

DUALAFFORD: LEARNING COLLABORATIVE VISUAL AFFORDANCE FOR DUAL-GRIPPER MANIPULATION

Yan Zhao^{1,6*} Ruihai Wu^{1,6*} Zhehuan Chen¹ Yourong Zhang¹ Qingnan Fan⁵
Kaichun Mo^{3,4} Hao Dong^{1,2,6†}

¹CFCS, CS Dept., PKU ²AIT, PKU ³Stanford University ⁴NVIDIA Research

⁵Tencent AI Lab ⁶BAAI

{zhaoyan790, wuruihai, acmlczh, minor_sixth, hao.dong}@pku.edu.cn
fqnchina@gmail.com kmo@nvidia.com

ABSTRACT

It is essential yet challenging for future home-assistant robots to understand and manipulate diverse 3D objects in daily human environments. Towards building scalable systems that can perform diverse manipulation tasks over various 3D shapes, recent works have advocated and demonstrated promising results learning visual actionable affordance, which labels every point over the input 3D geometry with an action likelihood of accomplishing the downstream task (*e.g.*, pushing or picking-up). However, these works only studied single-gripper manipulation tasks, yet many real-world tasks require two hands to achieve collaboratively. In this work, we propose a novel learning framework, DualAfford, to learn collaborative affordance for dual-gripper manipulation tasks. The core design of the approach is to reduce the quadratic problem for two grippers into two disentangled yet interconnected subtasks for efficient learning. Using the large-scale PartNet-Mobility and ShapeNet datasets, we set up four benchmark tasks for dual-gripper manipulation. Experiments prove the effectiveness and superiority of our method over baselines.

1 INTRODUCTION

We, humans, spend little or no effort perceiving and interacting with diverse 3D objects to accomplish everyday tasks in our daily lives. It is, however, an extremely challenging task for developing artificial intelligent robots to achieve similar capabilities due to the exceptionally rich 3D object space and high complexity manipulating with diverse 3D geometry for different downstream tasks. While researchers have recently made many great advances in 3D shape recognition (Chang et al., 2015; Wu et al., 2015), pose estimation (Wang et al., 2019; Xiang et al., 2017), and semantic understandings (Hu et al., 2018; Mo et al., 2019; Savva et al., 2015) from the vision community, as well as grasping (Mahler et al., 2019; Pinto & Gupta, 2016) and manipulating 3D objects (Chen et al., 2021; Xu et al., 2020) on the robotic fronts, there are still huge perception-interaction gaps (Batra et al., 2020; Gadre et al., 2021; Shen et al., 2021; Xiang et al., 2020) to close for enabling future home-assistant autonomous systems in the unstructured and complicated human environments.

One of the core challenges in bridging the gaps is figuring out good visual representations of 3D objects that are *generalizable* across diverse 3D shapes at a large scale and directly *consumable* by downstream planners and controllers for robotic manipulation. Recent works (Mo et al., 2021; Wu et al., 2022) have proposed a novel perception-interaction handshaking representation for 3D objects – *visual actionable affordance*, which essentially predicts an action likelihood for accomplishing the given downstream manipulation task at each point on the 3D input geometry. Such visual actionable affordance, trained across diverse 3D shape geometry (*e.g.*, refrigerators, microwaves) and for a specific downstream manipulation task (*e.g.*, pushing), is proven to *generalize* to novel unseen objects (*e.g.*, tables) and *benefits* downstream robotic executions (*e.g.*, more efficient exploration).

Though showing promising results, past works (Mo et al., 2021; Wu et al., 2022) are limited to single-gripper manipulation tasks. However, future home-assistant robots shall have two hands just like us humans, if not more, and many real-world tasks require two hands to achieve collaboratively. For example (Figure 1), to steadily pick up a heavy bucket, two grippers need to grasp it at two top

*Equal contribution

†Corresponding author; Project page: <https://hyperplane-lab.github.io/DualAfford>

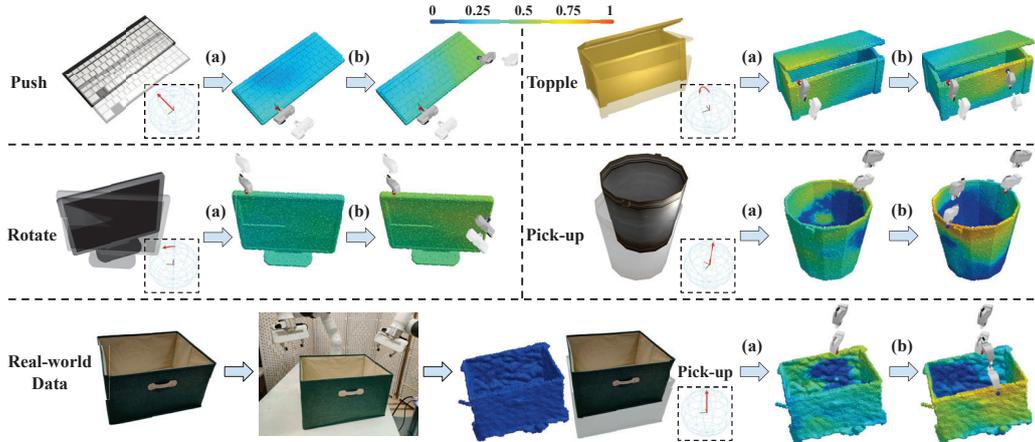


Figure 1: Given different shapes and manipulation tasks (*e.g.*, pushing the keyboard in the direction indicated by the red arrow), our proposed *DualAfford* framework predicts dual collaborative visual actionable affordance and gripper orientations. The prediction for the second gripper (b) is dependent on the first (a). We can directly apply our network to real-world data.

edges and move in the same direction; to rotate a display anticlockwise, one gripper points downward to hold it and the other gripper moves to the other side. Different manipulation patterns naturally emerge when the two grippers collaboratively attempt to accomplish different downstream tasks.

In this paper, we study the dual-gripper manipulation tasks and investigate learning collaborative visual actionable affordance. It is much more challenging to tackle dual-gripper manipulation tasks than single-gripper ones as the degree-of-freedom in action spaces is doubled and two affordance predictions are required due to the addition of the second gripper. Besides, the pair of affordance maps for the two grippers needs to be learned collaboratively. As we can observe from Figure 1, the affordance for the second gripper is dependent on the choice of the first gripper action. How to design the learning framework to learn such collaborative affordance is a non-trivial question.

We propose a novel method *DualAfford* to tackle the problem. At the core of our design, *DualAfford* disentangles the affordance learning problem of two grippers into two separate yet highly coupled subtasks, reducing the complexity of the intrinsically quadratic problem. More concretely, the first part of the network infers actionable locations for the first gripper where there exist second-gripper actions to cooperate, while the second part predicts the affordance for the second gripper conditioned on a given first-gripper action. The two parts of the system are trained as a holistic pipeline using the interaction data collected by manipulating diverse 3D shapes in a physical simulator.

