

Sparse multi-view hand-object reconstruction for unseen environments

Yik Lung Pang¹, Changjae Oh¹, Andrea Cavallaro^{1,2,3}

¹Centre for Intelligent Sensing, Queen Mary University of London

²Idiap Research Institute ³École Polytechnique Fédérale de Lausanne

{y.l.pang, c.oh}@qmul.ac.uk, a.cavallaro@idiap.ch

Abstract

Recent works in hand-object reconstruction mainly focus on the single-view and dense multi-view settings. On the one hand, single-view methods can leverage learned shape priors to generalise to unseen objects but are prone to inaccuracies due to occlusions. On the other hand, dense multi-view methods are very accurate but cannot easily adapt to unseen objects without further data collection. In contrast, sparse multi-view methods can take advantage of the additional views to tackle occlusion, while keeping the computational cost low compared to dense multi-view methods. In this paper, we consider the problem of hand-object reconstruction with unseen objects in the sparse multi-view setting. Given multiple RGB images of the hand and object captured at the same time, our model SVHO combines the predictions from each view into a unified reconstruction without optimisation across views. We train our model on a synthetic hand-object dataset and evaluate directly on a real world recorded hand-object dataset with unseen objects. We show that while reconstruction of unseen hands and objects from RGB is challenging, additional views can help improve the reconstruction quality.

1. Introduction

Reconstructing hand and object shapes from visual input has a lot of applications ranging from AR/VR to robotics [17, 21, 23, 29]. For example, in a human-to-robot handover scenario, the retrieved shapes help guide the robot to estimate possible grasps on the object while avoiding contact with the human hand. In this case, the accuracy of the reconstruction method is important, as a poorly estimated object shape will cause the grasp to fail, while a poorly estimated hand shape can cause injury to the human if the robot is unsure where the human hand is. Moreover, if the human attempts to handover an unseen object to the robot, the robot should be able to adapt to this new object on the fly without additional data collection to ensure the naturalness of the interaction.

Joint hand-object shape reconstruction from RGB has seen impressive progress in recent years [3–5, 7, 11, 13, 25, 30] due to the availability of high-quality parametric hand models [20] and neural implicit representations [18] for modelling complex 3D shapes. However, these methods are still rarely applied in the human-to-robot handover scenario either due to poor reconstruction quality or large data collection requirements.

Recent hand-object shape reconstruction methods can be separated into two categories, namely single-view or dense multi-view methods. Single-view methods can reconstruct hand-object shapes in a single forward pass by leveraging learned 3D shape priors from large hand-object interaction datasets [3–5, 7, 11, 13, 25, 30]. However, predicting the shape of unknown hands and objects from a single-view is challenging as the diversity of the hand poses and object shapes is vast. Moreover, when the object or hand is occluded or under motion blur, the reconstruction quality suffers greatly [4, 5, 30].

Using multiple views can alleviate the problem of occlusion by providing redundancy in the input. Neural implicit shape representation enables detailed reconstruction of unknown shapes from densely captured images of the object from multiple viewpoints [10, 19, 24, 31]. However, the reconstruction suffers where the object is occluded, which is common in the hand-object interaction scenario. Moreover, obtaining and optimising on dense multi-view images is time consuming and is not suitable in scenarios where the model has to adapt quickly to an unseen object.

Sparse multi-view methods provide a balanced approach between single-view and dense multi-view methods but has not been investigated in the hand-object reconstruction task. Compared to single-view methods, the additional views can help alleviate issues caused by occlusion or challenging poses. On the other hand, the data requirement is lower compared to dense multi-view methods, which enables applications where the model has to reconstruct unseen objects without further data collection or training.

