

D³Fields: Dynamic 3D Descriptor Fields for Zero-Shot Generalizable Robotic Manipulation

Yixuan Wang^{1*}, Zhuoran Li^{2,3*}, Mingtong Zhang¹, Katherine Driggs-Campbell¹,
Jiajun Wu², Li Fei-Fei², Yunzhu Li^{1,2}




Figure 1: **D³Fields Representation and Application to Various Manipulation Tasks.** D³Fields take in multi-view RGBD images and encode semantic features and instance masks using foundational models. The gray and colored points in the bottom left visualize background and semantic features mapped to RGB space using Principal Component Analysis (PCA), demonstrating consistency across instances. We use our representation for diverse tasks in a zero-shot manner. These tasks are defined by 2D goal images with diverse instances and styles. We address pick-and-place tasks such as shoe organization and tasks requiring dynamic modeling like collecting debris. We also demonstrate in the office table organization that our framework can accomplish 3D manipulation and compositional task specification.

Abstract

Scene representation has been a crucial design choice in robotic manipulation systems. An ideal representation should be 3D, dynamic, and semantic to meet the demands of diverse manipulation tasks. However, previous works often lack all three properties simultaneously. In this work, we introduce D³Fields — dynamic 3D descriptor fields. These fields capture the dynamics of the underlying 3D environment and encode both semantic features and instance masks. Specifically, we project arbitrary 3D points in the workspace onto multi-view 2D visual observations and interpolate features derived from foundational models. The resulting fused descriptor fields allow for flexible goal specifications using 2D images with varied contexts, styles, and instances. To evaluate the effectiveness of these descriptor fields, we apply our representation to a wide range of robotic manipulation tasks in a zero-shot manner. Through extensive evaluation in both real-world scenarios

*Denotes equal contribution. <https://robopil.github.io/d3fields/>

¹University of Illinois Urbana-Champaign ²Stanford University ³National University of Singapore

and simulations, we demonstrate that D^3 Fields are both generalizable and effective for zero-shot robotic manipulation tasks. In quantitative comparisons with state-of-the-art dense descriptors, such as Dense Object Nets and DINO, D^3 Fields exhibit significantly better generalization abilities and manipulation accuracy.

1 Introduction

The choice of scene representation is critical in robotic systems. An ideal representation should be simultaneously 3D, dynamic, and semantic to meet the needs of various robotic manipulation tasks in our daily lives. However, previous research into scene representations in robotics often does not encompass all three properties. Some representations exist in 3D space Manuelli et al. [2020], Li et al. [2021], Shi et al. [2022], Ze et al. [2023a], yet they overlook semantic information. Others focus on dynamic modeling Hafner et al. [2019b], Yan et al. [2021], Wang et al. [2023], Minderer et al. [2019], but only consider 2D data. Some other works are limited by only considering semantic information such as object instance and category Tremblay et al. [2018], Tyree et al. [2022], Jatavallabhula et al. [2023], Mazur et al. [2023], Liu et al. [2023b].

In this work, we aim to satisfy all three criteria by introducing D^3 Fields, unified descriptor fields that are 3D, dynamic, and semantic. D^3 Fields take in arbitrary points in the 3D world coordinate frame and output both geometric and semantic information related to these points. This includes the instance mask, dense semantic features, and the signed distance to the object surface. Notably, deriving these descriptor fields requires no training and is conducted in a zero-shot manner using large foundational vision models and vision-language models (VLMs). Specifically, we first use Grounding-DINO Liu et al. [2023a], Segment Anything (SAM) Kirillov et al. [2023], XMem Cheng and Schwing [2022], and DINOv2 Oquab et al. [2023] to extract information from multi-view 2D RGB images. We then project the 3D points back to each camera, interpolate to compute representations from each view, and fuse these data to derive the descriptors for the associated 3D points, as shown in Fig. 1 (left). By leveraging the dense semantic feature and instance mask of our representation, we can robustly track 3D points of the target object instance and train dynamics models. These learned dynamics models can then be incorporated into a Model-Predictive Control (MPC) framework to plan for manipulation tasks.

Notably, the derived representations allow for goal specification using 2D images sourced from the Internet, phones, or those generated by AI models. Such goal images have been challenging to manage with previous methods, because they contain varied styles, contexts, and object instances different from the robot’s workspace. Our proposed D^3 Fields can establish dense correspondences between the robot workspace and the target configurations. These correspondences give us the task objective, enabling us to plan the robot’s actions with the learned dynamics model within the MPC framework. This task execution process does not require any further training, offering a flexible and convenient interface for humans to instruct robots.

We evaluate our method across a wide range of household robotic manipulation tasks in a zero-shot manner. These tasks include organizing shoes, collecting debris, and organizing office desks, as shown in Fig. 1 (right). Furthermore, we offer detailed quantitative comparisons between our method and other state-of-the-art dense descriptor techniques. Our results indicate that our approach significantly outperforms in terms of generalizability and manipulation accuracy.

To summarize our contributions: (1) We introduce a novel representation, D^3 Fields, that is **3D**, **dynamic**, and **semantic**. (2) We present a novel and flexible goal specification method using 2D images that incorporate a range of styles, contexts, and instances. (3) Our proposed robotic manipulation framework supports zero-shot generalizable manipulation applicable to a broad spectrum of household tasks.

2 Related Works

2.1 Foundation Models for Robotics

Foundation models generally refer to those trained on broad data, often using self-supervision at scale, which can then be adapted (e.g., fine-tuned) to various downstream tasks. Large Language Models (LLMs) have showcased promising reasoning abilities for language. Robotics researchers have recently released a series of works that leverage LLMs, including SayCan Brohan et al. [2023] and Inner Monologue Huang et al. [2023b], to directly generate robot plans. Some later works have

used LLMs as a code generator: Code as Policies Liang et al. [2023] uses 2D object detectors as the perception API, whereas VoxPoser Huang et al. [2023a] creates a 3D value map. Yet, their perception modules fall short in modeling the precise geometry and dynamics of objects. Our D³Fields aim to address this by focusing on detailed 3D geometry and dynamics.