We evaluate the proposed method on four diverse dual-gripper manipulation tasks: pushing, rotating, toppling and picking-up. We set up a benchmark for experiments using shapes from PartNet-Mobility dataset (Mo et al., 2019; Xiang et al., 2020) and ShapeNet dataset (Chang et al., 2015). Quantitative comparisons against baseline methods prove the effectiveness of the proposed framework. Qualitative results further show that our method successfully learns interesting and reasonable dual-gripper collaborative manipulation patterns when solving different tasks. To summarize, in this paper,

- We propose a novel architecture *DualAfford* to learn collaborative visual actionable affordance for dual-gripper manipulation tasks over diverse 3D objects;
- We set up a benchmark built upon SAPIEN physical simulator (Xiang et al., 2020) using the PartNet-Mobility and ShapeNet datasets (Chang et al., 2015; Mo et al., 2019; Xiang et al., 2020) for four dual-gripper manipulation tasks;
- We show qualitative results and quantitative comparisons against three baselines to validate the effectiveness and superiority of the proposed approach.

2 RELATED WORK

Dual-gripper Manipulation. Many studies, from both computer vision and robotics communities, have been investigating dual-gripper or dual-arm manipulation (Chen et al., 2022; Simeonov et al., 2020; Weng et al., 2022; Chitnis et al., 2020; Xie et al., 2020; Liu & Kitani, 2021; Liu et al., 2022).

Vahrenkamp et al. (2009) presented two strategies for dual-arm planning: J+ and IK-RRT. Cohen et al. (2014) proposed a heuristic search-based approach using a manipulation lattice graph. Ha et al. (2020) presented a closed-loop and decentralized motion planner to avoid a collision. Multi-arm manipulation has also been investigated in various applications: grasping (Pavlichenko et al., 2018), pick-and-place (Shome & Bekris, 2019), and rearrangement (Shome et al., 2021; Hartmann et al., 2021). Our work pays more attention to learning object-centric visual actionable affordance heatmaps for dual-arm manipulation tasks, while previous works focus more on the planning and control sides. Gadre et al. (2021) learns affordances but for interactively part segmentation, they use one gripper to simply hold one articulated part, and use the other gripper to move the other articulated part.

Visual Affordance Prediction. Predicting affordance plays an important role in visual understanding and benefits downstream robotic manipulation tasks, which has been widely used in many previous works (Jiang et al., 2021b; Kokic et al., 2020; Mandikal & Grauman, 2021; Redmon & Angelova, 2015; Wang et al., 2021; Wu et al., 2022; 2023). For example, Kokic et al. (2017) used CNN to propose a binary map indicating contact locations for task-specific grasping. Jiang et al. (2021a) proposed the contact maps by exploiting the consistency between hand contact points and object contact regions. Following Where2Act (Mo et al., 2021), we use dense affordance maps to suggest action possibilities at every point on a 3D scan. In our work, we extend by learning two collaborative affordance maps for two grippers that are in deep cooperation for accomplishing downstream tasks.

3 PROBLEM FORMULATION

General Setting. We place a random 3D object from a random category on the ground, given its partially scanned point cloud observation $O \in \mathbb{R}^{N \times 3}$ and a specific task l , the network is required to propose two grippers actions $u_1 = (p_1, R_1)$ and $u_2 = (p_2, R_2)$, in which p is the contact point and R is the manipulation orientation. All inputs and outputs are represented in the camera base coordinate frame, with the z-axis aligned with the up direction and the x-axis points to the forward direction, which is in align with real robot’s camera coordinate system.

Task Formulation. We formulate four benchmark tasks: pushing, rotating, toppling and picking-up, which are widely used in manipulation benchmarks (Andrychowicz et al., 2017; Kumar et al., 2016; Mousavian et al., 2019; OpenAI et al., 2021) and commonly used as subroutines in object grasping and relocation (Chao et al., 2021; Mahler et al., 2019; Mandikal & Grauman, 2022; Rajeswaran et al., 2018; Zeng et al., 2020). We set different success judgments for difference tasks, and here we describe the pushing task as an example. Task $l \in \mathbb{R}^3$ is a unit vector denoting the object’s goal pushing direction. An object is successfully pushed if (1) its movement distance is over 0.05 unit-length, (2) the difference between its actual motion direction l' and goal direction l is within 30 degrees, (3) the object should be moved steadily, *i.e.*, the object can not be rotated or toppled by grippers.

4 METHOD

4.1 OVERVIEW OF *DualAfford* FRAMEWORK

Figure 2 presents the overview of our proposed *DualAfford* framework. Firstly, we collect large amount of interaction data to supervise the perception networks. Since it is costly to collect human annotations for dual-gripper manipulations, we use an interactive simulator named SAPIEN Xiang et al. (2020). We sample offline interactions by using either a random data sampling method or an optional reinforcement-learning (RL) augmented data sampling method described in Sec. 4.5.

We propose the novel Perception Module to learn collaborative visual actionable affordance and interaction policy for dual-gripper manipulation tasks over diverse objects. To reduce the complexity of the intrinsically quadratic problem of dual-gripper manipulation tasks, we disentangle the task into two separate yet highly coupled subtasks. Specifically, let N denote the point number of the point cloud, and θ_R denote the gripper orientation space on one point. If the network predicts the two gripper actions simultaneously, the combinatorial search space will be $O((\theta_R)^{(N \times N)})$. However, our Perception Module sequentially predicts two affordance maps and gripper actions in a conditional manner, which reduces the search space to $O((\theta_R)^{(N+N)})$. Therefore, we design two coupled submodules in the Perception Module: the First Gripper Module \mathcal{M}_1 (left) and the Second Gripper Module \mathcal{M}_2 (right), and each gripper module consists of three networks (Sec. 4.2).

The training and inference procedures, respectively indicated by the red and blue arrows in Figure 2, share the same architecture but with reverse dataflow directions. For inference, the dataflow direction is intuitive: \mathcal{M}_1 proposes u_1 , and then \mathcal{M}_2 proposes u_2 conditioned on u_1 . Although such dataflow guarantees the second gripper plays along with the first during inference, it cannot guarantee the first

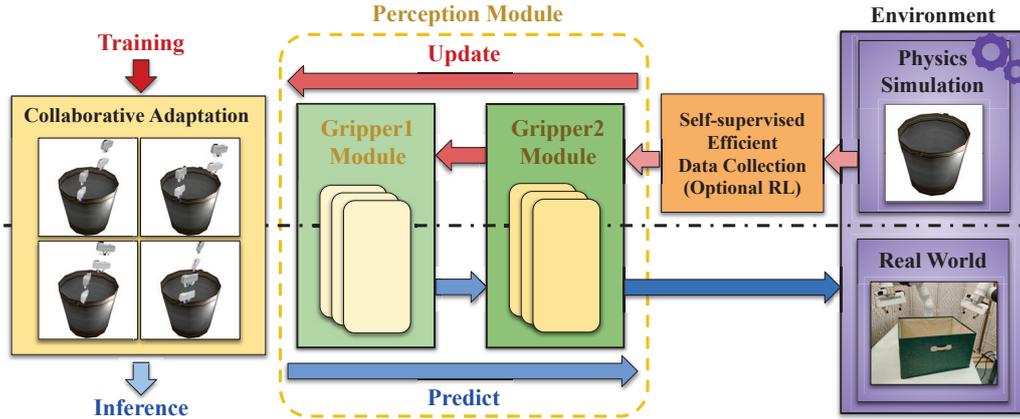


Figure 2: **Our proposed DualAfford framework**, first collects interaction data points in physics simulation, then uses them to train the Perception Module, which contains the First Gripper Module and the Second Gripper Module, and further enhances the cooperation between two grippers through the Collaborative Adaption procedure. The training and the inference procedures, as respectively indicated by the red and blue arrows, share the same architecture but with opposite dataflow directions.

