

# CenterGrasp: Object-Aware Implicit Representation Learning for Simultaneous Shape Reconstruction and 6-DoF Grasp Estimation

Eugenio Chisari<sup>1</sup>, Nick Heppert<sup>1</sup>, Tim Welschhold<sup>1</sup>, Wolfram Burgard<sup>2</sup>, Abhinav Valada<sup>1</sup>

**Abstract**—Reliable object grasping is a crucial capability for autonomous robots. However, many existing grasping approaches focus on general clutter removal without explicitly modeling objects and thus only relying on the visible local geometry. We introduce CenterGrasp, a novel framework that combines object awareness and holistic grasping. CenterGrasp learns a general object prior by encoding shapes and valid grasps in a continuous latent space. It consists of an RGB-D image encoder that leverages recent advances to detect objects and infer their pose and latent code, and a decoder to predict shape and grasps for each object in the scene. We perform extensive experiments on simulated as well as real-world cluttered scenes and demonstrate strong scene reconstruction and 6-DoF grasp-pose estimation performance. Compared to the state of the art, CenterGrasp achieves an improvement of 38.5 mm in shape reconstruction and 33 percentage points on average in grasp success. We make the code and trained models publicly available at <http://centergrasp.cs.uni-freiburg.de>.

**Index Terms**—Grasping, Deep Learning in Grasping and Manipulation, RGB-D Perception

## I. INTRODUCTION

GRASP-POSE estimation is a high-dimensional optimization problem as many different grasp candidates exist for a given object, with varying contact conditions and robot joint configurations. To simplify the problem, early works either assume prior knowledge of the object shape, e.g. via meshes or CAD models [1], [2], or they only consider top-down (i.e. 4-DoF) grasps [3], [4]. In contrast, in this paper, we address the general case of full 6-DoF object grasping, where the exact object instance and category are unknown a priori, and only a single RGB-D image of the scene is observed. Prior work has already shown impressive results in the general 6-DoF grasp-pose estimation problem [5]–[8]. Nevertheless, wider adoption beyond research labs is still limited. To close this gap, we identify two key features that the ideal object grasping pipeline should possess: *holistic grasping*, i.e. the ability to predict grasps all around the given object, even in regions occluded from the agent’s perspective, and *object awareness*, i.e. the ability to identify different objects in the scene and distinguish between them as well as the background. These properties would enable robots to grasp a specific target object, as opposed to just a random object from the clutter, as well




Fig. 1: Overview of CenterGrasp: (a) Input image, (b) Object heatmap prediction, (c) Subset of all valid and final grasp proposals overlaid on the observed scene, (d) Subset of all valid and final grasp proposals overlaid on the 3D scene reconstruction, (e) Selected grasps overlaid on the observed scene, (f) Selected grasps overlaid on the 3D scene reconstruction.

as to reason about the best grasp poses independently of the current camera point of view, e.g. grasping a mug from the handle even if it is currently not visible.

To achieve *holistic grasping*, prior works adopt a multi-camera setup around the scene [9], while others require the agent to first execute a full scan of the scene prior to estimating grasps [6]. While these strategies are effective, they restrict the applicability of the method, since the former requires carefully instrumented environments, while the latter drastically reduces the speed of operation. Additionally, there are common scenarios where both methods are not applicable, for example, when looking inside a cabinet to grasp a mug.

Most prior works are not object-aware, as they only consider the local geometry of the scene to predict grasps and discard additional information such as color images [6], [8]. Hence, these methods are limited to clutter removal tasks, where objects are grasped randomly from the scene. Recent work leveraged instance segmentation to achieve object awareness in their pipelines [7]. However, two-step pipelines increase complexity during inference and are thus more error-prone. Further, a combined model has the advantage to exploit mutually beneficial information.

In this work, we propose CenterGrasp, a novel, single-shot object-aware holistic grasp prediction approach. It consists of an image encoder that predicts an objectness heatmap, a pose,

Manuscript received: December, 13, 2023; Revised February, 29, 2024; Accepted April, 2, 2024.

This paper was recommended for publication by Editor M. Vincze upon evaluation of the Associate Editor and Reviewers’ comments. This work was funded by the Carl Zeiss Foundation with the ReScaLe project.

<sup>1</sup> Department of Computer Science, University of Freiburg, Germany

<sup>2</sup> Department of Engineering, University of Technology Nuremberg, Germany  
Digital Object Identifier (DOI): see top of this page.

and a latent code for each object, as well as an auto-decoder that reconstructs their respective 3D shape and predicts a set of grasp proposals. We train CenterGrasp only on synthetic images and demonstrate that it achieves zero-shot generalization to the real world. We extensively test CenterGrasp on both simulated as well as real-world scenes with two types of clutter: packed and pile. We evaluate on the joint task of simultaneous shape reconstruction and 6-DoF grasp estimation, and compare to the recent state-of-the-art method GIGA [8], demonstrating strong performance on this joint task.

This paper makes the following main contributions:

- CenterGrasp, a novel method for joint shape reconstruction and 6-DoF grasp estimation.
- An automated 6-DoF robotic grasping evaluation environment based on the SAPIEN [10] simulator.
- Extensive experimental evaluations comparing CenterGrasp with GIGA [8], a state-of-the-art baseline.
- The code (including our dataset generation and evaluation script), models, and videos publicly available at <http://centergrasp.cs.uni-freiburg.de>.

## II. RELATED WORK

There has been a plurality of 6-DoF grasping works and object detection methods. In the following, we highlight recent advances in both topics.

### A. 6-DoF Grasp Estimation

The 6-DoF grasp-pose estimation problem has attracted significant interest over the years [11]. We identify two main 6-DoF grasp estimation schemes: point cloud-based and voxel-based.

*Point cloud-based:* GPD [5] and PointNetGPD [12] are popular approaches where multiple grasp candidates are sampled using geometric heuristics from a point cloud, and a learned classifier is trained to evaluate their quality. 6-DoF GraspNet [13] uses a variational autoencoder to predict grasps proposals from a single object point cloud. AVN [14] processes RGB and depth separately to be robust against depth noise. GSNet [15] and AnyGrasp [16] adopt a per-point graspness metric to highlight graspable areas and guide the prediction. VoteGrasp [17], [18] uses contextual information to generate scene features that encode dependencies between objects. Liu *et al.* [19] propose a fusion network for joint foreground segmentation and grasp pose detection. Chen *et al.* [20] present a grasp heatmap model to guide the prediction. DGCAN [21] refines noisy depth features according to cross-modal relations with the RGB image. All the aforementioned methods are purely geometrical and do not include the notion of an object. To tackle this problem, Contact-GraspNet [7] proposes a two stages pipeline, where an unknown object segmentation model [22] is used to detect different object instances. In contrast, our heatmap prediction inherently distinguishes between multiple objects in the scene and thus enables object-aware grasping.

