

Slice3D: Multi-Slice, Occlusion-Revealing, Single View 3D Reconstruction

Yizhi Wang¹, Wallace Lira¹, Wenqi Wang², Ali Mahdavi-Amiri¹, Hao Zhang¹
¹Simon Fraser University ²Tsinghua University

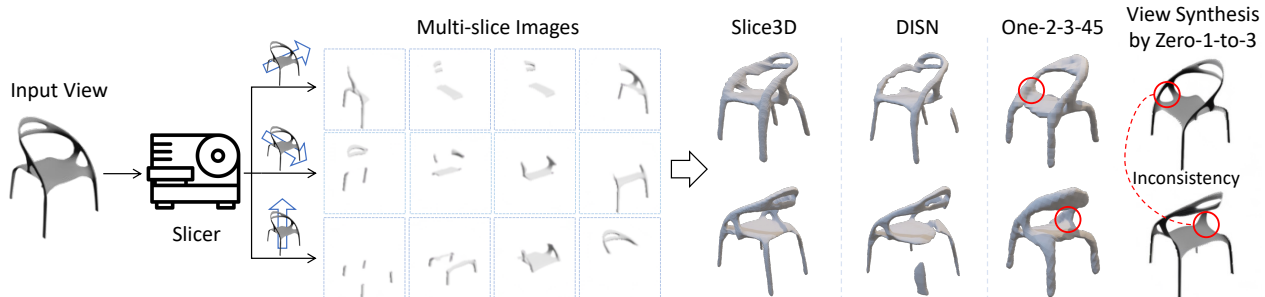


Figure 1. Our single-view 3D reconstruction method, Slice3D, predicts *multi-slice images* to reveal occluded parts without changing the camera (in contrast to *multi-view* synthesis), and then lifts the slices into a 3D model. This leads to better recovery of occluded parts than both encoder-decoder based methods such as DISN [65] and 3D reconstructions from diffusion-synthesized multi-view images (via Zero-1-to-3 [29]) such as One-2-3-45 [28], where the key challenge of the latter is multi-view inconsistency, as shown on the right.

Abstract

We introduce *multi-slice reasoning*, a new notion for *single-view 3D reconstruction* which challenges the current and prevailing belief that *multi-view synthesis* is the most natural conduit between *single-view* and *3D*. Our key observation is that *object slicing* is a more direct, and hence more advantageous, means to reveal occluded structures than altering camera views. Specifically, slicing can peel through any occluder without obstruction, and in the limit (i.e., with infinitely many slices), it is guaranteed to unveil all hidden object parts. We realize our idea by developing *Slice3D*, a novel method for *single-view 3D reconstruction* which first predicts *multi-slice images* from a single RGB input image and then integrates the slices into a 3D model using a coordinate-based transformer network to produce a signed distance function. The slice images can be regressed or generated, both through a U-Net based network. For the former, we inject a learnable slice indicator code to designate each decoded image into a spatial slice location, while the slice generator is a denoising diffusion model operating on the entirety of slice images stacked on the input channels. We conduct extensive evaluation against state-of-the-art alternatives to demonstrate superiority of our method, especially in recovering complex and severely occluded shape structures, amid ambiguities. All *Slice3D* results were produced by networks trained on a single Nvidia A40 GPU, with an inference time of less than 20 seconds.

1. Introduction

Single-view 3D reconstruction has been one of the extensively studied problems in computer vision. Despite the rapid advances on learning-based approaches to tackle this problem, including the latest attempts to leverage large foundational models and powerful diffusion-based generators, the most fundamental challenge still remains: how to faithfully reconstruct *occluded* parts from just one view?

Conventional encoder-decoder based methods, e.g., [9, 17, 25, 44, 53, 65], rely on direct 3D supervision to learn mappings from global or local image features to 3D, but such features from just a single view often do not have the capacity to handle severe occlusion. Also, recent evi-




Figure 2. Multi-view vs. multi-slice for occlusion revelation, with orange edges showing revealed parts and black edges showing occluded parts. In the limit, multi-view (i.e., infinitely many views) may still leave undisclosed parts (a), e.g., the bulb, while infinitely many slices would *guarantee complete structure revelation* (b). With limited slices (c) and views (d), four each, the multi-slice approach tends to reveal more shape structures.

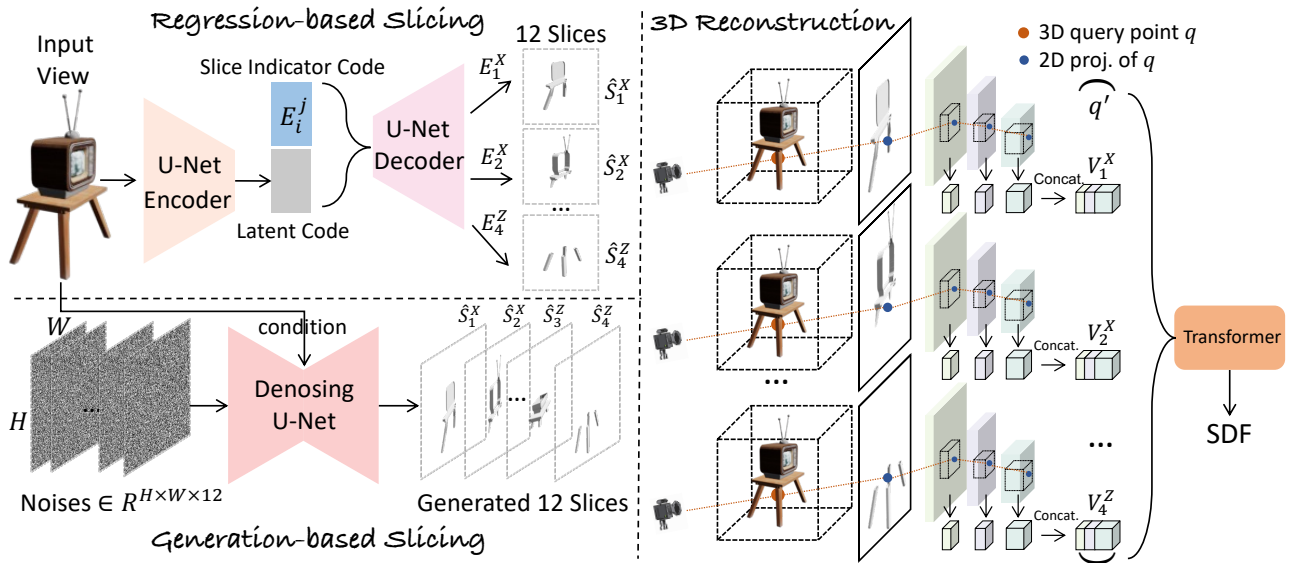


Figure 3. **Slice3D for single-view 3D reconstruction from multi-slice images.** The example shows 4 slice images per X, Y, Z direction, which can be obtained by either a regressor or a generator. Regression is carried out by a U-Net encoder-decoder, where the decoder takes an extra *slice indicator code* that is appended to the encoder output, to obtain a slice image. There is one indicator code for each of the $3 \times 4 = 12$ slices, that is optimized during training. Our generator is a denoising diffusion model which produces a concatenation of slice images following the distribution of the ground-truth (GT) images, when conditioned on the single-view input. Inference from random noises may output multiple plausible results. The 3D reconstruction stage inputs the predicted slice images and produces an implicit field. We first extract hierarchical CNN features from the input slices, then project a 3D query point onto the 2D planes of the feature maps, and finally gather all the query-related features and send them into a transformer [58] to predict the signed distance at the query point.