Meanwhile, foundational vision models, such as SAM Kirillov et al. [2023] and DINOv2 Oquab et al. [2023], have demonstrated impressive zero-shot generalization capabilities across various vision tasks. However, their focus is primarily on 2D vision tasks. Grounding these models in a dynamic 3D environment remains a challenge. The recent GROOT project showcases how to construct 3D object-centric representations using foundational models and exhibits notable few-shot generalization capabilities Zhu et al. [2023]. Still, GROOT does not emphasize learning about object dynamics or achieving zero-shot generalizable robotic manipulation.

2.2 Representation for Visual Robotic Manipulation

Scene representation has been a pivotal component in robotic manipulation systems. Some early work relies on 2D representations, such as bounding boxes Mülling et al. [2013], Duan et al. [2017]. Many recent methods construct particle representations of the environment and employ learned dynamics to capture the system’s underlying structure Li et al. [2019], Shi et al. [2022], Wang et al. [2023], Minderer et al. [2019], Wang et al. [2022], Gao and Tedrake [2021a], Manuelli et al. [2019], Gao and Tedrake [2021b]. They demonstrate impressive results in unstructured environments and with non-rigid objects. However, they are not semantic, which can hinder their ability to generalize to new tasks and scenarios. Some research opts for a fixed-dimension latent vector derived from high-dimensional sensory inputs as the representation Hafner et al. [2019a,b], Yan et al. [2021], Lin et al. [2022], Nair et al. [2023], Xiao et al. [2022], Radosavovic et al. [2022], Mandi et al. [2022], Stone et al. [2023], Li et al. [2021], but such a representation does not scale well to complex manipulation tasks that require high precision and explicit scene structures. Other approaches use 6 DoF object poses as their representation Tremblay et al. [2018], Tyree et al. [2022], Yoon et al. [2003], Zhu et al. [2014], though focusing primarily on grasping tasks instead of more dynamic ones. In this work, we aim to address these issues by introducing D³Fields, a representation that models dynamic 3D environments at varying semantic levels.

2.3 Neural Fields for Robotic Manipulation

Researchers have presented a variety of works using neural fields as a representation for robotic manipulation Zhu et al. [2021], Ichnowski et al. [2021], Wi et al. [2022a], Simeonov et al. [2022], Driess et al. [2022a], Jiang et al. [2021], Weng et al. [2023], Driess et al. [2022b], Shim et al. [2023], Yen-Chen et al. [2022], Tang et al. [2023], Zhou et al. [2023], [Mahi], Wi et al. [2022b,a]. Among them, Neural Descriptor Fields are the most relevant to ours Simeonov et al. [2022]. They build neural feature fields that generalize to different instances with several demonstrations; but they focus on learning geometric, not semantic features, which hinders cross-category generalization.

Recently, a series of works distilled neural feature fields using foundation models such as CLIP and DINO for supervision Radford et al. [2021], Caron et al. [2021]. LeRF distills neural feature fields to handle open-vocabulary 3D queries and develops task-oriented grasping based on it Kerr et al. [2023], Sharma et al. [2023]. Shen et al. [2023] use a similar distilled feature field for the grasping task. Both methods require dense camera views to train the neural field. GNFactor addresses this by introducing a voxel encoder Ze et al. [2023b]. However, distilling foundation models to create neural feature fields has drawbacks: (1) They often require dense camera views for a quality field. (2) Distilled neural fields need retraining for new scenes, limiting their generalization and making them ineffective for dynamic scenes. In contrast, our D³Fields do not need extra training for new scenes and can work with sparse views and dynamic settings.

3 Method

In this section, we introduce the problem formulation in Section 3.1 and define camera transformation and projection notations in Section 3.2. The construction of D³Fields is detailed in Section 3.3. Section 3.4 discusses tracking keypoints and learning dynamics, while Section 4.3 showcases how our representation enables zero-shot generalizable manipulation skills.

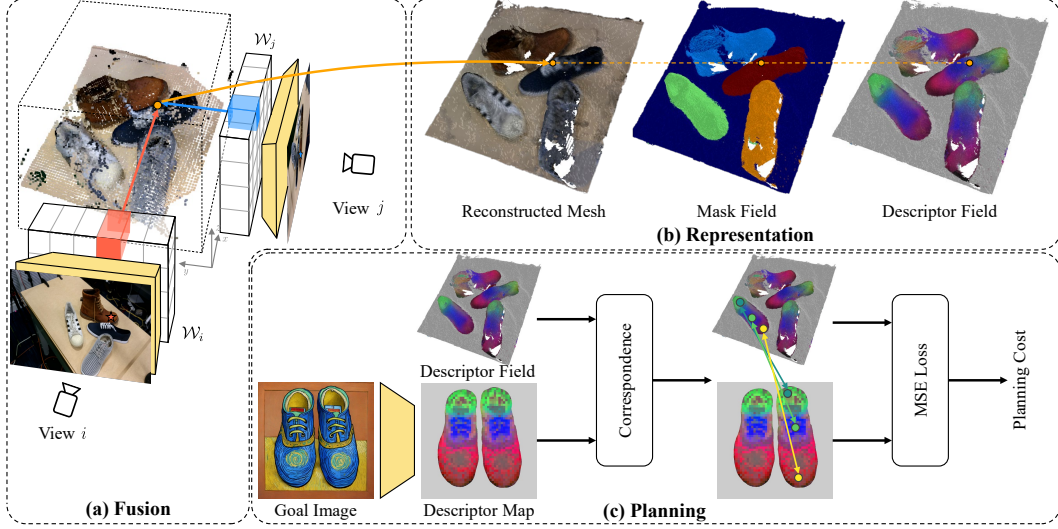


Figure 2: **Overview of the proposed framework.** (a) The fusion process fuses RGBD observations from multiple views. Each view is processed by foundation models to obtain the feature volume \mathcal{W} . Arbitrary 3D points are processed through projection and interpolation. (b) After fusing information from multiple views, we obtain an implicit distance function to reconstruct the mesh form. We also have instance masks and semantic features for evaluated 3D points, as shown by the mask field and descriptor field in the top right subfigure. (c) Given a 2D goal image, we use foundation models to extract the descriptor map. Then we correspond 3D features to 2D features and define the planning cost based on the correspondence.