*Voxel-based:* Instead of point clouds, some methods opt to encode the observed scene in a Truncated Signed Distance Field (TSDF) volume [23]. Volumetric Grasping Network (VGN) [6] executes a scan motion with the robot’s end-effector

and integrates all camera depth observations in a TSDF voxel-grid, which is then used as input to predict grasps candidates. While effective, the scan motion drastically reduces the speed of operation of the robot. To overcome this limitation, GIGA [8] proposes to exploit the synergy between geometric reasoning and grasp learning. The key idea is that a representation capable of encoding the shape information of the scene is also relevant to grasp prediction and vice-versa. Nevertheless, voxel-based methods can only be applied to a workspace of fixed dimensions, whose coordinates and calibration with respect to the robot are known a priori. Contrary to point cloud-based methods, this requires additional setup efforts and it prevents their use in more dynamic settings, such as mobile manipulation [24], [25].

### B. Object Detection and Representation

In order to fulfill both *object awareness* and *holistic grasping*, a method must be able to both detect objects in the scene, as well as produce an appropriate representation from which grasps can be extracted.

*Center-Based Detection:* The seminal work *Objects as Points* [26] led to the development of many center-based object detection methods [27], [28]. Recently, center-based detection methods have shown great results on category-level pose estimation and shape reconstruction [29]–[31] with extension to articulated objects [32]. We build upon these single-shot frameworks, mitigating the issues of two-stage pipelines for grasping [7], [33], and allowing us to predict a vast amount of grasp proposals while still retaining object-awareness.

*Implicit Object Representation:* A useful object representation needs to be able to encode relevant geometric information and to generalize to unseen object instances. The recently proposed NUNOCS [34] represents each object as a 6D transformation and 3D scaling of the category “mean shape”. This simple representation allows extracting dense correspondences and transferring relevant grasps across instances. While effective, this approach cannot handle large intra-category shape variations besides scaling. A different approach is to use the recently proposed OccupancyNet [35] as an implicit representation. GIGA [8] employs this model for scene reconstruction from partial observations. NDFs [9] adopt it to encode geometric and semantic descriptors in a 3D volume, which can then be queried to sample task-relevant keypoints [36] for unseen instances. They show impressive results in terms of generalization to unseen object poses and shape instances, but they require a setup of four cameras, as well as a separate model for each category, and a set of manually collected demonstrations.

*Grasp Distance Function:* DeepSDF [37] is a learned continuous Signed Distance Function representation that implicitly encodes a surface as the zero-level set of the learned model. Neural Grasp Distance Fields (NGDF) [38] extend the concept of the neural implicit distance function to grasping, where the learned model outputs scalar distance metrics for the current gripper pose to the closest grasp pose. Through differentiation, the joint positions of the robot can be iteratively updated, until the desired grasp pose is reached. While promising, this approach has only been evaluated with a single object in

the scene and training a separate model for each category. Moreover, it only provides the gradient towards one single grasp pose (the closest one). On the other hand, CenterGrasp is category-independent, it handles cluttered scenes, and it predicts an entire manifold of grasp candidates for each object in the scene. Most similar to our method, SceneGrasp [39] also learns a combined latent space for shape and grasp prediction. While interesting, this approach has only been trained on 6 categories of objects, and it lacks evaluation of the grasp quality, both in simulation and in the real world.

### III. TECHNICAL APPROACH

Given an RGB-D image as input, the goal is to simultaneously detect all the objects and predict their pose, shape, and valid grasps. We build upon the recently proposed CenterSnap [29] architecture, a single-shot 3D shape reconstruction and 6D pose estimation method, and adapt it to our grasp prediction setting. We adopt a point-based representation where we encode the complete 3D information of each object in the scene in the respective center point in the 2D image. At inference time, our image encoder (see Sec. III-A) detects the center point of each object as well as the respective 6D pose and latent shape vector. This information is then used by the proposed shape and grasps distance function (SGDF) decoder discussed in Sec. III-B to reconstruct both 3D shapes and valid grasps. Finally, as shown in Fig. 2, we transform the 3D shape of the object and the grasps back to the camera frame using our predicted poses.

The *object awareness* property is achieved thanks to the heatmap prediction, while the *holistic grasping* property is obtained from the full object shape and grasps reconstruction capabilities of the SGDF decoder. The SGDF decoder is trained first to learn a latent embedding space. After training, the optimized latent codes for each object in the dataset are saved, so that they can be used as labels for training the latent code prediction of the image encoder. In the following, we discuss both the encoder and decoder in detail as well as our data generation pipeline.

#### A. Image Encoder

*Model Architecture:* Following the CenterSnap [29] architecture, we generate a low resolution feature representation of both the RGB and depth images using a ResNet50 topology. We then concatenate both representations and feed them into a ResNet18-FPN backbone [40]. Finally, we use the resulting pyramid of features as input to three specialized heads: a heatmap head, a pose head, and a latent code head. All three network heads output a per-pixel prediction. The architecture of the three network heads follows the PanopticFPN [40] segmentation head. It takes FPN features as input and fuses information from all levels into a single output. Each feature level is upsampled to a common dimension and then summed elementwise. Each upsampling stage consists of a  $3 \times 3$  convolution, group norm, ReLU, and  $2 \times$  bilinear upsampling. The heatmap head predicts the probability of objectness for each pixel as a scalar value  $o \in [0, 1]$ . The pose head predicts the 6D pose of the object at the given pixel as a 12-dimensional vector, which is then reshaped

TABLE I: Overview of hyperparameters.

| (a) Decoder |         |                              | (b) Encoder |         |                            |
|-------------|---------|------------------------------|-------------|---------|----------------------------|
| Loss        | Scaling | Param. Value                 | Loss        | Scaling | Param. Value               |
| SDF         |         | $\delta_{\text{SDF}}$ 10.0   | Heatmap     |         | $\delta_{\text{Heat}}$ 100 |
| Grasp       |         | $\delta_{\text{Grasp}}$ 1.0  | Shape       |         | $\delta_{\text{Shape}}$ 1  |
| Code        |         | $\delta_{\text{Code}}$ 0.001 | Pose        |         | $\delta_{\text{Pose}}$ 5   |

into a  $3 \times 4$  matrix and projected to a valid homogeneous transformation via a Procrustes operation [41]. The code head predicts  $z \in \mathbb{R}^{D_l}$ , a  $D_l$ -dimensional vector representing the shape and possible grasps for the object at that pixel, denoted as latent code.