Figure 4. Our slice image prediction networks are trained using occlusion-revealing multi-slice images produced from a GT 3D object *mesh* along a predetermined set of slicing directions (the X, Y, Z axes of the object bounding box in the example above), while keeping the hollows as holes in the images, through which the *back faces* of the mesh can be visible. The 3D mesh may be textured, but the slice images are all rendered without texture.

dence suggests that many of these methods are essentially performing image classification or retrieval [55]. As such, occlusion recovery is by and large a “hit or miss”. Differentiable rendering enables 3D reasoning from images without 3D supervision [7, 38], but image-space rendering losses alone do not account for occluded structures.

Recently, more and more works model the reconstruction problem as a single-view conditioned *generative* task, aiming to learn a distribution of unseen structures. Typically, *multi-view* images are generated from single view via diffusion. This is followed by multi-view reconstruction, e.g., via neural radiance fields (NeRF) [6, 28, 42], to obtain 3D outputs. However, multi-view consistency is a reoccurring challenge (see Fig. 1), especially over disparate views.

We introduce a novel approach to single-view 3D reconstruction which is *occlusion-revealing* (Fig. 2) by first predicting *multi-slice images* from a single-view input and then integrating the slices into a 3D model; see Fig. 3. Fig. 4 shows slice images provided during training along a set of slicing directions. Each slice represents the projected image, in the input camera view, of one slice of the 3D *mesh* model. For the TV example shown, a completely occluded leg, in the camera view, is revealed in the slice images.

The motivations behind our multi-slice reconstruction are three-fold. First, slicing through a volume or 3D mesh is a natural way to reveal occluded parts. Second, with the slice images forming an intermediate representation, our approach breaks the difficult problem of mapping from single view to 3D into two comparatively simpler tasks: slice prediction followed by 3D reconstruction from slices.

More importantly, when contrasting our *multi-slice* representation with the popular *multi-view* approach for occlu-




Figure 5. Single-view reconstruction of severely occluded parts in a lamp (e.g., the bulb) by Slice3D vs. One-2-3-45 [28], a state-of-the-art multi-view approach. The multi-view images in [28] were obtained by Zero-1-to-3 [29], which was trained on renderings of 800K 3D models [10]. It is conceivable that with a limited number of views (12, the same as our slice count), the bulbs were barely visible due to occlusion by the lamp shades and bases.

sion revelation, we observe that slicing in the limit (i.e., with infinitely many slices) is *guaranteed* to unveil all hidden structures, while the same cannot be said about multi-view captures, as illustrated in Fig. 2 (a-b), for a 2D lamp example. In practice, when slice and view counts are both limited, the slicing approach is more occlusion-revealing in general since it can peel through any occluders without any obstruction, unlike the multi-view case; see (c-d). On the technical front, neither our slice prediction step nor the 3D reconstruction step is concerned with camera view changes or multi-view consistency, which facilitates the use of convolutional features. In Fig. 5, we show the above contrasting characteristics of the multi-slice vs. multi-view approaches manifested in a real 3D lamp reconstruction.

As outlined in Fig. 3, our slice images can be produced by either a regressive network or a generative one, both employing a U-Net architecture, with network losses measured against ground-truth (GT) images obtained by slicing 3D models from the training set; see Fig. 4. In the regressive network, we inject a learnable slice indicator code into the encoder-decoder setup to designate each decoded image into the proper spatial slice location. The slice generator is a denoising diffusion model operating on separate channels for different slice images. In the second stage, 3D reconstruction from the obtained slice images, we design a coordinate-based transformer to obtain an implicit field.

Our network is coined Slice3D. We conduct extensive experiments on ShapeNet [5] and Objverse [10] for evaluation. Quantitative and qualitative comparisons are made to representative methods including encoder-decoders capturing local image features [65], autoregressive latent models [36], generative model inversion [42], and three recent methods with diffusion + NeRF [6, 28, 33], demonstrating superiority of our method in terms of reconstruction quality, generalizability, and handling of ambiguities.

2. Related work

The earliest inspiration for our multi-slice 3D reconstruction idea was GANHopper by Lira et al. [26]. In the con-

text of unsupervised image-to-image translation, the key idea of their work is to *gradually* execute the translation between disparate domains by making small “hops,” rather than completing it in one go. We were then motivated to tackle the single-view 3D reconstruction problem as a domain translation task, hopping through the multi-slice images so as to bridge the gap between 2D and 3D.

Neural fields. Neural implicit functions [63] are learned by a neural network to predict such properties as color, occupancy, opacity, and surface normals. Earlier works such as IM-Net [8], occupancy networks [34], and DeepSDF [41] represent a 3D shape as a latent code and send it to a decoder formed by MLPs to predict occupancy or SDF values for any query point sampled in 3D. NeRF [35] and its variants [15] represent a 3D scene using a neural radiance field for novel view synthesis. Many works on neural 3D reconstruction [15, 63] take on single- or multi-view images as well as point clouds as input. In contrast, the reconstruction module in Slice3D is built on multi-slice images.

Single-view reconstruction from image features. An early work to exploit neural image features for 3D reconstruction was Pixel2Mesh [59], which uses a graph representation. Subsequently, neural implicit fields have been widely adopted for single-view reconstruction. IM-Net [8] first extracts a global latent code, from an input view, to learn a 3D implicit field. Auto-SDF [36] and ShapeFormer [66] employ vector quantization (VQ) [57] to improve the representation of latent features. More recent works [22, 37, 68] consider single-view reconstruction as a 3D generation task conditioned on the latent code of the input view. However, a global latent code is insufficient to capture finer details on a 3D object, often leading to overly smoothed or retrieved outcomes that do not faithfully represent the input. To alleviate this problem, both DISN [65] and D²IM-Net [25] opted to learn local image features, along with global ones, to improve the recovery of fine-grained topological (e.g., thin) features and surface details, respectively.