3.1 Problem Formulation

Given a 2D goal image \mathcal{I} , we denote the corresponding scene representation as \mathbf{s}_{goal} . Our goal is to find the action sequence $\{a^t\}$ to minimize the task objective:

$$\begin{aligned} \min_{\{a_t\}} \quad & c(\mathbf{s}^T, \mathbf{s}_{\text{goal}}), \\ \text{s.t.} \quad & \mathbf{s}^t = g(\mathbf{o}^t), \quad \mathbf{s}^{t+1} = f(\mathbf{s}^t, a^t), \end{aligned} \quad (1)$$

where $c(\cdot, \cdot)$ is the cost function measuring the distance between the terminal representation \mathbf{s}^T and the goal representation \mathbf{s}_{goal} . Representation extraction function $g(\cdot)$ takes in the current multi-view RGBD observations \mathbf{o}^t and outputs the current representation \mathbf{s}^t . $f(\cdot, \cdot)$ is the dynamics function that predicts the future representation \mathbf{s}^{t+1} , conditioned on the current representation \mathbf{s}^t and action a^t . The optimization aims to find the action sequence $\{a_t\}$ that minimizes the cost function $c(\mathbf{s}^T, \mathbf{s}_{\text{goal}})$.

3.2 Notation: Camera Transformation and Projection

We assume all cameras' intrinsic parameters \mathbf{K} and extrinsic parameters \mathbf{T} are known. The camera i extrinsic parameters are defined as follows.

$$\mathbf{T}_i = \begin{bmatrix} \mathbf{R}_i & \mathbf{t}_i \\ \mathbf{0}^T & 1 \end{bmatrix} \in \mathbb{SE}(3), \quad (2)$$

where Euclidean group $\mathbb{SE}(3) := \{\mathbf{R}, \mathbf{t} \mid \mathbf{R} \in \mathbb{SO}^3, \mathbf{t} \in \mathbb{R}^3\}$. For a 3D point \mathbf{x} in the world frame, we could obtain projected pixel \mathbf{u}_i and distance to camera \mathbf{r}_i as follows:

$$\mathbf{u}_i = \pi(\mathbf{K}_i(\mathbf{R}_i\mathbf{x} + \mathbf{t}_i)), \quad \mathbf{r}_i = [0, 0, 1]^T (\mathbf{R}_i\mathbf{x} + \mathbf{t}_i), \quad (3)$$

where π performs perspective projection, mapping a 3D vector $p = [x, y, z]^T$ to a 2D vector $q = [x/z, y/z]^T$.

3.3 D³Fields Representation

We fuse observation \mathbf{o}^t from multiple views to build the implicit 3D descriptor fields $\mathcal{F}^t(\cdot)$. For simplicity, we will represent \mathbf{o}^t as \mathbf{o} , and $\mathcal{F}^t(\cdot)$ as $\mathcal{F}(\cdot)$ in this subsection. The implicit 3D descriptor field $\mathcal{F}(\cdot)$ is defined as

$$(\mathbf{d}, \mathbf{f}, \mathbf{p}) = \mathcal{F}(\mathbf{x}), \quad (4)$$

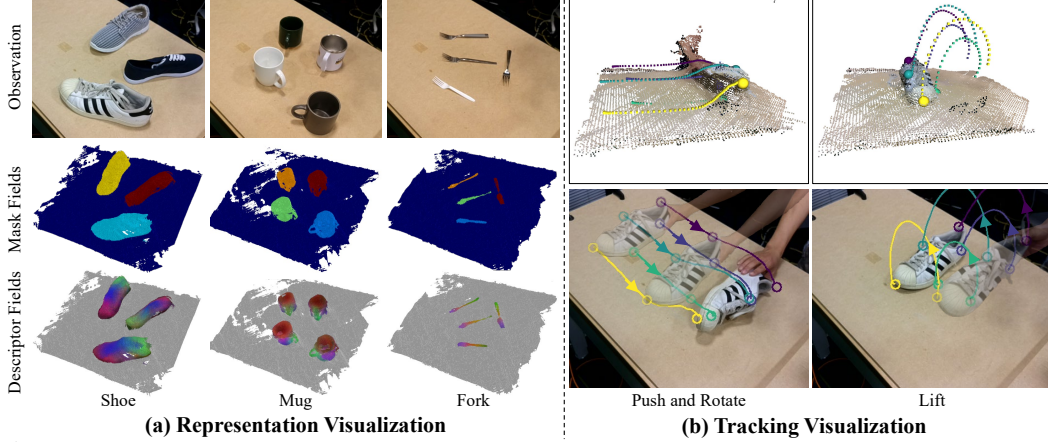


Figure 3: **Representation and Tracking Visualizations.** (a) To verify that the representation is both 3D and semantic, we visualize the representation across different object categories. Mask fields color 3D points based on their instance masks, which clearly differentiates between instances. Descriptor fields color 3D points by mapping features to RGB space using PCA. They display a consistent color pattern within a category, such as mug handles being colorized as green for different mug instances. (b) To demonstrate that our representation is dynamic, we apply it to tracking tasks and showcase two tracking examples, both of which involve 3D motions and partial observations in single views. The robust 3D tracking results confirm that our representation is 3D, dynamic, and semantic.

where \mathbf{x} is an arbitrary 3D point in the world frame, and $(\mathbf{d}, \mathbf{f}, \mathbf{p})$ is the corresponding geometric and semantic descriptor. $\mathbf{d} \in \mathbb{R}$ is the signed distance from \mathbf{x} to the surface. $\mathbf{f} \in \mathbb{R}^N$ represents the semantic information of N dimension. $\mathbf{p} \in \mathbb{R}^M$ denotes the instance probability distribution of M instances. M could be different across scenarios.