In this paper, we propose SVHO, a sparse multi-view method for hand-object reconstruction. Our proposed

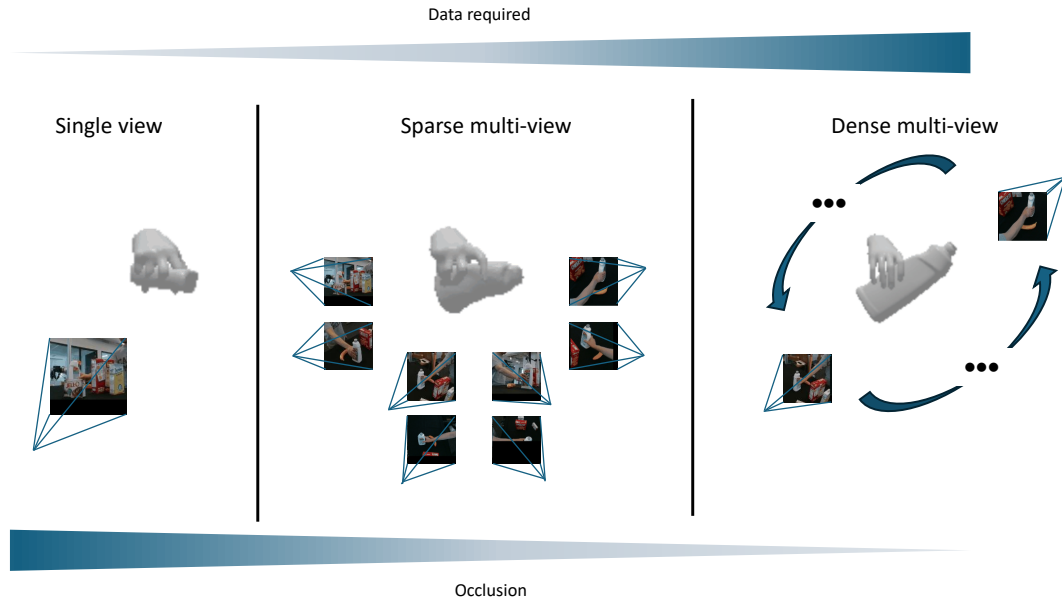


Figure 1. Single-view methods suffer from occlusion, while dense multi-view methods require a large amount of collected images. We propose to use sparse multi-view input to improve the reconstruction quality while keeping the data requirements low.

method takes as input the RGB images and the corresponding global hand poses from each view. The model predicts the hand and object shapes independently from each view and combines them to form a final reconstruction. We train our model entirely on the synthetic dataset ObMan [11] and evaluate on the real-world recorded dataset DexYCB [2] with unseen objects.

In summary, our contributions are:

- We propose a sparse multi-view setting for hand-object reconstruction
- We analyse our proposed sparse multi-view reconstruction method on up to 8 input views of unseen objects

2. Related works

2.1. 3D shape representations

Early works in hand-object shape estimation focused on using triangular meshes as a representation for the hand and object shape [11, 25]. The MANO model [20] provided a low dimensional representation that can be decoded back to a complete hand mesh efficiently. On the other hand, predicting meshes for unknown objects is challenging, especially for non genus 0 shapes [9]. Neural implicit shape representations [6, 18] have recently emerged as a popular choice for representing shapes in 3D due to their versatile capabilities of representing complex shapes, and have been adopted in hand object reconstruction works [3–5, 13, 30].

2.2. Single-view hand-object reconstruction

Works on hand-object shape reconstruction mainly follow the encoder-decoder architecture. An image encoder encodes the input image to the latent space, and then a 3D shape decoder maps the latent vectors to triangular meshes or signed distance fields (SDFs). A dual branch network was proposed in [11], each with a ResNet-18 [12] encoder pre-trained on ImageNet [22] to estimate features for the hand and object respectively. The hand features are passed through a fully connected layer to estimate the shape and pose parameters for the MANO [20] hand model and scale and translation parameters to help determine the size of the object. The object features are passed to AtlasNet [9] to estimate the object shape as an atlas of mesh surfaces in normalised vertices. A contact loss was also proposed to encourage contact between the hand and the object. However, there is no exchange of information between the hand and object shape estimation branch, which both contains important information for each other’s task (e.g. estimating contact surfaces).

