MVDD: Multi-View Depth Diffusion Models

Zhen Wang^{1,2†}, Qiangeng Xu¹, Feitong Tan¹, Menglei Chai¹, Shichen Liu¹, Rohit Pandey¹, Sean Fanello¹, Achuta Kadambi^{1,2}, and Yinda Zhang¹

¹ Google

² University of California, Los Angeles https://mvdepth.github.io/

Abstract. Denoising diffusion models have demonstrated outstanding results in 2D image generation, yet it remains a challenge to replicate its success in 3D shape generation. In this paper, we propose leveraging multi-view depth, which represents complex 3D shapes in a 2D data format that is easy to denoise. We pair this representation with a diffusion model, MVDD, that is capable of generating high-quality dense point clouds with 20K+ points with fine-grained details. To enforce 3D consistency in multi-view depth, we introduce an epipolar line segment attention that conditions the denoising step for a view on its neighboring views. Additionally, a depth fusion module is incorporated into diffusion steps to further ensure the alignment of depth maps. When augmented with surface reconstruction, MVDD can also produce highquality 3D meshes. Furthermore, MVDD stands out in other tasks such as depth completion, and can serve as a 3D prior, significantly boosting many downstream tasks, such as GAN inversion. State-of-the-art results from extensive experiments demonstrate MVDD's excellent ability in 3D shape generation, depth completion, and its potential as a 3D prior for downstream tasks.

Keywords: 3D Shape Generation · Diffusion Model

1 Introduction

3D shape generative models have made remarkable progress in the wave of AI-Generated Content. A powerful 3D generative model is expected to possess the following attributes: (i) *Scalability*. The model should be able to create objects with fine-grained details; (ii) *Faithfulness*. The generated 3D shapes should exhibit high fidelity and resemble the objects in the dataset; and (iii) *Versatility*. The model can be plugged in as a 3D prior in various downstream 3D tasks through easy adaptation. Selecting suitable probabilistic models becomes the key factor in achieving these criteria. Among popular generative methods such as GANs [10,20], VAEs [17], and normalizing flows [33], denoising diffusion models [11,45] explicitly model the data distribution; therefore, they are able to faithfully generate samples that reflect content diversity.

[†]Work done while the author was an intern at Google.



Fig. 1: Our proposed MVDD is versatile and can be utilized in various applications: (a) 3D shape generation: our model generates high-quality 3D shape with approximately 10X more points than diffusion-based point cloud generative models e.g., LION [56] and PVD [60] and contains diverse and fine-grained details. (b) Shape completion: we showcase shape completion results from partial inputs, highlighting the higher fidelity compared to PVD [60]. (c) Our model can serve as a powerful shape prior for downstream tasks such as 3D GAN inversion [3,39].

It is also important to choose suitable 3D representations for shape generation. While delivering high geometric quality and infinite resolution, implicit function-based models [5, 6, 19, 29, 31, 51, 52] tend to be computationally expensive. This is due to the fact that the number of inferences increases cubically with the resolution and the time-consuming post-process, e.g., marching cubes. On the other hand, studies [26, 56, 60] learn diffusion models on a point cloud by adding noise and denoising either directly on point positions or their latent embeddings. Due to the irregular data format of the point set, it requires over 10,000 epochs for these diffusion models to converge on a single ShapeNet [4] category, while the number of points that can be generated by these models typically hovers around 2048.

In this work, we investigate a multi-view depth representation and propose a novel diffusion model, namely MVDD, which generates 3D consistent multi-view depth maps for 3D shape generation. The benefits of using the multi-view depth representation with diffusion models come in three folds: 1) The representation is naturally supported by diffusion models. The 2D data format conveniently allows the direct adoption of powerful 2D diffusion architectures [21, 37, 42]; 2) Multi-view depth registers complex 3D surfaces onto 2D grids, essentially reducing the dimensionality of the 3D generation space to 2D. As a result, the generated 2D depth map can have higher resolution than volumetric implicit representations [29] and produce dense point clouds with a much higher point count; 3) Depth is a widely used representation; therefore, it is easy to use it as a 3D prior to support downstream applications.

While bearing this many advantages, one key challenge of using multi-view depths for 3D shape generation is cross-view consistency. Even with a welltrained diffusion model that learns the depth distribution from 3D consistent depth maps, the generated multi-view depth maps are not guaranteed to be consistent after ancestral sampling [26]. To tackle this challenge, our proposed MVDD conditions diffusion steps for each view on neighboring views, allowing different views to exchange information. This is realized by a novel epipolar "line segment" attention, which benefits from epipolar geometry. Differing from full attention [41] and epipolar attention [23], our epipolar "line segment" attention leverages the depth estimation in our diffusion process. Therefore, it only attends to features at the most relevant locations, making it both effective and efficient. However, even with relatively consistent multi-view maps, back-projected 3D points from each depth map are still not guaranteed to be perfectly aligned. resulting in "double layers" in the 3D shapes (see Fig. 7(c)). To address this issue, MVDD incorporates depth fusion in denoising steps to explicitly align depth from multiple views.

