

ZeroGrasp: Zero-Shot Shape Reconstruction Enabled Robotic Grasping

Shun Iwase^{1,2} Muhammad Zubair Irshad² Katherine Liu² Vitor Guizilini²
 Robert Lee³ Takuya Ikeda³ Ayako Amma³ Koichi Nishiwaki³ Kris Kitani¹
 Rareş Ambruş² Sergey Zakharov²
¹Carnegie Mellon University ²Toyota Research Institute ³Woven by Toyota

Abstract

Robotic grasping is a cornerstone capability of embodied systems. Many methods directly output grasps from partial information without modeling the geometry of the scene, leading to suboptimal motion and even collisions. To address these issues, we introduce ZeroGrasp, a novel framework that simultaneously performs 3D reconstruction and grasp pose prediction in near real-time. A key insight of our method is that occlusion reasoning and modeling the spatial relationships between objects is beneficial for both accurate reconstruction and grasping. We couple our method with a novel large-scale synthetic dataset, which comprises 1M photo-realistic images, high-resolution 3D reconstructions and 11.3B physically-valid grasp pose annotations for 12K objects from the Objaverse-LVIS dataset. We evaluate ZeroGrasp on the GraspNet-1B benchmark as well as through real-world robot experiments. ZeroGrasp achieves state-of-the-art performance and generalizes to novel real-world objects by leveraging synthetic data. <https://sh8.io/#/zerograsp>

1. Introduction

Safe and robust robotic grasping requires accurate geometric understanding of target objects, as well as their surroundings. However, most previous grasp detection methods [1–6] do not explicitly model the geometry of the target objects, which can lead to unexpected collisions and unstable contact with target objects. Although several methods [3, 7] leverage multi-view images to reconstruct the target objects in advance, this process introduces additional computational overhead and requires a more complex setup. Multi-view reconstruction is also often impractical for objects placed within confined spaces like shelves or boxes. Furthermore, the lack of large-scale datasets with ground-truth 3D shapes and grasp pose annotations further complicates accurate 3D reconstruction from a single RGB-D im-

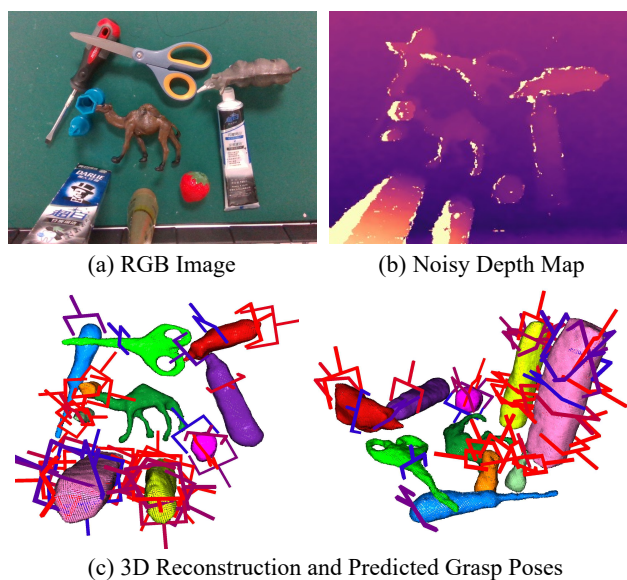


Figure 1. ZeroGrasp simultaneously reconstructs objects at high-resolution and predicts grasp poses from a single RGB-D image in near real-time (5FPS).

age. Recently, several works [8–10] demonstrate that sparse voxel representations outperform volumetric and NeRF-like implicit shape representations in terms of runtime, accuracy, and resolution, particularly for regression-based zero-shot 3D reconstruction.

To leverage reconstruction methods using sparse voxel representations for robotic grasping, it is essential to develop new approaches that can reason about both within a unified framework. To this end, we propose ZeroGrasp, a novel framework for near real-time 3D reconstruction and 6D grasp pose prediction. Our key hypothesis is that improved 3D reconstruction quality enhances grasp pose prediction, specifically by leveraging physics-based contact constraints and collision detection, which are essential for accurate grasping. Since robotic environments often involve

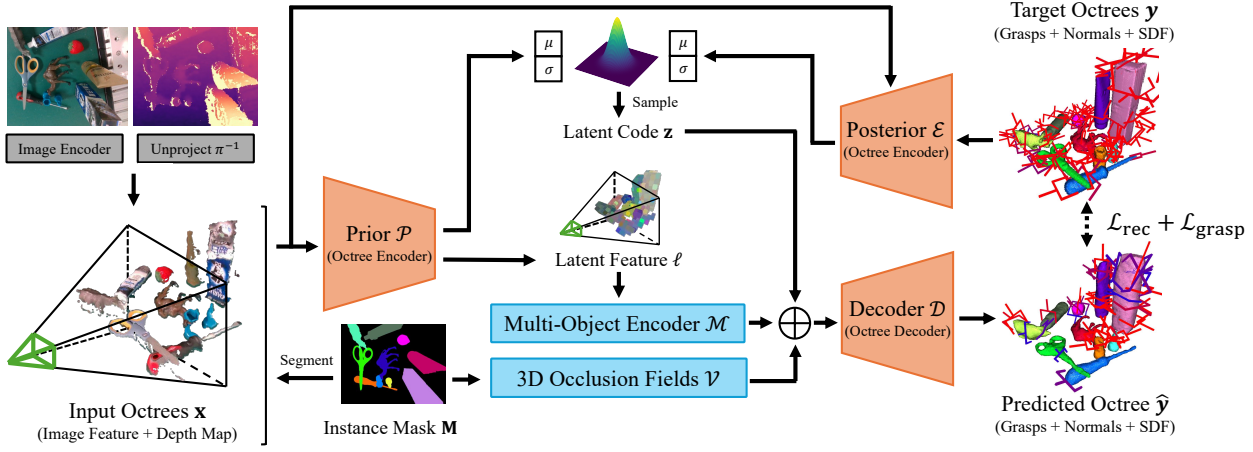


Figure 2. Overview of ZeroGrasp, a novel method for simultaneous 3D reconstruction and 6D grasp pose predictions from a single-view RGB-D image. The input octree \mathbf{x} is first fed into the octree-based CVAE (components with orange boxes). The multi-object encoder takes its latent feature ℓ to learn multi-object reasoning at the latent space. Further, 3D occlusion fields encode inter- and self-occlusion information via simple ray casting. The output features from the multi-object encoder and 3D occlusion fields are concatenated with the latent code \mathbf{z} , and 3D shapes and grasp poses are predicted by the decoder.

multiple objects with inter-object occlusions and close contacts, ZeroGrasp introduces two key components: a multi-object encoder and 3D occlusion fields. These components effectively model inter-object relationships and occlusions, thus, improving reconstruction quality. In addition, we design a simple refinement algorithm to improve grasp poses using the predicted reconstruction. Because the reconstruction is highly accurate, it provides reliable contact points and collision masks between the gripper and the target object, which we use to refine the grasp poses.

In addition to our proposed model, we also create a real-world dataset for evaluation, the ReOcS dataset, and a synthetic datasets for training, the ZeroGrasp-11B dataset. The ReOcS dataset is a real-world evaluation dataset of 3D reconstruction, with three splits representing different degrees of occlusion. We use this dataset to assess robustness to occlusions. The ZeroGrasp-11B dataset is a large-scale synthetic dataset designed to train models with zero-shot robotic grasping capability, containing high-quality and diverse 3D models from Objaverse-LVIS dataset [11], as shown in Table 1.