Later works explored the idea of sharing features between the hand and the object to improve the reconstruction quality. Relevant information between the hand and object can be shared via attention-guided graph convolution to improve both the estimation of the hand and object shape in an iterative manner [25]. However, the approach requires calculating the attention value of a fully connected graph formed from the hand and object mesh in multiple iterations, which can be time consuming. The speed and accu-

racy of the graph convolution architecture can be improved by introducing a dense mutual attention mechanism [27]. This removes the need to optimise the shapes in iterative steps. However, the features are shared in the global latent space, which can be difficult to optimise.

Several works have improved reconstruction results by estimating local shape features instead of a single global latent feature [4,5,30]. The estimation of the object shape can be conditioned on the estimated hand shape [30]. After estimating the hand shape using an off-the-shelf estimator, the object is reconstructed in the hand-wrist frame. Each query 3D position is projected back to the 2D image to retrieve the corresponding image features and the distance to each hand joint is encoded as additional features for the reconstruction of the object. Similar projection based feature retrieval was also proposed in [4,5], while also estimating object poses in addition to hand poses as condition to the shape reconstruction process. Most recently, a NeRF-based hand-object reconstruction model was proposed for single-view reconstruction which can be trained directly from multi-view images using photometric loss and does not require 3D ground-truth shape as supervision [7].

2.3. Multi-view hand-object reconstruction

Multi-view reconstruction of hand and object relies on the optimisation of the predicted shape and appearance based on densely collected images with predicted hand global pose [10,19,24,31]. NeRF [15] based methods were proposed to represent the hand and object shape [10,24]. The pose and shape of the hand and object are optimised by minimising the photometric loss between the rendered colour and the observed colour. These methods require a prerecorded video of the interaction and are not able to infer the shape in a new view. Moreover, the occluded parts of the object cannot be reconstructed as it is never observed by the camera. A diffusion-based prior was proposed to estimate the geometry of the parts occluded in the video [31]. A model was proposed to generalise to sparse novel views at test time [19]. However, the model has to be trained offline first with sparse view images of the same hand and object separately.

3. Proposed method

We aim to reconstruct the shape of the hand and the shape of the hand-held object from sparse multi-view images. We first predict the hand and object shape from each camera view independently in the canonical coordinate space. We then combine them to obtain a single final output.

We train autoencoders to encode hand and object shapes independently in the canonical coordinate space. We train a Patchwise VQ-VAE (P-VQ-VAE) [16] to autoencode the hand and object into 3D discrete latent cubes with cor-




Figure 2. We first encode hand and object shape independently using Patchwise VQ-VAE (P-VQ-VAE). This provides a compact representation to train our hand object shape prior.

responding shape codebooks. This allows us to treat the shape estimation problem as a classification problem with the classes being the codebook indices.

During test time, we obtain 2D features from the input image. We then form a 3D feature grid by projecting the 3D points in the canonical coordinate space to the image space using the global hand pose. We reconstruct the hand and object shapes in the canonical coordinate space and combine the classification probabilities of the different views. The final mesh is obtained using the trained decoder and the marching cubes algorithm [14].

3.1. Autoencoding hands and objects

We train autoencoders to map the 3D shape of the hand and the object into 3D discrete latent cubes $Z_h \in \mathbb{N}^{8 \times 8 \times 8}$ and $Z_o \in \mathbb{N}^{8 \times 8 \times 8}$, respectively. We use P-VQ-VAE [16] as the encoder architecture. The input SDF $\mathbb{R}^{128 \times 128 \times 128}$ is first divided into $8 \times 8 \times 8 = 512$ patches of $\mathbb{R}^{16 \times 16 \times 16}$. Each patch is encoded individually with 3D convolutional layers into a latent vector $z_e \in \mathbb{R}^V$ which combine to give the continuous latent cube of $Z' \in \mathbb{R}^{8 \times 8 \times 8 \times V}$. We then learn a codebook $B \in \mathbb{R}^{K \times V}$ with K elements of $e_k \in \mathbb{R}^V$, and discretize the continuous latent cube by mapping each patch’s latent vector $z_e \in \mathbb{R}^V$ to its nearest element e_i in the codebook [26].