*Training:* Our heatmap loss  $\mathcal{L}_{\text{heat}}$  is the mean squared error between the heatmap labels and the network output. Inspired by [7], we represent a 6D pose as a set of four 3D points  $v \in \mathbb{R}^{4 \times 3}$ . This formulation allows us to define the pose loss as the simple Euclidean distance between points, unifying both translation and rotation information with a homogeneous metric. Finally, the latent code loss  $\mathcal{L}_{\text{shape}}$  is calculated as the L1 loss between prediction and ground truth code maps. Both the pose and latent code losses are masked by the ground truth object binary masks and weighted at each pixel by the ground truth value of the object heatmap. This ensures that no loss is computed on background pixels, and emphasizes higher weight towards the center of the objects. Our final, total loss is the weighted sum of these three losses, averaged over all pixels:

$$\mathcal{L}_{\text{encoder}} = \delta_{\text{heat}} \mathcal{L}_{\text{heat}} + \delta_{\text{pose}} \mathcal{L}_{\text{pose}} + \delta_{\text{shape}} \mathcal{L}_{\text{shape}}. \quad (1)$$

We report the scaling parameters in Tab. Ib. We train the network for 100 epochs. To improve generalization, we normalize the RGB values to the dataset mean and standard deviation, and we add color jitter augmentation to the RGB images. We use the ADAM optimizer with default parameters and a learning rate of  $1e - 3$ .

#### B. Shape and Grasps Decoder

Our proposed shape and grasps distance function (SGDF) decoder maps a 3D coordinate and a latent code to both a shape distance and a grasp distance. The shape distance is a scalar value representing the signed distance between the sample coordinate and the closest point on the object’s surface. The grasp distance is the full 6D pose of the closest valid grasp to the sample coordinate, expressed in a reference frame aligned with the object frame and centered at the sample coordinate itself.

*Model Architecture:* Our decoder network is a deep multi-layer perceptron based on DeepSDF [37]. The input to our decoder is a latent code  $z \in \mathbb{R}^{D_l}$  and a coordinate in space  $x \in \mathbb{R}^3$ . The decoder outputs a scalar value  $s \in \mathbb{R}$  representing the distance to the closest surface as well as a grasp parameterized through a 9D vector raw grasp output  $\tilde{g} \in \mathbb{R}^9$ . The vector consists of three sub-vectors, an offset for the translation part of the grasp  $\delta t \in \mathbb{R}^3$  as well as two rotational components  $r_1, r_2 \in \mathbb{R}^3$ . We build a true, absolute rotation matrix defined in the object frame  $R \in \mathbb{R}^{3 \times 3}$  by applying the Gram-Schmidt procedure on the two rotational components. Thus, the final, full grasp pose is defined by constructing a transformation matrix using both

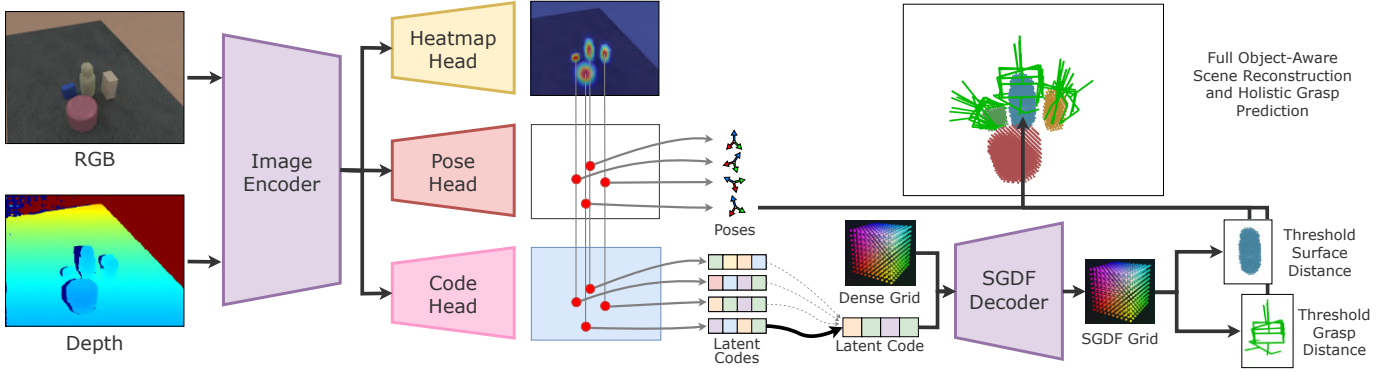


Fig. 2: Illustration of the CenterGrasp architecture. First, an RGB-D image is fed into the image encoder which outputs an object heatmap, a pose map, and a latent code map. Next, the object locations in the image are determined by extracting the peaks from the predicted heatmap. At these locations, each object pose and latent code is extracted accordingly. In the second step, the SGDF decoder infers the shape and grasps for each detected object. Finally, the object pose is used to transform the shape and grasp predictions from the canonical frame to the camera frame.




Fig. 3: Surface and Grasp Reconstruction. To highlight that CenterGrasp learns a continuous prior over the shape and grasp manifold, we randomly sample latent codes from the learned embedding space and reconstruct the surface as well as ten valid grasps for each object.

components  $\mathbf{g} = \mathbf{T}(\mathbf{x} + \delta\mathbf{t}, \mathbf{R}) \in SE(3)$  where the matrix is constructed through

$$\mathbf{T}(\mathbf{t}, \mathbf{R}) = \begin{bmatrix} \mathbf{R} & \mathbf{t} \\ \mathbf{0} & 1 \end{bmatrix}, \quad (2)$$

where  $\mathbf{t} \in \mathbb{R}^3$  and  $\mathbf{R} \in \mathbb{R}^{3 \times 3}$ .

*Training:* We train the decoder using multiple losses: first, we impose a loss on clamped SDF values

$$\mathcal{L}_{\text{SDF}} = \frac{\sum_{s_{\text{gt}}, s_{\text{pred}} \in \mathcal{S}} |\text{clamp}(s_{\text{gt}}, c) - \text{clamp}(s_{\text{pred}}, c)|}{|\mathcal{S}|}, \quad (3)$$

where  $\mathcal{S}$  is the set of all ground truth  $s_{\text{gt}}$  and predicted  $s_{\text{pred}}$  SDF values. The clamping function

$$\text{clamp}(x, c) = \min(\max(x, -c), c) \quad (4)$$

which clamps a value  $x$  to be between the specified range. In the experiments, we used  $c = 0.1$ .