Empowered by these modules, MVDD can generate high-quality 3D shapes, faithfully conduct depth completion, and distill 3D prior knowledge for down-stream tasks. We summarize our contributions as follows:

- To the best of our knowledge, we propose the first multi-view depth representation in the generative setting with a diffusion model MVDD. The representation reduces the dimension of the generation space and avoid unstructured data formats such as point set. Therefore, it is more scalable and suitable for diffusion frameworks and is easier to converge.
- We also propose an epipolar "line segment" attention and denoising depth fusion that could effectively enforce 3D consistency for multi-view depth maps.
- Through extensive experiments, we show the flexibility and versatility of MVDD in various tasks such as 3D shape generation and shape completion. Our method outperforms compared methods in both shape generation and shape completion by substantial margins.

2 Related Work

2.1 3D Shape Generative Models

Representations such as implicit functions, voxels, point clouds, and tetrahedron meshes have been used for 3D shape generation in previous studies.

Implicit-based models, such as AutoSDF [29], infer SDF from feature volumes. Since the computation for volumes grows cubically with resolution, the volume resolution is often limited. Voxel-based models, such as [30,60], face the same challenge. Other implicit-based models, such as 3D-LDM [31], IM-GAN [5], and Diffusion-sdf [6], generate latent codes and use auto-encoders to infer SDFs. The latent solution helps avoid the limitation of resolution but is prone to generate overly smoothed shapes. When combined with tetrahedron mesh, implicit

methods [9, 24] are able to generate compact implicit fields and achieve highquality shape generation. However, unlike multi-view depth, it is non-trivial for them to serve as a 3D prior in downstream tasks that do not use tetrahedron grids.

Point cloud-based methods avoid modeling empty space inherently. Previous explorations include SetVAE [15] and VG-VAE [1], which adopt VAEs for point latent sampling. GAN-based models [43, 50] employ adversarial loss to generate point clouds. Flow-based models [18, 54] use affine coupling layers to model point distributions. To enhance generation diversity, some studies leverage diffusion [11, 45] to generate 3D point cloud distributions. ShapeGF [2] applies the score-matching gradient flow to move the point set. DPM [26] and PVD [60] denoise Gaussian noise on point locations. LION [56] encodes the point set into latents and then conducts latent diffusion. Although these models excel in producing diverse shapes, the denoising scheme on unstructured point cloud data limits the number of points that can be generated. Our proposed model leverages multi-view depth representation, which can generate high-resolution point clouds, leading to 3D shapes with fine details.

2.2 Multi-View Diffusion Models

The infamous Janus problem [27, 36] and 3D inconsistency have plagued SDSbased [36] 3D content generation. MVDream [41] connects rendered images from different views via a 3D self-attention fashion to constrain multi-view consistency in the generated multi-view images. SyncDreamer [23] builds a cost volume that correlates the corresponding features across different views to synchronize the intermediate states of all the generated images at each step of the reverse process. EfficientDreamer [59] and TextMesh [48] concatenate canonical views either channel-wise or spatially into the diffusion models to enhance 3D consistency. Wonder3D [25] generates both RGB and normal maps for the task of single image to 3d. SweetDreamer [22] proposes aligned geometry priors by fine-tuning the 2D diffusion models to be viewpoint-aware and to produce view-specific coordinate maps. Our method differs from them in that we generate multi-view depth maps, instead of RGB images, and thus propose an efficient epipolar line segment attention tailored for depth maps to enforce 3D consistency.

3 Method

In this section, we introduce our Multi-View Depth Diffusion Models (MVDD). We first provide an overview of MVDD in Sec. 3.1, a model that aims to produce multi-view depth. After that, we illustrate how multi-view consistency is enforced among different views of depth maps in our model by using epipolar "line segment" attention (Sec. 3.1) and denoising depth fusion (Sec. 3.1). Finally, we introduce the training objectives in Sec. 3.2 and implementation details in Sec. 3.3.



Fig. 2: Our method collects ground truth from multi-view rendered depth maps (left). Starting with multiple 2D maps with randomly sampled noise, MVDD generates diverse multi-view depth through an iterative denoising diffusion process (right). To enforce multi-view 3D consistency, MVDD denoises each depth map with an efficient epipolar "line segment" attention (Sec. 3.1). Specifically, by leveraging the denoised value from the current step, MVDD only needs to attend to features on a line segment centered around the back-projected depth (the red dot), rather than the entire epipolar line. To further align the denoised multi-view depths, depth fusion (Sec. 3.1) is incorporated after the U-Net in a denoising step. The final multi-view depth can be fused together to obtain a high-quality dense point cloud, which can then be reconstructed into high quality 3D meshes with fine-grained details.