3D reconstruction via multi-view. Conventional 3D reconstruction approaches [14] involve multi-view reasoning. With recent advances in generative AI, a prevailing practice in single-view 3D reconstruction is to first employ generative models such as GANs [16, 23] and diffusion [11, 20, 45, 50] to generate several novel views from the input view, and then feed them to a multi-view 3D reconstruction module such as NeRF. NeRF-based surface reconstruction [39, 60, 67] can be further employed to output 3D meshes. EG3D [3] and NeRF-from-Image [42] build NeRFs or its variants upon the outputs of a style-GAN [23], while more recent works [4, 6, 18, 28–30, 33, 42, 43, 51, 62, 64] are favoring diffusion models such as Stable Diffusion (SD) [45]. To date, Zero-1-to-3 [29] represents the state of the art on single- to multi-view synthesis; it is built on SD and fine-tuned by 800K 3D models from Objaverse [10].

To alleviate the multi-view consistency problem, which is an reoccurring challenge for the image-space generative methods, recent methods [30, 31, 43, 52, 71] aim to enhance the consistency by performing spatial attention between different views, which dramatically increases the computational overhead. Nevertheless, the consistency of two dispartate views is still far from satisfactory. A deeper reason is that standard convolutional neural networks (CNNs) are not inherently rotation-invariant, making it difficult to build correspondences between significantly different views.

Slice representation. There have been techniques to use object slices to reconstruct a 3D shape, e.g., as inspired by Magnetic Resonance Imaging (MRI). Early works in computer graphics employ either parallel cross-sectional curves, e.g., [24], or non-parallel curve networks, e.g., [27], as inputs. Recently, with the advent of machine learning, OReX [48] and Cut-and-Approximate [40] develop neural reconstruction models to obtain 3D meshes, also from a sparse collection of 2D planar cross-sections, while Neural-Slice [21] uses slicing to obtain 3D shapes from 4D tetrahedral meshes. However, none of these techniques, nor others that we are aware of, address the problem of single-view 3D reconstruction since their inputs all consist of 3D data such as curve slices or point clouds. Moreover, our method learns to predict much *fewer* slice *images* and use them as input to a transformer-based reconstruction module.

Dip. Dip transform [2] is an ingenious way to acquire occluded 3D structures by measuring fluid displacement when an object is repeatedly oriented and dipped into a fluid. Similar to slicing, Dip was also designed to “reach” occluded parts that optical captures via lines of sight cannot. However, both the occlusion revelation (liquid penetration vs. slicing) and target applications (acquisition vs. single-view reconstruction) are completely different.

Multi-plane image (MPI). Introduced by Zhou et al. [72], an MPI representation for a scene consists of a set of fronto-parallel planes at fixed depths, with each image intended to capture scene appearance and visibility at a specific depth. Conceptually, both MPIs and our slice images provide a layered scene representation for occlusion revelation. However, MPIs [13, 19, 56, 69, 72] have only been employed for novel view synthesis, not 3D reconstruction, and the prediction networks were trained on *multi-view images*, e.g., from Youtube videos [72]. The generated MPIs only reveal a small amount of occluded scenes to account for view shifts. In contrast, our method directly cuts through a 3D object to reveal all the occluded parts at several depth ranges.

3. Method

Our method is based on a slicing representation; see Section 3.1 and Fig. 4. Given a single-view RGB image $I \in \mathbb{R}^{H \times W \times 3}$ of a 3D object, we first learn a slicing network

(Section 3.2) which takes I as input and outputs its corresponding multi-slice images. Then, an SDF prediction network is trained (Section 3.3) to leverage image features from the slices to produce an implicit field of the object.

3.1. Slicing Representation

As shown in Fig. 4, given a 3D object *mesh*, we slice it along the X , Y , and Z axes of its bounding box and render the sliced segments. The slice images can optionally include textures or have holes filled. By default, we remove the textures and keep the holes, which could reveal back faces of the mesh, and overall, it produces the best experimental results. The rendered slice images consist of $\{S_i^j | 1 \leq i \leq N_s, j \in \{X, Y, Z\}\}$, where $S_i^j \in \mathbb{R}^{H \times W \times 1}$, j denotes the axis and N_s is the number of slices along each axis. In our current implementation, we produce $N_s = 4$ slices per direction. Compared to planar cross-sections or tomography scans which consist of densely stacked planar images, our slicing representation is more compact with each slice image revealing more object semantics.

3.2. Slice Regression and Generation

Given an image I , observing a 3D shape from a single view, we aim to train a slicing network F_s that can produce the slice images from I : $F_s(I) = \{\hat{S}_i^j\}$. We have two different designs for such a network: regressive and generative. The former aims to regress GT slices, assuming deterministic occluded parts based on the partial observation of the input view. The latter aims to *generate* slices that are consistent with the input view, considering various possible occluded parts that do not necessarily match the GT slices.

3.2.1 Regression-based Slicing

We employ a U-Net [46] encoder-decoder to regress the slices from I . As shown in the top-left of Fig. 3, an indicator code $E_i^j \in \mathbb{R}^{N_e}$ is concatenated with the latent code of I to indicate which slice to output: $\hat{S}_i^j = F_d([F_e(I); E_i^j])$, where F_e and F_d are the encoder and decoder of F_s , respectively, and $[\cdot; \cdot]$ denotes concatenation.

The indicator codes $\{E_i^j\}$ are randomly initialized and jointly optimized with the whole network. The loss function \mathcal{L}_{reg} is the reconstruction error measured by the \mathcal{L}_1 -Norm and perceptual loss [70] (\mathcal{L}_p) for all slices:

$$\mathcal{L}_{reg} = \frac{\sum_{i,j} \|\hat{S}_i^j - S_i^j\|_1 + \mathcal{L}_p(\hat{S}_i^j, S_i^j)}{|S|}, \quad (1)$$

where $|S| = 3 \cdot N_s$ is the total number of slices.

3.2.2 Generation-based Slicing

Single-view occlusion revelation is clearly ill-posed, with potential ambiguities leading to multiple plausible 3D re-

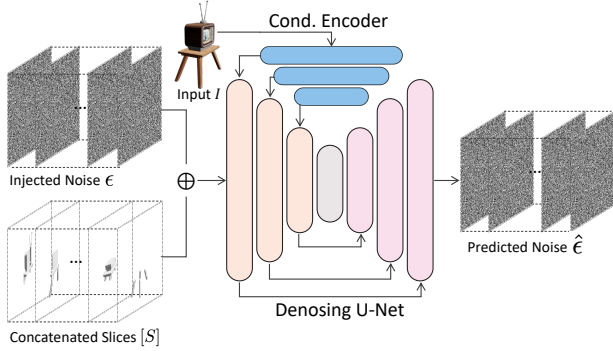


Figure 6. The pipeline of generation-based slicing via DDPM. \oplus denotes the injection of noise to slice images.

constructions. To this end, we consider slicing as an image-conditioned generation task and develop a generation-based slicing via Denoising Diffusion Probabilistic Models (DDPM) [20], whose detailed pipeline is shown in Fig. 6.