$$z = e_i, \text{ where } i = \operatorname{argmin}_k \|z_e - e_k\|_2, \quad (1)$$

where i is the index of the element e_i in the codebook nearest to the encoder output z_e , and $z \in \mathbb{R}^V$ is the mapped latent code. We save the extracted index i for each patch and reshape them back into the 3D cube $Z \in \mathbb{N}^{8 \times 8 \times 8}$. Note that we omit the subscript h and o as the same process is repeated for the hand and object.

By learning the codebook B we obtain a compact discrete latent representation of the hand and object shape. This representation preserves the local shape information of

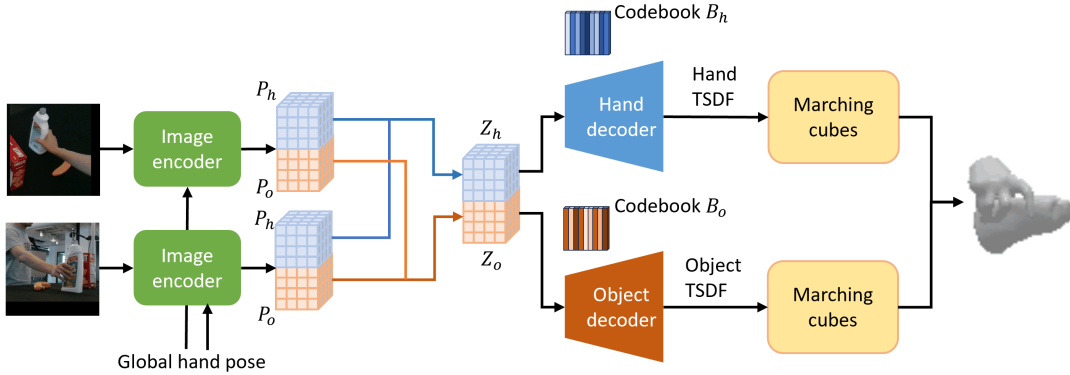


Figure 3. Our pipeline for hand object shape reconstruction from multi-view images. Predicted probabilities from individual views are averaged to get the final prediction.

the hand and object as each latent code corresponds to the local shape input only. Together, the latent cube is decoded jointly to ensure the consistency of the global shape. The continuous codes $Z' \in \mathbb{R}^{8 \times 8 \times 8 \times V}$ are first retrieved using the codebook B and discrete latent cube Z and enhanced with 3D convolutional layers across neighbouring patches to ensure spatial consistency.

We learn a decoder D to map latent codes e and 3D position $p = (x, y, z)$ to predicted SDF value \hat{s} . We use a 5-layer MLP decoder to be consistent with previous works [4, 5]. To retrieve the latent code e , the mixed codes are upsampled using trilinear interpolation to $\mathbb{R}^{128 \times 128 \times 128 \times V}$, and the corresponding code is retrieved using the query 3D position p . We first obtain the corresponding latent cube index $(i, j, k) = M(p)$ with the indexing function M . We then retrieve $e = Z'_{ijk}$ as the latent code for position p . We optimise the loss L_{AE} to train the autoencoders:

$$L_{AE} = \|s - \hat{s}\| + \|sg[z_e] - e\|_2^2 + \beta \|z_e - sg[e]\|_2^2, \quad (2)$$

where s and \hat{s} are the ground-truth and predicted signed distance field (SDF) values respectively, sg is the stop gradient operator, and e is the nearest element in codebook B to encoder output z_e . The first term minimises the difference between the reconstructed and input shape, while the second and third terms are the vector quantisation loss and commitment loss [26].

3.2. Reconstruction from multi-view images

After the autoencoders are trained, we run the encoder on the training set to extract the training pairs $\{I, Z_h, Z_o, \pi\}$, where I is the RGB image and π is the ground-truth transformation from the camera coordinate frame to the hand-wrist coordinate frame.