Second, we use a 5-point grasp pose loss similar to [7]. We use five 3D points on the gripper defined in the gripper frame, make them homogeneous, and stack them. We denote the resulting matrix as  $\mathbf{v} \in \mathbb{R}^{4 \times 5}$ . To accommodate for symmetry we flip the points along the x-y-plane, resulting in  $\tilde{\mathbf{v}} \in \mathbb{R}^{4 \times 5}$ . We then calculate the grasping loss at each point as

$$\mathcal{L}_{\text{Grasp}} = \frac{\sum_{\mathbf{g}_{\text{gt}}, \mathbf{g}_{\text{pred}} \in \mathcal{G}} \min(\|\mathbf{g}_{\text{gt}}\mathbf{v} - \mathbf{g}_{\text{pred}}\mathbf{v}\|, \|\mathbf{g}_{\text{gt}}\tilde{\mathbf{v}} - \mathbf{g}_{\text{pred}}\tilde{\mathbf{v}}\|)}{n_g}, \quad (5)$$

where we multiply the ground truth grasp  $\mathbf{g}_{\text{gt}}$  with the original and flipped gripper point matrices  $\mathbf{v}$  and  $\tilde{\mathbf{v}}$ , and use the minimal distance between either of them. The result is averaged over




Fig. 4: Generated Synthetic Data. To generate our training data, we render a random scene, consisting of a floor, table, and between 1 to 5 objects to yield an RGB image (left), a depth image with simulated sensor noise (center), and an object heatmap (right).

the number of grasps in the batch  $n_g$ . Third, we regularize the latent codes through a loss  $\mathcal{L}_{\text{Code}}$  computed as

$$\mathcal{L}_{\text{Code}} = \bar{\mathbf{Z}} \cdot \min(1, e/5), \quad (6)$$

where  $\bar{\mathbf{Z}}$  is the average norm of all embeddings and  $e$  is the current training epoch.

We calculate our final loss as

$$\mathcal{L}_{\text{decoder}} = \delta_{\text{SDF}}\mathcal{L}_{\text{SDF}} + \delta_{\text{Grasp}}\mathcal{L}_{\text{Grasp}} + \delta_{\text{Code}}\mathcal{L}_{\text{Code}}. \quad (7)$$

We report the scaling parameters in Tab. Ia. Throughout all of our experiments, we set the latent code size  $D_l$  to 32.

We train the network for 100 epochs. We apply dropout with probability 0.2 and weight normalization to regularize training. We use the ADAM optimizer with default parameters and a learning rate of  $1e-3$ . Fig. 3 shows the reconstructed output of our SGDF function for different sampled latent codes.

### C. Inference

Given a single RGB-D observation of the scene, our image encoder (see Sec. III-A) predicts a corresponding object heatmap, pose map, and latent code map. For each heatmap peak, both the 6D pose and latent code are extracted. For each object, we input to the SGDF decoder a  $64 \times 64 \times 64$  dense grid of sampled 3D coordinates together with the predicted latent code. We then extract the object surface points and a manifold of valid grasps as the respective iso-surfaces  $SGDF(\cdot) = 0$ , i.e. by discarding all points and grasps above a small threshold  $\epsilon$ . The remaining points and grasps are then transformed from their canonical frame to the camera frame using the predicted 6D object pose.

To achieve higher precision, this pose is further refined via point-to-plane ICP to match the predicted shape with the

observed point cloud. The list of predicted grasps is then filtered to discard all grasps that are in collision with the observed scene. Of the remaining grasp poses, we select the one that minimizes the torque around the grasp point, estimated using the gravitational force acting on the centroid of the predicted shape. These poses are then transformed from the camera frame into the robot base frame and the grasp is executed via an inverse kinematics controller.

#### D. Data Generation

CenterGrasp is trained solely on synthetic data and achieves zero-shot generalization to the real world. For a fair comparison, we train on the same 99,945 training scenes from GIGA [8], which consist of a list of object meshes and their respective poses. They are constructed from a set of 417 meshes as either packed or piled clutter.

*Shape and Grasp Distance Labels:* To generate grasp candidates for each object mesh, we first sample 1,000 surface points from the object. For each point, we align the gripper fingers to the point normal. Finally, we sample 24 different wrist rotations uniformly, i.e. we consecutively apply a 15 degrees rotation around the z-axis of the gripper, for a total of 24,000 grasp candidates per object. From these candidates, we compute the ground truth grasp labels by checking for collisions and evaluating their antipodality. On average, we find around 8,500 valid grasps per object. Similarly to NGDF [38], we find that a large discrete grasp set is a good approximation of the continuous manifold of valid grasps. For each object, we use `mesh_to_sdf*` to sample 100,000 points and their respective SDF values. For each point, we then find the closest grasp from the set of ground truth grasps of the given object and compute the respective grasp distance label.

*Image Labels:* To generate synthetic image observations (RGB and depth) as well as pose and instance segmentation labels from the precomputed GIGA scenes, we use the raytracing-based renderer from SAPIEN [10] and its realistic depth feature [42]. All the textures, materials, lights, and table shapes are randomized, see Fig. 4 for an example observation. For training, we render each scene from two random camera poses, resulting in roughly 200,000 RGB-D images and labels. To generate the object heatmaps, whose peaks represent each object in the image, we fit a Gaussian to the ground truth masks. For the pose map and code map labels, we use the instance masks to label each pixel of an object with its ground truth pose vector (12-dimensional) and its ground truth code vector ( $D_t$ -dimensional) respectively.

## IV. EXPERIMENTAL RESULTS

In this section, we first evaluate the proposed CenterGrasp on the joint task of simultaneous shape reconstruction and 6-DoF grasp estimation. We present a comparison with the state-of-the-art baseline GIGA [8], as well as an ablation study of our CenterGrasp approach without the ICP pose refinement step described in Sec. III-C. It is important to note that the ICP refinement is only possible thanks to the *object awareness*

property of our approach: predicting separate shapes for each object allows us to refine their poses individually. This is not possible with GIGA, since its prediction is at the full scene level. Moreover, we evaluate the output grasp quality on the GraspNet-1Billion [43] benchmark that evaluates grasp proposals via an analytic model. Finally, we demonstrate zero-shot transfer to a real robot setup.