More specifically, we denote a single view RGBD observation from camera i as $\mathbf{o}_i = (\mathcal{I}_i, \mathcal{R}_i)$, where RGB image $\mathcal{I}_i \in \mathbb{R}^{H \times W \times 3}$, and depth image $\mathcal{R}_i \in \mathbb{R}^{H \times W}$. For an arbitrary 3D point \mathbf{x} , we project it to image space using Eq. 3 and use bilinear interpolation to obtain the corresponding depth $r'_i = \mathcal{R}_i[\mathbf{u}_i]$. Then the descriptors from camera i are

$$\begin{aligned} \mathbf{d}_i &= r'_i - r_i, & \mathbf{d}'_i &= \max(\min(\mathbf{d}_i, \mu), -\mu), \\ \mathbf{f}_i &= \mathcal{W}_i^f[\mathbf{u}_i], & \mathbf{p}_i &= \mathcal{W}_i^p[\mathbf{u}_i], \end{aligned} \quad (5)$$

where DINOv2 Oquab et al. [2023] extracts the semantic feature volume $\mathcal{W}_i^f \in \mathbb{R}^{H \times W \times N}$ from RGB observation \mathcal{I}_i . $\mathcal{W}_i^p \in \mathbb{R}^{H \times W \times M}$ is the instance mask volume using Grounded-SAM Liu et al. [2023a], Kirillov et al. [2023]. μ is the truncation threshold for TSDF.

We fuse descriptors from all K views as follows:

$$v_i = H(\mathbf{d}_i + \mu), \quad w_i = \exp\left(\frac{\min(\mu - |\mathbf{d}_i|, 0)}{\mu}\right), \quad (6)$$

and then

$$\mathbf{d} = \frac{\sum_{i=1}^K v_i \mathbf{d}'_i}{\delta + \sum_{i=1}^K v_i}, \quad \mathbf{f} = \frac{\sum_{i=1}^K v_i w_i \mathbf{f}_i}{\delta + \sum_{i=1}^K v_i}, \quad \mathbf{m} = \frac{\sum_{i=1}^K v_i w_i \mathbf{m}_i}{\delta + \sum_{i=1}^K v_i}, \quad (7)$$

where H is the unit step function and δ is a small value to avoid numeric errors. $v_i = 0$ when \mathbf{x} is not observable in camera i , because if \mathbf{x} is occluded in camera i , it should not contribute to the descriptor of \mathbf{x} . In addition, we could only have a confident estimation when \mathbf{x} is close to the surface. Therefore, w_i will decay as $|\mathbf{d}_i|$ increases. For \mathbf{x} that is far away, \mathbf{f} and \mathbf{m} will degrade to 0^T .

We convert the implicit field function $\mathcal{F}(\cdot)$ to a set of keypoints \mathbf{s} . First, we create voxels $\mathbf{x} \in \mathbb{R}^{W \times L \times H \times 3}$ in the workspace and evaluate $(\mathbf{d}, \mathbf{f}, \mathbf{p}) = \mathcal{F}(\mathbf{x})$. We filter out $\mathbf{x}_i \in \mathbf{x}$ where \mathbf{d}_i is large or \mathbf{p}_i has a low probability to avoid empty space and the background. After obtaining filtered points \mathbf{x}' , we use farthest point sampling to find surface points $\mathbf{s} \in \mathbb{R}^{3 \times n_s}$ of an instance.

3.4 Keypoints Tracking and Dynamics Training

This section will present how to use the dynamic implicit 3D descriptor field $\mathcal{F}(\cdot)$ to track keypoints and train dynamics. Without losing generalization, consider the tracking of a single instance $\mathbf{s}^t \in$

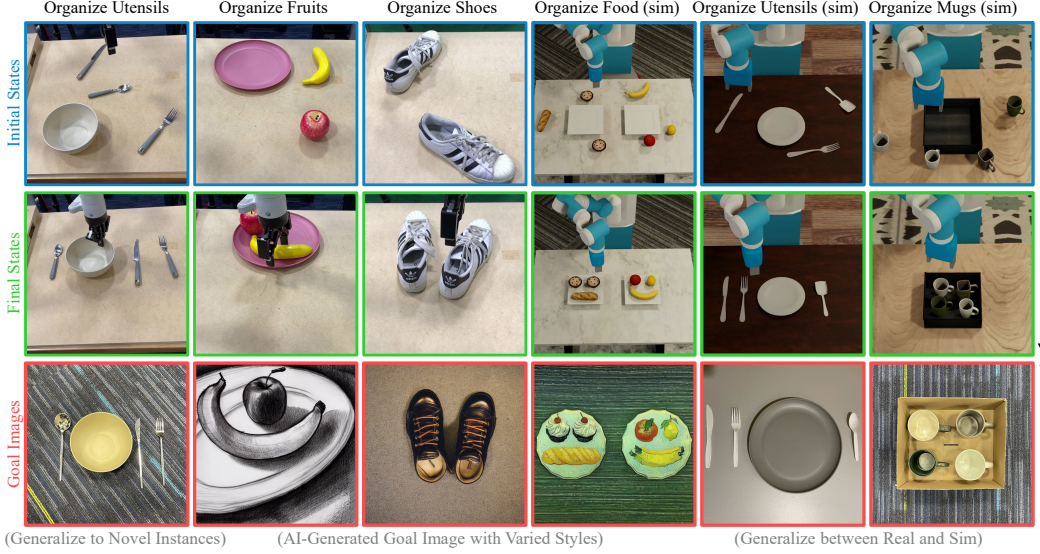


Figure 4: **Qualitative results.** We qualitatively evaluate our proposed framework on household manipulation tasks, both in the real world and in simulation, encompassing tasks such as organizing utensils, fruits, shoes, food, and mugs. The figure highlights that our representation can generalize across varied instances, styles, and contexts. For instance, in the organizing fruits example, the goal image, unlike the workspace, is styled as a sketch drawing. Because our representation can map bananas with varied styles and appearances to similar features, the banana in the workspace can correspond to the banana in the sketch. This allows the task to be successfully completed. This wide range of tasks showcases the generalization capabilities and manipulation precision of our framework.