We first train our model on the single-view reconstruction task. We use ResNet-18 [12] to encode the input im-

age into 2D features. Spatially aligned 3D features are retrieved by projecting the query 3D point back to the 2D image space using the hand pose. We treat the reconstruction task as a classification task with the codebook indices as the classes [16]. We use 3D convolutional layers to process the spatially aligned 3D features and output $\hat{P} \in \mathbb{R}^{8 \times 8 \times 8 \times K}$ which corresponds to the predicted probabilities of each codebook index at each location in 3D space. We train the model using a weighted cross-entropy where a weight of 0.25 is given to the index corresponding to an empty space and a weight of 0.75 is given to all other indices equally as the dataset is biased towards empty space.

To reconstruct the hand and object shapes from multiple views C with input images $I = \{I_1, \dots, I_{|C|}\}$, we first obtain independent reconstructions from each camera view c . We aggregate the results across different views as P by averaging the predicted probabilities of the selected views.

$$P = \frac{1}{|C'|} \sum_{c \in C'} \hat{P}_c, \quad (3)$$

where $C' \subseteq C$ is the subset of views selected for reconstruction. After P is obtained, we select the codebook index i at each location with the highest predicted probability to form the predicted latent cube \hat{Z} . We retrieve the continuous latent cube \hat{Z}' by indexing the codebook B and we use the trained decoder D to obtain the predicted SDF values \hat{s} . Finally, we obtain the reconstructed mesh $R_{C'}$ using marching cubes [14] from the predicted SDF values.

4. Experiments

4.1. Datasets

ObMan [11] We train our model on the synthetic only dataset ObMan containing 8 object categories (bottles, bowls, cans, jars, knives, cellphones, cameras and remote




Figure 4. Autoencoder reconstruction of 3D hand and objects

controls) from ShapeNet [1]. Each sample contains an annotated hand grasp with the corresponding ground-truth hand and object mesh. We follow AlignSDF [5] and discard the samples with too many double-sided triangles, resulting in a training set of 87190 samples and a testing set of 6285 samples.

DexYCB [2] We use DexYCB as a testing dataset that includes multi-view video recordings of humans performing grasping of objects. The dataset contains 20 unique objects from the YCB-Video dataset [28]. Following AlignSDF [5], we only consider the samples in the S0 split where the minimum distance between the hand and the object mesh is less than 5mm. For each object, we randomly sample 100 frames, where each frame contains 8 RGB images from the 8 camera views, the ground-truth hand and object mesh, and the annotated hand pose. This results in a total of 2000 sampled frames, with 16000 images.

4.2. Implementation details

Our goal is to mimic the scenarios in human-robot interactions where the robot has to adapt quickly to unseen objects grasped by humans. Therefore, we train our model entirely on the synthetic ObMan dataset, and we evaluate the performance on the DexYCB dataset where the objects used are unseen.

We first train our P-VQ-VAE to autoencode hand and objects independently on the training set of ObMan. We use a codebook size of 512 with a latent code of length 128, and we sample the hand and object SDF as a truncated SDF (TSDF) of size $128 \times 128 \times 128$. Due to instability in the codebook training, we perform codebook restart [8] on codes that are not used every 25 batches. We double this duration every time the codebook is restarted. Once the model is trained, we run the encoders on the training set to obtain the latent cube Z for each sample.

We then pair these sampled latent cubes Z with their corresponding RGB image and global hand pose to form the training set for the image reconstruction model. Our image reconstruction model outputs the probabilities for each en-

Table 1. Results for autoencoding hand and object shapes (CD - Chamfer distance; FS - F-score)

$CD_h \downarrow$	$FS_h@1 \uparrow$	$FS_h@5 \uparrow$	$CD_o \downarrow$	$FS_o@5 \uparrow$	$FS_o@10 \uparrow$
0.015	0.798	0.999	0.437	0.822	0.953

try in the codebook at each location in the latent cube. We choose the predicted codebook index with the highest probability to perform the reconstruction. We retrieve the continuous latent cube Z' from the codebook and use marching cubes to obtain the final mesh at the 0 level set.

