

# D-SCo: Dual-Stream Conditional Diffusion for Monocular Hand-Held Object Reconstruction

Bowen Fu<sup>1,2</sup>, Gu Wang<sup>1</sup>, Chenyangguang Zhang<sup>1</sup>, Yan Di<sup>2</sup>,  
Ziqin Huang<sup>1</sup>, Zhiying Leng<sup>2</sup>, Fabian Manhardt<sup>3</sup>, Xiangyang Ji<sup>1\*</sup>  
and Federico Tombari<sup>2,3\*</sup>

<sup>1</sup>Tsinghua University    <sup>2</sup>Technical University of Munich    <sup>3</sup>Google  
{fbw19, zcyg22, huang-zq20}@mails.tsinghua.edu.cn,  
{yan.di@, zhiying.leng@, tombari@in.}tum.de,  
fabianmanhardt@google.com, {wanggu1, xyji}@tsinghua.edu.cn

**Abstract.** Reconstructing hand-held objects from a single RGB image is a challenging task in computer vision. In contrast to prior works that utilize deterministic modeling paradigms, we employ a point cloud denoising diffusion model to account for the probabilistic nature of this problem. In the core, we introduce centroid-fixed Dual-Stream Conditional diffusion for monocular hand-held object reconstruction (**D-SCo**), tackling two predominant challenges. First, to avoid the object centroid from deviating, we utilize a novel hand-constrained centroid fixing paradigm, enhancing the stability of diffusion and reverse processes and the precision of feature projection. Second, we introduce a dual-stream denoiser to semantically and geometrically model hand-object interactions with a novel unified hand-object semantic embedding, enhancing the reconstruction performance of the hand-occluded region of the object. Experiments on the synthetic ObMan dataset and three real-world datasets HO3D, MOW and DexYCB demonstrate that our approach can surpass all other state-of-the-art methods.

**Keywords:** Point Cloud Diffusion · Hand-Held Object Reconstruction

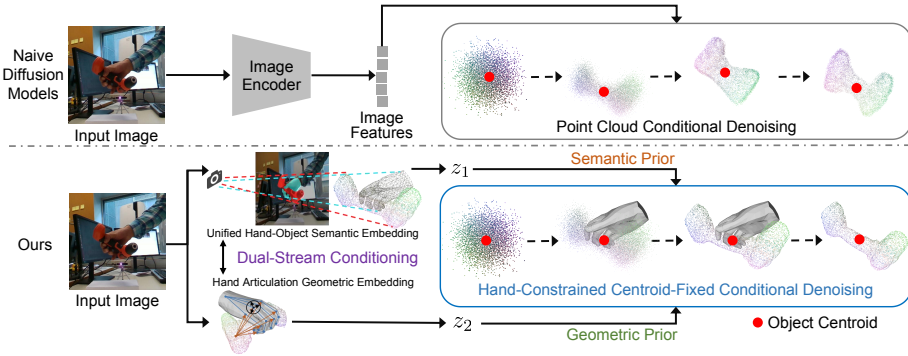
## 1 Introduction

Modeling hand-object interaction has a wide range of applications across various domains, including AR/VR [38, 55, 56] and human-robot interaction [20, 23, 50, 77]. Significant progress has been recently made for this task [4, 38, 77], where monocular hand-held object reconstruction draws particular attention [31, 37, 78]. Reconstructing hand-held objects from a single RGB image is a highly challenging and ill-posed problem. Suffering from the lack of real-world data and the ambiguity caused by hand- and self-occlusion, the performance of single-view hand-held object reconstruction remains limited [31, 78, 80].

In essence, most existing works either rely on Signed Distance Fields (SDFs) [37, 78] or Directed Distance Fields (DDFs) [80] to represent object shapes. Despite

---

\* corresponding author



**Fig. 1:** Comparison between D-SCo and naive diffusion models for hand-held object reconstruction. **Naive diffusion models** are conditioned only on image features without controlling object centroid deviation or modeling the uncertainty induced by hand occlusion. **D-SCo**, however, keeps the object centroid fixed under the constraint of the hand, making the diffusion model focus on shape reconstruction, and utilizes a dual-stream architecture to individually process semantic and geometric priors to learn a suitable representation for their own domain, tackling the aforementioned problems.

being effective under favorable conditions, such techniques generally tend to result in over-smoothed and undetailed reconstruction [11, 52]. Moreover, prior work [31, 37, 78, 80] typically utilize a deterministic modeling paradigm, making it difficult to reason about the uncertainties introduced by hand- or self-occlusion. Recently, probabilistic point cloud denoising diffusion models [42, 44, 83] have shown to be effective for the task of single-view object reconstruction. Compared to works employing surface-based representations like SDFs and DDFs, diffusion-driven methods for reconstructing point clouds enjoy better capabilities in overcoming artifacts such as noise, sparsity, and irregularities [11]. Moreover, the probabilistic nature of denoising diffusion models is particularly beneficial for modeling uncertainties and underconstrained problems.

Thus, in this work, we propose to leverage these probabilistic point cloud denoising diffusion models to conduct hand-held object reconstruction from a single RGB image. However, directly employing diffusion models in single-view hand-held object reconstruction faces two main problems: **First**, in existing diffusion models [43, 74, 75], the centroid of a partially denoised point cloud is not controlled and can thus deviate, for example, to the back side of the hand or even intersect with the hand, leading to physically implausible results. Additionally, the centroid deviation causes the misalignment of the semantic features [16] and can thus have adverse effects on the object reconstruction quality. **Second**, diffusion models for 3D reconstruction are typically conditioned only on single-stream 2D image features [44, 74], without modeling geometric hand-object interaction or addressing the uncertainty induced by hand occlusion.

To solve the aforementioned problems, we present D-SCo, a centroid-fixed dual-stream conditional point cloud denoising diffusion model for single-view hand-held object reconstruction. As shown in Fig. 1, we compare our D-SCo with naive diffusion models. First, a hand-constrained centroid fixing scheme is proposed to ensure that the centroid of the partially denoised point cloud does not diverge during the diffusion as well as the reverse process. In particular, we use a small and efficient neural network to estimate the hand-constrained object centroid and use it as a guide during the reverse process. Hence, the model only needs to consider the simpler shape diffusion task rather than having to simultaneously account for shape and position. Second, to best leverage the hand-object interaction prior, we introduce a dual-stream denoiser, which individually processes semantic and geometric priors. In particular, we utilize a unified hand-object semantic embedding to compensate for hand occlusion in the 2D feature extraction stage. Our contributions can be summarized as follows.

- We present D-SCo, the first conditional point cloud diffusion model for 3D reconstruction of hand-held objects from a single RGB image.
- We introduce a novel hand-constrained centroid fixing scheme, utilizing the hand vertices prior to prevent the centroid of the partially denoised point cloud from diverging during diffusion and reverse processes.
- We propose a dual-stream denoiser to semantically and geometrically model hand-object interactions. Our novel unified hand-object semantic embedding serves as a strong prior for reconstructing occluded regions of the object.

Experiments on the synthetic ObMan [31] dataset and three real-world datasets, *i.e.*, HO3D [28], MOW [4] and DexYCB [6], demonstrate that D-SCo can surpass all existing methods by a large margin.