$\mathbb{R}^{3 \times n_s}$. For clarity, we denote \mathbf{f} and \mathbf{d} from $\mathcal{F}(\cdot)$ as $\mathcal{F}_f(\cdot)$ and $\mathcal{F}_d(\cdot)$. We formulate the tracking problem as an optimization problem:

$$\min_{\mathbf{s}^{t+1}} \|\mathcal{F}_f(\mathbf{s}^{t+1}) - \mathcal{F}_f(\mathbf{s}^0)\|_2. \quad (8)$$

Since $\mathcal{F}(\cdot)$ is differentiable, we could use a gradient-based optimizer. This method could be naturally extended to multiple-instance scenarios. We found that relying solely on features for tracking is unstable. We added rigid constraints and distance regularization for a more stable tracking.

Keypoint tracking enables dynamics model training on real data. We instantiate the dynamics model $f(\cdot, \cdot)$ as graph neural networks (GNNs). We follow Li et al. [2020] to predict object dynamics. Please refer to Li et al. [2019, 2020] for more details on how to train the GNN-based dynamics model. The trained dynamics will be used for trajectory optimization in Section 3.5.

3.5 Zero-Shot Generalizable Robotic Manipulation

As described in Section 3.3, we denote initial tracked points and features as \mathbf{s}^0 and \mathbf{f}^0 . We estimate $\mathbf{s}_{\text{goal}} \in \mathbb{R}^{2 \times n_s}$ of goal image $\mathcal{I}_{\text{goal}}$ as follows:

$$\alpha_{ij} = \exp\left(\|\mathcal{W}_{\text{goal}}^f[\mathbf{u}_i] - \mathbf{f}_j^0\|_2\right),$$

$$w_{ij} = \frac{\exp(s\alpha_{ij})}{\sum_{i=1}^{H \times W} \exp(s\alpha_{ij})}, \quad (9)$$

then we have $\mathbf{s}_{\text{goal},j} = \sum_{i=1}^{H \times W} w_{ij} \mathbf{u}_i$, where $\mathcal{W}_{\text{goal}}^f$ is the feature volume extracted from $\mathcal{I}_{\text{goal}}$ using DINOv2. s is the hyperparameter to determine whether the heatmap w_{ij} is more smooth or concentrating. Although Eq. 9 only shows a single instance case, it could be naturally extended to multiple instances by using instance mask information.

However, \mathbf{s}_{goal} is in the image space, while \mathbf{s}^t is in the 3D space. We bridge this gap by introducing a reference camera with approximate intrinsic and extrinsic parameters \mathbf{K}' and \mathbf{T}' . Instead of rendering images in the reference view, we focus on projecting 3D keypoints into 2D images and define the task cost function in image space as follows:

$$c(\mathbf{s}^t, \mathbf{s}_{\text{goal}}) = \|\pi(\mathbf{K}'(\mathbf{R}'\mathbf{s}^t + \mathbf{t}')) - \mathbf{s}_{\text{goal}}\|_2^2. \quad (10)$$




Figure 5: **Quantitative Evaluation.** We perform real-world quantitative evaluations by measuring final goal-achieving performance and keypoints correspondence accuracy. (a) We use IoU to measure goal-achieving performance. Results indicate that our method aligns with the goal configurations much better than DON and DINO across various object categories and scenarios. (b) We measure the keypoints correspondence accuracy according to the fraction of points with accurate matches, with correct matches determined by a distance threshold. Our method is consistently better at aligning with the goal image, regardless of the chosen threshold.

4 Experiments

In this section, we evaluate our representation across various manipulation tasks with varying goal image styles, instances, and contexts. We visualize D³Fields and showcase tracking results in Section 4.2. Then, we highlight our framework’s zero-shot generalizability in both real-world and simulated tasks in Section 4.3. Finally, a quantitative comparison with baselines in Section 4.4 underscores our framework’s generalization and manipulation precision.

4.1 Experiment Setup

In the real world, we employ four OAK-PRO D cameras to gather RGBD observations and use the Kinova[®] Gen3 for action execution. In simulation, we utilize OmniGibson and deploy Fetch for mobile manipulation tasks Li et al. [2022]. Our evaluations span a variety of tasks, including organizing shoes, collecting debris, tidying the office table, arranging utensils, and more.

We implement the baseline methods using Dense Object Nets (DON) and DINO for feature extraction Florence et al. [2018], Caron et al. [2021]. We quantitatively evaluate these methods on five object classes for single-instance manipulation tasks in the real world. The results and analysis are presented in Section 4.4.

4.2 Descriptor Fields Visualization and Keypoints Tracking

D³Fields provide a good 3D semantic representation, as shown in Fig. 3(a). We first visualize the mask fields by coloring 3D points according to their most likely instance, and our visualization shows a clear 3D instance segmentation. Additionally, we map the semantic features to RGB space using PCA, as with DINOv2 Oquab et al. [2023]. Visualization of the descriptor fields reveals that D³Fields retain a dense semantic understanding of objects. In the provided shoe example, even though various shoes have distinct appearances and poses, they exhibit similar color patterns: shoe heels are represented in green, and shoe toes in red. We observed similar patterns when evaluating the model on mugs and forks.

As discussed before, D³Fields can also capture scene dynamics. We evaluate it by tracking the object keypoints. We show two examples of 3D keypoint tracking in Fig. 3(b). In the first example, a shoe is pushed and then flipped. Although only a portion of the shoe is visible from the view, our framework tracks it reliably. In another example, a shoe is lifted and then set down. Despite parts of the shoe being out of the camera’s view, we can robustly track it in 3D.

4.3 Zero-Shot Generalizable Manipulation

We conduct a qualitative evaluation of D³Fields in common household robotic manipulation tasks in a zero-shot manner, with partial results displayed in Fig. 1 and Fig. 4. The following capabilities of our framework are observed:

Generalization to AI-Generated Goal Images. In Fig. 1, the goal image, rendered in a Van Gogh style, depicts shoes distinct from those in the workspace. Since D³Fields encode semantic information, capturing shoes with varied appearances under similar descriptors, our framework can manipulate shoes based on AI-generated goal images.

Compositional Goal Images and 3D Manipulation. Using the office desk organization example in Fig. 1, the robot first arranges the mouse and pen according to the goal image. It then repositions the mug from the box to the mug pad, referencing a goal image of the upright mug.