4.3. Metrics

Following gSDF [4], we first align the predicted and ground-truth hand mesh with the global hand pose. The object meshes are transformed also using the global hand pose. We use Chamfer distance (cm^2) and F-score to evaluate the reconstruction quality of the predicted hand and object meshes.

4.4. Results

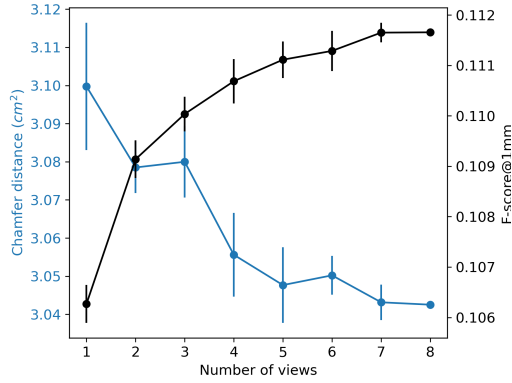
4.4.1 3D shape autoencoding

We evaluate our P-VQ-VAE autoencoder on the testing set of ObMan. Our model takes as input the ground-truth TSDF and outputs the reconstructed TSDF of the same size. We found that our autoencoder can accurately reconstruct hands and objects of various poses and shapes (Table 1).


Figure 4 shows example visualisations of hand and object reconstructions from the object classes bottles, can, remote and knife. Since the latent codes extracted by the shape encoder will act as a ground-truth for the image reconstruction method, it is important to ensure the quality of the autoencoder’s reconstruction as it acts as the upper-bound for our image reconstruction model.

4.4.2 Multi-view RGB reconstruction

We evaluate our multi-view hand-object shape reconstruction model on the DexYCB dataset. To analyse the impact



(a) Hand reconstruction



(b) Object reconstruction

Figure 5. Average F-score and standard deviation across 6 runs for (a) hand and (b) object reconstruction when varying the number of input views

of the number of views on the reconstruction quality, we vary the number of input views to the model from 1 to 8. For each sample, we randomly select the corresponding number of views from the 8 available views. We repeat the experiment for each number of views 6 times and we report the mean and standard deviation across the runs.

Figure 5 shows the average Chamfer distance and F-score of hand and object with varying numbers of views. For the hand, the F-score increases consistently as the number of views increases. For the object, the F-score peaks at 2 views and decreases as the number of views increases further. The Chamfer distance follows the same trend but in reverse as lower Chamfer distance is better. Since the images are centered around the wrist, the model can effectively take advantage of the multiple views to improve the features for hand reconstruction since the position of the hand in the image is known. However, as the number of views increases, other objects in the scene can distract the model from recognising the target object being grasped. Thus, more views can actually negatively impact the reconstruction quality of the object. This issue can be mitigated




Figure 6. Average F-score across 6 runs per object when varying the number of input views

by introducing a hand-object segmentation model, to allow our reconstruction model to focus on the target object only, and improve the object reconstruction quality with the increasing number of views.

We further analyse the F-score per object with varying numbers of views in Figure 6. We observe that the reconstruction quality of each object changes differently to the increasing number of views. Objects that have good reconstruction quality at a single-view benefits the most when the number of views increases, while objects that have lower reconstruction quality become worse when the number of views increases.

We also observe that objects that stand out from the background can benefit the most from the increasing number of views, for example mustard bottle, power drill and foam brick. This further supports the need for hand-object segmentation masks to allow our reconstruction model to fully take advantage of the multiple views when reconstructing the object.

We show additional qualitative results in Figure 7. For the gelatine box in the first example, we observe that the first view used I_1 has a cluttered background with colours similar to that of the gelatine box. As the model receives the second view I_2 and third view I_3 , the reconstruction improves as the object can be clearly separated from the dark background. However, further increasing the number of views causes the reconstruction to deteriorate, as more and more objects appear in the background. Similar issues can be observed in the fourth view I_4 and the sixth view I_6 of the large marker and the banana as the target object blends with the background or other objects in the scene. Self occlusion by the hand on the object can also cause the reconstruction to deteriorate especially for smaller objects.