We evaluate both the baseline and our methods, showing that our approach — trained on the GraspNet-1B dataset [1] alone, as well as on a combination of the GraspNet-1B dataset and ZeroGrasp-11B — achieves state-of-the-art performance on the GraspNet-1B benchmark. Our ablation studies further show that the proposed components enhance both reconstruction and grasp pose prediction quality. Finally, we conduct real-robot evaluations to demonstrate its generalizability in real-world scenarios.

Our contributions are summarized as follows:

- We propose ZeroGrasp, a novel framework for simultaneous 3D reconstruction and 6D grasp pose prediction using an octree-based conditional variational autoencoder

(CVAE). ZeroGrasp achieves the best performance on the GraspNet-1B benchmark and real-robot evaluation.

- We introduce a multi-object encoder and 3D occlusion fields to model inter-object relationships and occlusions.
- We propose a simple grasp pose refinement algorithm that improves grasp accuracy using the reconstructed 3D shape.
- We create two datasets, 1) the ReOcS dataset for evaluating 3D reconstruction under occlusions, and 2) the ZeroGrasp-11B dataset for zero-shot robotic grasping.

2. Related Works

Regression-based 3D reconstruction. Regression-based 3D reconstruction from a single-view RGB-D image [8, 20–47] have been a major focus of research in 3D computer vision. These methods explore different 3D representations, including dense voxel grids [23, 31, 39, 48], sparse voxel grids [8, 9, 49] (e.g. octree [9], VDB [49], hash table [8], and etc.), and implicit representations [20, 33, 34, 38]. Nevertheless, dense voxel grid and implicit representations face limitations in output resolution due to expensive memory and computational costs. On the other hand, several works [9, 20, 21, 49] show that sparse voxel representations such as an octree and VDB [50] enable high-resolution 3D reconstruction with faster runtime thanks to its efficient hierarchical structure. Alternatively, single-view reconstruction through novel view synthesis achieves impressive results. Recent works such as GeNVS [51], Zero-1-to-3 [52], 3DiM [53], and InstantMesh [54] leverage diffusion models to render multi-view images given a canonical camera pose. However, these approaches are slow (often over 10 seconds) and inter-object occlusions degrade the performance significantly. Further, integrating grasp pose prediction is non-trivial. Thus, we adopt an octree as a shape representation

Table 1. **Dataset comparisons.** We create a large-scale grasp detection dataset for zero-shot robotic grasping using 12K 3D models from Objaverse-LVIS dataset [11]. Our ZeroGrasp-11B dataset includes 1 million RGB-D images and physics-based dense 6D grasp annotations of cluttered scenes.

Dataset	# Images	# 3D Models	# Grasps	# Cat.	Type	Modality	Grasp Alg.	Grasp Rep.
Cornell [2]	1K	0.2K	8K	16	Real	RGB-D	Manual	Planar
Jacquard [12]	54K	11K	1.1M	N/A	Sim.	RGB-D	Physics	Planar
Zhang <i>et al.</i> [13]	4.7K	$\approx 15\text{K}$	100K	N/A	Real	RGB	Manual	Planar
VR-Grasping-101 [14]	10K	0.1K	4.8M	7	Sim.	RGB-D	Manual	6D
GraspNet-1Billion [1]	97K	0.1K	<u>1.2B</u>	30-35	Real	RGB-D	Analytical	6D
ACRONYM [15]	N/A	9K	17.7M	<u>262</u>	Sim.	N/A	Physics	6D
REGRAD [16]	<u>900K</u>	50K	100M	55	Sim.	N/A	Physics	6D
HouseCat6D [17]	23.5K	0.2K	10M	10	Real	RGB-D+P	Physics	6D
Grasp-Anything-6D [18]	1M	N/A	200M	N/A	Synth.	RGB + ZoeDepth [19]	Analytical	6D
ZeroGrasp-11B (Ours)	1M	<u>12K</u>	11.3B	606	Sim.	RGB-D	Physics	6D

and design our framework based on octree-based U-Net [9].

Regression-based Grasp Pose Prediction. Traditional grasp pose prediction methods typically assume prior knowledge of 3D objects and often rely on simplified analytical models based on force closure principles [55, 56]. Recently, tremendous progress has been made in learning-based approaches [1, 6, 57, 58] which have allowed models to predict 6D grasp poses directly from RGB(-D) images and point clouds. This has enabled the regression of grasp poses in highly cluttered scenes without explicitly modeling object geometries. However, this could result in unstable grasping and unexpected collision, as accurately learning collision avoidance and precise contact points remains challenging. Although some methods [42, 59, 60] explore 3D reconstruction to improve grasp prediction, their choices of shape representations and network architectures often limit its full potential.

Zero-shot robotic grasping. Zero-shot robotic grasping refers to the ability to grasp unseen target objects without prior knowledge. To achieve this, there are mainly two directions — (1) optimizing grasp poses at test time based on contact points using reconstructed or ground-truth 3D shapes [3, 61], and (2) augmenting or synthesizing large-scale grasp data to improve generalization [1, 15, 62]. For instance, Ma *et al.* [3] propose a contact-based optimization algorithm to refine initial grasp poses by using a reconstructed 3D scene from multi-view RGB-D images. Existing large-scale grasp pose datasets such as ACRONYM [15], GraspNet-1B [1], and EGAD [62] explore the second approach. Nevertheless, they are limited to object diversity or missing annotations like RGB-D images. Inspired by these two approaches, we aim to improve generalization to unseen objects with a simple and efficient grasp pose refinement algorithm that utilizes predicted reconstructions. Further, we create a large-scale synthetic dataset for grasp pose detection. Our dataset comprises

high-quality and diverse objects, as well as 1M photorealistic RGB images and physics-based grasp pose annotations.

3. Proposed Method

Our goal is to build an efficient and generalizable model for simultaneous 3D shape reconstruction and grasp pose prediction from a single RGB-D observation, and to demonstrate that the predicted reconstructions can be used to refine grasp poses via contact-based constraints and collision detection. In this section, we describe the network architecture and grasp pose refinement algorithm.

3D shape representation. We adopt an octree as a shape representation where attributes such as image features, the signed distance function (SDF), normals, and grasp poses are defined at the deepest level of the octree. For instance, we represent an input octree as a tuple of voxel centers \mathbf{p} at the final depth, associated with image features \mathbf{f} ,

$$\mathbf{x} = (\mathbf{p}, \mathbf{f}), \quad \mathbf{p} \in \mathbb{R}^{N \times 3}, \mathbf{f} \in \mathbb{R}^{N \times D}. \quad (1)$$

where N is the number of voxels. Unlike point clouds, an octree structure enables efficient depth-first search and recursive subdivision to octants, making it ideal for high-resolution shape reconstruction and dense grasp pose prediction in a memory and computationally efficient manner.