The key challenge of applying DDPM lies in how to generate *consistent* slices. Rather than diffusing each single slice, we propose to concatenate all slices into an aggregated structure $[S]$, along either the color dimension ($[S] \in \mathbb{R}^{H \times W \times |S|}$) or a spatial dimension such as height ($[S] \in \mathbb{R}^{(|S| \cdot H) \times W \times 1}$). The concatenation of slices instructs DDPM to learn the correlation of slices, yielding consistent generation results. By default, we concatenate along the color dimension. The comparison of these two concatenation ways can be found in the supplementary.

Following DDPM, we first inject a noise $\epsilon \sim \mathcal{N}(0, 1)$ to $[S]$ to obtain $[\bar{S}]$. The noised slices $[\bar{S}]$ are then sent to a U-Net based denoising network F_{dn} that outputs the predicted noise conditioned on I : $\hat{\epsilon} = F_{dn}([\bar{S}], I)$. In practice, the injection and prediction of noise ϵ involve an iteration step t , which we omit here for brevity. To better condition I for denoising diffusion, we employ a pre-trained VGG to extract multi-layer features from I and add them into the corresponding layers in the encoder of F_{dn} (see the blue part of Fig. 6). The loss function is the denoising loss: $\mathcal{L}_{gen} = \|\hat{\epsilon} - \epsilon\|_2^2$. After training, we can sample random noises to generate a set of slice images $\{\hat{S}_i^j\}$.

Compared to regression-based slicing, generation-based slicing can produce multiple plausible outputs rather than a single deterministic result, but requires much longer inference time due to the nature of diffusion models.

3.3. Learning Implicit Field from Slices

We learn an implicit field from slice images, by projecting a 3D query point onto each slice image, extracting query-specific features, and then aggregating them to predict a signed distance. For regression-based slicing, we train the SDF prediction network on the predicted slices from the previous section thus the whole pipeline can be trained end-to-end. For generation-based slicing, we train the SDF pre-

diction network on GT slice images since the generated slices may not match the GT signed distances.

Specifically, for a query point $q \in \mathbb{R}^3$, we project it to a location $\tilde{q} \in \mathbb{R}^2$ on the 2D image plane according to the camera parameters Θ . For slice \hat{S}_i^j (regression) or S_i^j (generation), we first extract its feature maps, then retrieve features from each feature map corresponding to location \tilde{q} , and finally concatenate them to obtain the image feature vector $V_i^j \in \mathbb{R}^{1 \times 1 \times N_c}$. The retrieval is accomplished via bi-linear interpolation over the features grids that encompass the location of \tilde{q} . We design a transformer network [58] F_t which learns to aggregate the feature vectors from all slices to predict the signed distance \hat{d} : $\hat{d} = F_t([q', V_1^X, \dots, V_{N_c}^Z])$, where $q' \in \mathbb{R}^{N_c}$ is linearly projected from q to be the same dimension with V_i^j . The optimization objective is to minimize the \mathcal{L}_1 distance between predicted and GT signed distance: $\mathcal{L}_{imp} = \|\hat{d} - d\|_1$. In regression-based slicing, \mathcal{L}_{imp} and \mathcal{L}_{reg} are jointly optimized.

DISN [65] extracts from only the input view I to obtain a feature vector V . For query points located around occluded parts of the object (e.g., the query point in Fig. 3), V is irrelevant to q since it is extracted from the front parts that block q , often leading to erroneous SDF prediction. In contrast, one or more of our slice images could reveal the occluded parts, yielding more accurate features fed to the transformer for SDF prediction. After learning the implicit field, we apply Marching Cubes [32] to obtain triangle meshes. Estimating the extrinsics (camera poses) of Θ can output 3D shapes in canonical poses, which is optional and applied only on ShapeNet [5]; See more details in Supplemental.

4. Results and Evaluation

Dataset. We utilize two well-known datasets, ShapeNet [5] and Objaverse [10], to train and test our network. For ShapeNet, our method is applied to Chairs and Airplanes separately, using the train/validation/test split as in SHREC16 [47]. Objaverse is a challenging dataset, where many objects may not belong to a clear category. In our experiment, we randomly selected $\approx 40K$ shapes from Objaverse to keep resource and time costs manageable, while ensuring diversity of the collection. We follow a 95%-5% train-test split over the 40K selected shapes, ensuring that every training model we use is also from the training set of Zero-1-2-3 [29]. Note that Zero-1-2-3 used a train-test split of 99%-1% over the entire 800K models in Objaverse.

Rendering Slice Images. We employ the *bisect* function in Blender [1] to cut an object mesh into slices and then render them. By default, we keep internal cavities and holes produced by slicing. Since many 3D models in Objaverse are not in, or have no, canonical poses and feature random rotations, slicing them along the default X , Y , and Z directions is not suitable. To this end, we opted for *camera-aligned* slicing instead. Specifically, we set the camera position, in-

Datasets Metrics	ShapeNet Chairs			ShapeNet Airplanes		
	CD↓	F1↑	HD↓	CD↓	F1↑	HD↓
DISN [65]	3.25	2.34	7.84	3.48	2.81	7.62
Auto-SDF [36]	25.0	0.77	15.8	-	-	-
NeRF-Img [42]	19.5	1.27	13.4	16.0	1.41	14.0
One-2-3-45 [28]	12.5	1.01	11.7	23.4	0.75	12.6
Ours	2.80	2.43	6.75	2.99	2.83	6.45
DISN* [65]	3.02	2.42	7.86	3.19	2.96	7.56
SSDNeRF* [6]	8.98	2.06	10.9	10.4	2.00	10.4
Ours*	2.67	2.53	6.70	2.81	2.96	6.38

Table 1. Quantitative results of single-view 3D reconstruction on ShapeNet Chairs and Airplanes. ‘*’ denotes using ground-truth (GT) camera poses. Slice3D (ours) uses regression-based slicing.

roduce a random rotation to the object, position it in front of the camera, and utilize the camera’s axes for slicing. Under this setting, estimating the rotation of objects is unnecessary; see more details in the supplementary material.

Implementation Details. When preparing training data, we use the Vega FEM Library [49] for SDF computation, same as DISN [65] and AutoSDF [36]. For each object, the diagonal length of the object’s 3D bounding box is normalized to 1. We set H and W to 128, N_e to 128, N_s to 4, and N_p to 1,024. We train our model on *only one* Nvidia A40 GPU.

Evaluation Metrics. We use Chamfer \mathcal{L}_2 Distance (CD), F-score%1 (F1) [54], and Hausdorff Distance (HD) to evaluate reconstruction results. Besides these 3D metrics, we also evaluate using 2D metrics, PSNR, SSIM [61], and LPIPS [70], on 24 rendered views of reconstructed meshes.