Figure 7. Qualitative reconstruction results from multi-view RGB. The top row shows the 8 views used in the reconstruction. The bottom row shows (from left to right) the reconstruction obtained using 1 to 8 views.

5. Conclusion

In this paper, we tackled the challenging problem of sparse multi-view reconstruction of unseen (hand-held) objects. We showed that our model can take advantage of the additional information from multiple views for reconstruction. However, in cluttered scenes, increasing the number of views may negatively impact the reconstruction quality of the object. Our future work includes the introduction of a hand-object segmentation model as a pre-processing step for the model to better take advantage of the multiple views.

Acknowledgement

This project made use of time on Tier 2 HPC facility JADE2, funded by EPSRC (EP/T022205/1).

References

- [1] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015. 5
- [2] Yu-Wei Chao, Wei Yang, Yu Xiang, Pavlo Molchanov, Ankur Handa, Jonathan Tremblay, Yashraj S Narang, Karl Van Wyk, Umar Iqbal, Stan Birchfield, et al. Dexycb: A benchmark for capturing hand grasping of objects. In *CVPR*, pages 9044–9053, 2021. 2, 5
- [3] Yujin Chen, Zhigang Tu, Di Kang, Ruizhi Chen, Linchao Bao, Zhengyou Zhang, and Junsong Yuan. Joint hand-object 3d reconstruction from a single image with cross-branch feature fusion. *IEEE Transactions on Image Processing*, 30:4008–4021, 2021. 1, 2
- [4] Zerui Chen, Shizhe Chen, Cordelia Schmid, and Ivan Laptev. gsdf: Geometry-driven signed distance functions for 3d hand-object reconstruction. In *CVPR*, pages 12890–12900, 2023. 1, 2, 3, 4, 5
- [5] Zerui Chen, Yana Hasson, Cordelia Schmid, and Ivan Laptev. Alignsdf: Pose-aligned signed distance fields for hand-object reconstruction. In *ECCV*, pages 231–248. Springer, 2022. 1, 2, 3, 4, 5
- [6] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *CVPR*, pages 5939–5948, 2019. 2
- [7] Hongsuk Choi, Nikhil Chavan-Dafle, Jiacheng Yuan, Volkan Isler, and Hyunsoo Park. Handnerf: Learning to reconstruct hand-object interaction scene from a single rgb image. In *ICRA*, 2024. 1, 3