Grasp pose representation. We represent grasp poses using a general two-finger parallel gripper model, as used in GraspNet [1]. Specifically, our grasp poses consist of the following components: view graspness score $\mathbf{s} \in \mathbb{R}^M$, which indicates the robustness of grasp positions [57]; quality $\mathbf{q} \in \mathbb{R}^M$, computed using the force closure algorithm [55]; view direction $\mathbf{v} \in \mathbb{R}^{3M}$; angle $\mathbf{a} \in \mathbb{R}^M$; width $\mathbf{w} \in \mathbb{R}^M$; and depth $\mathbf{d} \in \mathbb{R}^M$:

$$\mathbf{g} = [\mathbf{s} \quad \mathbf{q} \quad \mathbf{v} \quad \mathbf{a} \quad \mathbf{w} \quad \mathbf{d}], \quad (2)$$

where M denotes the number of total grasps in the target octree, and the closest grasp poses within a 5 mm radius




Figure 3. **3D occlusion fields** localize occlusion information by casting rays from the camera to the voxel centers around the target object and performing depth tests. Specifically, if a ray intersects the target object, a self-occlusion flag o_{self} is set to 1. If it intersects non-target objects, an inter-object occlusion flag o_{inter} is set to 1.

are assigned to each point. If none exists, we set its corresponding graspness to 0. In GraspNet-1B and ZeroGrasp-11B datasets, each point is annotated with a dense set of grasp poses covering all combinations of views, angles, and depths ($300 \times 12 \times 4$). With the grasp poses \mathbf{g} , the target octree is defined as

$$\mathbf{y} = (\mathbf{p}^{gt}, \mathbf{f}^{gt}) = (\mathbf{p}^{gt}, [\phi \quad \mathbf{n} \quad \mathbf{g}]), \quad (3)$$

where $\phi \in \mathbb{R}^M$ is the SDF, and $\mathbf{n} \in \mathbb{R}^{M \times 3}$ is normal vectors of the target octree.

3.1. Architecture

Given input octrees \mathbf{x} , composed of per-instance partial point clouds derived from depth maps and instance masks, along with their corresponding image features, we aim to predict 3D reconstructions and grasp poses $\hat{\mathbf{y}}$ represented as octrees. ZeroGrasp is built upon an octree-based U-Net [9] and conditional variational autoencoder (CVAE) [63] to model shape reconstruction uncertainty and grasp pose prediction, while maintaining near real-time inference. We present two key innovations to improve its accuracy and generalization. Specifically, we introduce (1) **multi-object encoder** to model spatial relations between objects via a 3D transformer in the latent space, enabling collision-free 3D reconstructions and grasp poses, and (2) **3D occlusion fields**, a novel 3D occlusion representation which captures inter-object occlusions to enhance shape reconstruction in occluded regions.

Octree feature extraction. An RGB image $\mathbf{I} \in \mathbb{R}^{H \times W \times 3}$ is encoded to extract an image feature \mathbf{W} . We fine-tune SAM 2 [64] to generate 2D instance masks $\mathbf{M} \in \mathbb{R}^{H \times W}$ and \mathbf{M}_i represents an i -th object mask. The image feature map is then unprojected into 3D space by $(\mathbf{q}_i, \mathbf{w}_i) = \pi^{-1}(\mathbf{W}, \mathbf{D}, \mathbf{K}, \mathbf{M}_i)$ where \mathbf{q}_i and \mathbf{w}_i denote

3D point cloud and its corresponding features of an i -th object, respectively. Here, π is the unprojection function, $\mathbf{D} \in \mathbb{R}^{H \times W}$ is the depth map and $\mathbf{K} \in \mathbb{R}^{3 \times 3}$ denotes the camera intrinsics. The 3D point cloud features are converted to an octree $\mathbf{x}_i = (\mathbf{p}_i, \mathbf{f}_i) = \mathcal{G}(\mathbf{q}_i, \mathbf{w}_i)$ where \mathcal{G} is the conversion function from the point cloud and its features to an octree.

Octree-based CVAE. To improve the shape reconstruction quality, ZeroGrasp utilizes probabilistic modeling through an octree-based conditional variational autoencoder (CVAE) to address the inherent uncertainty in single-view shape reconstruction, which is crucial for improving both reconstruction and grasp pose prediction quality. Inspired by [63], our Octree-based CVAE consists of an encoder \mathcal{E} , prior \mathcal{P} , and decoder \mathcal{D} to learn latent representations of 3D shapes and grasp poses together as diagonal Gaussian. Concretely, the encoder $\mathcal{E}(\mathbf{z}_i | \mathbf{x}_i, \mathbf{y}_i)$ learns to predict the latent code \mathbf{z}_i based on the predicted and ground-truth octrees \mathbf{x}_i and \mathbf{y}_i . The prior $\mathcal{P}(\ell_i, \mathbf{z}_i | \mathbf{x}_i)$ takes the octree \mathbf{x}_i as input and computes the latent feature $\ell_i \in \mathbb{R}^{N'_i \times D'}$ and code $\mathbf{z}_i \in \mathbb{R}^{D'}$ where N'_i and D' are the number of points and the dimension of the latent feature. Internally, the latent code is sampled from predicted mean and variance via a reparameterization trick [65]. The decoder $\mathcal{D}(\mathbf{y}_i | \ell_i, \mathbf{z}_i, \mathbf{x}_i)$ predicts a 3D reconstruction along with grasp poses. To save computational cost, the decoder predicts occupancy at each depth, discarding grid cells with a probability below 0.5. Only in the final layer does the decoder predict the SDF, normal vectors and grasp poses as well as occupancy. During training, KL divergence between the encoder and prior is minimized such that their distributions are matched.

Multi-object encoder. The prior \mathcal{P} computes features per object, lacking the capability of modeling global spatial arrangements for collision-free reconstruction and grasp pose prediction. To address this, we incorporate a transformer in the latent space, composed of K standard Transformer blocks with self-attention and RoPE [66] positional encoding, following in [10]. The multi-object encoder \mathcal{M} takes voxel centers $\mathbf{r}_i \in \mathbb{R}^{N'_i \times 3}$ and its features $\ell_i \in \mathbb{R}^{N'_i \times D'}$ of all the objects at the latent space are updated as

$$[\ell_1 \cdots \ell_L] \leftarrow \mathcal{M}([\mathbf{r}_1, \ell_1) \cdots (\mathbf{r}_L, \ell_L)]), \quad (4)$$

where L represents the total number of objects.

3D occlusion fields. Our key insight is that the multi-object encoder primarily learns to avoid collisions between objects and grasp poses in a cluttered scene, as collision modeling requires only local context, making it easier to handle. In contrast, occlusion modeling requires a comprehensive understanding of the global context to accurately capture visibility relationships, since occluders and

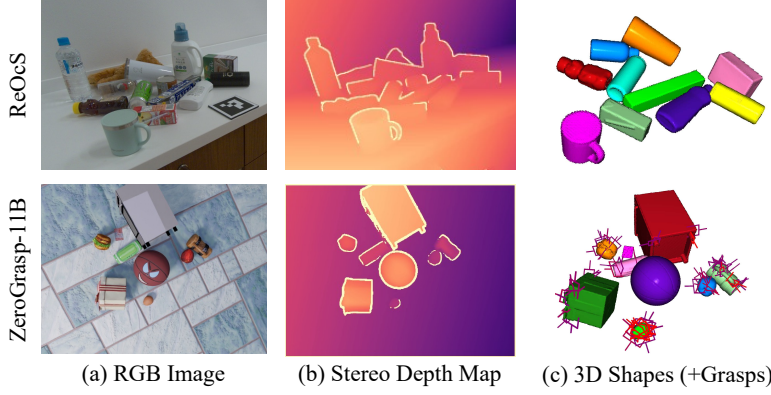


Figure 4. Example RGB images, stereo depth maps, 3D shapes and grasp poses from the ReOcs and ZeroGrasp-11B datasets. The grasp poses of the ZeroGrasp-11B dataset are subsampled by grasp-NMS [1] for better visibility of the 3D shapes and grasps. More examples are found in the supplementary material.

occludees can be positioned far apart. To mitigate this issue, we design 3D occlusion fields that localizes visibility information to voxels via simplified octree-based volume rendering. Concretely, we subdivide a voxel at the latent space into B^3 smaller blocks (B blocks per axis), which are projected into the image space. As shown in Figure 3, if a block lies within the instance mask corresponding to the target object, a self-occlusion flag o_{self} is set to 1. If the block lies within the instance mask of neighbor objects, inter-object occlusion flag o_{inter} is set to 1. After computing the flags of all the blocks, we construct the 3D occlusion fields $\mathcal{V}_i \in \mathbb{R}^{N' \times B^3 \times 2}$ by concatenating the two flags of the i -th object. Finally, we encode it by three layers of 3D CNNs that downsample the resolution by a factor of two at each layer to obtain an occlusion feature $\mathbf{o}_i \in \mathbb{R}^{N' \times D''}$ at the latent space, and update the latent feature by $\ell_i \leftarrow [\ell_i \ \mathbf{o}_i]$ to account for occlusions as well as collisions.