4.1. Comparisons to State Of The Art (SOTA)

We compare our method with SOTA methods including DISN [65], AutoSDF [36], NeRF-From-Img [42], SSDNeRF [6], RealFusion [33], and One-2-3-45 [28]. Note that when showing visual results, we always present two views to remove any view bias towards the input.

ShapeNet. Figs. 7 and 8 demonstrate the effectiveness of our approach in reconstructing obscured sections of objects while precisely restoring visible ones. In contrast, DISN [65] can barely reconstruct the occluded parts. AutoSDF [36] and NeRF-From-Image [42] rely on latent features for reconstruction, which can yield rather clean results but they do not necessarily respect the input image, as the chair results from Fig. 7 evidently showed. SSDNeRF [6] tends to produce erroneous geometric details and the surfaces are not as smooth as from other methods, especially for the airplane examples in Fig. 8. Generally speaking, One-2-3-45 [28] yields rough and bulky geometries that clearly lack details. Additionally, our method outperforms the other methods in the reconstruction of unconventional inputs, as shown by the two airplanes in Fig. 8.

Tab. 1 shows quantitative comparisons to the above methods on the two ShapeNet categories. When our method operates without using GT camera poses and relies on esti-

Method	CD↓	F1↑	HD↓	PSNR↑	SSIM↑	LPIPS↓
DISN [65]	28.2	1.47	23.0	21.9	8.79	15.2
One-2-3-45 [28]	35.4	0.87	29.8	19.8	8.55	19.8
Ours (G)	25.0	1.51	16.4	22.3	8.84	13.4
Ours (R)	16.6	1.53	13.4	22.6	8.88	12.8

Table 2. Quantitative results of single-view 3D reconstruction on the Objaverse dataset. ‘R’ and ‘G’ denote regression-based and generation-based slicing in Slice3D, respectively.

mated poses, its performance sees only a marginal degradation. Notably, for Chair, our method with estimated poses even outperforms DISN and SSDNeRF with GT poses. It is also notable that One-2-3-45 tends to output dilated shapes. Although it falls behind DISN in quantitative results, it usually generates more complete shapes and more accurate topologies than DISN; see Figs. 7 and 8.

Objaverse. Since AutoSDF [36], NeRF-From-Image [42], and SSDNeRF [6] are all category-specific methods, they cannot be trained on Objaverse which consists of many 3D objects with no clear categories. Thus, we mainly compare our method with DISN [65], RealFusion [33], and One-2-3-45 [28]; see Fig. 9 for a qualitative comparison.

Note that RealFusion [33] and One-2-3-45 [28] both leverage a significantly larger volume of training data compared to our method, as they both rely on a pre-trained Stable Diffusion [45] model. Despite this disparity in the scale of the training sets, our method still outperforms these two methods, as well as DISN; see Fig. 9. Quantitatively, Tab. 2 demonstrates our method’s across-the-board advantage over DISN and One-2-3-45 (more so on Objaverse than on ShapeNet), in both 3D and 2D metrics, especially CD, HD, and LPIPS. These results also show that our method possesses the capacity to generalize across common shapes rather than being confined to specific categories.

Real-world Image Input. To further stress-test the methods, we perform Google image search on object keywords (e.g., “chair”, “airplane”, etc.) and take the top-1 image with occlusions as input. In addition, we use some examples from Google Scanned Objects (GSO) [12] dataset. Fig. 10 shows that our method performs the best in recovering occluded parts while preserving fine details, e.g., the turbines of the airplane, the heel, and the opening of the boot.

Multi-slice vs. Multi-view Reconstruction. Fig. 11 compares between multi-slice and multi-view reconstructions. We demonstrate two different sets of multi-slices and multi-views generated from our method and One-2-3-45 [28]. Our generated slice images depict two distinct configurations of chair legs, maintaining high consistency across slices. In contrast, the multiple views synthesized in One-2-3-45 do encounter inconsistency issues, compromising an accurate 3D reconstruction. We show three slices in each set and more can be found in the supplementary material.

Training and Inference Costs. Unlike RealFusion [33] and One-2-3-45 [28], our method does not depend on a large




Figure 7. Visual comparison between single-view 3D reconstruction methods on two ShapeNet chairs. DISN and our method (based on regressive slicing) utilize the same estimated camera parameters. Two different views are displayed to remove view bias.




Figure 8. Visual comparison between single-view 3D reconstruction methods on two somewhat unconventional airplanes from ShapeNet.

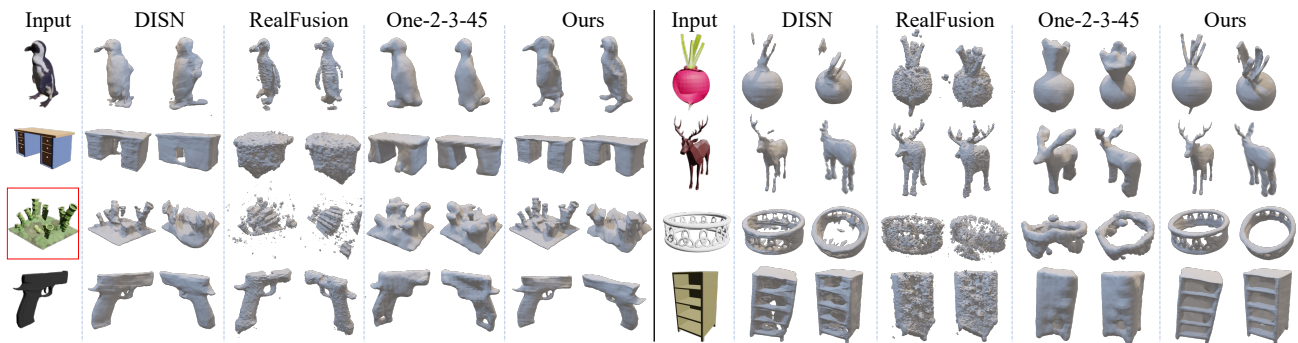


Figure 9. Visual comparison between cross-category reconstruction methods on Objaverse. Our method better recovers overall geometries, fine details, and is more faithful to the inputs. Except for the seaweed (squared), all the other models are in the *training* set of Zero-1-to-3.




Figure 10. On real-world image inputs. Top two from Google image search and the others from GSO [12]. Zoom in for details.

pre-trained model [45]. NeRF-based methods tend to have long inference times, as many of them require per-test optimization. Our method has much faster inference times: **10 or 20 seconds** with slice regression or generation, com-

pared to 90 minutes for RealFusion and 45 seconds for One-2-3-45 and SSDNeRF, while delivering superior results. With similar inference times to DISN, Auto-SDF, and NeRF-From-Img, our method has better reconstruction quality. See supplementary material for more details.