gripper’s action is suitable for the second to collaborate with. To tackle this problem, for training, we employ the reverse dataflow: \mathcal{M}_2 is trained first, and then \mathcal{M}_1 is trained with the awareness of the trained \mathcal{M}_2 . Specifically, given diverse u_1 in training dataset, \mathcal{M}_2 is first trained to propose u_2 collaborative with them. Then, with the trained \mathcal{M}_2 able to propose u_2 collaborative with different u_1 , \mathcal{M}_1 learns to propose u_1 that are easy for \mathcal{M}_2 to propose successful collaborations. In this way, both \mathcal{M}_1 and \mathcal{M}_2 are able to propose actions easy for the other to collaborate with.

Although such design encourages two grippers to cooperate, the two gripper modules are separately trained using only offline collected data, and their proposed actions are never truly executed as a whole, so they are not explicitly taught if their collaboration is successful. To further enhance their cooperation, we introduce the Collaborative Adaptation procedure (Sec. 4.4), in which we execute two grippers’ actions simultaneously in simulator, using the outcomes to provide training supervision.

4.2 PERCEPTION MODULE AND INFERENCE

To reduce the complexity of the intrinsically quadratic problem and relieve the learning burden of our networks, we disentangle the dual-gripper learning problem into two separate yet coupled subtasks. We design a conditional perception pipeline containing two submodules shown in Figure 3, in which u_2 is proposed conditioned on u_1 during inference, while \mathcal{M}_1 is trained conditioned on the trained \mathcal{M}_2 during training. There are three networks in each gripper module: Affordance Network \mathcal{A} , Proposal Network \mathcal{P} and Critic Network \mathcal{C} . First, as the gripper action can be decomposed into a contact point and a gripper orientation, we design Affordance Network and Proposal Network to respectively predict them. Also, to evaluate whether an action of the gripper is suitable for collaboration, we design Critic Network for this purpose. Below we describe the design of each module.

Backbone Feature Extractors. The networks in Perception Module may receive four kinds of input entities or intermediate results: point cloud O , task l , contact point p , and gripper orientation R . In different submodules, the backbone feature extractors share the same architectures. We use a segmentation-version PointNet++ (Qi et al., 2017) to extract per-point feature $f_s \in \mathbb{R}^{128}$ from O , and employ three MLP networks to respectively encode l , p , and R into $f_l \in \mathbb{R}^{32}$, $f_p \in \mathbb{R}^{32}$, and $f_R \in \mathbb{R}^{32}$.

4.2.1 THE FIRST GRIPPER MODULE

The First Gripper Module contains three sequential networks. Given an object and a task configuration, the Affordance Network \mathcal{A}_1 indicates where to interact by predicting affordance map, the Proposal Network \mathcal{P}_1 suggests how to interact by predicting manipulation orientations, and the Critic Network \mathcal{C}_1 evaluates the per-action success likelihood.

Affordance Network. This network \mathcal{A}_1 predicts an affordance score $a_1 \in [0, 1]$ for each point p , indicating the success likelihood when the first gripper interacts with the point, with the assumption that there exists an expert second gripper collaborating with it. Aggregating the affordance scores, we acquire an affordance map A_1 over the partial observation, from which we can filter out low-rated

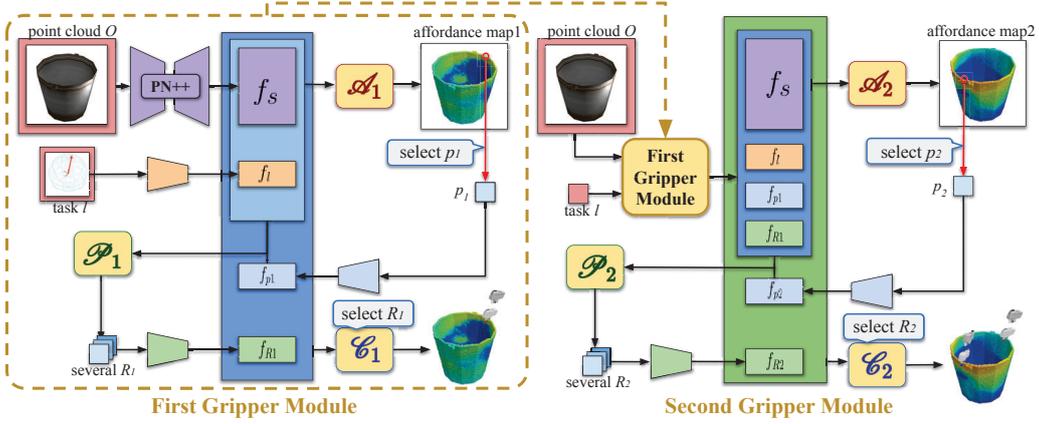


Figure 3: **Architecture details of the Perception Module.** Given a 3D partial scan and a specific task, our network sequentially predicts the first and second grippers’ affordance maps and manipulation actions in a conditional manner. Each gripper module is composed of 1) an Affordance Network \mathcal{A} indicating where to interact; 2) a Proposal Network \mathcal{P} suggesting how to interact; 3) a Critic Network \mathcal{C} evaluating the success likelihood of an interaction.

proposals and select a contact point p_1 for the first gripper. This network is implemented as a single-layer MLP that receives the feature concatenation of f_s , f_l and f_{p_1} .

Proposal Network. This network \mathcal{P}_1 models the distribution of the gripper’s orientation R_1 on the given point p_1 . It is implemented as a conditional variational autoencoder (Sohn et al., 2015), where an encoder maps the gripper orientation into a Gaussian noise $z \in \mathbb{R}^{32}$, a decoder reconstructs it from z . Implemented as MLPs, they both take the feature concatenation of f_s , f_l , and f_{p_1} as the condition.

Critic Network. This network \mathcal{C}_1 rates the success likelihood of each manipulation orientation on each point by predicting a scalar $c_1 \in [0, 1]$. A higher c_1 indicates a higher potential for the second gripper to collaboratively achieve the given task. It is implemented as a single-layer MLP that consumes the feature concatenation of f_s , f_l , f_{p_1} and f_{R_1} .