Training. Similar to the standard VAEs [63, 65], we train our model by maximizing the evidence lower bound (ELBO). Additionally, we opt for economic supervision [67] to learn grasp pose prediction efficiently. Therefore, the loss function is defined as

$$\mathcal{L}_{\text{rec}} = \omega_{\text{occ}} \sum_h \mathcal{L}_{\text{occ}}^h + \omega_{\text{nm}} \mathcal{L}_{\text{nm}} + \omega_{\text{SDF}} \mathcal{L}_{\text{SDF}}, \quad (5)$$

$$\mathcal{L}_{\text{grasp}} = \omega_s \mathcal{L}_s + \omega_q \mathcal{L}_q + \omega_a \mathcal{L}_a + \omega_w \mathcal{L}_w + \omega_d \mathcal{L}_d, \quad (6)$$

$$\mathcal{L}_{\text{KL}} = \omega_{\text{KL}} D_{\text{KL}}(\mathcal{E}(\mathbf{z}_i | \mathbf{x}_i, \mathbf{y}_i) \| \mathcal{P}(\ell_i, \mathbf{z}_i | \mathbf{x}_i)), \quad (7)$$

$$\mathcal{L} = \mathcal{L}_{\text{rec}} + \mathcal{L}_{\text{grasp}} + \mathcal{L}_{\text{KL}}, \quad (8)$$

where $\mathcal{L}_{\text{occ}}^h$ computes the mean of the binary cross entropy (BCE) function of occupancy at each depth h , and \mathcal{L}_{nm} , and \mathcal{L}_{SDF} represent the averaged L1 distances of surface normal and SDF, respectively, at the final depth of the octree. \mathcal{L}_s , \mathcal{L}_q , \mathcal{L}_a , \mathcal{L}_w , and \mathcal{L}_d computes the averaged L1 distance of graspness of all the possible views, and cross entropy for




Figure 5. **Contact-based constraints** are used to effectively refine grasp poses. We first obtain contact points c_L and c_R . Next, the contact distance $D(c_{L|R})$, and the depth is computed by $Z(c_{L|R})$ are computed. Finally, the width and height of the grasp is refined based on Eq. (10) and Eq. (11).

quality, angle, width, and depth respectively. Finally, the term \mathcal{L}_{KL} measures the KL divergence between the encoder \mathcal{E} and the prior \mathcal{P} . Each ω term is a weight parameter to align the scale of different loss terms.

3.2. Grasp Pose Refinement

We find that a strong advantage of 3D reconstruction is its ability to utilize the reconstruction to refine predicted grasp poses. While Ma *et al.* [3] propose a contact-based optimization algorithm, it requires an accurate truncated signed distance field (TSDF) reconstructed from multi-view images and its runtime is relatively slow. In contrast, we introduce a simple refinement algorithm that applies contact-based constraints and collision detection on the 3D reconstruction. Specifically, we first detect contact points by finding the closest points on the reconstruction to the left and right fingers of the gripper. We then adjust the predicted width and depth so that both fingertips have contact. Finally, we perform collision detection with the reconstruction to discard grasp poses with collisions. In the following, we explain the details of these two refinement processes.

Contact-based constraints. Accurate contacts are essential for successful grasping, as they ensure stability and control during manipulation. While our network predicts width and depth of the gripper, we observe that even small errors can result in unstable grasping. To address this issue, we refine a grasp pose by adjusting the fingertip locations of the gripper to align with the nearest contact points of the left and right fingers c_L and c_R on the reconstruction. Based on the contact points the width w is refined as

$$\Delta w = \min(D(c_L), D(c_R)), \quad (9)$$

$$w \leftarrow w + 2(\max(\gamma_{\min}, \min(\Delta w, \gamma_{\max})) - \Delta w), \quad (10)$$

so that the contact distance Δw remains within the range γ_{\min} to γ_{\max} . Note that $D(c)$ denotes the contact distance




Figure 6. Comparisons of 3D reconstruction methods using sparse voxel representations on the ReOcS dataset. Except OctMAE [10], an RGB-D image and predicted instance mask are given as input, and the methods output per-object reconstructions. For OctMAE, we visualize its results with normal maps since it is designed to predict a scene-level reconstruction. For a fair comparison, all the models are trained only on the ZeroGrasp-11B dataset. The red rectangles highlight the regions with major differences.

Table 2. Quantitative evaluation of 3D reconstruction on the GraspNet-1B [1], and ReOcS datasets with different difficulties. Chamfer distance (CD), F1-Score@10mm (F1), and normal consistency (NC) are reported in the unit of mm. Seg. denotes an output 3D reconstruction is segmented or not.

Method	Seg.	GraspNet-1B [1]			ReOcS (Ours)								
					Easy			Normal			Hard		
		CD↓	F1↑	NC↑	CD↓	F1↑	NC↑	CD↓	F1↑	NC↑	CD↓	F1↑	NC↑
Minkowski [8]	✓	6.84	81.45	77.89	5.59	85.40	84.74	6.05	82.15	82.68	9.11	77.10	80.86
OCNN [43]	✓	7.23	<u>82.22</u>	<u>78.44</u>	<u>5.26</u>	85.43	85.66	5.96	82.33	<u>84.25</u>	8.69	77.58	<u>82.08</u>
OctMAE [10]		<u>7.57</u>	78.38	75.19	5.53	<u>87.62</u>	86.90	<u>5.93</u>	<u>83.98</u>	83.45	<u>6.76</u>	<u>80.24</u>	80.58
Ours	✓	6.05	84.08	78.46	4.76	88.71	<u>86.74</u>	5.54	84.67	85.13	6.73	80.86	82.95

from \mathbf{c} . We further adjust the depth \mathbf{d} by

$$\mathbf{d} \leftarrow \max(Z(\mathbf{c}_L), Z(\mathbf{c}_R)), \quad (11)$$

where $Z(\mathbf{c})$ compute depth of the contact point \mathbf{c} . These simple refinement steps help ensure stable grasps.

Collision detection. We implement a simple model-free collision detector using the two-finger gripper, following GS-Net [57]. Although the previous method uses partial point cloud obtained from a depth map, it fails to discard predicted grasp poses that result in collisions with occluded regions. To overcome this limitation, we instead leverage the reconstructed shapes, which allows more precise collision detection. To justify this approach, we perform extensive analysis in our experiments and show the advantages.

4. Datasets

We create two datasets for evaluation and training — 1) the ReOcS dataset is designed to evaluate the quality of 3D reconstruction under varying occlusion levels, and 2) the ZeroGrasp-11B dataset is intended for training baselines and our model for zero-shot robotic grasping. Figure 4 highlights several examples of the datasets.