## 2 Related Work

**Single-view Object Reconstruction.** Object reconstruction is a long-standing problem in the computer vision community. Traditionally, varieties of works focus on reconstruction from multi-view geometric techniques [29, 62–64]. Recently, advanced by the powerful representation capacity of deep learning, single-view-based methods for object reconstruction have shown promising results [12, 74, 75], despite the highly ill-posed setting. Early learning-based methods propose to learn category-specific networks [35, 36, 70], while more recent works try to learn a generalizable model across multiple categories for either meshes [25, 27], voxels [12, 24, 73], point clouds [21, 39], or implicit representations such as NeRFs [45] and SDFs [10, 52]. In this work, we focus on the difficult yet very important problem of learning hand-held object reconstruction from a single RGB image [31], putting a particular emphasis on modeling the impact of the hand occlusion.

**Single-View Hand-Held Object Reconstruction.** Reconstructing hand-held objects is very challenging due to intricate hand- and self-occlusions. Earlier works opt for the simplified task of 6DoF object pose estimation [28, 30, 67, 69, 72], assuming known object templates. To better utilize hand-object interactions,

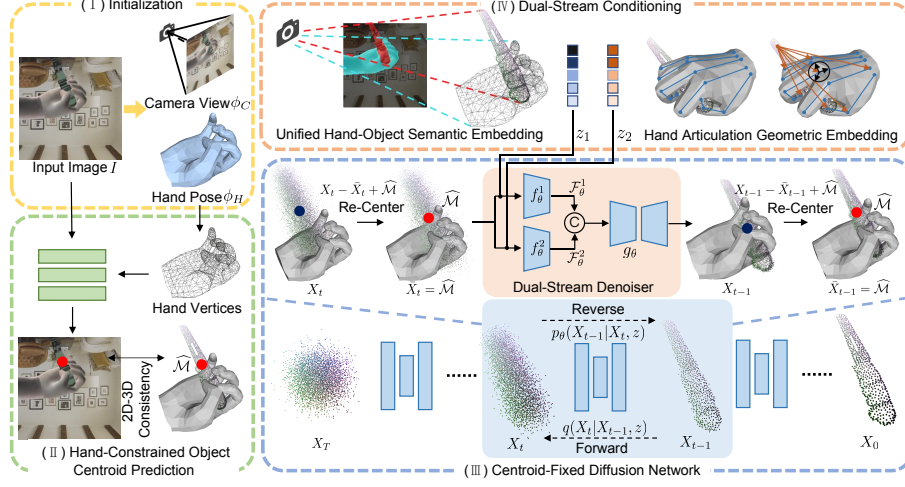
some works jointly reason about hand and object poses by means of implicit feature fusion [7, 40, 65, 69], utilizing explicit geometric constraints [2, 4, 13, 26], or enforcing of physical constraints [53, 71]. Increasing attention has recently been paid to model-free hand-held object reconstruction, as it is a more applicable setting. Exemplarily, while [31] explicitly predicts the object mesh, [8, 9, 37, 80] leverage implicit 3D representations such as SDFs, DDFs for reconstructing the shape of the hand and object. Recently, the concurrent work [81] instead utilizes NeRF as a representation. However, all these methods follow a deterministic approach, oftentimes resulting in low-quality reconstruction results for occluded and invisible parts, especially in the ill-posed single-view setting. We instead draw inspiration from recent advances in probabilistic 3D generation [16, 33, 44, 83] and additionally leverage semantic and geometric hand-object priors to achieve object reconstructions with high fidelity.

**Diffusion Models for 3D Reconstruction.** Denoising diffusion models (DDMs) [33, 49] have recently attracted increasing attention in 3D reconstruction. While Luo *et al.* [43] and LION [79] employ latent variables for the diffusion process in point cloud generation, PVD [83] directly applies DDMs to point clouds, leading to a unified framework for both unconditional 3D generation and conditional completion of partial shapes. Built upon PVD, PC<sup>2</sup> [44] conducts single-view reconstruction by conditioning on projected image features. In this work, we employ point cloud diffusion models [44, 83] for hand-held object reconstruction, as we observed that these models are more robust towards producing fragmented or distorted surfaces under ambiguous views than SDF-driven approaches [78].

**Hand Pose estimation.** Hand pose estimation methods from RGB(-D) images can be primarily categorized into model-free and model-based methods. Model-free approaches typically detect 2D keypoints, which are then lifted to 3D joint positions [34, 47, 48, 51, 58, 59, 86], whereas model-based approaches commonly exploit statistical models such as MANO [60] to work in a low-dimensional parameter spaces [1, 61, 66, 82, 85]. Compared with model-free methods, the latter line of works exhibits better robustness to occlusions as well as domain discrepancies [78]. Consequently, in this work, we rely on an off-the-shelf model-based approach [31, 61] to obtain the hand poses, which are then further leveraged within object shape inference.

### 3 Method

Given a single RGB image  $I$  capturing a hand-held object, we aim at reconstructing its shape as a 3D point cloud. As shown in Fig. 2, we first utilize off-the-shelf methods [31, 61] to predict hand parameters  $\phi_H$  and camera view  $\phi_C$ . Thereby,  $\phi_H$  is defined with respect to the MANO [60] hand model, having 45DoF joint parameters, and  $\phi_C$  represents the 6DoF pose of the hand wrist in the world reference system. Given the estimated hand vertices as obtained from the hand parameters, we employ a small yet efficient network to predict the object centroid  $\mathcal{M}$ . Eventually, we leverage the estimated centroid as part of the point cloud diffusion process to enable robust object point cloud reconstruc-



**Fig. 2: Architecture of D-SCo.** (I) Given a single-view RGB image, we first predict the hand pose  $\phi_H$  and camera view  $\phi_C$  by an off-the-shelf network. (II) The object centroid  $\hat{\mathcal{M}}$  is then estimated by our simple yet efficient hand-constrained centroid prediction network. (III) We further introduce a centroid-fixed diffusion network, which always keeps the centroid of partially denoised point cloud fixed at the predicted centroid  $\hat{\mathcal{M}}$  during the reverse process. (IV) A dual-stream denoiser is proposed to individually process and then aggregate semantic and geometric hand-object interaction priors as condition. A unified hand-object semantic embedding is introduced to serve as a strong prior of hand-occlusion.

tion. Note that in contrast to existing diffusion methods, our model can thus fully focus on shape reconstruction, leading to improved performance. Furthermore, we introduce a dual-stream denoiser to first independently process and then aggregate semantic and geometric priors with a novel unified hand-object embedding, which helps the reconstruction of the hand-occluded regions of the object. Due to the probabilistic nature of diffusion models together with our explicit modeling of hand-object interaction, D-SCo shows superior performance in handling uncertainties arising from hand- and self-occlusion.