#### A. Simulation Experiments

We build our evaluation tasks on SAPIEN [10], a realistic and physics-rich simulator. Our scenes consist of a Franka Emika Panda arm, a table, and a number of objects placed on top of it. While GIGA [8] evaluated the grasping performance with a floating gripper that can be placed freely in space, we evaluate with a complete 7-DoF robot arm, which more closely mirrors a real-world scenario. To ensure a fair comparison with GIGA [8], we first use the same object meshes and randomized scenes from their test set. Additionally, we evaluate the out-of-distribution capability of all methods by testing on the YCB object set [44]. Following prior works [6], [8], we generate two different types of cluttered scenes: packed clutter, where objects are randomly placed upright next to each other, and pile clutter, where objects are randomly dropped on the table and lie on top of each other. From the combination of two object sets and two clutter types, our full evaluation consists of four different environments. Each baseline is evaluated on 200 random scenes from each environment. The number of objects in each scene is sampled from a Poisson distribution with an expected value equal to 4.

*Shape Reconstruction:* To evaluate the shape reconstruction quality, we use the ground truth meshes and their poses in the scene to extract a ground truth point cloud. We then compare it to both the point cloud predicted by CenterGrasp as well as to a point cloud sampled from GIGA’s predicted meshes. Fig. 5 shows a qualitative comparison of the reconstruction from both methods. GIGA yields adequate results in the in-distribution evaluation (i.e. using GIGA objects), but its reconstruction quality drastically decreases in the out-of-distribution settings. On the other hand, CenterGrasp demonstrates good reconstruction quality in all environments, including real-world evaluation.

In Tab. II, we report the quantitative results of our evaluation, which consists of two shape reconstruction metrics calculated between predictions and ground truth: the L2 bi-directional Chamfer Distance (CD), and the 3D Intersection over Union (IoU). Averaging across environments, CenterGrasp achieves an improvement of 38.5 mm in CD and 0.38 in the IoU score. This numerical evaluation is consistent with the qualitative results from Fig. 5, highlighting the better generalization capabilities of CenterGrasp compared to the state of the art. These results demonstrate the advantages of learning to predict geometries at the object level rather than at the full scene level. Furthermore, our ablation of CenterGrasp without ICP demonstrates the importance of pose refinement. The results show that CenterGrasp without ICP achieves similar results to GIGA when evaluated with YCB objects, and worse results when evaluated with GIGA objects.

*Grasp Pose Prediction:* We perform comprehensive evaluations of the grasping capabilities of CenterGrasp on our simulation

\*[https://github.com/marian42/mesh\\_to\\_sdf](https://github.com/marian42/mesh_to_sdf)

TABLE II: Evaluation of the shape and grasp-pose prediction in four different simulated environments. Results are averaged over three random seeds.  $CD = L2$ -Chamfer Distance ( $\downarrow$ ) in mm,  $IoU = 3D$  Intersection over Union ( $\uparrow$ ),  $SR = Success Rate$  ( $\uparrow$ ),  $DR = Declutter Rate$  ( $\uparrow$ ).

|             | GIGA Objects Packed |             |             |             | GIGA Objects Pile |             |             |             | YCB Objects Packed |             |             |             | YCB Objects Pile |             |             |             |
|-------------|---------------------|-------------|-------------|-------------|-------------------|-------------|-------------|-------------|--------------------|-------------|-------------|-------------|------------------|-------------|-------------|-------------|
|             | CD                  | IoU         | SR          | DR          | CD                | IoU         | SR          | DR          | CD                 | IoU         | SR          | DR          | CD               | IoU         | SR          | DR          |
| GIGA [8]    | 35.5                | 0.18        | 0.52        | 0.59        | 71.5              | 0.10        | 0.41        | 0.18        | 63.3               | 0.13        | 0.49        | 0.41        | 83.2             | 0.11        | 0.44        | 0.27        |
| CG w/o ICP  | 55.5                | 0.10        | 0.23        | 0.16        | 55.3              | 0.20        | 0.17        | 0.12        | 63.2               | 0.13        | 0.14        | 0.14        | 87.9             | 0.14        | 0.18        | 0.14        |
| CenterGrasp | <b>23.2</b>         | <b>0.34</b> | <b>0.84</b> | <b>0.87</b> | <b>16.1</b>       | <b>0.66</b> | <b>0.84</b> | <b>0.71</b> | <b>18.1</b>        | <b>0.61</b> | <b>0.83</b> | <b>0.81</b> | <b>42.0</b>      | <b>0.47</b> | <b>0.67</b> | <b>0.70</b> |

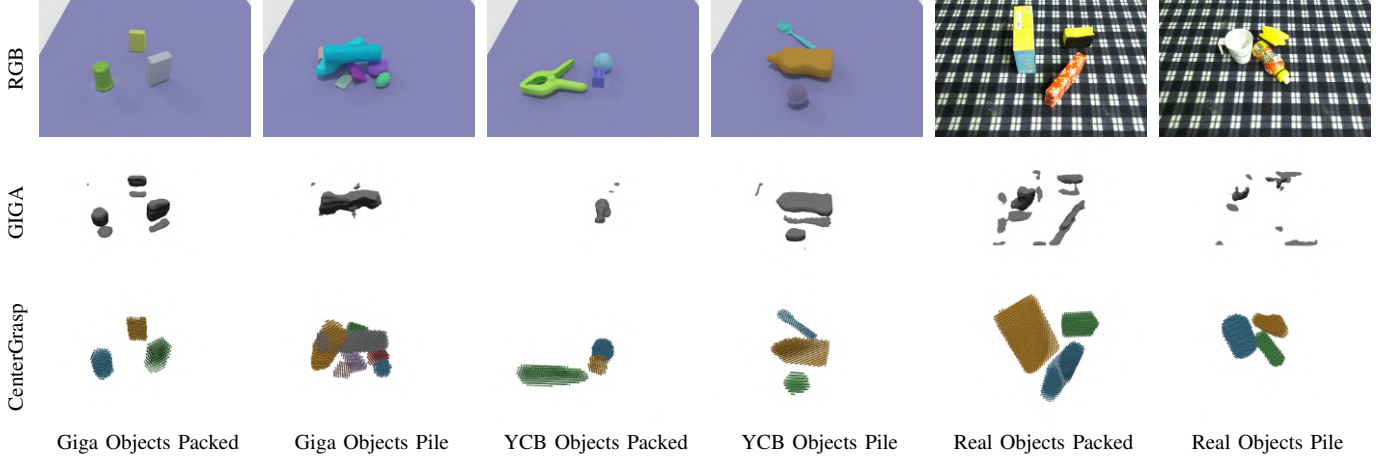


Fig. 5: A comparison between the mesh reconstruction of GIGA [8] and the point cloud reconstruction of CenterGrasp. GIGA yields adequate results in the in-distribution evaluation (i.e. using GIGA objects), but its reconstruction quality drastically decreases in the out-of-distribution settings. On the other hand, CenterGrasp demonstrates good reconstruction quality in all environments, including real-world evaluations.