4.1. ReOcS Dataset

The ReOcS dataset contains 1,125 RGB-D images and ground-truth instance masks, 6D poses, and 3D shapes. To

obtain accurate depth maps of metallic and transparent objects, we use a learning-based stereo depth estimation algorithm [71]. There are three splits — easy, normal and hard — based on the extent of occlusions. We use this dataset to compare the robustness of baselines and our method under different occlusion conditions. For the details, please refer to the supplementary material.

4.2. ZeroGrasp-11B Dataset

As shown in Table 1, the ZeroGrasp-11B dataset leverages 12K 3D models and create 1M photorealistic RGB images, ground-truth and stereo depth maps of 25,000 scenes with BlenderProc [72]. In addition, it provides ground-truth 3D reconstructions and 6D object poses. While Grasp-Anything-6D [18] has 6D annotations of a larger number of objects, 3D models are missing, which is crucial for reconstruction. Further, its synthesized images and predicted depth maps have no guarantee that they are physically valid, and grasp pose annotations are sparse and generated from planar grasp poses. We solve these issues with the ZeroGrasp-11B dataset to enable zero-shot robotic grasping. In the following, we describe the procedure of grasp pose generation.

Grasp pose generation. Following [6], we begin by randomly sampling N_s surface points on ground-truth 3D reconstructions. N_s is determined by $N_s = \mathcal{A}/\rho$ with

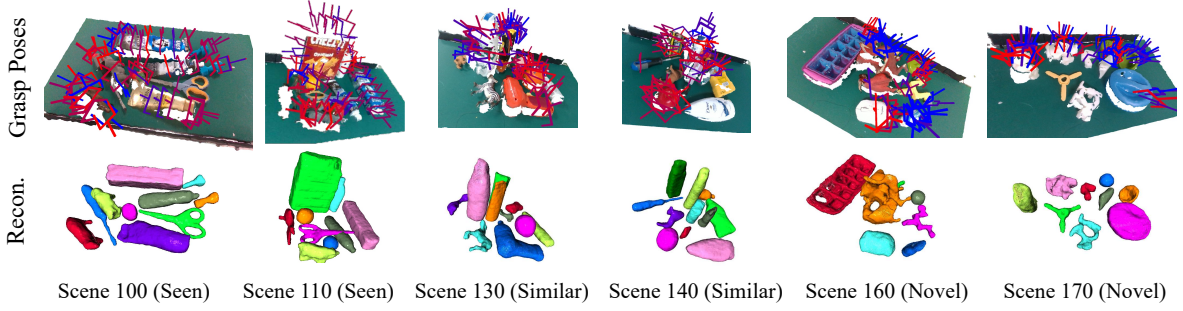


Figure 7. Qualitative results on grasp pose prediction of ZeroGrasp. Following GSNet [57], we show the best 50 grasp predictions after grasp-NMS [1] from six different scenes (two scenes per split). Red and blue grasps denote high and low grasp quality scores, respectively.

Table 3. Quantitative evaluation of grasp pose prediction on the GraspNet-1Billion benchmark. Similar to the other baseline methods, we report the average precision (AP), $AP_{0.4}$, and $AP_{0.8}$. Note that 0.4 and 0.8 denote the friction coefficients, and the lower the more difficult. G and R in the output column indicate whether the reconstructed posture prediction are predicted or not.

Method	Output		Seen			Similar			Novel		
	G	R	AP	$AP_{0.8}$	$AP_{0.4}$	AP	$AP_{0.8}$	$AP_{0.4}$	AP	$AP_{0.8}$	$AP_{0.4}$
GG-CNN [5]	✓		15.48	21.84	10.25	13.26	18.37	4.62	5.52	5.93	1.86
Chu <i>et al.</i> [68]	✓		15.97	23.66	10.80	15.41	20.21	7.06	7.64	8.69	2.52
CenterGrasp [†] [59]	✓	✓	16.46	20.24	11.74	9.52	11.92	5.71	1.60	1.89	1.12
GPD [69]	✓		22.87	28.53	12.84	21.33	27.83	9.64	8.24	8.89	2.67
Lian <i>et al.</i> [4]	✓		25.96	33.01	15.37	22.68	29.15	10.76	9.23	9.89	2.74
GraspNet [1]	✓		27.56	33.43	16.59	26.11	34.18	14.23	10.55	11.25	3.98
GSNet [57]	✓		67.12	78.46	60.90	54.81	66.72	46.17	24.31	30.52	14.23
Ma <i>et al.</i> [70]	✓		63.83	74.25	58.66	58.46	70.05	51.32	24.63	31.05	12.85
HGGD	✓		64.45	72.81	61.16	53.59	64.12	45.91	24.59	30.46	<u>15.58</u>
EconomicGrasp [67]	✓		68.21	79.60	63.54	61.19	73.60	53.77	25.48	31.46	13.85
Ours	✓	✓	<u>70.53</u>	<u>82.28</u>	<u>64.26</u>	<u>62.51</u>	<u>74.26</u>	<u>54.97</u>	<u>26.46</u>	<u>33.13</u>	15.11
Ours+FT	✓	✓	72.43	83.12	65.57	65.45	78.32	55.48	28.49	34.21	15.80

Table 4. Ablations on the network input, architecture, and refinement algorithm. For reconstruction and grasp pose prediction, we report the metrics of the hard split of the ReOcS dataset and GraspNet-1B dataset, respectively.

Method	Reconstruction			Grasp Pose		
	CD↓	F1↑	NC↑	Seen	Similar	Novel
Baseline (OCNN [9])	8.69	77.58	82.08	41.27	36.48	17.46
No CVAE	7.67	78.79	82.35	70.23	60.31	26.28
No Multi-Obj. Encoder	7.09	79.62	82.60	69.52	61.03	26.17
No 3D Occlusion Fields	7.54	78.81	81.94	67.34	58.45	25.00
No Contact Constraints	6.73	80.86	82.95	65.67	55.34	24.92
No Collision Detection	6.73	80.86	82.95	49.35	44.28	21.03
+ Depth Map	6.73	80.86	82.95	59.93	51.58	24.07
Ours	6.73	80.86	82.95	70.53	62.51	26.46

A denoting the surface area and ρ as a density parameter. For each surface point, we synthesize candidate grasps with all combinations of views, orientations around the point’s normal vector, and depth respectively, following GraspNet-1B [1]. Next, we conduct collision detection to eliminate any grasps with collision and compute the grasp quality q for the remaining candidates. The quality metric [55] is computed based on the normal vectors \mathbf{n}_L and \mathbf{n}_R of the contact points \mathbf{c}_L and \mathbf{c}_R by $q = \min(\mathbf{n}_L \cdot \mathbf{c}_{LR}, \mathbf{n}_R \cdot \mathbf{c}_{LR})$, where $\mathbf{c}_{LR} = (\mathbf{c}_L - \mathbf{c}_R) / (\|\mathbf{c}_L\| \|\mathbf{c}_R\|)$. Finally, we physically validate the generated grasps with IsaacGym [73]. To make the Obverse 3D models compatible with simulation, we decompose them into convex hulls using V-HACD [74]. Figure 4 shows the grasp poses before and after the collision and physics-based filtering process.

5. Experiments

Implementation details. Our proposed method, ZeroGrasp, adopts a ResNeXt [75] architecture, pretrained on the ImageNet dataset [76], as an image encoder, and all the parameters except the last layer are fixed during training.