4.2.2 THE SECOND GRIPPER MODULE

Conditioned on the first gripper action $u_1 = (p_1, R_1)$ proposed by \mathcal{M}_1 , \mathcal{M}_2 first generates a point-level collaborative affordance A_2 for the second gripper and samples a contact point p_2 . Then, \mathcal{M}_2 proposes multiple candidate orientations, among which we can choose a suitable one as R_2 . The design philosophy and implementations of \mathcal{M}_2 are the same as \mathcal{M}_1 , except that all three networks (\mathcal{A}_2 , \mathcal{P}_2 and \mathcal{C}_2) take the first gripper’s action u_1 , i.e., p_1 and R_1 , as the additional input.

4.3 TRAINING AND LOSSES

As shown in Figure 2, during inference (indicated by blue arrows), the first gripper predicts actions without seeing how the second gripper will collaborate. To enable the first gripper to propose actions easy for the second to collaborate with, we train the Perception Module in the dataflow direction indicated by red arrows, as described in Sec. 4.1. We adopt the Critic Network \mathcal{C}_1 of the first gripper as a bridge to connect two gripper modules. \mathcal{C}_1 scores whether an action of the first gripper is easy for \mathcal{M}_2 to propose collaborative actions. With the trained \mathcal{C}_1 , \mathcal{M}_1 will propose actions with the assumption that there exists an expert gripper to cooperate with. Therefore, \mathcal{M}_1 and \mathcal{M}_2 will both learn to collaborate with each other.

Critic Loss. It is relatively easy to train the second Critic Network \mathcal{C}_2 . Given the interaction data with the corresponding ground-truth interaction result r , where $r = 1$ means positive and $r = 0$ means negative, we can train \mathcal{C}_2 using the standard binary cross-entropy loss. For simplicity, we use f^m to denote each network’s input feature concatenation, as mentioned in Sec.4.2:

$$\mathcal{L}_{\mathcal{C}_2} = r_j \log(\mathcal{C}_2(f_{p_2}^m)) + (1 - r_j) \log(1 - \mathcal{C}_2(f_{p_2}^m)). \quad (1)$$

However, for the first Critic Network \mathcal{C}_1 , since we only know the first gripper’s action $u_1 = (p_1, R_1)$, we can not directly obtain the ground-truth interaction outcome of a single action u_1 . To tackle this problem, given the first gripper’s action, we evaluate it by estimating the potential for the second

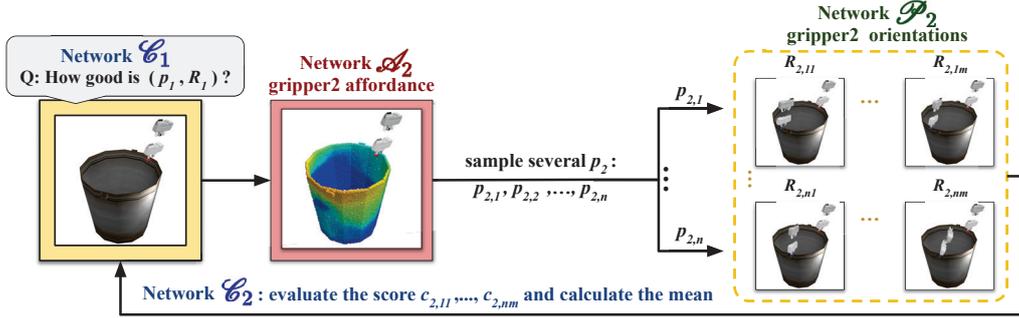


Figure 4: To train \mathcal{E}_1 that evaluates how the first action can collaborate with the trained Second Gripper Module \mathcal{M}_2 , we comprehensively use the trained \mathcal{A}_2 , \mathcal{P}_2 and \mathcal{E}_2 of \mathcal{M}_2 to provide supervision.

gripper to collaboratively accomplish the given task. As shown in Figure 4, we comprehensively use the trained \mathcal{A}_2 , \mathcal{P}_2 and \mathcal{E}_2 of the Second Gripper Module \mathcal{M}_2 . Specifically, to acquire the ground-truth action score \hat{c} for the first gripper, we first use \mathcal{A}_2 to predict the collaborative affordance map A_2 and sample n contact points: $p_{2,1}, \dots, p_{2,n}$, then we use \mathcal{P}_2 to sample m interaction orientations on each contact point i : $R_{2,i,1}, \dots, R_{2,i,m}$. Finally, we use \mathcal{E}_2 to rate the scores of these actions: $c_{2,11}, \dots, c_{2,nm}$ and calculate their average value. Thus we acquire the ground-truth score of \mathcal{E}_1 , and we apply \mathcal{L}_1 loss to measure the error between the prediction and the ground-truth:

$$\hat{c}_{p_1} = \frac{1}{nm} \sum_{j=1}^n \sum_{k=1}^m \mathcal{E}_2(f_{p_{2j}}^{in}, \mathcal{P}_2(f_{p_{2j}}^{in}, z_{jk})); \quad \mathcal{L}_{\mathcal{E}_1} = |\mathcal{E}_1(f_{p_1}^{in}) - \hat{c}_{p_1}|. \quad (2)$$

Proposal Loss. \mathcal{P}_1 and \mathcal{P}_2 are implemented as cVAE (Sohn et al., 2015). For the i -th gripper, we apply geodesic distance loss to measure the error between the reconstructed gripper orientation R_i and ground-truth \hat{R}_i , and KL Divergence to measure the difference between two distributions:

$$\mathcal{L}_{\mathcal{P}_i} = \mathcal{L}_{geo}(R_i, \hat{R}_i) + D_{KL}(q(z|\hat{R}_i, f^{in}) || \mathcal{N}(0, 1)). \quad (3)$$

Affordance Loss. Similar to Where2Act (Mo et al., 2021), for each point, we adopt the 'affordance' score as the expected success rate when executing action proposals generated by the Proposal Network \mathcal{P} , which can be directly evaluated by the Critic Network \mathcal{E} . Specifically, to acquire the ground-truth affordance score \hat{a} for the i -th gripper, we sample n gripper orientations on the point p_i using \mathcal{P}_i , and calculate their average action scores rated by \mathcal{E}_i . We apply \mathcal{L}_1 loss to measure the error between the prediction and the ground-truth affordance score on a certain point:

$$\hat{a}_{p_i} = \frac{1}{n} \sum_{j=1}^n \mathcal{E}_i(f_{p_i}^{in}, \mathcal{P}_i(f_{p_i}^{in}, z_j)); \quad \mathcal{L}_{\mathcal{A}_i} = |\mathcal{A}_i(f_{p_i}^{in}) - \hat{a}_{p_i}|. \quad (4)$$

4.4 COLLABORATIVE ADAPTATION PROCEDURE

Although the above training procedure can enable two gripper modules to propose affordance and actions collaboratively, their collaboration is limited, because they are trained in a separate and sequential way using only offline collected data, without any real and simultaneous executions of proposed actions. To further enhance the collaboration between the two gripper modules, we introduce the Collaborative Adaptation procedure, in which the two modules are trained in a simultaneous manner using online executed and collected data, with loss functions the same as in Sec. 4.3. In this procedure, the proposed dual-gripper actions are simultaneously executed in the simulator, using interaction outcomes to update the two gripper modules. In this way, the two gripper modules can better understand whether their proposed actions are successful or not as they are aware of interaction results, and thus the two separately trained modules are integrated into one collaborative system.