### 3.1 Conditional Point Cloud Denoising Diffusion

We formulate single-view hand-held object reconstruction as conditional point cloud denoising diffusion, consisting of two Markov chains called the diffusion process and the reverse diffusion process. Suppose that we have a target point cloud with  $N$  points  $X_0 \in \mathbb{R}^{3N}$  from the conditional distribution  $q(X|z)$ , where  $z = I_{\phi_C, \phi_H}$  is the input RGB image with the corresponding camera view  $\phi_C$  and the hand pose  $\phi_H$ . For the diffusion process, we gradually add Gaussian noise

to the target point cloud at different levels  $t \in \{1, \dots, T\}$  as

$$q(X_t|X_{t-1}, z) = \mathcal{N}(X_t|z; \sqrt{1 - \beta_t}X_{t-1}|z, \beta_t\mathbf{I}). \quad (1)$$

Notice that with a fixed variance schedule  $\{\beta_t\}_{t=0}^T$ ,  $X_t$  can be simply expressed by  $X_0$  according to

$$q(X_t|X_0, z) = \mathcal{N}(X_t|z; \sqrt{\bar{\alpha}_t}X_0|z, (1 - \bar{\alpha}_t)\mathbf{I}), \quad (2)$$

where  $\alpha_t = 1 - \beta_t$ ,  $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$ . Therefore, we can reparameterize  $X_t$  as a linear combination of  $X_0$  and a noise variable  $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$  as follows

$$X_t = \sqrt{\bar{\alpha}_t}X_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon. \quad (3)$$

Starting with a point cloud sample  $X_T$  from random Gaussian noise, the reverse process iteratively samples from  $q(X_{t-1}|X_t, z)$  to remove the added noise from the diffusion process. To approximate this reverse process, we train a point cloud conditional denoiser  $\mathcal{D}_\theta(X_t, t, z)$  to learn the distribution  $q(X_{t-1}|X_t, z)$  by

$$\begin{aligned} p_\theta(X_{t-1}|X_t, z) &= \mathcal{N}(X_{t-1}; \mu_\theta(X_t, t, z), \sigma_t^2\mathbf{I}), \\ \mu_\theta(X_t, t, z) &= \frac{1}{\sqrt{\alpha_t}}(X_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}}\epsilon_\theta(X_t, t, z)), \end{aligned} \quad (4)$$

where  $\mu_\theta$  is the estimated mean.

### 3.2 Hand-Constrained Centroid-Fixed Conditional Diffusion

Existing works [78, 80] commonly obtain the camera view  $\phi_C$  and the hand pose  $\phi_H$  from standard hand pose estimation methods [31, 61]. Without knowing the object pose during inference, we can only attempt to denoise the points in the hand wrist coordinate system. However, directly using a diffusion model to jointly learn the object’s shape and pose can become very challenging [16]. Thus, we propose a hand-constrained centroid fixing scheme to ease learning.

**Centroid-Fixed Diffusion.** We first reduce the problem of object pose estimation to centroid prediction, which defines a new object coordinate system. Essentially, while the origin of the object coordinate system is located at the centroid of the object, the axis orientation is shared with the hand wrist frame. Notice that during training, we can directly use the ground-truth object centroid  $\mathcal{M}$  to constrain the object point cloud. Therefore, we harness  $\mathcal{M}$  to stabilize the diffusion process and make the point cloud fixed at  $\mathcal{M}$ . In particular, during the centroid-fixed diffusion, we re-center the object point cloud to  $\mathcal{M}$  and guarantee that the added noise has zero-mean via

$$\begin{aligned} X_0 &\sim q(X_0), X_0 \leftarrow X_0 - \bar{X}_0 + \mathcal{M}, \\ \epsilon &\sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \epsilon \leftarrow \epsilon - \bar{\epsilon}, \end{aligned} \quad (5)$$

where  $\bar{X}_0$  denotes the centroid of  $X_0$ . Noteworthy, by keeping the object centroid unchanged, the misalignment error within the semantic feature projection due

to centroid movements could also be alleviated. Thus, the training behavior becomes more stable as well.

**Centroid Prediction for Reverse Diffusion.** During the reverse diffusion process, we propose the use of a simple yet effective network  $\mathcal{G}$  to estimate the translation of the object w.r.t. the hand wrist coordinate system. Given the input RGB image  $I$  along with the corresponding hand pose  $\phi_H$  and camera view  $\phi_C$ , the predicted object centroid is obtained as

$$\widehat{\mathcal{M}} = \mathcal{G}(I_{\phi_C, \phi_H}). \quad (6)$$

As shown in Fig. 2 (II),  $\mathcal{G}$  first encodes the hand vertices  $X_H$  using a PointNet-like [54] model to constrain the object centroid in 3D space. Subsequently, the global hand features are combined with the image features, extracted by a pre-trained ResNet-18 [32], and processed by two parallel Multilayer Perceptrons (MLPs) to respectively output the 3D and 2D object centroid.

During the reverse process, we instead start from  $X_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ ,  $X_T \leftarrow X_T - \bar{X}_T + \widehat{\mathcal{M}}$ . We then re-center the predicted noise and restrict the centroid of the denoised point cloud at  $\widehat{\mathcal{M}}$  to always remain locked at each step  $t \in \{T-1, \dots, 0\}$  according to

$$\begin{aligned} \epsilon_\theta &\leftarrow \epsilon_\theta - \bar{\epsilon}_\theta, \\ X_t &\sim p_\theta(X_t | X_{t+1}, z), X_t \leftarrow X_t - \bar{X}_t + \widehat{\mathcal{M}}. \end{aligned} \quad (7)$$

After the denoising process, we can directly obtain the reconstructed object point cloud by taking  $X_0$ .

Noteworthy, estimating object orientation is more challenging than predicting only the translation and will significantly increase the complexity of the network. Therefore, we only constrain the object centroid since previous works [41, 44] have shown diffusion models are able to reconstruct the object well even if it is not in a canonical orientation.

### 3.3 Dual-Stream Conditional Point Cloud Denoising

The essence of the conditional point cloud denoiser  $\mathcal{D}_\theta(X_t, t, z)$  is the modeling of the condition  $z$  so to fully utilize the information of the input image along with the corresponding camera view and hand pose. Our key insight is that the hand can offer both semantic and geometric priors to facilitate object reconstruction. Therefore, we utilize a dual-stream architecture to individually process these priors to avoid mutual interference.

**Unified Hand-Object Semantic Embedding.** Related work has shown that using 2D image features is a vital component for object reconstruction [74, 75]. Unlike prior works that extract a global embedding for the image [12, 74, 75], we instead extract unique deep image features for each point of the partially denoised point cloud at each diffusion step in a projective manner. Specifically, we first extract image features  $\mathcal{F} \in \mathbb{R}^{H \times W \times C}$  using a standard 2D network, such as ResNet [32] or ViT [18], with  $C$  being the number of output feature

channels. We then project the point cloud onto the image by the efficient point cloud rasterizer [57]  $\mathcal{R}(\phi_C, X_t)$ . Subsequently, the pointwise semantic features at timestep  $t$  can be represented as  $X_t^O = \pi(\mathcal{R}(\phi_C, X_t), \mathcal{F})$  with  $X_t^O \in \mathbb{R}^{N \times C}$ , where  $\pi(\cdot)$  denotes the back-projection to 3D space. In this way, every individual point is pixel-aligned with the deep features corresponding to the pixel onto which the point is rasterized.

However, due to the inevitable occlusion induced by the hand, it is very challenging to reconstruct the occluded part of the object as only a single view is provided. On the other side, the hand information should not be fully ignored as the hand naturally serves as a prior for estimating the shape of the hand-occluded object region. Therefore, we apply  $\pi(\mathcal{R}(\cdot))$  to both object and hand points to obtain  $X_t^{HO} = \pi(\mathcal{R}(\phi_C, [X_t, X_H]), \mathcal{F})$  with  $X_t^{HO} \in \mathbb{R}^{(N+N_H) \times C}$ , serving as a unified hand-object semantic embedding. Thereby,  $N_H$  denotes the number of the hand mesh vertices  $X_H$ , which can be obtained from the estimated hand pose  $\phi_H$  together with the MANO hand model [60]. Noteworthy, we also apply an extra one-hot encoding to indicate whether the points belong to the object or the hand, making  $X_t^{HO} \in \mathbb{R}^{(N+N_H) \times (C+1)}$ .