Similar to EconomicGrasp [67], we use the predicted view graspness s to determine a view direction. For training, we use AdamW [77] with a learning rate of 0.001, batch size of 16 on NVIDIA A100. The weights of the loss function is provided in the supplemental. We set the dimensions of the input image feature D , the latent feature D' , and the 3D occlusion fields \mathcal{V} to 32, 192, and 16 respectively. For the 3D occlusion fields, we use 8 for the block resolution B . Following Ma *et al.*, the ranges of the contact distance γ_{\min} and γ_{\max} are defined to 0.005m and 0.02m, respectively. To generate grasp poses, we use 0.005m² as the sampling density ρ .

Metrics. Similar to OctMAE [10], we use the Chamfer distance (CD), F-1 score, and normal consistency (NC) to evaluate the quality of 3D reconstruction. To evaluate the quality of grasp pose prediction, we use average precision (AP), a standard metric of the GraspNet-1B benchmark, which evaluates average precision based on the top-k ranked grasps in a scene. The AP_μ metric measures the precision with the friction of μ by evaluating grasps with friction μ over m different thresholds. The final AP score is computed as the mean of AP_μ , using friction values μ from 0.2 to 1.2

at intervals of 0.2.

5.1. Main Results

3D reconstruction. As shown in Table 2, our method outperforms the other single-view reconstruction methods. We choose the three methods using sparse voxel representations due to its superior efficiency and accuracy in a zero-shot setup, reported in Iwase *et al.* [10]. We train the baseline and our methods on the ZeroGrasp-11B dataset and evaluate them on the GraspNet-1B and ReOcS dataset to test generalization to real-world images. Our qualitative evaluation in Figure 6 demonstrates the robustness of ZeroGrasp to real-world images and inter-object occlusions.

Grasp pose prediction. Table 3 demonstrates the comparison against state-of-the-art methods for grasp pose prediction on the RealSense data of the GraspNet-1Billion benchmark. The baselines and our model are trained on the training split of the GraspNet-1Billion dataset for 20 epochs. Notably, our method achieves the state-of-the-art performance across all the AP metrics. In the Ours+FT setup, our model is initially pre-trained on the ZeroGrasp-11B dataset, then fine-tuned on the GraspNet-1Billion dataset for 2 epochs. As a result, fine-tuning improves 1.9%, 2.94%, and 2.03% in the seen, similar and novel splits. This result supports the importance of large-scale grasp pose datasets for zero-shot robotic grasping. Figure 7 shows qualitative results of ZeroGrasp. Unlike the previous methods, ZeroGrasp enables accurate grasp pose prediction even in occluded or truncated regions.

5.2. Ablations

Table 4 shows our ablation studies to validate the effectiveness of each component. We provide detailed analyses from the perspectives of the two tasks addressed in our work.

3D reconstruction. We observe a consistent drop in performance across all reconstruction metrics when each of CVAE, the multi-object encoder, and 3D occlusion fields is individually excluded. This highlights the importance of multi-object reasoning to achieve higher reconstruction quality. As shown in Figure 6, our visualizations further demonstrate that these components contribute to better reconstruction, especially in regions with inter-object occlusions and close contacts between objects.

Grasp pose prediction. As illustrated in Table 4, most of the components contribute to improved grasp pose detection. In particular, collision detection and contact-based constraints provide a significant boost to grasp pose quality. Our comparison of collision detection using a depth map (partial point clouds) as in GSNet [57] and our predicted reconstruction (59.93 vs 70.53) reveals that reconstruction-based collision detection is more effective. Furthermore, the




Figure 8. Example scenes of our real-robot evaluation.

substantial performance drop without 3D occlusion fields underscores the importance of reasoning about inter-object occlusions.

5.3. Real-Robot Evaluation

We validate the feasibility and generalizability of the baseline (OCNN [9]) and our method, trained only on our synthetic dataset, through real-world evaluations. Our robotic setup uses Franka Emika Panda robot and Robotiq 2F-85 hand. As shown in Figure 8, we set up 5 scenes with 3 to 4 objects. Each object is picked up in repeated trials, with a maximum of 3 attempts per object. Our success rate, measured by the ratio of objects successfully picked up, is 56.25% for the baseline and 75% for our method, highlighting the strong generalization of our approach in real-world scenarios. We describe more details about the robotic setup and show qualitative results in the supplementary material.

6. Conclusion

In this paper, we propose ZeroGrasp, a novel approach for simultaneous 3D reconstruction and grasp pose prediction. By integrating five key components, ZeroGrasp enhances both shape reconstruction and grasp prediction quality. Our extensive analysis confirms the effectiveness of these components. In addition, we strongly believe that ZeroGrasp-11B dataset facilitates future research in zero-shot robotic grasping. Despite its promising generalization capabilities, ZeroGrasp has some limitations. First, our method does not support incremental or multi-view 3D reconstruction [78, 79], which is beneficial when using a wrist-mounted camera on an end effector. Second, it does not account for placement poses that could leverage predicted 3D reconstructions. While this paper focuses on single-view 3D reconstruction and grasp pose prediction, exploring these directions would be valuable.

Acknowledgment

We thank Tianyi Ko for help with the real-world robot experiments. This research was supported by Toyota Research Institute.