4.5 OFFLINE DATA COLLECTION

Instead of acquiring costly human annotations, we use SAPIEN (Xiang et al., 2020) to sample large amount of offline interaction data. For each interaction trial with each object, we sample two gripper actions u_1, u_2 , and test the interaction result r . We define a trial to be positive when: (1) the two grippers successfully achieve the task, e.g., pushing a display over a threshold length without rotating it; (2) the task can be accomplished only by the collaboration of two grippers, i.e., when we replay each gripper action without the other, the task can not be achieved. We represent each interaction data as $(O, l, p_1, p_2, R_1, R_2) \rightarrow r$, and balance the number of positive and negative interactions. Here we introduce two data collection methods: random and RL augmented data sampling.

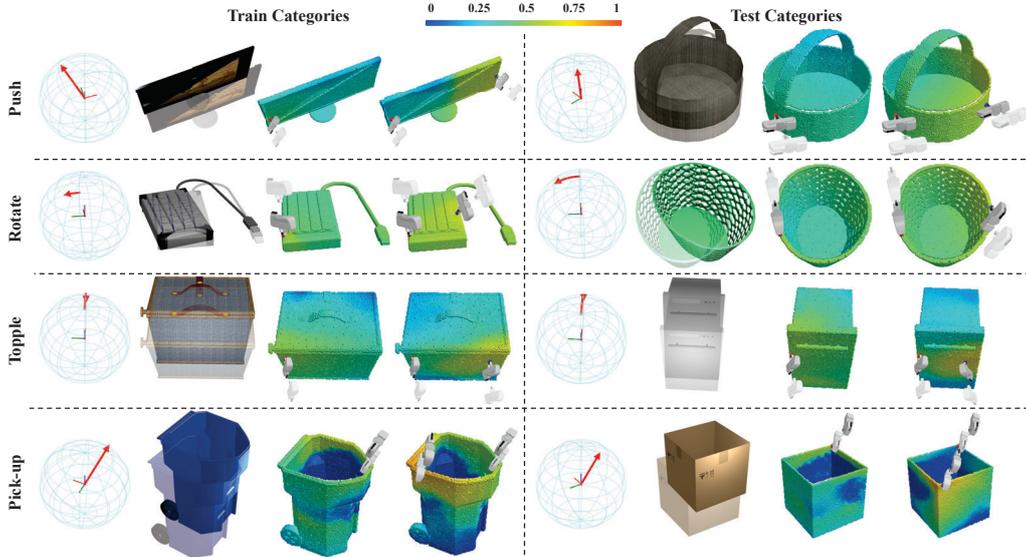


Figure 5: Qualitative results of Affordance Networks. In each block, we respectively show (1) task represented by a red arrow, (2) object which should be moved from transparent to solid, (3) the first affordance map predicted by \mathcal{A}_1 , (4) the second affordance map predicted by \mathcal{A}_2 conditioned on the first action. Left shapes are from training categories, while right shapes are from unseen categories.

Random Data Sampling. We can efficiently sample interaction data by parallelizing simulation environments across multiple CPUs. For each data point, we first randomly sample two contact points on the object point cloud, then we randomly sample two interaction orientations from the hemisphere above the tangent plane around the point, and finally test the interaction result.

RL Augmented Data Sampling. For tasks with complexity, such as picking-up, it is nearly impossible for a random policy to collect positive data. To tackle this problem, we propose the RL method. We first leverage Where2Act (Mo et al., 2021) to propose a prior affordance map, highlighting where to grasp. After sampling two contact points, we use SAC (Haarnoja et al., 2018) with the manually designed dense reward functions to efficiently predict interaction orientations.

5 EXPERIMENTS

5.1 RESULTS AND ANALYSIS

We perform large-scale experiments under four dual-gripper manipulation tasks, and set up three baselines for comparisons. Results prove the effectiveness and superiority of our proposed approach.

5.2 ENVIRONMENT SETTINGS AND DATASET

We follow the environment settings of Where2Act (Mo et al., 2021) except that we use two Franka Panda Flying grippers. We conduct our experiments on SAPIEN (Xiang et al., 2020) simulator with the large-scale PartNet-Mobility (Mo et al., 2019) and ShapeNet (Chang et al., 2015) dataset. To analyze whether the learned representations can generalize to novel unseen categories, we reserve some categories only for testing. See Supplementary Sec. C for more details.

5.3 EVALUATION METRICS, BASELINES AND ABLATION

Evaluation Metrics. To quantitatively evaluate the action proposal quality, we run interaction trials in simulation and report sample-success-rate (Mo et al., 2021), which measures the percentage of successful interactions among all interaction trials proposed by the networks.

Baselines. We compare our approach with three baselines and one ablated version: (1) A random approach that randomly selects the contact points and gripper orientations. (2) A heuristic approach in which we acquire the ground-truth object poses and hand-engineer a set of rules for different tasks. For example, for the picking-up task, we set the two contact points on the objects’ left and right top edges and set the two gripper orientations the same as the given picking-up direction. (3) M-Where2Act: a dual-gripper Where2Act (Mo et al., 2021) approach. While Where2Act initially

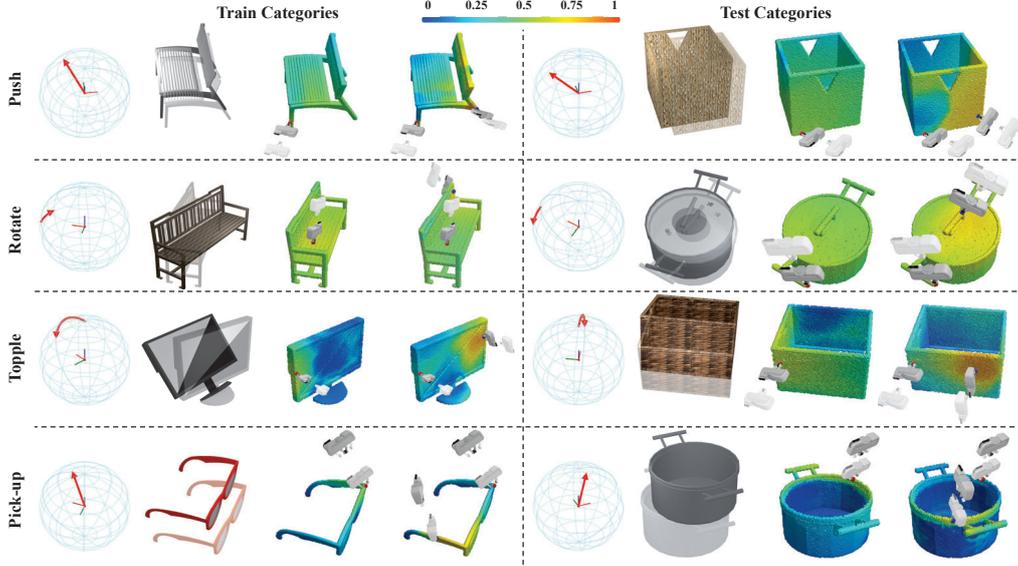


Figure 6: The per-point action scores predicted by Critic Networks \mathcal{C}_1 and \mathcal{C}_2 . In each result block, from left to right, we show the task, the input shape, the per-point success likelihood predicted by \mathcal{C}_1 given the first gripper orientation, and the per-point success likelihood predicted by \mathcal{C}_2 given the second gripper orientation, conditioned on the first gripper’s action.

considers interactions for a single gripper, we adapt it as a baseline by modifying each module in Where2Act to consider the dual grippers as a combination, and assign a task l to it as well. (4) Ours w/o CA: an ablated version of our method that removes the Collaborative Adaptation procedure.