Compared with existing works [78, 80], our proposed unified hand-object semantic embedding  $X_t^{HO}$  holds information of the hand-induced occlusion, which is crucial for the reconstruction of the occluded part of the object, thus increasing the robustness (see Tab. 1 (b)).


**Hand Articulation Geometric Embedding.** Furthermore, the object shape is also highly constrained by the hand and very geometrically related to the hand articulation. Inspired by [8, 78, 80], we explicitly encode hand-object interaction by transforming the partially denoised point cloud at every step to each hand joint frame.

Given the partially denoised point cloud  $X_t$  and hand parameters  $\phi_H$ , we first compute the rotation  $R_j$  and translation  $T_j$  of each joint  $j$  with respect to the hand wrist using forward kinematics given the hand model. We then transform  $X_t$  from the hand wrist coordinate system to each hand joint frame via  $X_t^j = R_j X_t + T_j \in \mathbb{R}^{N \times 3}$  to encode the hand articulation onto object points. Finally, the hand articulation embedding  $X_t^A \in \mathbb{R}^{N \times J}$  is calculated via concatenation and flattening of  $X_t^j$ , with  $J$  equaling 45 in our experiments as 15 hand joints are utilized.

**Dual-Stream Denoiser.** The naive way to utilize the aforementioned semantic and geometric priors is to directly concatenate them as the condition  $z = [X_t^{HO}, X_t^A]$  and produce the denoised point cloud controlled by  $z$  in a single-stream manner. However, forcibly integrating these embeddings from different domains may result in interference and cause a drop in performance (see Tab. 4).

Hence, we instead employ a dual-stream denoiser to separately process  $X_t^{HO}$  and  $X_t^A$ . As shown in Fig. 2, given  $z_1 = X_t^{HO}$  and  $z_2 = X_t^A$ , we first feed the object and hand points along with their corresponding semantic embedding  $[X_t, z_1]$  to one branch of the dual-stream denoiser  $f_\theta^1$  and obtain  $\mathcal{F}_\theta^1 \in \mathbb{R}^{(N+N_H) \times S}$  as feature representation guided by the semantic prior, with  $S$  denoting the number of latent feature channels. Similarly, we feed the object point cloud along





**Fig. 3:** Qualitative results on the ObMan [31] dataset. For each method and ground truth, we show the reconstruction results in the camera view (column 1) and a novel view (column 2).

with its corresponding geometric embedding  $[X_t, z_2]$  to  $f_\theta^2$  using the identical architecture to obtain  $\mathcal{F}_\theta^2 \in \mathbb{R}^{N \times S}$  as our geometric feature. The final noise is then predicted from the concatenation of the semantic and geometric features with  $\epsilon_\theta = g_\theta([\mathcal{F}_\theta^1, \mathcal{F}_\theta^2]) \in \mathbb{R}^{N \times 3}$ , where  $g_\theta$  consists of stacked MLPs.

In this way, each branch can learn a specialized representation suitable for its own domain and then serve as conditioning to the diffusion model, contributing to the reconstruction of object shape from different domains. The detailed architecture of the dual-stream denoiser is provided in supplementary material.

### 3.4 Training Objectives

**Diffusion model.** For optimization, we use the common MSE loss between model prediction  $\epsilon_\theta(X_t, t, z)$  and applied noise  $\epsilon$  with

$$\mathcal{L}_{denoise} = \|\epsilon - \epsilon_\theta(X_t, t, z)\|, \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (8)$$

We further regularize the object shape using a projective mask loss  $\mathcal{L}_{mask}$ . To this end, at each timestep  $t$  the sampled and predicted point cloud are projected onto the image via the aforementioned rasterizer, and the L1 loss between them is computed as

$$\mathcal{L}_{mask} = \|\mathcal{R}(X_t) - \mathcal{R}(\hat{X}_t)\|_1, \quad (9)$$

where  $\hat{\bullet}$  denotes predicted results. The overall loss is then a weighted sum of both terms with

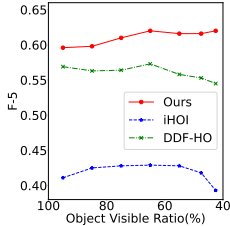
$$\mathcal{L}_{overall} = \mathcal{L}_{denoise} + \eta_1 \mathcal{L}_{mask}, \quad (10)$$

where the weight  $\eta_1$  controls the strength of the projective 2D regularization.

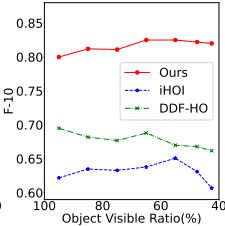
**Centroid prediction network.** Notice that the centroid prediction network is separately trained from  $\mathcal{D}_\theta$ . Thereby, the 3D object centroid  $\mathcal{M}_{3d}$  in the hand wrist frame and the 2D object centroid  $\mathcal{M}_{2d}$  in the normalized device coordinate (NDC) space are both supervised along with a 2D-3D projection loss. Overall,

**Table 1:** Comparison with the state-of-the-art on ObMan [31]. (a) F-score of 5mm and 10mm, Chamfer Distance (mm) metrics are utilized for evaluation. (b) Robustness against hand occlusion. We analyze the patterns of the F-5 and F-10 metrics as a function of the object visibility ratio on the test set of ObMan.

Method	F-5 $\uparrow$	F-10 $\uparrow$	CD $\downarrow$
HO [31]	0.23	0.56	0.64
GF [37]	0.30	0.51	1.39
AlignSDF [9]	0.40	0.64	-
iHOI [78]	0.42	0.63	1.02
gSDF [8]	0.44	0.66	-
DDF-HO [80]	0.55	0.67	0.14
Ours	<b>0.61</b>	<b>0.81</b>	<b>0.11</b>



(a)



(b)

our total loss is defined as

$$\mathcal{L}_{centroid} = \|\mathcal{M}_{3d} - \widehat{\mathcal{M}}_{3d}\| + \lambda_1 \|\mathcal{M}_{2d} - \widehat{\mathcal{M}}_{2d}\| + \lambda_2 \|\mathcal{P}(\widehat{\mathcal{M}}_{3d}) - \widehat{\mathcal{M}}_{2d}\|, \quad (11)$$

where  $\lambda_1$  and  $\lambda_2$  are hyperparameters,  $\mathcal{P}$  is the transformation from the hand wrist frame to the NDC space, and  $\widehat{\bullet}$  refers to predicted results.

## 4 Experiments

### 4.1 Datasets and Setup

**Datasets.** We evaluate our method on four common benchmark datasets, including the synthetic ObMan [31] dataset and three real-world datasets, namely HO3D [28], MOW [4] and DexYCB [6]. **ObMan** [31] consists of 2,772 objects from 8 categories taken from ShapeNet [5], with 21K plausible grasps generated using GraspIt [46]. The dataset contains 141K frames for training and 6K frames for testing, which are rendered using Blender on top of random backgrounds. **HO3D** [28] includes 77,558 frames from 68 sequences, containing 10 different objects from the YCB dataset [3], which are manipulated by 10 different users. The hand and object pose annotations are obtained from multi-camera optimization procedures. **MOW** [4] consists of 442 images of 121 object templates, with in-the-wild source data collected from EPIC Kitchens [14, 15] and the 100 Days of Hands [65] datasets. We use the same training and testing split as iHOI [78] and DDF-HO [80] for all three aforementioned datasets. **DexYCB** [6] is a large-scale real-world hand-object dataset. Following [8, 9, 76], we focus on right-hand samples and use the official s0 split. 29,656 training samples and 5,928 testing samples are downsampled following the setup of gSDF [8].