setup. For each scene, the robot executes the predicted grasp, lifts the object, and moves it to a drop area. The attempt is considered successful if the objects are successfully moved to the drop area. The episode is terminated when all objects in the scene have been cleared or after three consecutive failed attempts. In Tab. II, we report two object grasping metrics: Success Rate (SR) and Declutter Rate (DR). Success rate is defined as the number of successful grasps divided by the total grasps attempted, whereas declutter rate is defined as the number of successful grasps divided by the total number of objects in the evaluation scenes. The latter metric is important to highlight cases where objects are still present in the scene, but no grasp is predicted by the model.

In this experiment, CenterGrasp achieves 33 percentage points higher SR and 41 percentage points higher DR compared to GIGA on average, demonstrating its ability to clear drastically more objects from the scene while requiring fewer attempts. Furthermore, our ablation study shows a significant decrease in the success rate when using CenterGrasp without ICP, which achieves significantly worse results than both CenterGrasp and GIGA. Compared to the full pipeline, CenterGrasp without ICP achieves over 60 percentage points less SR and DR, highlighting the importance of precise object poses for the success of the entire pipeline.

### B. GraspNet-1Billion benchmark

While the target task of our work is simultaneous shape reconstruction and grasp estimation, we additionally evaluate the grasp quality of our method on the GraspNet-1Billion benchmark [43], to allow a direct comparison with dedicated

grasp-pose estimation approaches. This benchmark evaluates the top 50 grasp proposals using an analytic model of grasp quality. Following [43], we report the Average Precision (AP) for two different friction coefficients ( $AP_{0.4}$  and  $AP_{0.8}$ ), as well as the total AP averaged across the full range of coefficients. We train and evaluate CenterGrasp on the Kinect camera data. The results are shown in Tab. III, where we compare against the baselines originally reported in the benchmark. For the seen objects case, CenterGrasp performs better than the two lowest scoring baseline [45], [46] and worse than the two best scoring methods [12], [43]. For the unseen and novel object cases, a similar trend is observable for the  $AP_{0.4}$  metric. We can attribute this underperformance to multiple factors. First, CenterGrasp is designed to detect objects in the scene, extract their poses and predict their full shape and grasp manifold, while the compared methods only focus on the grasp estimation task considered in this benchmark. Second, the benchmark reports the average metric over the best 50 grasps: since CenterGrasp does not directly predict a grasp score, ranking the poses in the grasp manifold is not straightforward. Third, we observe that the heatmap objectness prediction struggles in heavily cluttered scenes in the benchmark, leading to failure cases downstream. Fourth, the small amount of objects in the GraspNet-1B dataset (only 87), paired with a high variability in their shape, prevents the SGDF decoder to generalize well to unseen objects.

### C. Real Robot Experiments

We carry out evaluations of the grasp prediction quality on a real robot setup consisting of a 7-DoF Franka Emika Panda

TABLE III: Evaluation of 6-Dof grasp pose prediction on the GraspNet-1Billion benchmark [43]. We report AP for two friction coefficients (0.4, 0.8) as well as the total AP averaged across the full range of coefficients. *AP* = Average Precision ( $\uparrow$ ), *G* = Grasp, *R* = Reconstruction, *D* = Detection.

| Method           |   |   |   | Seen         |                   |                   | Unseen       |                   |                   | Novel        |                   |                   |
|------------------|---|---|---|--------------|-------------------|-------------------|--------------|-------------------|-------------------|--------------|-------------------|-------------------|
|                  | G | R | D | AP           | AP <sub>0.8</sub> | AP <sub>0.4</sub> | AP           | AP <sub>0.8</sub> | AP <sub>0.4</sub> | AP           | AP <sub>0.8</sub> | AP <sub>0.4</sub> |
| GG-CNN [45]      | ✓ | - | - | 16.89        | 22.47             | 11.23             | 15.05        | 19.76             | 6.19              | 7.38         | 8.78              | 1.32              |
| Chu et al. [46]  | ✓ | - | - | 17.59        | 24.67             | 12.74             | 17.36        | 21.64             | 8.86              | 8.04         | 9.34              | 1.76              |
| GPD [5]          | ✓ | - | - | 24.38        | 30.16             | 13.46             | 23.18        | 28.64             | 11.32             | 9.58         | 10.14             | 3.16              |
| PointnetGPD [12] | ✓ | - | - | 27.59        | 34.21             | 17.83             | 24.38        | 30.84             | 12.83             | 10.66        | 11.24             | 3.21              |
| GraspNet-1B [43] | ✓ | - | - | <b>29.88</b> | <b>36.19</b>      | <b>19.31</b>      | <b>27.84</b> | <b>33.19</b>      | <b>16.62</b>      | <b>11.51</b> | <b>12.92</b>      | <b>3.56</b>       |
| CenterGrasp      | ✓ | ✓ | ✓ | 22.35        | 27.16             | 15.35             | 13.78        | 17.01             | 8.22              | 5.72         | 6.93              | 2.64              |



Fig. 6: The simple (left) and non-convex (right) household objects used in our real-world experiments.

TABLE IV: Evaluation of 6-Dof grasp pose prediction in the real world. *SR* = Success Rate ( $\uparrow$ ), *DR* = Declutter Rate ( $\uparrow$ ), *G* = Grasp, *R* = Reconstruction, *D* = Detection.

| Method      |   |   |   | Packed      |             | Pile        |             | NonConvex   |             |
|-------------|---|---|---|-------------|-------------|-------------|-------------|-------------|-------------|
|             | G | R | D | SR          | DR          | SR          | DR          | SR          | DR          |
| GPD [5]     | ✓ | - | - | 0.48        | 0.79        | 0.41        | 0.67        | 0.38        | 0.59        |
| GIGA [8]    | ✓ | ✓ | - | 0.36        | 0.64        | 0.20        | 0.37        | 0.26        | 0.44        |
| CenterGrasp | ✓ | ✓ | ✓ | <b>0.57</b> | <b>0.85</b> | <b>0.53</b> | <b>0.83</b> | <b>0.49</b> | <b>0.70</b> |

arm equipped with a parallel jaw-gripper and a wrist-mounted ZED2 depth camera. We collect a set of 20 simple and a set of 6 non-convex household objects, shown in Fig. 6. We evaluate on three settings: packed and pile (similar to the simulation experiments) and non-convex. We build 30 different scenes with three objects each, for a total of 90 possible grasps for all three scenarios. To minimize variance, we replicate the scenes across all the methods as accurately as possible. For each scene, we capture an RGB-D image and query the model to predict a grasp. We then execute the predicted grasp and remove the object from the scene. If a grasp attempt fails two times for the same object, that object is removed manually from the scene. For each baseline and type of scene, we report the success rate (SR) and declutter rate (DR) in Tab. IV. Similar to the simulation results, we observe that CenterGrasp demonstrates the best performance across all the evaluations, achieving on average 27 percentage points higher SR and 20 percentage points higher DR compared to GIGA.