Figure 5 presents the dual affordance maps predicted by our Affordance Networks \mathcal{A}_1 and \mathcal{A}_2 , as well as the proposed grippers interacting with the high-rated points. We can observe that: (1) the affordance maps reasonably highlight where to interact (*e.g.*, for picking-up, the grippers can only grasp the top edge); (2) the affordance maps embody the cooperation between the two grippers (*e.g.*, to collaboratively push a display, the two affordance maps sequentially highlight its left and right half part, so that the display can be pushed steadily.) Besides, we find that our method has the ability to generalize to novel unseen categories.

In Figure 6, we additionally visualize the results of Critic Networks \mathcal{C}_1 and \mathcal{C}_2 . Given different gripper orientations, the Critic Networks propose the per-point action scores over the whole point cloud. We can observe that our network is aware of the shape geometries, gripper orientations and tasks. For example, in the Rotate-Train-Categories block, the first map highlights a part of chair surface since the first gripper is downward, and the second map accordingly highlights the chair back on the other side given the second-gripper orientation, which collaboratively ensures the chair is rotated clockwise. It is noteworthy that in the first map the chair surface has higher scores than the arm, because the chair tends to skid when selecting the arm as a fulcrum for rotation.

Figure 7 (a) visualizes the diverse collaborative actions proposed by Proposal networks \mathcal{P}_1 and \mathcal{P}_2 on an example display. Our networks can propose different orientations on the same points.

Table 1: Baseline comparison on the sample-success-rate metric.

	Train Categories				Test Categories			
	pushing	rotating	toppling	picking-up	pushing	rotating	toppling	picking-up
Random	7.40	10.40	6.40	3.00	3.20	9.00	3.00	6.00
Heuristic	32.40	24.20	54.00	31.93	25.80	21.80	38.00	37.90
M-Where2Act	28.00	15.67	36.60	5.00	23.40	10.67	25.60	13.80
Ours w/o CA	35.87	17.53	56.00	28.87	34.67	15.33	39.67	38.33
Ours	48.76	33.73	65.53	40.33	42.93	35.07	41.80	54.33

Table 1 presents the sample-success-rate of different methods over the four challenging tasks. We can see that our method outperforms three baselines over all comparisons.

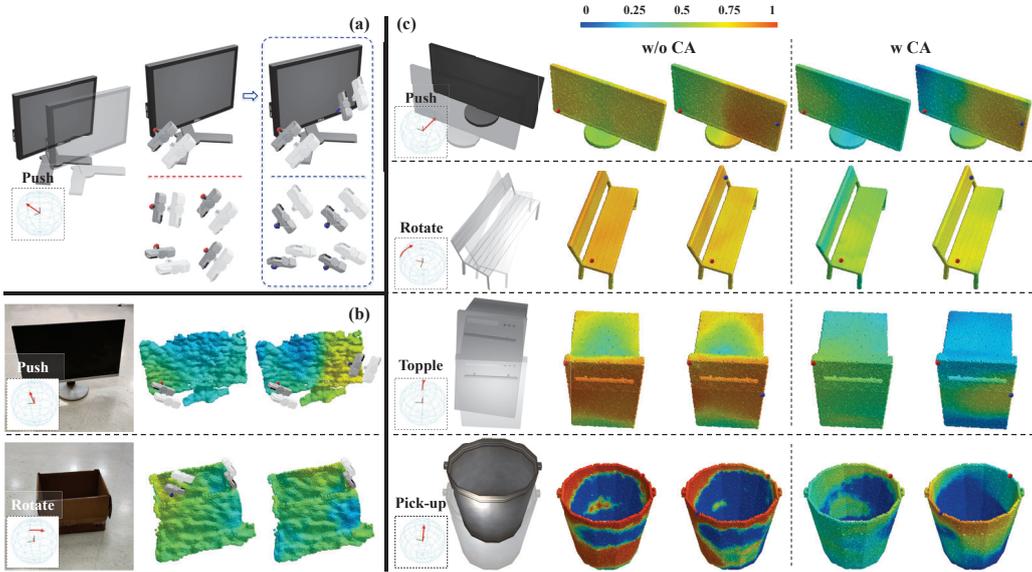


Figure 7: (a) The diverse and collaborative actions proposed by the Proposal Networks \mathcal{P}_1 and \mathcal{P}_2 . (b) The promising results testing on real-world data. (c) The actionable affordance maps of the ablated version that removes the Collaborative Adaptation procedure (left) and ours (right).

For **the heuristic baseline**, it gains relatively high numbers since it proposes actions with the ground-truth object poses and orientations. However, the inter- and intra-category shape geometries are exceptionally diverse, and we can not expect the hand-engineered rules to work for all shapes. Moreover, in the real world, this approach needs more effort to acquire ground-truth information.

For **M-Where2Act**, it learns the dual contact points and orientations as a combination and has worse performance. In comparison, our method disentangles the collaboration learning problem and reduces the complexity. Besides, M-Where2Act consumes nearly quadratic time to give proposals for the reason that it has to query the affordance scores of all the $n * n$ pair combinations of n points.

For **Ours w/o CA**, this ablated version of our method shows that the Collaborative Adaption procedure helps boost the performance. Figure 7 (c) visualizes the affordance maps without (left) and with (right) Collaborative Adaption procedure. We find that the affordance maps become more refined. For example, to push the display, the affordance scores of the base become lower since it is difficult to interact with; to collaboratively topple the dishwasher, in the second affordance map, the left front surface receives lower scores while the right maintains relatively higher.

Table 2 shows the success rate of the Random Data Sampling and RL Augmented Data Sampling method. The RL method significantly improves data collection efficiency on each object category.

Figure 1 and Figure 7 (b) show qualitative results that our networks can directly transfer the learned affordance to real-world data. We show more real-robot experiments in supplementary Sec. A.

Table 2: The success rate of data collection in the picking-up task.

	Train Categories						Test Categories			
	eyeglasses	bucket	trash can	pliers	basket	display	box	kitchen pot	scissors	laptop
Random-Sampling	0.06	0.12	0.04	< 0.01	0.09	0.03	0.03	0.06	< 0.01	< 0.01
RL-Sampling	6.12	9.65	5.26	5.79	6.41	9.78	5.12	9.18	6.38	7.13

6 CONCLUSION

In this paper, we proposed a novel framework *DualAfford* for learning collaborative actionable affordance for dual-gripper manipulation over diverse 3D shapes. We set up large-scale benchmarks for four dual-gripper manipulation tasks using the PartNet-Mobility and ShapeNet datasets. Results proved the effectiveness of the approach and its superiority over the three baselines.