References

- [1] H.-S. Fang, C. Wang, M. Gou, and C. Lu, “Graspnet-1billion: A large-scale benchmark for general object grasping,” in *CVPR*, 2020. 1, 2, 3, 5, 6, 7
- [2] Y. Jiang, S. Moseson, and A. Saxena, “Efficient grasping from rgb-d images: Learning using a new rectangle representation,” in *ICRA*, 2011. 3
- [3] M. Haoxiang, S. Modi, B. Gao, and H. Di, “Generalizing 6-dof grasp detection via domain prior knowledge,” in *CVPR*, 2024. 1, 3, 5
- [4] H. Liang, X. Ma, S. Li, M. Görner, S. Tang, B. Fang, F. Sun, and J. Zhang, “PointNetGPD: Detecting grasp configurations from point sets,” in *ICRA*, 2019. 7
- [5] D. Morrison, P. Corke, and J. Leitner, “Closing the Loop for Robotic Grasping: A Real-time, Generative Grasp Synthesis Approach,” in *RSS*, 2018. 7
- [6] A. Mousavian, C. Eppner, and D. Fox, “6-dof graspnet: Variational grasp generation for object manipulation,” in *ICCV*, 2019. 1, 3, 6
- [7] W. Shen, G. Yang, A. Yu, J. Wong, L. P. Kaelbling, and P. Isola, “Distilled feature fields enable few-shot language-guided manipulation,” in *CoRL*, 2023. 1
- [8] C. Choy, J. Gwak, and S. Savarese, “4D Spatio-Temporal ConvNets: Minkowski Convolutional Neural Networks,” in *CVPR*, 2019. 1, 2, 6
- [9] P.-S. Wang, Y. Liu, Y.-X. Guo, C.-Y. Sun, and X. Tong, “O-CNN: Octree-Based Convolutional Neural Networks for 3D Shape Analysis,” *SIGGRAPH*, 2017. 2, 3, 4, 7, 8
- [10] L. K. Iwase, Shun and, V. Guizilini, A. Gaidon, K. Kitani, R. Ambrus, and S. Zakharov, “Zero-shot multi-object scene completion,” in *ECCV*, 2024. 1, 4, 6, 7, 8
- [11] M. Deitke, D. Schwenk, J. Salvador, L. Weihs, O. Michel, E. VanderBilt, L. Schmidt, K. Ehsani, A. Kembhavi, and A. Farhadi, “Objaverse: A Universe of Annotated 3D Objects,” *CVPR*, 2022. 2, 3
- [12] A. Depierre, E. Dellandréa, and L. Chen, “Jacquard: A large scale dataset for robotic grasp detection,” *IROS*, 2018. 3
- [13] H. Zhang, X. Lan, S. Bai, X. Zhou, Z. Tian, and N. Zheng, “Roi-based robotic grasp detection for object overlapping scenes,” in *IROS*, 2019. 3
- [14] X. Yan, J. Hsu, M. Khansari, Y. Bai, A. Pathak, A. Gupta, J. Davidson, and H. Lee, “Learning 6-dof grasping interaction via deep geometry-aware 3d representations,” in *ICRA*, 2018. 3
- [15] C. Eppner, A. Mousavian, and D. Fox, “ACRONYM: A large-scale grasp dataset based on simulation,” in *ICRA*, 2021. 3
- [16] H. Zhang, D. Yang, H. Wang, B. Zhao, X. Lan, J. Ding, and N. Zheng, “Regrad: A large-scale relational grasp dataset for safe and object-specific robotic grasping in clutter,” *RA-L*, 2022. 3
- [17] H. Jung, G. Zhai, S.-C. Wu, P. Ruhkamp, H. Schieber, P. Wang, G. Rizzoli, H. Zhao, S. D. Meier, D. Roth, N. Navab, *et al.*, “Housecat6d—a large-scale multi-modal category level 6d object perception dataset with household objects in realistic scenarios,” *CVPR*, 2024. 3
- [18] T. Nguyen, M. N. Vu, B. Huang, A. Vuong, Q. Vuong, N. Le, T. Vo, and A. Nguyen, “Language-driven 6-dof grasp detection using negative prompt guidance,” in *ECCV*, 2024. 3, 6
- [19] S. F. Bhat, R. Birkel, D. Wofk, P. Wonka, and M. Müller, “Zoedepth: Zero-shot transfer by combining relative and metric depth,” 2023. 3
- [20] Z. Huang, S. Stojanov, A. Thai, V. Jampani, and J. M. Rehg, “ZeroShape: Regression-based Zero-shot Shape Reconstruction,” *CVPR*, 2023. 2
- [21] X. Ren, J. Huang, X. Zeng, K. Museth, S. Fidler, and F. Williams, “Xcube: Large-scale 3d generative modeling using sparse voxel hierarchies,” in *CVPR*, 2024. 2
- [22] L. Mescheder, M. Oechsle, M. Niemeyer, S. Nowozin, and A. Geiger, “Occupancy Networks: Learning 3D Reconstruction in Function Space,” in *CVPR*, 2019.
- [23] S. Peng, M. Niemeyer, L. Mescheder, M. Pollefeys, and A. Geiger, “Convolutional Occupancy Networks,” in *ECCV*, 2020. 2
- [24] M. Z. Irshad, T. Kollar, M. Laskey, K. Stone, and Z. Kira, “Centersnap: Single-shot multi-object 3d shape reconstruction and categorical 6d pose and size estimation,” 2022.
- [25] M. Z. Irshad, S. Zakharov, R. Ambrus, T. Kollar, Z. Kira, and A. Gaidon, “Shapo: Implicit representations for multi-object shape appearance and pose optimization,” 2022.
- [26] A. Bozic, P. Palafox, J. Thies, A. Dai, and M. Nießner, “TransformerFusion: Monocular rgb scene reconstruction using transformers,” in *NeurIPS*, 2021.
- [27] A. Dai, D. Ritchie, M. Bokeloh, S. Reed, J. Sturm, and M. Nießner, “ScanComplete: Large-Scale Scene Completion and Semantic Segmentation for 3D Scans,” in *CVPR*, 2018.
- [28] H.-X. Chen, J. Huang, T.-J. Mu, and S.-M. Hu, “CIRCLE: Convolutional Implicit Reconstruction And Completion For Large-Scale Indoor Scene,” in *ECCV*, 2022.
- [29] M. Lunayach, S. Zakharov, D. Chen, R. Ambrus, Z. Kira, and M. Z. Irshad, “Fsd: Fast self-supervised single rgb-d to categorical 3d objects,” in *Int. Conf. on Robotics and Automation*, IEEE, 2024.
- [30] J. Huang, Z. Gojcic, M. Atzmon, O. Litany, S. Fidler,