**Metrics.** Following [78, 80], we report the Chamfer Distance (CD) in mm and F-score at thresholds of 5mm (F-5) and 10mm (F-10).

**Table 2:** Results on real-world datasets HO3D [28] and MOW [4]. We report results for both finetuning (top) and zero-shot transfer (bottom).

Method	Finetuning	HO3D			MOW		
		F-5 $\uparrow$	F-10 $\uparrow$	CD $\downarrow$	F-5 $\uparrow$	F-10 $\uparrow$	CD $\downarrow$
HO [31]	✓	0.11	0.22	4.19	0.03	0.06	49.8
GF [37]	✓	0.12	0.24	4.96	0.06	0.13	40.1
iHOI [78]	✓	0.28	0.50	1.53	0.13	0.24	23.1
DDF-HO [80] <sup>1</sup>	✓	0.27	0.40	0.86	0.17	0.24	1.59
Ours	✓	<b>0.41</b>	<b>0.63</b>	<b>0.34</b>	<b>0.30</b>	<b>0.47</b>	<b>1.35</b>
iHOI [78]	✗	0.14	0.27	4.36	0.09	0.17	8.43
DDF-HO [80]	✗	0.24	0.36	<b>0.73</b>	0.14	0.19	1.89
Ours	✗	<b>0.27</b>	<b>0.41</b>	1.19	<b>0.24</b>	<b>0.36</b>	<b>1.44</b>

## 4.2 Evaluation on ObMan Dataset

We first present quantitative results on the large-scale synthetic ObMan dataset in Tab. 1 (a). As shown, our approach surpasses all state-of-the-art methods by a large margin for F-score and Chamfer Distance. Specifically, compared with the current best-performing method DDF-HO [80], we achieve a relative improvement of 10.9 % and 20.9 % for F-5 and F-10 metrics, respectively. Furthermore, we reduce the CD by 21.4 % compared with DDF-HO and 89.2 % compared with iHOI, demonstrating that our reconstructed shape possesses significantly fewer outliers.

We also visualize the reconstructed objects along with the hands in Fig. 3. Since we focus on object reconstruction, we use ground-truth hand poses for all images. The object surfaces are reconstructed using alpha shapes [19] implemented with Open3D [84] to qualitatively compare with previous works. Noteworthy, the metrics are not affected as they are computed from the point clouds. While SDF- or DDF-based methods, including iHOI [78], gSDF [8] and DDF-HO [80], tend to result in either over-smoothed and less-detailed, or fragmented reconstructions, our approach is able to generate geometrically coherent point clouds with plausible details, even for thin objects and heavily occluded parts. Specifically, the bottle in Row 2 suffers a 54.4% occlusion rate. Nonetheless, D-SCo shows strong capabilities of inferring the occluded part of the object.

## 4.3 Evaluation on Real-World Datasets

Aside from the synthetic ObMan dataset, we also conduct experiments on the three real-world datasets HO3D [28], MOW [4] and DexYCB [6] to demonstrate our approach’s generalization capabilities for real-world scenarios. The model is respectively finetuned on the HO3D and MOW, starting from the ObMan pre-trained model as initialization. As shown in Tab. 2 (top) and Tab. 3, our approach achieves state-of-the-art results on HO3D, MOW as well as DexYCB. Concretely,

<sup>1</sup> The results are provided by the authors of DDF-HO [80], with exact same training split as iHOI [78] and ours.

**Table 3:** Comparison with the state-of-the-art on DexYCB [6].

Metric	HO [31]	GF [37]	AlignSDF [9]	gSDF [8]	Ours
F-5 $\uparrow$	0.38	0.39	0.41	0.44	<b>0.63</b>
F-10 $\uparrow$	0.64	0.66	0.68	0.71	<b>0.82</b>
CD $\downarrow$	0.42	0.45	0.39	0.34	<b>0.13</b>

regarding F-5 and F-10 metrics, we exhibit a noticeable average improvement of 54.7 %, 86.2 %, and 29.4 %, respectively. In contrast to the F-score metric, the Chamfer Distance is known to be more vulnerable to outliers [8, 68, 78]. Thus, the significant reduction of the CD metric illustrates that our approach is significantly more robust, producing fewer outliers. The qualitative results in Fig. 4 similarly demonstrate the superiority of our approach.


**Zero-shot transfer to HO3D and MOW.** To further evaluate our zero-shot transfer abilities, we also directly apply our model, trained on ObMan, to HO3D and MOW without conducting any additional finetuning. The results in Tab. 2 (bottom) show that our method again achieves a remarkable improvement on HO3D (13 % and 14 % in F-5 and F-10) and MOW (71 % and 89 % in F-5 and F-10), demonstrating D-SCo’s strength in synthetic-to-real generalization.

#### 4.4 Ablation Studies

**Effectiveness of dual-stream denoiser.** We demonstrate the effectiveness of dual-stream denoiser in Tab. 4. Instead of individually processing semantic and geometric embeddings, we simply concatenate the embeddings as conditioning to the denoiser. The decrease in performance (C0 *vs.* B0) shows that independently processing semantic and geometric information within our dual-stream denoiser enhances reconstruction quality.

**Effectiveness of semantic and geometric condition.** In Tab. 4, we show the impact of the semantic and geometric embeddings. The proposed unified hand-object semantic embedding  $X_t^{HO}$  utilizes the semantic information of the hand to supplement the occluded semantic information of objects, which is crucial to model hand-induced occlusion. Adding the embedding leads again to improved performance (D2 *vs.* D0, C0 *vs.* D1). Moreover, the hand articulation geometric embedding  $X_t^A$  explicitly models hand-object interaction geometrically and results in improved performance (D1 *vs.* D0, C0 *vs.* D2).

We further noticed that the way of encoding the hand-object conditional information has a significant impact on performance. In particular, we have explored two alternative strategies to model the hand-object interaction. First, inspired by [17], we implement a GCN-based hand embedding, which utilizes a graph convolutional network [22] for feature extraction. We again apply the GCN hand embedding to the partially denoised point cloud at each step. Although the GCN hand embedding implicitly encodes the hand articulation, it does not encode the hand-object interaction. Therefore, simply applying such an embedding as conditioning does not further performance (See D3 *vs.* C0). The



**Fig. 4:** Qualitative results on HO3D [28] (top) and MOW [4] (middle and bottom) datasets. For each method and ground truth, we show the reconstruction results in the camera view (column 1) and a novel view (column 2). We also show the unoccluded objects for our method and ground truth.

hand articulation geometric embedding, on the other hand, applies the unique articulation-aware embedding to each point of the partially denoised point cloud, which explicitly encodes hand-object interaction, thus significantly benefiting the object reconstruction performance. Second, we also utilize standard PointNet [54] to encode 3D hand vertices into a global hand embedding, which is then served to the partially denoised point cloud. Again, without modeling the hand-object interaction, the global hand embedding falls short of providing sufficient semantic information about the hand-occluded object. Consequently, the predictions end up being inferior to our proposed  $X_t^{HO}$  (D4 *vs.* C0).