Acknowledgements. Yan Zhao, Ruihai Wu, Zhehuan Chen, Yourong Zhang, Hao Dong were supported by National Natural Science Foundation of China (No. 62136001).

Ethics Statements. Our project helps build future home-assistant robots. The visual affordance predictions can be visualized and analyzed before use, which is less vulnerable to potential dangers. The training data used may introduce data bias, but this is a general concern for similar methods. We do not see any particular major harm or issue our work may raise up.

REFERENCES

- Marcin Andrychowicz, Filip Wolski, Alex Ray, Jonas Schneider, Rachel Fong, Peter Welinder, Bob McGrew, Josh Tobin, OpenAI Pieter Abbeel, and Wojciech Zaremba. Hindsight experience replay. *Advances in neural information processing systems*, 30, 2017.
- Dhruv Batra, Angel X Chang, Sonia Chernova, Andrew J Davison, Jia Deng, Vladlen Koltun, Sergey Levine, Jitendra Malik, Igor Mordatch, Roozbeh Mottaghi, et al. Rearrangement: A challenge for embodied ai. *arXiv preprint arXiv:2011.01975*, 2020.
- Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015.
- Yu-Wei Chao, Wei Yang, Yu Xiang, Pavlo Molchanov, Ankur Handa, Jonathan Tremblay, Yashraj S. Narang, Karl Van Wyk, Umar Iqbal, Stan Birchfield, Jan Kautz, and Dieter Fox. DexYCB: A benchmark for capturing hand grasping of objects. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- Tao Chen, Jie Xu, and Pulkit Agrawal. A system for general in-hand object re-orientation. *Conference on Robot Learning*, 2021.
- Yuanpei Chen, Yaodong Yang, Tianhao Wu, Shengjie Wang, Xidong Feng, Jiechuang Jiang, Stephen Marcus McAleer, Hao Dong, Zongqing Lu, and Song-Chun Zhu. Towards human-level bimanual dexterous manipulation with reinforcement learning, 2022.
- Rohan Chitnis, Shubham Tulsiani, Saurabh Gupta, and Abhinav Gupta. Efficient bimanual manipulation using learned task schemas. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1149–1155. IEEE, 2020.
- Sachin Chitta, Ioan Sucan, and Steve Cousins. Moveit! *IEEE Robotics & Automation Magazine*, 19(1):18–19, 2012.
- Benjamin Cohen, Sachin Chitta, and Maxim Likhachev. Single- and dual-arm motion planning with heuristic search. *The International Journal of Robotics Research*, 33(2):305–320, 2014. doi: 10.1177/0278364913507983. URL <https://doi.org/10.1177/0278364913507983>.
- Samir Yitzhak Gadre, Kiana Ehsani, and Shuran Song. Act the part: Learning interaction strategies for articulated object part discovery. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 15752–15761, October 2021.
- Huy Ha, Jingxi Xu, and Shuran Song. Learning a decentralized multi-arm motion planner. In *Conference on Robotic Learning (CoRL)*, 2020.
- Tuomas Haarnoja, Aurick Zhou, Kristian Hartikainen, George Tucker, Sehoon Ha, Jie Tan, Vikash Kumar, Henry Zhu, Abhishek Gupta, Pieter Abbeel, and Sergey Levine. Soft actor-critic algorithms and applications. *CoRR*, abs/1812.05905, 2018. URL <http://arxiv.org/abs/1812.05905>.
- Valentin Noah Hartmann, Andreas Orthey, Danny Driess, Ozgur S. Oguz, and Marc Toussaint. Long-horizon multi-robot rearrangement planning for construction assembly. *CoRR*, abs/2106.02489, 2021. URL <https://arxiv.org/abs/2106.02489>.

- Ruizhen Hu, Manolis Savva, and Oliver van Kaick. Functionality representations and applications for shape analysis. In *Computer Graphics Forum*, volume 37, pp. 603–624. Wiley Online Library, 2018.
- Hanwen Jiang, Shaowei Liu, Jiashun Wang, and Xiaolong Wang. Hand-object contact consistency reasoning for human grasps generation. *CoRR*, abs/2104.03304, 2021a. URL <https://arxiv.org/abs/2104.03304>.
- Zhenyu Jiang, Yifeng Zhu, Maxwell Svetlik, Kuan Fang, and Yuke Zhu. Synergies between affordance and geometry: 6-dof grasp detection via implicit representations. *CoRR*, abs/2104.01542, 2021b. URL <https://arxiv.org/abs/2104.01542>.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Mia Kokic, Johannes A Stork, Joshua A Haustein, and Danica Kragic. Affordance detection for task-specific grasping using deep learning. In *2017 IEEE-RAS 17th International Conference on Humanoid Robotics (Humanoids)*, pp. 91–98. IEEE, 2017.
- Mia Kokic, Danica Kragic, and Jeannette Bohg. Learning task-oriented grasping from human activity datasets. *IEEE Robotics and Automation Letters*, 5(2):3352–3359, 2020. doi: 10.1109/LRA.2020.2975706.
- Vikash Kumar, Abhishek Gupta, Emanuel Todorov, and Sergey Levine. Learning dexterous manipulation policies from experience and imitation. *arXiv preprint arXiv:1611.05095*, 2016.
- Junjia Liu, Yiting Chen, Zhipeng Dong, Shixiong Wang, Sylvain Calinon, Miao Li, and Fei Chen. Robot cooking with stir-fry: Bimanual non-prehensile manipulation of semi-fluid objects. *IEEE Robotics and Automation Letters*, 7(2):5159–5166, 2022.
- Xingyu Liu and Kris M Kitani. V-mao: Generative modeling for multi-arm manipulation of articulated objects. In *5th Annual Conference on Robot Learning*, 2021.
- Jeffrey Mahler, Matthew Matl, Vishal Satish, Michael Danielczuk, Bill DeRose, Stephen McKinley, and Ken Goldberg. Learning ambidextrous robot grasping policies. *Science Robotics*, 4(26): eaau4984, 2019.
- Priyanka Mandikal and Kristen Grauman. Learning dexterous grasping with object-centric visual affordances. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 6169–6176, 2021. doi: 10.1109/ICRA48506.2021.9561802.
- Priyanka Mandikal and Kristen Grauman. Dexvip: Learning dexterous grasping with human hand pose priors from video. *CoRR*, abs/2202.00164, 2022. URL <https://arxiv.org/abs/2202.00164>.
- Kaichun Mo, Shilin Zhu, Angel X Chang, Li Yi, Subarna Tripathi, Leonidas J Guibas, and Hao Su. Partnet: A large-scale benchmark for fine-grained and hierarchical part-level 3d object understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 909–918, 2019.
- Kaichun Mo, Leonidas J. Guibas, Mustafa Mukadam, Abhinav Gupta, and Shubham Tulsiani. Where2act: From pixels to actions for articulated 3d objects. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 6813–6823, October 2021.
- Arsalan Mousavian, Clemens Eppner, and Dieter Fox. 6-dof graspnet: Variational grasp generation for object manipulation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 2901–2910, 2019.
- OpenAI OpenAI, Matthias Plappert, Raul Sampedro, Tao Xu, Ilge Akkaya, Vineet Kosaraju, Peter Welinder, Ruben D’Sa, Arthur Petron, Henrique P d O Pinto, et al. Asymmetric self-play for automatic goal discovery in robotic manipulation. *arXiv preprint arXiv:2101.04882*, 2021.