4.2. Ablation Study

Generative vs. Regressive Slicing. As shown in the first row of Fig. 11, for generation-based slicing, we can produce multiple plausible results for a given input view. Tab. 2 shows that regressive slicing outperforms generative slicing when compared to GT. However, it is notable that the 3D outputs may differ from GT when input views have ambiguous occluded parts. Regressive training struggles to handle such one-to-many mappings, leading to generation outperforming regression in reconstructing occluded parts.

Number of Slices. We test the performance of our method

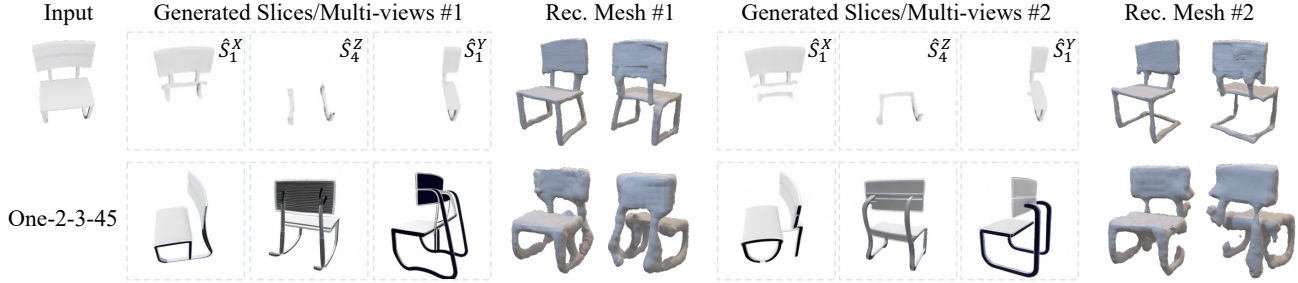


Figure 11. Multi-slice vs. multi-view reconstructions amid ambiguities in the chair legs. Both One-2-3-45 (bottom) and Slice3D (top) can produce multiple results. Our results are both plausible from consistent slices, while One-2-3-45 suffers from multi-view inconsistencies.

Setting	CD↓	F1↑	HD↓	Setting	CD↓	F1↑	HD↓
One-hot ind.	2.68	2.53	6.82	only X	2.88	2.39	7.74
DISN+Depth	2.99	2.43	7.81	only Y	2.75	2.45	6.84
$N_s = 2$	2.90	2.46	7.31	only Z	3.08	2.38	7.60
$N_s = 4$	2.67	2.53	6.70	Filled	2.73	2.50	7.37
$N_s = 8$	2.56	2.56	6.93	w/ textures	2.82	2.50	7.70

Table 3. Ablation studies on ShapeNet Chairs.

with regression-based slicing by varying N_s between 2, 4, and 8, on ShapeNet Chairs. From Tab. 3, it is evident that larger N_s leads to better reconstruction quality. However, $N_s = 8$ will dramatically increase the computational cost of our method since its complexity is $O(N_s^2)$, as pair-wised attentions are performed on the slices in the transformer. As a trade-off between reconstruction quality and computational complexity, we set $N_s = 4$ by default.

Tab. 3 also covers several other ablation studies.

Learned v.s. One-hot Slice Indicator. Learned indicator codes (the default setting) achieve slightly better performance than using one-hot codes to indicate slice locations. This can be seen by comparing numbers from the first half row of Tab. 3 to those in the last row of Tab. 1.

Slices v.s. Depth Images. Since our slice images offer layered information, it is logical to perform a comparison against depth images as supplementary information. We compared the performance of DISN when concatenating estimated depth images with the input view. Results are shown in the “DISN+Depth” half row, which indicates that depth images can only provide marginal improvement.

Three Axes v.s. One Axis. We also investigated the option of slicing along only one axis. Not surprisingly, Tab. 3 shows that this underperforms against using three axes.

Solid v.s. Hollow Slices. Tab. 3 also demonstrates that not filling the holes improves the performance as more information about the inner/outer surfaces is revealed this way.

Using Textured Slices. We can optionally keep object textures in the slice images. Experimental results show that removing the textures in the slice images can achieve better reconstruction. The reasons could be: 1) learning to produce slice images without textures could help the model better distinguish geometries from textures; 2) for some 3D

models the color of inner and outer surfaces could be different which may complicate the learning process.

5. Discussion and Future Work

Our multi-slice idea represents a significant departure from conventional approaches to single-view 3D reconstruction. Through conceptual illustrations (Fig. 5) and extensive experiments, we demonstrate its compelling properties in recovering occluded object structures. Note again that all the Slice3D results on Objaverse [10] were obtained by networks trained on a very *small subset* (40K or 5%) of the 800K 3D models therein, whereas One-2-3-45 [28] utilizes Zero-1-to-3 [29] for multi-view synthesis and the latter was fine-tuned on the entirety of the 800K. Yet our results (Table 2) show that Slice3D outperforms One-2-3-45 across the board, suggesting that our model generalizes better, likely owing to its superior occlusion-revealing abilities.

Our current implementation only uses four slices per direction, which is quite coarse. Since there is no camera rotation, parts of an object far away from the camera are in smaller scale, e.g., the rear leg of the chair in Fig. 10 and of the buck in Fig. 9. This can hinder our network’s ability to extract detailed features. More refined and gradual slicing over those parts may be needed to improve matters.

The foremost limitation of slicing however, is that realistically, it can only be executed at scale on *digital* 3D models, not physical ones. On the other hand, the multi-view approach can be applied to physical 3D objects by capturing photos only, without ever needing 3D supervision, as in RealFusion [33]. However, this is not without a high cost: 90 minutes by RealFusion vs. 45 seconds by One-2-3-45 (requiring 3D fine-tuning for multi-view synthesis) and less than 20 seconds by Slice3D. We envision synthesizing 3D models, not many but carefully crafted few, for the multi-slice approach as an intriguing direction to explore.

Acknowledgements. We thank all the anonymous reviewers and area chairs for their valuable comments. We thank Daniel Cohen-Or and Xingguang Yan for helpful discussions. This work was supported in part by NSERC.