**Effectiveness of hand-constrained centroid fixing.** In Tab. 4, we illustrate the importance of the centroid fixing scheme. Our centroid fixing operation improves the stability of the diffusion and the reverse processes, thus enhancing the performance (D0 *vs.* E0). Further, without the hand-constrained centroid prediction network, the diffusion model has to simultaneously learn the centroid deviation and object shape, which leads to clearly worse results (E1 *vs.* D0). Additionally, in F0 we report results when using the actual ground-truth object centroid, and in F1 when using the ground-truth object pose. Without our centroid fixing scheme, the results remain inferior even when using the ground-truth object centroid on ObMan (F0 *vs.* D0). This further demonstrates the power of our proposed centroid fixing paradigm.

**Effectiveness of  $\mathcal{L}_{mask}$ .** We compare the quantitative results w/ or w/o  $\mathcal{L}_{mask}$ . Supervising the diffusion process in both 2D and 3D domains further boosts the object reconstruction performance (A0 *vs.* B0). Noteworthy, we can surpass existing methods on both ObMan and HO3D datasets even without  $\mathcal{L}_{mask}$  (B0).

**Robustness against occlusion.** To illustrate the robustness of our approach against hand occlusion, we split the test set of ObMan into groups according to the visible ratio of the object and compute the mean F-score for each group. As shown in Tab. 1 (b), when the object is undergoing strong occlusion by the hand (object visible ratio < 50%), iHOI suffers a significant decline in both F-5 and F-10 metrics. Similarly, DDF-HO also experiences an apparent decrease

**Table 4:** Ablation study on ObMan [31] and HO3D [28] datasets.

Row	Method	ObMan			HO3D		
		F-5 $\uparrow$	F-10 $\uparrow$	CD $\downarrow$	F-5 $\uparrow$	F-10 $\uparrow$	CD $\downarrow$
A0	Ours	0.61	0.81	0.11	0.41	0.63	0.34
A1	Ours Oracle	0.67	0.86	0.09	0.51	0.76	0.23
B0	A0 $\rightarrow$ w/o $\mathcal{L}_{mask}$	0.57	0.76	0.23	0.36	0.56	0.61
C0	B0 $\rightarrow$ w/o dual-stream denoiser	0.54	0.74	0.27	0.34	0.53	0.76
D0	C0 $\rightarrow$ w/o $X_t^{HO}$ & $X_t^A$	0.48	0.67	0.41	0.28	0.46	0.96
D1	C0 $\rightarrow$ w/o $X_t^{HO}$	0.51	0.69	0.37	0.33	0.50	0.81
D2	C0 $\rightarrow$ w/o $X_t^A$	0.51	0.69	0.38	0.30	0.48	0.89
D3	C0 $\rightarrow$ w/ GCN hand embedding	0.52	0.71	0.30	0.34	0.53	0.82
D4	C0 $\rightarrow$ w/ global hand embedding	0.52	0.71	0.30	0.31	0.49	0.86
E0	D0 $\rightarrow$ w/o centroid fixing	0.44	0.61	0.65	0.27	0.45	1.00
E1	D0 $\rightarrow$ w/o centroid prediction network	0.32	0.45	2.48	0.23	0.36	1.31
F0	E0 $\rightarrow$ Test with GT object centroid	0.45	0.67	0.36	0.29	0.47	0.93
F1	E0 $\rightarrow$ Test with GT object pose	0.50	0.70	0.34	0.31	0.49	0.84

with  $< 60\%$  visible ratio. Nevertheless, the performance of our approach instead remains high. Aided by our hand-constrained centroid fixing and the modeling of the hand-induced occlusion, D-SCo exhibits strong robustness against occlusion.

**Oracle Experiments.** In line with other diffusion models [42, 44] for 3D object reconstruction, in Tab. 4 (A1) we report oracle results for D-SCo. To this end, we predict five possible shapes for each object starting from different sampled Gaussian noises and choose the best sample for each input image with respect to the F-score. Essentially, the oracle results are supposed to demonstrate the probabilistic nature of our method and represent an upper bound. The obtained results underline the ability of D-SCo to overcome the ill-posed essence of the problem. Thanks to the probabilistic formulation in diffusion models, our approach is able to generate multiple plausible shapes (See supplementary material), illustrating our capability of modeling the uncertainty induced by hand- and/or self-occlusion.

## 5 Conclusion

This paper presents D-SCo, a novel centroid-fixed dual-stream conditional diffusion model for single-view hand-held object reconstruction. D-SCo does not require any object templates, category priors, or depth information, and excels at modeling uncertainties induced by hand- and self-occlusion. In the core, we propose a hand-constrained centroid-fixing paradigm, utilizing estimated hand vertices to prevent the centroid of the partially denoised point cloud from diverging during diffusion and reverse processes. Further, a dual-stream denoiser is introduced to semantically and geometrically model hand-object interaction, with a novel unified hand-object semantic embedding enhancing the robustness against occlusion.

## Acknowledgements

This work was supported in part by the China Scholarship Council and in part by the Shuimu-Zhiyuan Tsinghua Scholar Program.

## References

1. Boukhayma, A., Bem, R.d., Torr, P.H.: 3d hand shape and pose from images in the wild. In: CVPR. pp. 10843–10852 (2019)
2. Brahmabhatt, S., Tang, C., Twigg, C.D., Kemp, C.C., Hays, J.: Contactpose: A dataset of grasps with object contact and hand pose. In: ECCV. pp. 361–378. Springer (2020)
3. Calli, B., Singh, A., Walsman, A., Srinivasa, S., Abbeel, P., Dollar, A.M.: The ycb object and model set: Towards common benchmarks for manipulation research. In: 2015 international conference on advanced robotics (ICAR). pp. 510–517. IEEE (2015)
4. Cao, Z., Radosavovic, I., Kanazawa, A., Malik, J.: Reconstructing hand-object interactions in the wild. In: ICCV. pp. 12417–12426 (2021)
5. Chang, A.X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., et al.: Shapenet: An information-rich 3d model repository. arXiv preprint arXiv:1512.03012 (2015)
6. Chao, Y.W., Yang, W., Xiang, Y., Molchanov, P., Handa, A., Tremblay, J., Narang, Y.S., Van Wyk, K., Iqbal, U., Birchfield, S., Kautz, J., Fox, D.: DexYCB: A benchmark for capturing hand grasping of objects. In: CVPR (2021)
7. Chen, Y., Tu, Z., Kang, D., Chen, R., Bao, L., Zhang, Z., Yuan, J.: Joint hand-object 3d reconstruction from a single image with cross-branch feature fusion. IEEE TIP **30**, 4008–4021 (2021)
8. Chen, Z., Chen, S., Schmid, C., Laptev, I.: gsdf: Geometry-driven signed distance functions for 3d hand-object reconstruction. In: CVPR. pp. 12890–12900 (2023)
9. Chen, Z., Hasson, Y., Schmid, C., Laptev, I.: Alignsdf: Pose-aligned signed distance fields for hand-object reconstruction. In: ECCV. pp. 231–248. Springer (2022)
10. Chen, Z., Zhang, H.: Learning implicit fields for generative shape modeling. In: CVPR. pp. 5939–5948 (2019)
11. Choe, J., Joung, B., Rameau, F., Park, J., Kweon, I.S.: Deep point cloud reconstruction. In: ICLR (2021)
12. Choy, C.B., Xu, D., Gwak, J., Chen, K., Savarese, S.: 3d-r2n2: A unified approach for single and multi-view 3d object reconstruction. In: ECCV. pp. 628–644. Springer (2016)
13. Corona, E., Pumarola, A., Alenya, G., Moreno-Noguer, F., Rogez, G.: Ganhand: Predicting human grasp affordances in multi-object scenes. In: CVPR. pp. 5031–5041 (2020)
14. Damen, D., Doughty, H., Farinella, G.M., Fidler, S., Furnari, A., Kazakos, E., Moltisanti, D., Munro, J., Perrett, T., Price, W., Wray, M.: The epic-kitchens dataset: Collection, challenges and baselines. IEEE TPAMI **43**(11), 4125–4141 (2021). <https://doi.org/10.1109/TPAMI.2020.2991965>
15. Damen, D., Doughty, H., Farinella, G.M., Fidler, S., Furnari, A., Kazakos, E., Moltisanti, D., Munro, J., Perrett, T., Price, W., et al.: Scaling egocentric vision: The epic-kitchens dataset. In: ECCV. pp. 720–736 (2018)