- and F. Williams, "Neural Kernel Surface Reconstruction," in *CVPR*, 2023.
- [31] Y. Li, Z. Yu, C. Choy, C. Xiao, J. M. Alvarez, S. Fidler, C. Feng, and A. Anandkumar, "VoxFormer: Sparse Voxel Transformer for Camera-based 3D Semantic Scene Completion," in *CVPR*, 2023. 2
- [32] N. Heppert, M. Z. Irshad, S. Zakharov, K. Liu, R. A. Ambrus, J. Bohg, A. Valada, and T. Kollar, "Carto: Category and joint agnostic reconstruction of articulated objects," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 21201–21210, 2023.
- [33] C.-Y. Wu, J. Johnson, J. Malik, C. Feichtenhofer, and G. Gkioxari, "Multiview Compressive Coding for 3D Reconstruction," in *CVPR*, 2023. 2
- [34] A. Boulch and R. Marlet, "POCO: Point Convolution for Surface Reconstruction," in *CVPR*, 2022. 2
- [35] T. Shen, J. Gao, K. Yin, M.-Y. Liu, and S. Fidler, "Deep Marching Tetrahedra: a Hybrid Representation for High-Resolution 3D Shape Synthesis," in *NeurIPS*, 2021.
- [36] Z. Liu, Y. Feng, M. J. Black, D. Nowrouzezahrai, L. Paull, and W. Liu, "MeshDiffusion: Score-based Generative 3D Mesh Modeling," in *ICLR*, 2023.
- [37] J. J. Park, P. Florence, J. Straub, R. Newcombe, and S. Lovegrove, "DeepSDF: Learning Continuous Signed Distance Functions for Shape Representation," in *CVPR*, 2019.
- [38] X. Yu, Y. Rao, Z. Wang, Z. Liu, J. Lu, and J. Zhou, "PoinTr: Diverse Point Cloud Completion with Geometry-Aware Transformers," in *ICCV*, 2021. 2
- [39] X. Yan, L. Lin, N. J. Mitra, D. Lischinski, D. Cohen-Or, and H. Huang, "ShapeFormer: Transformer-based Shape Completion via Sparse Representation," in *CVPR*, 2022. 2
- [40] P. Mittal, Y.-C. Cheng, M. Singh, and S. Tulsiani, "AutoSDF: Shape Priors for 3D Completion, Reconstruction and Generation," in *CVPR*, 2022.
- [41] Y.-C. Cheng, H.-Y. Lee, S. Tulyakov, A. G. Schwing, and L.-Y. Gui, "SDFusion: Multimodal 3d shape completion, reconstruction, and generation," in *CVPR*, 2023.
- [42] J. Varley, C. DeChant, A. Richardson, J. Ruales, and P. Allen, "Shape completion enabled robotic grasping," in *IROS*, 2017. 3
- [43] P.-S. Wang, Y. Liu, and X. Tong, "Deep Octree-based CNNs with Output-Guided Skip Connections for 3D Shape and Scene Completion," in *CVPRW*, 2020. 6
- [44] S. Song, F. Yu, A. Zeng, A. X. Chang, M. Savva, and T. Funkhouser, "Semantic Scene Completion from a Single Depth Image," *CVPR*, 2017.
- [45] D. Zhang, C. Choi, I. Park, and Y. M. Kim, "Probabilistic Implicit Scene Completion," in *ICLR*, 2022.
- [46] S. S. Mohammadi, N. F. Duarte, D. Dimou, Y. Wang, M. Taiana, P. Morerio, A. Dehban, P. Moreno, A. Bernardino, A. Del Bue, and J. Santos-Victor, "3DSGrasp: 3D Shape-Completion for Robotic Grasp," in *ICRA*, 2023.
- [47] P. Zhang, W. Liu, Y. Lei, H. Lu, and X. Yang, "Cascaded Context Pyramid for Full-Resolution 3D Semantic Scene Completion," in *ICCV*, 2019. 2
- [48] J. Li, K. Han, P. Wang, Y. Liu, and X. Yuan, "Anisotropic Convolutional Networks for 3D Semantic Scene Completion," in *CVPR*, 2020. 2
- [49] F. Williams, J. Huang, J. Swartz, G. Klar, V. Thakkar, M. Cong, X. Ren, R. Li, C. Fuji-Tsang, S. Fidler, E. Sifakis, and K. Museth, "fvdb: A deep-learning framework for sparse, large-scale, and high-performance spatial intelligence," *SIGGRAPH*, 2024. 2
- [50] K. Museth, "VDB: High-resolution sparse volumes with dynamic topology," 2013. 2
- [51] E. R. Chan, K. Nagano, M. A. Chan, A. W. Bergman, J. J. Park, A. Levy, M. Aittala, S. D. Mello, T. Karras, and G. Wetzstein, "GeNVS: Generative novel view synthesis with 3D-aware diffusion models," in *CoRR*, 2023. 2
- [52] R. Liu, R. Wu, B. V. Hoorick, P. Tokmakov, S. Zakharov, and C. Vondrick, "Zero-1-to-3: Zero-shot One Image to 3D Object," in *CVPR*, 2023. 2
- [53] D. Watson, W. Chan, R. Martin-Brualla, J. Ho, A. Tagliasacchi, and M. Norouzi, "Novel View Synthesis with Diffusion Models," *CoRR*, 2022. 2
- [54] J. Xu, W. Cheng, Y. Gao, X. Wang, S. Gao, and Y. Shan, "Instantmesh: Efficient 3d mesh generation from a single image with sparse-view large reconstruction models," *CoRR*, 2024. 2
- [55] V.-D. Nguyen, "Constructing force-closure grasps," in *Proceedings. 1986 IEEE International Conference on Robotics and Automation*, 1986. 3, 7
- [56] A. Bicchi and V. Kumar, "Robotic grasping and contact: a review," in *ICRA*, 2000. 3
- [57] C. Wang, H.-S. Fang, M. Gou, H. Fang, J. Gao, and C. Lu, "Graspness discovery in clutters for fast and accurate grasp detection," in *ICCV*, 2021. 3, 6, 7, 8
- [58] J. Mahler, J. Liang, S. Niyaz, M. Laskey, R. Doan, X. Liu, J. A. Ojea, and K. Goldberg, "Dex-net 2.0: Deep learning to plan robust grasps with synthetic point clouds and analytic grasp metrics," in *RSS*, 2017. 3
- [59] E. Chisari, N. Heppert, T. Welschehold, W. Burgard, and A. Valada, "Centergrasp: Object-aware implicit representation learning for simultaneous shape reconstruction and 6-dof grasp estimation," *RA-L*, 2024. 3, 7

- [60] Z. Jiang, Y. Zhu, M. Svetlik, K. Fang, and Y. Zhu, “Synergies between affordance and geometry: 6-dof grasp detection via implicit representations,” *RSS*, 2021. 3
- [61] P. Grady, C. Tang, C. D. Twigg, M. Vo, S. Brahmabhatt, and C. C. Kemp, “ContactOpt: Optimizing contact to improve grasps,” in *CVPR*, 2021. 3
- [62] D. Morrison, P. Corke, and J. Leitner, “Egad! an evolved grasping analysis dataset for diversity and reproducibility in robotic manipulation,” *RA-L*, 2020. 3
- [63] D. Rempe, T. Birdal, A. Hertzmann, J. Yang, S. Sridhar, and L. J. Guibas, “Humor: 3d human motion model for robust pose estimation,” in *ICCV*, 2021. 4, 5
- [64] N. Ravi, V. Gabeur, Y.-T. Hu, R. Hu, C. Ryali, T. Ma, H. Khedr, R. Rädle, C. Rolland, L. Gustafson, E. Mintun, J. Pan, K. V. Alwala, N. Carion, C.-Y. Wu, R. Girshick, P. Dollár, and C. Feichtenhofer, “Sam 2: Segment anything in images and videos,” *CoRR*, 2024. 4
- [65] D. P. Kingma and M. Welling, “Auto-encoding variational bayes,” *ICLR*, 2014. 4, 5
- [66] J. Su, Y. Lu, S. Pan, B. Wen, and Y. Liu, “RoFormer: Enhanced Transformer with Rotary Position Embedding,” in *ICLR*, 2020. 4
- [67] X.-M. Wu, J.-F. Cai, J.-J. Jiang, D. Zheng, Y.-L. Wei, and W.-S. Zheng, “An economic framework for 6-dof grasp detection,” in *ECCV*, 2025. 5, 7
- [68] F. Chu, R. Xu, and P. A. Vela, “Real-world multiobject, multigrasp detection,” in *RA-L*, 2018. 7
- [69] A. ten Pas, M. Gualtieri, K. Saenko, and R. W. Platt, “Grasp pose detection in point clouds,” *The International Journal of Robotics Research (IJRR)*, 2017. 7
- [70] M. Haoxiang and D. Huang, “Towards scale balanced 6-dof grasp detection in cluttered scenes,” in *CoRL*, 2022. 7
- [71] K. Shankar, M. Tjersland, J. Ma, K. Stone, and M. Bajracharya, “A learned stereo depth system for robotic manipulation in homes,” *RA-L*, 2021. 6
- [72] M. Denninger, D. Winkelbauer, M. Sundermeyer, W. Boerdijk, M. Knauer, K. H. Strobl, M. Humt, and R. Triebel, “BlenderProc2: A Procedural Pipeline for Photorealistic Rendering,” *Journal of Open Source Software*, 2023. 6
- [73] V. Makoviychuk, L. Wawrzyniak, Y. Guo, M. Lu, K. Storey, M. Macklin, D. Hoeller, N. Rudin, A. Allshire, A. Handa, and G. State, “Isaac gym: High performance gpu-based physics simulation for robot learning,” *CoRR*, 2021. 7
- [74] K. Mamou, “V-hacd: Volumetric hierarchical approximate convex decomposition.” <https://github.com/kmamou/v-hacd>, 2016. 7
- [75] S. Xie, R. Girshick, P. Dollár, Z. Tu, and K. He, “Aggregated Residual Transformations for Deep Neural Networks,” *CVPR*, 2017. 7
- [76] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, “Imagenet: A large-scale hierarchical image database,” in *CVPR*, 2009. 7
- [77] I. Loshchilov and F. Hutter, “Decoupled weight decay regularization,” in *ICLR*, 2019. 7
- [78] R. Mur-Artal and J. D. Tardós, “ORB-SLAM2: an open-source SLAM system for monocular, stereo and RGB-D cameras,” *IEEE Transactions on Robotics*, 2017. 8
- [79] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, “Nerf: Representing scenes as neural radiance fields for view synthesis,” in *ECCV*, 2020. 8