Generalization across Instances and Materials. Granular objects, unlike rigid ones, have more complex dynamics. Our framework effectively handles these materials, as shown in the debris collection in Fig. 1. Fig. 4 further showcases our framework’s instance-level generalization, where the goal image displays instances different from the workspace.

Generalization across Simulation and Real World. We evaluated our framework on household tasks in the simulator, as shown in the utensil organization and mug organization examples in Fig. 4. Given goal images taken from the real world, our framework can also manipulate objects to the goal configurations. Our framework demonstrates generalization capabilities between simulation and the real world.

4.4 Quantitative Comparisons with Baselines

In Fig. 5(a), we measure performance using the IoU between the goal image mask and the final state mask after manipulation, with higher values indicating better alignment. Evaluating across five object classes, our method consistently outperforms the baselines, underscoring its generalization and manipulation accuracy. While DINO struggles with distinguishing object components, leading to imprecise results, it still works better than DON. Although DON performs well on familiar object classes and configurations, it lacks generalization in novel scenarios.

In Fig. 5(b), we present the correspondence results. We manually label corresponding keypoints on both the goal image and the final manipulation result to evaluate the correspondence accuracy. We calculate the fraction of accurately matched points based on a distance threshold. Our method consistently outperforms the baselines, regardless of the threshold. DINO ranks second, while DON lags behind. Consistent with Fig. 5(a), our method excels in generalization and accuracy, DINO is broadly applicable but less precise, and DON struggles with generalization.

5 Conclusion

In this work, we introduce D³Fields, which implicitly encode 3D semantic features and 3D instance masks, and model the underlying dynamics. Our emphasis is on zero-shot generalizable robotic manipulation tasks specified by 2D goal images of varying styles, contexts, and instances. Our framework excels in executing a diverse array of household manipulation tasks in both simulated and real-world scenarios. Its performance greatly surpasses baseline methods such as Dense Object Nets and DINO in terms of generalization capabilities and manipulation accuracy.

References

- Anthony Brohan, Yevgen Chebotar, Chelsea Finn, Karol Hausman, Alexander Herzog, Daniel Ho, Julian Ibarz, Alex Irpan, Eric Jang, Ryan Julian, et al. Do as i can, not as i say: Grounding language in robotic affordances. In *Conference on Robot Learning*, pages 287–318. PMLR, 2023.
- Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9650–9660, 2021.
- Ho Kei Cheng and Alexander G. Schwing. XMem: Long-term video object segmentation with an atkinson-shiffrin memory model. In *ECCV*, 2022.
- Danny Driess, Jung-Su Ha, Marc Toussaint, and Russ Tedrake. Learning models as functionals of signed-distance fields for manipulation planning. In *Conference on Robot Learning*, pages 245–255. PMLR, 2022a.

- Danny Driess, Ingmar Schubert, Pete Florence, Yunzhu Li, and Marc Toussaint. Reinforcement learning with neural radiance fields. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022b.
- Yan Duan, Marcin Andrychowicz, Bradly Stadie, OpenAI Jonathan Ho, Jonas Schneider, Ilya Sutskever, Pieter Abbeel, and Wojciech Zaremba. One-shot imitation learning. *Advances in neural information processing systems*, 30, 2017.
- Peter R. Florence, Lucas Manuelli, and Russ Tedrake. Dense object nets: Learning dense visual object descriptors by and for robotic manipulation. In Aude Billard, Anca Dragan, Jan Peters, and Jun Morimoto, editors, *Proceedings of The 2nd Conference on Robot Learning*, volume 87 of *Proceedings of Machine Learning Research*, pages 373–385. PMLR, 29–31 Oct 2018.
- Wei Gao and Russ Tedrake. kpm 2.0: Feedback control for category-level robotic manipulation. *IEEE Robotics and Automation Letters*, 6(2):2962–2969, 2021a.
- Wei Gao and Russ Tedrake. kpm-sc: Generalizable manipulation planning using keypoint affordance and shape completion. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6527–6533. IEEE, 2021b.
- Danijar Hafner, Timothy Lillicrap, Jimmy Ba, and Mohammad Norouzi. Dream to control: Learning behaviors by latent imagination. *arXiv preprint arXiv:1912.01603*, 2019a.
- Danijar Hafner, Timothy Lillicrap, Ian Fischer, Ruben Villegas, David Ha, Honglak Lee, and James Davidson. Learning latent dynamics for planning from pixels. In *International conference on machine learning*, pages 2555–2565. PMLR, 2019b.
- Wenlong Huang, Chen Wang, Ruohan Zhang, Yunzhu Li, Jiajun Wu, and Li Fei-Fei. Voxposer: Composable 3d value maps for robotic manipulation with language models. In *7th Annual Conference on Robot Learning*, 2023a.
- Wenlong Huang, Fei Xia, Ted Xiao, Harris Chan, Jacky Liang, Pete Florence, Andy Zeng, Jonathan Tompson, Igor Mordatch, Yevgen Chebotar, et al. Inner monologue: Embodied reasoning through planning with language models. In *Conference on Robot Learning*, pages 1769–1782. PMLR, 2023b.
- Jeffrey Ichnowski, Yahav Avigal, Justin Kerr, and Ken Goldberg. Dex-nerf: Using a neural radiance field to grasp transparent objects. In *5th Annual Conference on Robot Learning*, 2021.
- Krishna Murthy Jatavallabhula, Alihusein Kuwajerwala, Qiao Gu, Mohd Omama, Ganesh Iyer, Soroush Saryazdi, Tao Chen, Alaa Maalouf, Shuang Li, Nikhil Varma Keetha, Ayush Tewari, Joshua Tenenbaum, Celso de Melo, Madhava Krishna, Liam Paull, Florian Shkurti, and Antonio Torralba. ConceptFusion: Open-set multimodal 3D mapping. In *Proceedings of Robotics: Science and Systems*, Daegu, Republic of Korea, July 2023. doi: 10.15607/RSS.2023.XIX.066.
- Zhenyu Jiang, Yifeng Zhu, Maxwell Svetlik, Kuan Fang, and Yuke Zhu. Synergies Between Affordance and Geometry: 6-DoF Grasp Detection via Implicit Representations. In *Proceedings of Robotics: Science and Systems*, Virtual, July 2021. doi: 10.15607/RSS.2021.XVII.024.
- Justin Kerr, Chung Min Kim, Ken Goldberg, Angjoo Kanazawa, and Matthew Tancik. Lrf: Language embedded radiance fields. In *International Conference on Computer Vision (ICCV)*, 2023.
- Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. Segment anything. *arXiv:2304.02643*, 2023.
- Chengshu Li, Ruohan Zhang, Josiah Wong, Cem Gokmen, Sanjana Srivastava, Roberto Martín-Martín, Chen Wang, Gabriel Levine, Michael Lingelbach, Jiankai Sun, Mona Anvari, Minjune Hwang, Manasi Sharma, Arman Aydin, Dhruva Bansal, Samuel Hunter, Kyu-Young Kim, Alan Lou, Caleb R Matthews, Ivan Villa-Renteria, Jerry Huayang Tang, Claire Tang, Fei Xia, Silvio Savarese, Hyowon Gweon, Karen Liu, Jiajun Wu, and Li Fei-Fei. BEHAVIOR-1k: A benchmark for embodied AI with 1,000 everyday activities and realistic simulation. In *6th Annual Conference on Robot Learning*, 2022.