References

- [1] Bisect - blender manual. <https://docs.blender.org/manual/en/2.80/modeling/meshes/editing/subdividing/bisect.html>. Accessed: 2023-11-15. **5**
- [2] Kfir Aberman, Oren Katzir, Qiang Zhou, Zegang Luo, Andrei Sharf, Chen Greif, Baoquan Chen, and Daniel Cohen-Or. Dip transform for 3D shape reconstruction. *ACM Transactions on Graphics*, 36(4):Article 79, 2017. **4**
- [3] Eric R Chan, Connor Z Lin, Matthew A Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas J Guibas, Jonathan Tremblay, Sameh Khamis, et al. Efficient geometry-aware 3d generative adversarial networks. In *CVPR*, pages 16123–16133, 2022. **3**
- [4] Eric R Chan, Koki Nagano, Matthew A Chan, Alexander W Bergman, Jeong Joon Park, Axel Levy, Miika Aittala, Shalini De Mello, Tero Karras, and Gordon Wetzstein. Generative novel view synthesis with 3d-aware diffusion models. *arXiv preprint arXiv:2304.02602*, 2023. **3**
- [5] 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. **3, 5**
- [6] Hansheng Chen, Jiatao Gu, Anpei Chen, Wei Tian, Zhuowen Tu, Lingjie Liu, and Hao Su. Single-stage diffusion NeRF: A unified approach to 3d generation and reconstruction. *arXiv preprint arXiv:2304.06714*, 2023. **2, 3, 6**
- [7] Wenzheng Chen, Jun Gao, Huan Ling, Edward J. Smith, Jaakko Lehtinen, Alec Jacobson, and Sanja Fidler. Learning to predict 3d objects with an interpolation-based differentiable renderer. *Advances in Neural Information Processing Systems*, 2019. **2**
- [8] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *CVPR*, pages 5939–5948, 2019. **3**
- [9] Zhiqin Chen, Andrea Tagliasacchi, and Hao Zhang. Bsp-net: Generating compact meshes via binary space partitioning. In *CVPR*, 2020. **1**
- [10] Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of annotated 3d objects. In *CVPR*, pages 13142–13153, 2023. **3, 5, 8**
- [11] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794, 2021. **3**
- [12] Laura Downs, Anthony Francis, Nate Koenig, Brandon Kinman, Ryan Hickman, Krista Reymann, Thomas B McHugh, and Vincent Vanhoucke. Google scanned objects: A high-quality dataset of 3d scanned household items. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 2553–2560. IEEE, 2022. **6, 7**
- [13] John Flynn, Michael Broxton, Paul Debevec, Matthew Duvall, Graham Fyffe, Ryan Overbeck, Noah Snavely, and Richard Tucker. Deepview: View synthesis with learned gradient descent. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2367–2376, 2019. **4**
- [14] Yasutaka Furukawa, Carlos Hernández, et al. Multi-view stereo: A tutorial. *Foundations and Trends® in Computer Graphics and Vision*, 9(1-2):1–148, 2015. **3**
- [15] Kyle Gao, Yina Gao, Hongjie He, Dening Lu, Linlin Xu, and Jonathan Li. Nerf: Neural radiance field in 3d vision, a comprehensive review. *arXiv preprint arXiv:2210.00379*, 2022. **3**
- [16] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014. **3**
- [17] Thibault Groueix, Matthew Fisher, Vladimir G. Kim, Bryan Russell, and Mathieu Aubry. Atlasnet: A papier-mâché approach to learning 3d surface generation. In *CVPR*, 2018. **1**
- [18] Jiatao Gu, Alex Trevithick, Kai-En Lin, Joshua M Susskind, Christian Theobalt, Lingjie Liu, and Ravi Ramamoorthi. Nerfdiff: Single-image view synthesis with nerf-guided distillation from 3d-aware diffusion. In *International Conference on Machine Learning*, pages 11808–11826. PMLR, 2023. **3**
- [19] Yuxuan Han, Ruicheng Wang, and Jiaolong Yang. Single-view view synthesis in the wild with learned adaptive multipane images. In *ACM SIGGRAPH 2022 Conference Proceedings*, pages 1–8, 2022. **4**
- [20] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. **3, 5**
- [21] Chenbo Jiang, Jie Yang, Shwai He, Yu-Kun Lai, and Lin Gao. Neurslice: Neural 3d triangle mesh reconstruction via slicing 4d tetrahedral meshes. 2023. **4**
- [22] Heewoo Jun and Alex Nichol. Shap-e: Generating conditional 3d implicit functions. *arXiv preprint arXiv:2305.02463*, 2023. **3**
- [23] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, pages 4401–4410, 2019. **3**
- [24] Shay Kels and Nira Dyn. Reconstruction of 3d objects from 2d cross-sections with the 4-point subdivision scheme adapted to sets. *Computers & Graphics*, 35:741–746, 2011. **4**
- [25] Manyi Li and Hao Zhang. D²IM-Net: Learning detail disentangled implicit fields from single images. In *CVPR*, 2021. **1, 3**
- [26] Wallace Lira, Johannes Merz, Daniel Ritchies, Daniel Cohen-Or, and Hao Zhang. GANHopper: Multi-hop gan for unsupervised image-to-image translation. In *ECCV*, 2020. **3**
- [27] Lu Liu, Chandrajit Bajaj, Joseph O. Deasy, Daniel A. Low, and Tao Ju. Surface reconstruction from non-parallel curve networks. *Computer Graphics Forum*, 27:155–163, 2008. **4**
- [28] Minghua Liu, Chao Xu, Haian Jin, Linghao Chen, Zexiang Xu, Hao Su, et al. One-2-3-45: Any single image to 3d mesh in 45 seconds without per-shape optimization. *arXiv preprint arXiv:2306.16928*, 2023. **1, 2, 3, 6, 8**

