthesis—a survey," *IEEE Transactions on Robotics*, vol. 30, no. 2, pp. 289–309, 2014.
- [9] F. Xiang, Y. Qin, K. Mo, Y. Xia, H. Zhu, F. Liu, M. Liu, H. Jiang, Y. Yuan, H. Wang, L. Yi, A. X. Chang, L. J. Guibas, and H. Su, "SAPIEN: A simulated part-based interactive environment," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [10] K. Mo, S. Zhu, A. X. Chang, L. Yi, S. Tripathi, L. J. Guibas, and H. Su, "PartNet: A large-scale benchmark for fine-grained and hierarchical part-level 3D object understanding," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [11] K. Mo, L. J. Guibas, M. Mukadam, A. Gupta, and S. Tulsiani, "Where2act: From pixels to actions for articulated 3d objects," *CoRR*, vol. abs/2101.02692, 2021. [Online]. Available: <https://arxiv.org/abs/2101.02692>
- [12] R. Wu, Y. Zhao, K. Mo, Z. Guo, Y. Wang, T. Wu, Q. Fan, X. Chen, L. Guibas, and H. Dong, "VAT-Mart: Learning Visual Action Trajectory Proposals for Manipulating 3D ARTiculated Objects," Apr. 2022, arXiv:2106.14440 [cs]. [Online]. Available: <http://arxiv.org/abs/2106.14440>
- [13] M. Sundermeyer, A. Mousavian, R. Triebel, and D. Fox, "Contact-graspnet: Efficient 6-dof grasp generation in cluttered scenes," *CoRR*, vol. abs/2103.14127, 2021. [Online]. Available: <https://arxiv.org/abs/2103.14127>
- [14] X. Liu and K. M. Kitani, "V-MAO: Generative Modeling for Multi-Arm Manipulation of Articulated Objects," Nov. 2021, arXiv:2111.03987 [cs]. [Online]. Available: <http://arxiv.org/abs/2111.03987>
- [15] H. Geng, H. Xu, C. Zhao, C. Xu, L. Yi, S. Huang, and H. Wang, "Gapartnet: Cross-category domain-generalizable object perception and manipulation via generalizable and actionable parts," 2023.
- [16] H. Zhang, B. Eisner, and D. Held, "Flowbot++: Learning generalized articulated objects manipulation via articulation projection," 2023.
- [17] E. Klingbeil, A. Saxena, and A. Ng, "Learning to open new doors," *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 2751–2757, 2010.
- [18] M. Arduengo, C. Torras, and L. Sentis, "A versatile framework for robust and adaptive door operation with a mobile manipulator robot," *ArXiv*, vol. abs/1902.09051, 2019.
- [19] T. Rühr, J. Sturm, D. Pangercic, M. Beetz, and D. Cremers, "A generalized framework for opening doors and drawers in kitchen environments," *2012 IEEE International Conference on Robotics and Automation*, pp. 3852–3858, 2012.
- [20] H. Xiong, R. Mendonca, K. Shaw, and D. Pathak, "Adaptive mobile manipulation for articulated objects in the open world," 2024.
- [21] H. Geng, Z. Li, Y. Geng, J. Chen, H. Dong, and H. Wang, "Partmanip: Learning cross-category generalizable part manipulation policy from point cloud observations," 2023.
- [22] C. Bao, H. Xu, Y. Qin, and X. Wang, "DexArt: Benchmarking Generalizable Dexterous Manipulation with Articulated Objects," May 2023, arXiv:2305.05706 [cs]. [Online]. Available: <http://arxiv.org/abs/2305.05706>
- [23] A. Curtis, L. Kaelbling, and S. Jain, "Task-directed exploration in continuous pomdps for robotic manipulation of articulated objects," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 3721–3728.
- [24] Z. Xu, Z. He, and S. Song, "Universal Manipulation Policy Network for Articulated Objects," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 2447–2454, Apr. 2022, conference Name: IEEE Robotics and Automation Letters.
- [25] B. Eisner, H. Zhang, and D. Held, "Flowbot3d: Learning 3d articulation flow to manipulate articulated objects," 2023.
- [26] C. Ning, R. Wu, H. Lu, K. Mo, and H. Dong, "Where2explore: Few-shot affordance learning for unseen novel categories of articulated objects," 2023.
- [27] C. Goldfeder, M. Ciocarlie, H. Dang, and P. K. Allen, "The columbia grasp database," in *2009 IEEE international conference on robotics and automation*. IEEE, 2009, pp. 1710–1716.
- [28] D. Kappler, J. Bohg, and S. Schaal, "Leveraging big data for grasp planning," in *2015 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2015, pp. 4304–4311.
- [29] 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," *arXiv preprint arXiv:1703.09312*, 2017.
- [30] A. Mousavian, C. Eppner, and D. Fox, "6-dof graspnet: Variational grasp generation for object manipulation," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 2901–2910.
- [31] D. Morrison, P. Corke, and J. Leitner, "Egad! an evolved grasping analysis dataset for diversity and reproducibility in robotic manipulation," *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pp. 4368–4375, 2020.
- [32] H.-S. Fang, C. Wang, M. Gou, and C. Lu, "Graspnet-1billion: A large-scale benchmark for general object grasping," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 11 444–11 453.
- [33] C. Eppner, A. Mousavian, and D. Fox, "A billion ways to grasp: An evaluation of grasp sampling schemes on a dense, physics-based grasp data set," in *Robotics Research: The 19th International Symposium ISRR*. Springer, 2022, pp. 890–905.
- [34] —, "Acronym: A large-scale grasp dataset based on simulation," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 6222–6227.
- [35] H.-S. Fang, C. Wang, H. Fang, M. Gou, J. Liu, H. Yan, W. Liu, Y. Xie, and C. Lu, "Anygrasp: Robust and efficient grasp perception in spatial and temporal domains," 2023. [Online]. Available: <https://arxiv.org/abs/2212.08333>
- [36] E. Coumans and Y. Bai, "Pybullet, a python module for physics simulation for games, robotics and machine learning," 2016.
- [37] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, "Pointnet++: Deep hierarchical feature learning on point sets in a metric space," *CoRR*, vol. abs/1706.02413, 2017. [Online]. Available: <http://arxiv.org/abs/1706.02413>
- [38] C. Choy, J. Park, and V. Koltun, "Fully convolutional geometric features," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 8958–8966.
- [39] J. Zhou, M. Wang, W. Mao, M. Gong, and X. Liu, "Siamesepointnet: A siamese point network architecture for learning 3d shape descriptor," in *Computer Graphics Forum*, vol. 39, no. 1. Wiley Online Library, 2020, pp. 309–321.
- [40] S. Xie, J. Gu, D. Guo, C. R. Qi, L. Guibas, and O. Litany, "Pointcontrast: Unsupervised pre-training for 3d point cloud understanding," in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*. Springer, 2020, pp. 574–591.
- [41] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo, P. Dollár, and R. Girshick, "Segment anything," 2023.