- Yunzhu Li, Jiajun Wu, Russ Tedrake, Joshua B Tenenbaum, and Antonio Torralba. Learning particle dynamics for manipulating rigid bodies, deformable objects, and fluids. In *ICLR*, 2019.
- Yunzhu Li, Toru Lin, Kexin Yi, Daniel Bear, Daniel L.K. Yamins, Jiajun Wu, Joshua B. Tenenbaum, and Antonio Torralba. Visual grounding of learned physical models. In *International Conference on Machine Learning*, 2020.
- Yunzhu Li, Shuang Li, Vincent Sitzmann, Pulkit Agrawal, and Antonio Torralba. 3d neural scene representations for visuomotor control. *arXiv preprint arXiv:2107.04004*, 2021.
- Jacky Liang, Wenlong Huang, Fei Xia, Peng Xu, Karol Hausman, Brian Ichter, Pete Florence, and Andy Zeng. Code as policies: Language model programs for embodied control. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9493–9500. IEEE, 2023.
- Xingyu Lin, Carl Qi, Yunchu Zhang, Yunzhu Li, Zhiao Huang, Katerina Fragkiadaki, Chuang Gan, and David Held. Planning with spatial-temporal abstraction from point clouds for deformable object manipulation. In *6th Annual Conference on Robot Learning*, 2022.
- Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. *arXiv preprint arXiv:2303.05499*, 2023a.
- Weiyu Liu, Yilun Du, Tucker Hermans, Sonia Chernova, and Chris Paxton. Structdiffusion: Language-guided creation of physically-valid structures using unseen objects. In *RSS 2023*, 2023b.
- Nur Muhammad (Mahi)Shafiullah, Chris Paxton, Lerrel Pinto, Soumith Chintala, and Arthur Szlam. CLIP-Fields: Weakly Supervised Semantic Fields for Robotic Memory. In *Proceedings of Robotics: Science and Systems*, Daegu, Republic of Korea, July 2023. doi: 10.15607/RSS.2023.XIX.074.
- Zhao Mandi, Homanga Bharadhwaj, Vincent Moens, Shuran Song, Aravind Rajeswaran, and Vikash Kumar. Cacti: A framework for scalable multi-task multi-scene visual imitation learning. *arXiv preprint arXiv:2212.05711*, 2022.
- Lucas Manuelli, Wei Gao, Peter Florence, and Russ Tedrake. kpm: Keypoint affordances for category-level robotic manipulation. In *The International Symposium of Robotics Research*, pages 132–157. Springer, 2019.
- Lucas Manuelli, Yunzhu Li, Pete Florence, and Russ Tedrake. Keypoints into the future: Self-supervised correspondence in model-based reinforcement learning. In *Conference on Robot Learning (CoRL)*, 2020.
- Kirill Mazur, Edgar Sucar, and Andrew J Davison. Feature-realistic neural fusion for real-time, open set scene understanding. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 8201–8207. IEEE, 2023.
- Matthias Minderer, Chen Sun, Ruben Villegas, Forrester Cole, Kevin P Murphy, and Honglak Lee. Unsupervised learning of object structure and dynamics from videos. *Advances in Neural Information Processing Systems*, 32, 2019.
- Katharina Mülling, Jens Kober, Oliver Kroemer, and Jan Peters. Learning to select and generalize striking movements in robot table tennis. *The International Journal of Robotics Research*, 32(3): 263–279, 2013.
- Suraj Nair, Aravind Rajeswaran, Vikash Kumar, Chelsea Finn, and Abhinav Gupta. R3m: A universal visual representation for robot manipulation. In *Conference on Robot Learning*, pages 892–909. PMLR, 2023.
- Maxime Oquab, Timothée Darcet, Theo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Russell Howes, Po-Yao Huang, Hu Xu, Vasu Sharma, Shang-Wen Li, Wojciech Galuba, Mike Rabbat, Mido Assran, Nicolas Ballas, Gabriel Synnaeve, Ishan Misra, Herve Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. Dinov2: Learning robust visual features without supervision, 2023.

- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR, 18–24 Jul 2021.
- Ilija Radosavovic, Tete Xiao, Stephen James, Pieter Abbeel, Jitendra Malik, and Trevor Darrell. Real-world robot learning with masked visual pre-training. *CoRL*, 2022.
- Satvik Sharma, Adam Rashid, Chung Min Kim, Justin Kerr, Lawrence Yunliang Chen, Angjoo Kanazawa, and Ken Goldberg. Language embedded radiance fields for zero-shot task-oriented grasping. In *7th Annual Conference on Robot Learning*, 2023.
- William Shen, Ge Yang, Alan Yu, Jensen Wong, Leslie Pack Kaelbling, and Phillip Isola. Distilled feature fields enable few-shot manipulation. In *7th Annual Conference on Robot Learning*, 2023.
- Haochen Shi, Huazhe Xu, Zhiao Huang, Yunzhu Li, and Jiajun Wu. Robocraft: Learning to see, simulate, and shape elasto-plastic objects with graph networks. *arXiv preprint arXiv:2205.02909*, 2022.
- Dongseok Shim, Seungjae Lee, and H. Jin Kim. SNeRL: Semantic-aware neural radiance fields for reinforcement learning. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett, editors, *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 31489–31503. PMLR, 23–29 Jul 2023.
- Anthony Simeonov, Yilun Du, Andrea Tagliasacchi, Joshua B Tenenbaum, Alberto Rodriguez, Pulkit Agrawal, and Vincent Sitzmann. Neural descriptor fields: Se (3)-equivariant object representations for manipulation. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 6394–6400. IEEE, 2022.
- Austin Stone, Ted Xiao, Yao Lu, Keerthana Gopalakrishnan, Kuang-Huei Lee, Quan Vuong, Paul Wohlhart, Brianna Zitkovich, Fei Xia, Chelsea Finn, and Karol Hausman. Open-world object manipulation using pre-trained vision-language models, 2023.
- Zhenggang Tang, Balakumar Sundaralingam, Jonathan Tremblay, Bowen Wen, Ye Yuan, Stephen Tyree, Charles Loop, Alexander Schwing, and Stan Birchfield. RGB-only reconstruction of tabletop scenes for collision-free manipulator control. In *ICRA*, 2023.
- Jonathan Tremblay, Thang To, Balakumar Sundaralingam, Yu Xiang, Dieter Fox, and Stan Birchfield. Deep object pose estimation for semantic robotic grasping of household objects. In *Conference on Robot Learning*, pages 306–316. PMLR, 2018.
- Stephen Tyree, Jonathan Tremblay, Thang To, Jia Cheng, Terry Mosier, Jeffrey Smith, and Stan Birchfield. 6-dof pose estimation of household objects for robotic manipulation: An accessible dataset and benchmark. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 13081–13088. IEEE, 2022.
- Weiyao Wang, Andrew S Morgan, Aaron M Dollar, and Gregory D Hager. Dynamical scene representation and control with keypoint-conditioned neural radiance field. In *2022 IEEE 18th International Conference on Automation Science and Engineering (CASE)*, pages 1138–1143. IEEE, 2022.
- Yixuan Wang, Yunzhu Li, Katherine Driggs-Campbell, Li Fei-Fei, and Jiajun Wu. Dynamic-Resolution Model Learning for Object Pile Manipulation. In *Proceedings of Robotics: Science and Systems*, Daegu, Republic of Korea, July 2023. doi: 10.15607/RSS.2023.XIX.047.
- Thomas Weng, David Held, Franziska Meier, and Mustafa Mukadam. Neural grasp distance fields for robot manipulation. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2023.
- Youngsun Wi, Pete Florence, Andy Zeng, and Nima Fazeli. VirDo: Visio-tactile implicit representations of deformable objects. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 3583–3590. IEEE, 2022a.

- Youngsun Wi, Andy Zeng, Pete Florence, and Nima Fazeli. VirDo++: Real-world, visuo-tactile dynamics and perception of deformable objects. *arXiv preprint arXiv:2210.03701*, 2022b.
- Tete Xiao, Ilija Radosavovic, Trevor Darrell, and Jitendra Malik. Masked visual pre-training for motor control. *arXiv:2203.06173*, 2022.
- Wilson Yan, Ashwin Vangipuram, Pieter Abbeel, and Lerrel Pinto. Learning predictive representations for deformable objects using contrastive estimation. In Jens Kober, Fabio Ramos, and Claire Tomlin, editors, *Proceedings of the 2020 Conference on Robot Learning*, volume 155 of *Proceedings of Machine Learning Research*, pages 564–574. PMLR, 16–18 Nov 2021.
- Lin Yen-Chen, Pete Florence, Jonathan T. Barron, Tsung-Yi Lin, Alberto Rodriguez, and Phillip Isola. NeRF-Supervision: Learning dense object descriptors from neural radiance fields. In *IEEE Conference on Robotics and Automation (ICRA)*, 2022.
- Youngrock Yoon, Guilherme N DeSouza, and Avinash C Kak. Real-time tracking and pose estimation for industrial objects using geometric features. In *2003 IEEE International conference on robotics and automation (cat. no. 03CH37422)*, volume 3, pages 3473–3478. IEEE, 2003.
- Yanjie Ze, Nicklas Hansen, Yinbo Chen, Mohit Jain, and Xiaolong Wang. Visual reinforcement learning with self-supervised 3d representations. *IEEE Robotics and Automation Letters (RA-L)*, 2023a.
- Yanjie Ze, Ge Yan, Yueh-Hua Wu, Annabella Macaluso, Yuying Ge, Jianglong Ye, Nicklas Hansen, Li Erran Li, and Xiaolong Wang. Multi-task real robot learning with generalizable neural feature fields. In *7th Annual Conference on Robot Learning*, 2023b.
- Allan Zhou, Moo Jin Kim, Lirui Wang, Pete Florence, and Chelsea Finn. Nerf in the palm of your hand: Corrective augmentation for robotics via novel-view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17907–17917, 2023.
- Luyang Zhu, Arsalan Mousavian, Yu Xiang, Hammad Mazhar, Jozef van Eenbergen, Shoubhik Debnath, and Dieter Fox. Rgb-d local implicit function for depth completion of transparent objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4649–4658, 2021.
- Menglong Zhu, Konstantinos G Derpanis, Yinfei Yang, Samarth Brahmabhatt, Mabel Zhang, Cody Phillips, Matthieu Lecce, and Kostas Daniilidis. Single image 3d object detection and pose estimation for grasping. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3936–3943. IEEE, 2014.
- Yifeng Zhu, Zhenyu Jiang, Peter Stone, and Yuke Zhu. Learning generalizable manipulation policies with object-centric 3d representations. In *7th Annual Conference on Robot Learning*, 2023.