- [29] Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. Zero-1-to-3: Zero-shot one image to 3d object. *arXiv preprint arXiv:2303.11328*, 2023. 1, 3, 5, 8
- [30] Yuan Liu, Cheng Lin, Zijiao Zeng, Xiaoxiao Long, Lingjie Liu, Taku Komura, and Wenping Wang. Syncdreamer: Generating multiview-consistent images from a single-view image, 2023. 3, 4
- [31] Xiaoxiao Long, Yuan-Chen Guo, Cheng Lin, Yuan Liu, Zhiyang Dou, Lingjie Liu, Yuexin Ma, Song-Hai Zhang, Marc Habermann, Christian Theobalt, et al. Wonder3d: Single image to 3d using cross-domain diffusion. *arXiv preprint arXiv:2310.15008*, 2023. 4
- [32] 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. 5
- [33] Luke Melas-Kyriazi, Iro Laina, Christian Rupprecht, and Andrea Vedaldi. Realfusion: 360deg reconstruction of any object from a single image. In *CVPR*, pages 8446–8455, 2023. 3, 6, 8
- [34] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *CVPR*, pages 4460–4470, 2019. 3
- [35] 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
- [36] Paritosh Mittal, Yen-Chi Cheng, Maneesh Singh, and Shubham Tulsiani. Autosdf: Shape priors for 3d completion, reconstruction and generation. In *CVPR*, pages 306–315, 2022. 3, 6
- [37] Alex Nichol, Heewoo Jun, Prafulla Dhariwal, Pamela Mishkin, and Mark Chen. Point-e: A system for generating 3d point clouds from complex prompts. *arXiv preprint arXiv:2212.08751*, 2022. 3
- [38] M. Niemeyer, L. Mescheder, M. Oechsle, and A. Geiger. Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In *CVPR*, 2020. 2
- [39] Michael Oechsle, Songyou Peng, and Andreas Geiger. Unisurf: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. In *ICCV*, pages 5589–5599, 2021. 3
- [40] Azimkhon Ostonov. Cut-and-approximate: 3d shape reconstruction from planar cross-sections with deep reinforcement learning. *arXiv preprint arXiv:2210.12509*, 2022. 4
- [41] 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. 3
- [42] Dario Pavllo, David Joseph Tan, Marie-Julie Rakotosaona, and Federico Tombari. Shape, pose, and appearance from a single image via bootstrapped radiance field inversion. In *CVPR*, pages 4391–4401, 2023. 2, 3, 6
- [43] Guocheng Qian, Jinjie Mai, Abdullah Hamdi, Jian Ren, Aliaksandr Siarohin, Bing Li, Hsin-Ying Lee, Ivan Skokhodov, Peter Wonka, Sergey Tulyakov, et al. Magic123: One image to high-quality 3d object generation using both 2d and 3d diffusion priors. *arXiv preprint arXiv:2306.17843*, 2023. 3, 4
- [44] Stephan R. Richter and Stefan Roth. Matryoshka networks: Predicting 3d geometry via nested shape layers. In *CVPR*, 2018. 1
- [45] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, pages 10684–10695, 2022. 3, 6, 7
- [46] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015. 4
- [47] Manolis Savva, Fisher Yu, Hao Su, M Aono, B Chen, D Cohen-Or, W Deng, Hang Su, Song Bai, Xiang Bai, et al. Shrec16 track: largescale 3d shape retrieval from shapenet core55. In *Proceedings of the eurographics workshop on 3D object retrieval*, 2016. 5
- [48] Haim Sawdayee, Amir Vaxman, and Amit H. Bermano. Ores: Object reconstruction from planar cross-sections using neural fields. In *CVPR*, pages 20854–20862, 2023. 4
- [49] Fun Shing Sin, Daniel Schroeder, and Jernej Barbič. Vega: non-linear fem deformable object simulator. In *Computer Graphics Forum*, pages 36–48. Wiley Online Library, 2013. 6
- [50] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pages 2256–2265. PMLR, 2015. 3
- [51] Junshu Tang, Tengfei Wang, Bo Zhang, Ting Zhang, Ran Yi, Lizhuang Ma, and Dong Chen. Make-it-3d: High-fidelity 3d creation from a single image with diffusion prior. *arXiv preprint arXiv:2303.14184*, 2023. 3
- [52] Shitao Tang, Fuyang Zhang, Jiacheng Chen, Peng Wang, and Yasutaka Furukawa. Mvdifffusion: Enabling holistic multi-view image generation with correspondence-aware diffusion. *arXiv preprint arXiv:2307.01097*, 2023. 4
- [53] Maxim Tatarchenko, Alexey Dosovitskiy, and Thomas Brox. Octree generating networks: Efficient convolutional architectures for high-resolution 3d outputs. In *ICCV*, 2017. 1
- [54] Maxim Tatarchenko, Stephan R Richter, René Ranftl, Zhuwen Li, Vladlen Koltun, and Thomas Brox. What do single-view 3d reconstruction networks learn? In *CVPR*, pages 3405–3414, 2019. 6
- [55] Maxim Tatarchenko, Stephan R. Richter, René Ranftl, Zhuwen Li, Vladlen Koltun, and Thomas Brox. What do single-view 3d reconstruction networks learn? In *CVPR*, 2019. 2
- [56] Richard Tucker and Noah Snavely. Single-view view synthesis with multiplane images. In *Proceedings of the IEEE/CVF*

- Conference on Computer Vision and Pattern Recognition*, pages 551–560, 2020. 4
- [57] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017. 3
- [58] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 2, 5
- [59] Nanyang Wang, Yinda Zhang, Zhuwen Li, Yanwei Fu, Wei Liu, and Yu-Gang Jiang. Pixel2mesh: Generating 3d mesh models from single rgb images. In *ECCV*, pages 52–67, 2018. 3
- [60] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *arXiv preprint arXiv:2106.10689*, 2021. 3
- [61] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 6
- [62] Daniel Watson, William Chan, Ricardo Martin-Brualla, Jonathan Ho, Andrea Tagliasacchi, and Mohammad Norouzi. Novel view synthesis with diffusion models. *arXiv preprint arXiv:2210.04628*, 2022. 3
- [63] Yiheng Xie, Towaki Takikawa, Shunsuke Saito, Or Litany, Shiqin Yan, Numair Khan, Federico Tombari, James Tompkin, Vincent Sitzmann, and Srinath Sridhar. Neural fields in visual computing and beyond. *Comput. Graph. Forum*, 2022. 3
- [64] DeJia Xu, Yifan Jiang, Peihao Wang, Zhiwen Fan, Yi Wang, and Zhangyang Wang. Neurallift-360: Lifting an in-the-wild 2d photo to a 3d object with 360deg views. In *CVPR*, pages 4479–4489, 2023. 3
- [65] Qiangeng Xu, Weiyue Wang, Duygu Ceylan, Radomir Mech, and Ulrich Neumann. Disn: Deep implicit surface network for high-quality single-view 3d reconstruction. *Advances in neural information processing systems*, 32, 2019. 1, 3, 5, 6
- [66] Xingguang Yan, Liqiang Lin, Niloy J Mitra, Dani Lischinski, Daniel Cohen-Or, and Hui Huang. Shapeformer: Transformer-based shape completion via sparse representation. In *CVPR*, pages 6239–6249, 2022. 3
- [67] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. *Advances in Neural Information Processing Systems*, 34:4805–4815, 2021. 3
- [68] Biao Zhang, Jiapeng Tang, Matthias Niessner, and Peter Wonka. 3dshape2vecset: A 3d shape representation for neural fields and generative diffusion models. *arXiv preprint arXiv:2301.11445*, 2023. 3
- [69] Mingfang Zhang, Jinglu Wang, Xiao Li, Yifei Huang, Yoichi Sato, and Yan Lu. Structural multiplane image: Bridging neural view synthesis and 3d reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16707–16716, 2023. 4
- [70] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, pages 586–595, 2018. 4, 6
- [71] Minda Zhao, Chaoyi Zhao, Xinyue Liang, Lincheng Li, Zeng Zhao, Zhipeng Hu, Changjie Fan, and Xin Yu. Efficientdreamer: High-fidelity and robust 3d creation via orthogonal-view diffusion prior. *arXiv preprint arXiv:2308.13223*, 2023. 4
- [72] Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. Stereo magnification: Learning view synthesis using multiplane images. *arXiv preprint arXiv:1805.09817*, 2018. 4