## V. DISCUSSION AND LIMITATIONS

The effective scene understanding of CenterGrasp not only enhances grasping capabilities but also unlocks various additional opportunities. For instance, the concept of *object awareness* enables human interaction, allowing users to select desired objects by utilizing object masks extracted from the heatmap, as shown in Fig. 8. Additionally, the ability of CenterGrasp to perform *holistic grasping* ensures accurate detection, reconstruction, and successful grasping of objects, even when they are partially (self-) occluded (Fig. 7), or behind



Fig. 7: Object-aware and holistic scene understanding. Given the input image (left), CenterGrasp is able to predict a full per-object reconstruction even in the presence of partial occlusions (center-right).



Fig. 8: The *object awareness* property allows human interaction. User can select the desired object to be grasped.




Fig. 9: The *holistic grasping* property allows CenterGrasp to predict grasps even behind the visible region of the camera (viewing from left to right).

the visible region of the camera (Fig. 9). For example, with the object in Fig. 9, we predict a total of 1267 grasps, of which 61% are on the visible region and 39% on the invisible region of the object. Although our approach demonstrates significant performance improvements compared to the state-of-the-art baseline, it does have some limitations. First, precise pose predictions are critical to the success of the pipeline, which in practice implies the adoption of an ICP refinement step. Second, two separate models need to be trained, the image encoder, which operates at the scene level, and the SGDF decoder, which operates at the object level. Both aspects represent an exciting opportunity for future research.

## VI. CONCLUSION

In this work, we presented CenterGrasp, a simultaneous shape reconstruction and grasp-pose prediction architecture. We designed our framework around two key concepts: *object awareness* and *holistic grasping*, with the aim to combine 3D scene understanding with object grasping. We compared the proposed approach with a recent state-of-the-art method and showed that it achieves exceptional performance both in simulation and real-world robotic experiments. By combining accurate perception and successful grasping, our framework paves the way for dexterous robotic systems capable of effectively interacting with objects in complex real-world environments. We made the code and trained models publicly available to facilitate future research.