16. Di, Y., Zhang, C., Wang, P., Zhai, G., Zhang, R., Manhardt, F., Busam, B., Ji, X., Tombari, F.: Ccd-3dr: Consistent conditioning in diffusion for single-image 3d reconstruction. arXiv preprint arXiv:2308.07837 (2023)
17. Doosti, B., Naha, S., Mirbagheri, M., Crandall, D.J.: Hope-net: A graph-based model for hand-object pose estimation. In: CVPR. pp. 6608–6617 (2020)
18. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An image is worth 16x16 words: Transformers for image recognition at scale. ICLR (2021)
19. Edelsbrunner, H., Kirkpatrick, D., Seidel, R.: On the shape of a set of points in the plane. IEEE Transactions on information theory **29**(4), 551–559 (1983)
20. Edsinger, A., Kemp, C.C.: Human-robot interaction for cooperative manipulation: Handing objects to one another. In: RO-MAN 2007-The 16th IEEE International Symposium on Robot and Human Interactive Communication. pp. 1167–1172. IEEE (2007)
21. Fan, H., Su, H., Guibas, L.J.: A point set generation network for 3d object reconstruction from a single image. In: CVPR. pp. 605–613 (2017)
22. Gao, H., Ji, S.: Graph u-nets. In: ICML. pp. 2083–2092. PMLR (2019)
23. Gao, Q., Chen, Y., Ju, Z., Liang, Y.: Dynamic hand gesture recognition based on 3d hand pose estimation for human–robot interaction. IEEE Sensors Journal **22**(18), 17421–17430 (2021)
24. Girdhar, R., Fouhey, D.F., Rodriguez, M., Gupta, A.: Learning a predictable and generative vector representation for objects. In: ECCV. pp. 484–499. Springer (2016)
25. Gkioxari, G., Malik, J., Johnson, J.: Mesh r-cnn. In: ICCV. pp. 9785–9795 (2019)
26. Grady, P., Tang, C., Twigg, C.D., Vo, M., Brahmbhatt, S., Kemp, C.C.: Contactopt: Optimizing contact to improve grasps. In: CVPR. pp. 1471–1481 (2021)
27. Groueix, T., Fisher, M., Kim, V.G., Russell, B.C., Aubry, M.: A papier-mâché approach to learning 3d surface generation. In: CVPR. pp. 216–224 (2018)
28. Hampali, S., Rad, M., Oberweger, M., Lepetit, V.: Honnotate: A method for 3d annotation of hand and object poses. In: CVPR. pp. 3196–3206 (2020)
29. Hartley, R., Zisserman, A.: Multiple view geometry in computer vision. Cambridge university press (2003)
30. Hasson, Y., Tekin, B., Bogo, F., Laptev, I., Pollefeys, M., Schmid, C.: Leveraging photometric consistency over time for sparsely supervised hand-object reconstruction. In: CVPR. pp. 571–580 (2020)
31. Hasson, Y., Varol, G., Tzionas, D., Kalevatykh, I., Black, M.J., Laptev, I., Schmid, C.: Learning joint reconstruction of hands and manipulated objects. In: CVPR. pp. 11807–11816 (2019)
32. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR. pp. 770–778 (2016)
33. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. NeurIPS **33**, 6840–6851 (2020)
34. Iqbal, U., Molchanov, P., Gall, T.B.J., Kautz, J.: Hand pose estimation via latent 2.5 d heatmap regression. In: ECCV. pp. 118–134 (2018)
35. Kanazawa, A., Tulsiani, S., Efros, A.A., Malik, J.: Learning category-specific mesh reconstruction from image collections. In: ECCV. pp. 371–386 (2018)
36. Kar, A., Tulsiani, S., Carreira, J., Malik, J.: Category-specific object reconstruction from a single image. In: CVPR. pp. 1966–1974 (2015)



37. Karunratanakul, K., Yang, J., Zhang, Y., Black, M.J., Muandet, K., Tang, S.: Grasping field: Learning implicit representations for human grasps. In: 3DV. pp. 333–344. IEEE (2020)
38. Leng, Z., Chen, J., Shum, H.P., Li, F.W., Liang, X.: Stable hand pose estimation under tremor via graph neural network. In: 2021 IEEE Virtual Reality and 3D User Interfaces (VR). pp. 226–234. IEEE (2021)
39. Lin, C.H., Kong, C., Lucey, S.: Learning efficient point cloud generation for dense 3d object reconstruction. In: AAAI. vol. 32 (2018)
40. Liu, S., Jiang, H., Xu, J., Liu, S., Wang, X.: Semi-supervised 3d hand-object poses estimation with interactions in time. In: CVPR. pp. 14687–14697 (2021)
41. Liu, Z., Tang, H., Lin, Y., Han, S.: Point-voxel cnn for efficient 3d deep learning. *NeurIPS* **32** (2019)
42. Luo, S., Hu, W.: Diffusion probabilistic models for 3d point cloud generation. In: CVPR. pp. 2837–2845 (2021)
43. Luo, S., Hu, W.: Diffusion probabilistic models for 3d point cloud generation. In: CVPR. pp. 2837–2845 (June 2021)
44. Melas-Kyriazi, L., Rupprecht, C., Vedaldi, A.: PC<sup>2</sup>: Projection-conditioned point cloud diffusion for single-image 3D reconstruction. In: CVPR. pp. 12923–12932 (June 2023)
45. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: ECCV (2020)
46. Miller, A.T., Allen, P.K.: Graspit! a versatile simulator for robotic grasping. *IEEE Robotics & Automation Magazine* **11**(4), 110–122 (2004)
47. Mueller, F., Bernard, F., Sotnychenko, O., Mehta, D., Sridhar, S., Casas, D., Theobalt, C.: Generated hands for real-time 3d hand tracking from monocular rgb. In: CVPR. pp. 49–59 (2018)
48. Mueller, F., Davis, M., Bernard, F., Sotnychenko, O., Verschoor, M., Otaduy, M.A., Casas, D., Theobalt, C.: Real-time pose and shape reconstruction of two interacting hands with a single depth camera. *ACM TOG* **38**(4), 1–13 (2019)
49. Nichol, A.Q., Dhariwal, P.: Improved denoising diffusion probabilistic models. In: ICML. pp. 8162–8171. PMLR (2021)
50. Ortenzi, V., Cosgun, A., Pardi, T., Chan, W.P., Croft, E., Kulić, D.: Object handovers: a review for robotics. *IEEE Transactions on Robotics* **37**(6), 1855–1873 (2021)
51. Panteleris, P., Oikonomidis, I., Argyros, A.: Using a single rgb frame for real time 3d hand pose estimation in the wild. In: WACV. pp. 436–445. IEEE (2018)
52. Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: DeepSDF: Learning continuous signed distance functions for shape representation. In: CVPR. pp. 165–174 (2019)
53. Pham, T.H., Kyriazis, N., Argyros, A.A., Kheddar, A.: Hand-object contact force estimation from markerless visual tracking. *IEEE TPAMI* **40**(12), 2883–2896 (2017)
54. Qi, C.R., Su, H., Mo, K., Guibas, L.J.: Pointnet: Deep learning on point sets for 3d classification and segmentation. In: CVPR. pp. 652–660 (2017)
55. Qian, X., He, F., Hu, X., Wang, T., Ramani, K.: Arnnote: An augmented reality interface for collecting custom dataset of 3d hand-object interaction pose estimation. In: Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology. pp. 1–14 (2022)