- [8] Prafulla Dhariwal, Heewoo Jun, Christine Payne, Jong Wook Kim, Alec Radford, and Ilya Sutskever. Jukebox: A generative model for music. *arXiv preprint arXiv:2005.00341*, 2020. [5](#)
- [9] Thibault Groueix, Matthew Fisher, Vladimir G Kim, Bryan C Russell, and Mathieu Aubry. A papier-mâché approach to learning 3d surface generation. In *CVPR*, pages 216–224, 2018. [2](#)
- [10] Shreyas Hampali, Tomas Hodan, Luan Tran, Lingni Ma, Cem Keskin, and Vincent Lepetit. In-hand 3d object scanning from an rgb sequence. In *CVPR*, pages 17079–17088, 2023. [1](#), [3](#)
- [11] Yana Hasson, Gul Varol, Dimitrios Tzionas, Igor Kalevatykh, Michael J Black, Ivan Laptev, and Cordelia Schmid. Learning joint reconstruction of hands and manipulated objects. In *CVPR*, pages 11807–11816, 2019. [1](#), [2](#), [4](#)
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. [2](#), [4](#)
- [13] Korrawe Karunratanakul, Jinlong Yang, Yan Zhang, Michael J Black, Krikamol Muandet, and Siyu Tang. Grasping field: Learning implicit representations for human grasps. In *3DV*, pages 333–344. IEEE, 2020. [1](#), [2](#)
- [14] William E Lorensen and Harvey E Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *Seminal graphics: pioneering efforts that shaped the field*, pages 347–353. 1998. [3](#), [4](#)
- [15] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. [3](#)
- [16] Paritosh Mittal, Yen-Chi Cheng, Maneesh Singh, and Shubham Tulsiani. AutoSDF: Shape priors for 3d completion, reconstruction and generation. In *CVPR*, 2022. [3](#), [4](#)
- [17] Yik Lung Pang, Alessio Xompero, Changjae Oh, and Andrea Cavallaro. Towards safe human-to-robot handovers of unknown containers. In *RO-MAN*, pages 51–58, 2021. [1](#)
- [18] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *CVPR*, pages 165–174, 2019. [1](#), [2](#)
- [19] Wentian Qu, Zhaopeng Cui, Yinda Zhang, Chenyu Meng, Cuixia Ma, Xiaoming Deng, and Hongan Wang. Novel-view synthesis and pose estimation for hand-object interaction from sparse views. In *ICCV*, pages 15100–15111, 2023. [1](#), [3](#)
- [20] Javier Romero, Dimitrios Tzionas, and Michael J Black. Embodied hands: Modeling and capturing hands and bodies together. *arXiv preprint arXiv:2201.02610*, 2022. [1](#), [2](#)
- [21] Patrick Rosenberger, Akansel Cosgun, Rhys Newbury, Jun Kwan, Valerio Ortenzi, Peter Corke, and Manfred Grafinger. Object-independent human-to-robot handovers using real time robotic vision. *IEEE Robotics and Automation Letters*, 6(1):17–23, 2020. [1](#)
- [22] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015. [2](#)
- [23] Ricardo Sanchez-Matilla, Konstantinos Chatzilygeroudis, Apostolos Modas, Nuno Ferreira Duarte, Alessio Xompero, Pascal Frossard, Aude Billard, and Andrea Cavallaro. Benchmark for human-to-robot handovers of unseen containers with unknown filling. *IEEE Robotics and Automation Letters*, 5(2):1642–1649, 2020. [1](#)
- [24] Anilkumar Swamy, Vincent Leroy, Philippe Weinzaepfel, Fabien Baradel, Salma Galaoui, Romain Brégier, Matthieu Armando, Jean-Sebastien Franco, and Grégory Rogez. Showme: Benchmarking object-agnostic hand-object 3d reconstruction. In *ICCV*, pages 1935–1944, 2023. [1](#), [3](#)
- [25] Tze Ho Elden Tse, Kwang In Kim, Ales Leonardis, and Hyung Jin Chang. Collaborative learning for hand and object reconstruction with attention-guided graph convolution. In *CVPR*, pages 1664–1674, 2022. [1](#), [2](#)
- [26] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *NeurIPS*, 30, 2017. [3](#), [4](#)
- [27] Rong Wang, Wei Mao, and Hongdong Li. Interacting hand-object pose estimation via dense mutual attention. In *WACV*, pages 5735–5745, 2023. [3](#)
- [28] Yu Xiang, Tanner Schmidt, Venkatraman Narayanan, and Dieter Fox. Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes. In *Robotics: Science and Systems (RSS)*, 2018. [5](#)
- [29] Wei Yang, Chris Paxton, Arsalan Mousavian, Yu-Wei Chao, Maya Cakmak, and Dieter Fox. Reactive human-to-robot handovers of arbitrary objects. In *ICRA*, pages 3118–3124. IEEE, 2021. [1](#)
- [30] Yufei Ye, Abhinav Gupta, and Shubham Tulsiani. What’s in your hands? 3d reconstruction of generic objects in hands. In *CVPR*, pages 3895–3905, 2022. [1](#), [2](#), [3](#)
- [31] Yufei Ye, Poorvi Hebbar, Abhinav Gupta, and Shubham Tulsiani. Diffusion-guided reconstruction of everyday hand-object interaction clips. In *ICCV*, pages 19717–19728, 2023. [1](#), [3](#)