- Dmytro Pavlichenko, Diego Rodriguez, Max Schwarz, Christian Lenz, Arul Selvam Periyasamy, and Sven Behnke. Autonomous dual-arm manipulation of familiar objects. In *2018 IEEE-RAS 18th International Conference on Humanoid Robots (Humanoids)*, pp. 1–9, 2018. doi: 10.1109/HUMANOIDS.2018.8624922.
- Lerrel Pinto and Abhinav Gupta. Supersizing self-supervision: Learning to grasp from 50k tries and 700 robot hours. In *2016 IEEE international conference on robotics and automation (ICRA)*, pp. 3406–3413. IEEE, 2016.
- Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems*, 30, 2017.
- Morgan Quigley, Ken Conley, Brian Gerkey, Josh Faust, Tully Foote, Jeremy Leibs, Rob Wheeler, Andrew Y Ng, et al. Ros: an open-source robot operating system. In *ICRA workshop on open source software*, volume 3, pp. 5. Kobe, Japan, 2009.
- Aravind Rajeswaran, Vikash Kumar, Abhishek Gupta, Giulia Vezzani, John Schulman, Emanuel Todorov, and Sergey Levine. Learning Complex Dexterous Manipulation with Deep Reinforcement Learning and Demonstrations. In *Proceedings of Robotics: Science and Systems (RSS)*, 2018.
- Joseph Redmon and Anelia Angelova. Real-time grasp detection using convolutional neural networks. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1316–1322, 2015. doi: 10.1109/ICRA.2015.7139361.
- Manolis Savva, Angel X Chang, and Pat Hanrahan. Semantically-enriched 3d models for common-sense knowledge. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 24–31, 2015.
- Bokui Shen, Fei Xia, Chengshu Li, Roberto Martín-Martín, Linxi Fan, Guanzhi Wang, Claudia Pérez-D’Arpino, Shyamal Buch, Sanjana Srivastava, Lyne P. Tchapmi, Micael E. Tchapmi, Kent Vainio, Josiah Wong, Li Fei-Fei, and Silvio Savarese. igibson 1.0: a simulation environment for interactive tasks in large realistic scenes. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. accepted. IEEE, 2021.
- Rahul Shome and Kostas E. Bekris. Anytime multi-arm task and motion planning for pick-and-place of individual objects via handoffs. In *2019 International Symposium on Multi-Robot and Multi-Agent Systems (MRS)*, pp. 37–43, 2019. doi: 10.1109/MRS.2019.8901083.
- Rahul Shome, Kiril Solovey, Jingjin Yu, Kostas Bekris, and Dan Halperin. Fast, high-quality two-arm rearrangement in synchronous, monotone tabletop setups. *IEEE Transactions on Automation Science and Engineering*, 18(3):888–901, 2021. doi: 10.1109/TASE.2021.3055144.
- Anthony Simeonov, Yilun Du, Beomjoon Kim, Francois R. Hogan, Joshua Tenenbaum, Pulkit Agrawal, and Alberto Rodriguez. A long horizon planning framework for manipulating rigid pointcloud objects. In *Conference on Robot Learning (CoRL)*, 2020. URL <https://anthonysimeonov.github.io/rpo-planning-framework/>.
- Kihyuk Sohn, Honglak Lee, and Xinchen Yan. Learning structured output representation using deep conditional generative models. *Advances in neural information processing systems*, 28:3483–3491, 2015.
- Nikolaus Vahrenkamp, Dmitry Berenson, Tamim Asfour, James Kuffner, and Rüdiger Dillmann. Humanoid motion planning for dual-arm manipulation and re-grasping tasks. In *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 2464–2470, 2009. doi: 10.1109/IROS.2009.5354625.
- He Wang, Srinath Sridhar, Jingwei Huang, Julien Valentin, Shuran Song, and Leonidas J Guibas. Normalized object coordinate space for category-level 6d object pose and size estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2642–2651, 2019.

- Yian Wang, Ruihai Wu, Kaichun Mo, Jiaqi Ke, Qingnan Fan, Leonidas Guibas, and Hao Dong. Adaafford: Learning to adapt manipulation affordance for 3d articulated objects via few-shot interactions. *arXiv preprint arXiv:2112.00246*, 2021.
- Thomas Weng, Sujay Man Bajracharya, Yufei Wang, Khush Agrawal, and David Held. Fabricflownet: Bimanual cloth manipulation with a flow-based policy. In *Conference on Robot Learning*, pp. 192–202. PMLR, 2022.
- Ruihai Wu, Yan Zhao, Kaichun Mo, Zizheng Guo, Yian Wang, Tianhao Wu, Qingnan Fan, Xuelin Chen, Leonidas Guibas, and Hao Dong. VAT-mart: Learning visual action trajectory proposals for manipulating 3d ARTiculated objects. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=iEx3PiooLy>.
- Ruihai Wu, Chuanruo Ning, and Hao Dong. Learning foresightful dense visual affordance for deformable object manipulation. *arXiv preprint arXiv:2303.11057*, 2023.
- Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1912–1920, 2015.
- Fanbo Xiang, Yuzhe Qin, Kaichun Mo, Yikuan Xia, Hao Zhu, Fangchen Liu, Minghua Liu, Hanxiao Jiang, Yifu Yuan, He Wang, et al. Sapien: A simulated part-based interactive environment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11097–11107, 2020.
- Yu Xiang, Tanner Schmidt, Venkatraman Narayanan, and Dieter Fox. Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes. *arXiv preprint arXiv:1711.00199*, 2017.
- Fan Xie, Alexander Chowdhury, M De Paolis Kaluza, Linfeng Zhao, Lawson Wong, and Rose Yu. Deep imitation learning for bimanual robotic manipulation. *Advances in neural information processing systems*, 33:2327–2337, 2020.
- Zhenjia Xu, Zhanpeng He, Jiajun Wu, and Shuran Song. Learning 3d dynamic scene representations for robot manipulation. In *Conference on Robotic Learning (CoRL)*, 2020.
- Andy Zeng, Pete Florence, Jonathan Tompson, Stefan Welker, Jonathan Chien, Maria Attarian, Travis Armstrong, Ivan Krasin, Dan Duong, Vikas Sindhwani, and Johnny Lee. Transporter networks: Rearranging the visual world for robotic manipulation. *Conference on Robot Learning (CoRL)*, 2020.
- Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the continuity of rotation representations in neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5745–5753, 2019.