## REFERENCES

- [1] A. Bicchi and V. Kumar, "Robotic grasping and contact: a review," in *Int. Conf. on Robotics and Automation*, 2000, pp. 348–353.
- [2] J. Bohg, A. Morales, T. Asfour, and D. Kragic, "Data-driven grasp synthesis—a survey," *IEEE Transactions on Robotics*, vol. 30, no. 2, pp. 289–309, 2014.
- [3] 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 *Robotics: Science and Systems*, 2017.
- [4] D. Morrison, P. Corke, and J. Leitner, "Learning robust, real-time, reactive robotic grasping," *The International Journal of Robotics Research*, vol. 39, no. 2-3, pp. 183–201, 2020.
- [5] A. Ten Pas, M. Gualtieri, K. Saenko, and R. Platt, "Grasp pose detection in point clouds," *Int. Journal of Robotics Research*, vol. 36, no. 13-14, pp. 1455–1473, 2017.
- [6] M. Breyer, J. J. Chung, L. Ott, S. Roland, and N. Juan, "Volumetric grasping network: Real-time 6 dof grasp detection in clutter," in *Conference on Robot Learning*, 2020.
- [7] M. Sundermeyer, A. Mousavian, R. Triebel, and D. Fox, "Contact-graspnet: Efficient 6-dof grasp generation in cluttered scenes," in *Int. Conf. on Robotics and Automation*, 2021.
- [8] Z. Jiang, Y. Zhu, M. Svetlik, K. Fang, and Y. Zhu, "Synergies between affordance and geometry: 6-dof grasp detection via implicit representations," *Robotics: Science and Systems*, 2021.
- [9] A. Simeonov, Y. Du, A. Tagliasacchi, J. B. Tenenbaum, A. Rodriguez, P. Agrawal, and V. Sitzmann, "Neural descriptor fields:  $Se(3)$ -equivariant object representations for manipulation," in *Int. Conf. on Robotics and Automation*, 2022.
- [10] F. Xiang, Y. Qin, K. Mo, Y. Xia, H. Zhu, F. Liu, M. Liu, H. Jiang, Y. Yuan, H. Wang, L. Yi, A. X. Chang, L. J. Guibas, and H. Su, "SAPIEN: A simulated part-based interactive environment," in *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition*, 2020.
- [11] R. Newbury, M. Gu, L. Chumbley, A. Mousavian, C. Eppner, J. Leitner, J. Bohg, A. Morales, T. Asfour, D. Kragic, D. Fox, and A. Cosgun, "Deep learning approaches to grasp synthesis: A review," *IEEE Transactions on Robotics*, vol. 39, no. 5, pp. 3994–4015, 2023.
- [12] 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 *Int. Conf. on Robotics and Automation*, 2019, pp. 3629–3635.
- [13] A. Mousavian, C. Eppner, and D. Fox, "6-dof graspnet: Variational grasp generation for object manipulation," in *Int. Conf. on Computer Vision*, 2019, pp. 2901–2910.
- [14] M. Gou, H.-S. Fang, Z. Zhu, S. Xu, C. Wang, and C. Lu, "Rgb matters: Learning 7-dof grasp poses on monocular rgb-d images," in *Int. Conf. on Robotics and Automation*. IEEE, 2021, pp. 13 459–13 466.
- [15] 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 *Int. Conf. on Computer Vision*, 2021, pp. 15 964–15 973.
- [16] H.-S. Fang, C. Wang, H. Fang, M. Gou, J. Liu, H. Yan, W. Liu, Y. Xie, and C. Lu, "Anygrasp: Robust and efficient grasp perception in spatial and temporal domains," 2023.
- [17] D.-C. Hoang, J. A. Stork, and T. Stoyanov, "Context-aware grasp generation in cluttered scenes," in *Int. Conf. on Robotics and Automation*. IEEE, 2022, pp. 1492–1498.
- [18] D.-C. Hoang, A.-N. Nguyen, V.-D. Vu, D.-Q. Vu, V.-T. Nguyen, T.-U. Nguyen, C.-T. Tran, K.-T. Phan, and N.-T. Ho, "Grasp configuration synthesis from 3d point clouds with attention mechanism," *Journal of Intelligent & Robotic Systems*, vol. 109, no. 3, p. 71, 2023.
- [19] X. Liu, Y. Zhang, H. Cao, D. Shan, and J. Zhao, "Joint segmentation and grasp pose detection with multi-modal feature fusion network," in *Int. Conf. on Robotics and Automation*, 2023, pp. 1751–1756.
- [20] S. Chen, W. Tang, P. Xie, W. Yang, and G. Wang, "Efficient heatmap-guided 6-dof grasp detection in cluttered scenes," *IEEE Robotics and Automation Letters*, 2023.
- [21] R. Qin, H. Ma, B. Gao, and D. Huang, "Rgb-d grasp detection via depth guided learning with cross-modal attention," in *Int. Conf. on Robotics and Automation*, 2023, pp. 8003–8009.
- [22] C. Xie, Y. Xiang, A. Mousavian, and D. Fox, "Unseen object instance segmentation for robotic environments," *IEEE Transactions on Robotics*, 2021.
- [23] B. Curless and M. Levoy, "A volumetric method for building complex models from range images," in *Proc. of the Annual Conf. on Computer Graphics and Interactive Techniques*, 1996, pp. 303–312.
- [24] D. Honerkamp, T. Welschehold, and A. Valada, " $N^2m^2$ : Learning navigation for arbitrary mobile manipulation motions in unseen and dynamic environments," *IEEE Transactions on Robotics*, 2023.
- [25] F. Schmalstieg, D. Honerkamp, T. Welschehold, and A. Valada, "Learning long-horizon robot exploration strategies for multi-object search in continuous action spaces," in *The International Symposium of Robotics Research*, 2022, pp. 52–66.
- [26] X. Zhou, D. Wang, and P. Krähenbühl, "Objects as Points," *arXiv preprint arXiv:1904.07850*, 2019.
- [27] K. Duan, S. Bai, L. Xie, H. Qi, Q. Huang, and Q. Tian, "CenterNet: Keypoint Triplets for Object Detection," in *Int. Conf. on Computer Vision*, 2019, pp. 6568–6577.
- [28] N. Heppert, T. Migimatsu, B. Yi, C. Chen, and J. Bohg, "Category-Independent Articulated Object Tracking with Factor Graphs," in *Int. Conf. on Intelligent Robots and Systems*, 2022, pp. 3800–3807.
- [29] 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," in *Int. Conf. on Robotics and Automation*, 2022, pp. 10 632–10 640.
- [30] 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," in *Europ. Conf. on Computer Vision*, 2022, pp. 275–292.
- [31] 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*, 2024.
- [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. of the IEEE Conf. on Computer Vision and Pattern Recognition*, 2023, pp. 21 201–21 210.
- [33] L. Manuelli, W. Gao, P. Florence, and R. Tedrake, *KPAM: KeyPoint Affordances for Category-Level Robotic Manipulation*, 2022, pp. 132–157.
- [34] B. Wen, W. Lian, K. Bekris, and S. Schaal, "Catgrasp: Learning category-level task-relevant grasping in clutter from simulation," in *Int. Conf. on Robotics and Automation*, 2022, pp. 6401–6408.
- [35] L. Mescheder, M. Oechsle, M. Niemeyer, S. Nowozin, and A. Geiger, "Occupancy networks: Learning 3d reconstruction in function space," in *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition*, 2019, pp. 4460–4470.
- [36] J. O. von Hartz, E. Chisari, T. Welschehold, W. Burgard, J. Boedecker, and A. Valada, "The treachery of images: Bayesian scene keypoints for deep policy learning in robotic manipulation," *IEEE Robotics and Automation Letters*, vol. 8, no. 11, pp. 6931–6938, 2023.
- [37] J. J. Park, P. Florence, J. Straub, R. Newcombe, and S. Lovegrove, "Deepsdf: Learning continuous signed distance functions for shape representation," in *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition*, 2019.
- [38] T. Weng, D. Held, F. Meier, and M. Mukadam, "Neural Grasp Distance Fields for Robot Manipulation," in *Int. Conf. on Robotics and Automation*, 2023, pp. 1814–1821.
- [39] S. Agrawal, N. Chavan-Dafle, I. Kasahara, S. Engin, J. Huh, and V. Isler, "Real-time simultaneous multi-object 3d shape reconstruction, 6dof pose estimation and dense grasp prediction," in *Int. Conf. on Intelligent Robots and Systems*, 2023.
- [40] A. Kirillov, R. Girshick, K. He, and P. Dollár, "Panoptic feature pyramid networks," in *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition*, 2019, pp. 6399–6408.
- [41] R. Bréger, "Deep regression on manifolds: a 3D rotation case study," in *International Conference on 3D Vision*, 2021.
- [42] X. Zhang, R. Chen, A. Li, F. Xiang, Y. Qin, J. Gu, Z. Ling, M. Liu, P. Zeng, S. Han, Z. Huang, T. Mu, J. Xu, and H. Su, "Close the optical sensing domain gap by physics-grounded active stereo sensor simulation," *IEEE Transactions on Robotics*, pp. 1–19, 2023.
- [43] H.-S. Fang, C. Wang, M. Gou, and C. Lu, "Graspnet-1billion: A large-scale benchmark for general object grasping," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 11 444–11 453.
- [44] B. Calli, A. Walsman, A. Singh, S. Srinivasa, P. Abbeel, and A. M. Dollar, "Benchmarking in manipulation research: Using the yale-cmu-berkeley object and model set," *IEEE Robotics & Automation Magazine*, vol. 22, no. 3, pp. 36–52, 2015.
- [45] D. Morrison, P. Corke, and J. Leitner, "Closing the loop for robotic grasping: A real-time, generative grasp synthesis approach," *Robotics: Science and Systems XIV*, pp. 1–10, 2018.
- [46] F.-J. Chu, R. Xu, and P. A. Vela, "Real-world multiobject, multigrasp detection," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 3355–3362, 2018.