56. Rantamaa, H.R., Kangas, J., Kumar, S.K., Mehtonen, H., Järnstedt, J., Raisamo, R.: Comparison of a vr stylus with a controller, hand tracking, and a mouse for object manipulation and medical marking tasks in virtual reality. *Applied Sciences* **13**(4), 2251 (2023)
57. Ravi, N., Reizenstein, J., Novotny, D., Gordon, T., Lo, W.Y., Johnson, J., Gkioxari, G.: Accelerating 3d deep learning with pytorch3d. *arXiv preprint arXiv:2007.08501* (2020)
58. Rogez, G., Khademi, M., Supančič III, J., Montiel, J.M.M., Ramanan, D.: 3d hand pose detection in egocentric rgb-d images. In: *ECCVW*. pp. 356–371. Springer (2015)
59. Rogez, G., Supancic, J.S., Ramanan, D.: Understanding everyday hands in action from rgb-d images. In: *ICCV*. pp. 3889–3897 (2015)
60. Romero, J., Tzionas, D., Black, M.J.: Embodied hands: Modeling and capturing hands and bodies together. *ACM TOG* **36**(6) (2017)
61. Rong, Y., Shiratori, T., Joo, H.: Frankmocap: A monocular 3d whole-body pose estimation system via regression and integration. In: *ICCVW* (2021)
62. Schönberger, J.L., Frahm, J.M.: Structure-from-motion revisited. In: *CVPR* (2016)
63. Schönberger, J.L., Zheng, E., Pollefeys, M., Frahm, J.M.: Pixelwise view selection for unstructured multi-view stereo. In: *ECCV* (2016)
64. Seitz, S.M., Curless, B., Diebel, J., Scharstein, D., Szeliski, R.: A comparison and evaluation of multi-view stereo reconstruction algorithms. In: *CVPR*. vol. 1, pp. 519–528. IEEE (2006)
65. Shan, D., Geng, J., Shu, M., Fouhey, D.F.: Understanding human hands in contact at internet scale. In: *CVPR*. pp. 9869–9878 (2020)
66. Sridhar, S., Mueller, F., Oulasvirta, A., Theobalt, C.: Fast and robust hand tracking using detection-guided optimization. In: *CVPR*. pp. 3213–3221 (2015)
67. Sridhar, S., Mueller, F., Zollhöfer, M., Casas, D., Oulasvirta, A., Theobalt, C.: Real-time joint tracking of a hand manipulating an object from rgb-d input. In: *ECCV*. pp. 294–310. Springer (2016)
68. Tatarchenko, M., Richter, S.R., Ranftl, R., Li, Z., Koltun, V., Brox, T.: What do single-view 3d reconstruction networks learn? In: *CVPR*. pp. 3405–3414 (2019)
69. Tekin, B., Bogo, F., Pollefeys, M.: H+o: Unified egocentric recognition of 3d hand-object poses and interactions. In: *CVPR* (June 2019)
70. Tulsiani, S., Kar, A., Carreira, J., Malik, J.: Learning category-specific deformable 3d models for object reconstruction. *IEEE TPAMI* **39**(4), 719–731 (2016)
71. Tzionas, D., Ballan, L., Srikantha, A., Aponte, P., Pollefeys, M., Gall, J.: Capturing hands in action using discriminative salient points and physics simulation. *IJCV* **118**, 172–193 (2016)
72. Tzionas, D., Gall, J.: 3d object reconstruction from hand-object interactions. In: *ICCV*. pp. 729–737 (2015)
73. Wu, J., Zhang, C., Xue, T., Freeman, B., Tenenbaum, J.: Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. *NeurIPS* **29** (2016)
74. Xie, H., Yao, H., Sun, X., Zhou, S., Zhang, S.: Pix2vox: Context-aware 3d reconstruction from single and multi-view images. In: *ICCV*. pp. 2690–2698 (2019)
75. Xie, H., Yao, H., Zhang, S., Zhou, S., Sun, W.: Pix2vox++: Multi-scale context-aware 3d object reconstruction from single and multiple images. *IJCV* **128**(12), 2919–2935 (2020)
76. Yang, L., Li, K., Zhan, X., Lv, J., Xu, W., Li, J., Lu, C.: Artiboost: Boosting articulated 3d hand-object pose estimation via online exploration and synthesis. In: *CVPR*. pp. 2750–2760 (2022)

77. Yang, L., Zhan, X., Li, K., Xu, W., Li, J., Lu, C.: Cpf: Learning a contact potential field to model the hand-object interaction. In: ICCV. pp. 11097–11106 (2021)
78. Ye, Y., Gupta, A., Tulsiani, S.: What’s in your hands? 3D reconstruction of generic objects in hands. In: CVPR. pp. 3895–3905 (June 2022)
79. Zeng, X., Vahdat, A., Williams, F., Gojcic, Z., Litany, O., Fidler, S., Kreis, K.: Lion: Latent point diffusion models for 3d shape generation. In: NeurIPS (2022)
80. Zhang, C., Di, Y., Zhang, R., Zhai, G., Manhardt, F., Tombari, F., Ji, X.: Ddf-ho: Hand-held object reconstruction via conditional directed distance field. *NeurIPS* **36** (2024)
81. Zhang, C., Jiao, G., Di, Y., Huang, Z., Wang, G., Zhang, R., Fu, B., Tombari, F., Ji, X.: Moho: Learning single-view hand-held object reconstruction with multi-view occlusion-aware supervision. *arXiv preprint arXiv:2310.11696* (2024)
82. Zhang, X., Li, Q., Mo, H., Zhang, W., Zheng, W.: End-to-end hand mesh recovery from a monocular rgb image. In: ICCV. pp. 2354–2364 (2019)
83. Zhou, L., Du, Y., Wu, J.: 3d shape generation and completion through point-voxel diffusion. In: ICCV. pp. 5826–5835 (October 2021)
84. Zhou, Q.Y., Park, J., Koltun, V.: Open3d: A modern library for 3d data processing. *arXiv preprint arXiv:1801.09847* (2018)
85. Zhou, Y., Habermann, M., Xu, W., Habibi, I., Theobalt, C., Xu, F.: Monocular real-time hand shape and motion capture using multi-modal data. In: CVPR. pp. 5346–5355 (2020)
86. Zimmermann, C., Brox, T.: Learning to estimate 3d hand pose from single rgb images. In: ICCV. pp. 4903–4911 (2017)