3.1 Multi-View Depth Diffusion

Our method represents a 3D shape \mathcal{X} by using its multi-view depth maps $\mathbf{x} \in \mathbb{R}^{N \times H \times W} = {\mathbf{x}^v | v = 1, 2, ..., N}$, where v is the index of the view, N is the total number of views, and H and W are the depth map resolution. To generate a 3D shape that is both realistic and faithful to the diversity distribution, we adopt the diffusion process [11,44] that gradually denoise N depth maps. These depth maps can be fused to obtain a dense point cloud, which can optionally be used to reconstruct [13,35] a high-quality mesh model. We illustrate the entire pipeline in Fig. 2

In the diffusion process, we first create the ground truth multi-view depth diffusion distribution $q(\mathbf{x}_{0:T})$ in a *forward process*. In this process, we gradually add Gaussian noise to each ground truth depth map \mathbf{x}_0^v for T steps, obtaining N depth maps of pure Gaussian noise $\mathbf{x}_T = {\mathbf{x}_T^v | v = 1, 2, ..., N}$. The joint distributions can be factored into a product of per-view Markov chains:

$$q(\mathbf{x}_{0:T}) = q(\mathbf{x}_0) \prod_{t=1}^{T} q(\mathbf{x}_t | \mathbf{x}_{t-1})$$

= $q(\mathbf{x}_0^{1:N}) \prod_{v=1}^{N} \prod_{t=1}^{T} q(\mathbf{x}_t^v | \mathbf{x}_{t-1}^v),$ (1)

$$q(\mathbf{x}_t^v | \mathbf{x}_{t-1}^v) := \mathcal{N}(\mathbf{x}_t^v; \sqrt{1 - \beta_t} \mathbf{x}_{t-1}^v, \beta_t \mathbf{I}),$$
(2)

where β_t is the step t variance schedule at step t shared across views.

We then learn a diffusion denoising model to predict the distribution of a reverse process $p_{\theta}(\mathbf{x}_{0:T})$ to iteratively denoise the \mathbf{x}_T back to the ground truth

 \mathbf{x}_0 . The joint distribution can be formulated as:

$$p_{\theta}(\mathbf{x}_{0:T}) = p(\mathbf{x}_{T}) \prod_{t=1}^{T} p_{\theta}(\mathbf{x}_{t-1} | \mathbf{x}_{t})$$
$$= p(\mathbf{x}_{T}^{1:N}) \prod_{v=1}^{N} \prod_{t=1}^{T} p_{\theta}(\mathbf{x}_{t-1}^{v} | \mathbf{x}_{t}^{v}),$$
(3)

$$p_{\theta}(\mathbf{x}_{t-1}^{v}|\mathbf{x}_{t}^{v}) := \mathcal{N}(\mathbf{x}_{t-1}^{v}; \mu_{\theta}(\mathbf{x}_{t}^{v}, t), \beta_{t}\mathbf{I}),$$
(4)

where $\mu_{\theta}(\mathbf{x}_{t}^{v}, t)$ estimates the mode of depth map distribution for view v at step t-1.

However, following Eq. (3) and Eq. (4), diffusion process denoises each view independently. Starting from N maps of pure random noise, a well-trained model of this kind would generate realistic depth maps $\mathbf{x}_0^{1:N}$, which however could not be fused into an intact shape due to no 3D consistency across views. Therefore, we propose to condition denoising steps for each view on its R neighboring views $\mathbf{x}_t^{r_1:r_R}$ and replace Eq. (3) and Eq. (4) with:

$$p_{\theta}(\mathbf{x}_{0:T}) = p(\mathbf{x}_{T}^{1:N}) \prod_{t=1}^{T} \prod_{v=1}^{N} p_{\theta}(\mathbf{x}_{t-1}^{v} | \mathbf{x}_{t}^{v}, \mathbf{x}_{t}^{r_{1}:r_{R}}),$$
(5)

$$p_{\theta}(\mathbf{x}_{t-1}^{v}|\mathbf{x}_{t}^{v},\mathbf{x}_{t}^{r_{1}:r_{R}}) := \mathcal{N}(\mathbf{x}_{t-1}^{v};\mu_{\theta}(\mathbf{x}_{t}^{v},\mathbf{x}_{t}^{r_{1}:r_{R}},t),\beta_{t}\mathbf{I}).$$
(6)

MVDD achieves this through an efficient epipolar 'line segment' attention (Sec. 3.1). Additionally, even though the denoising process is multi-view conditioned, back-projected depth maps are still not guaranteed to be perfectly aligned in 3D space. Inspired by multi-view stereo methods [8, 40, 53], MVDD conducts denoising depth fusion (Sec. 3.1) in each diffusion step (Eq. (6)).

Epipolar Line Segment Attention To promote consistency across all depth maps, we introduce an attention module named epipolar "line segment" attention. With the depth value of current step, MVDD leverages this information and attends only to features from visible locations on other views. To be specific, we sample on the line segment centered by the back-projected depth, rather than on the entire epipolar line [41, 49]. This design allows the proposed attention to obtain more relevant cross-view features, making it excel in both efficiency and effectiveness. The attention is defined as:

$$Q \in \mathbb{R}^{(B \times N \times H \times W) \times 1 \times F},$$

$$K, V \in \mathbb{R}^{(B \times N \times H \times W) \times (R \times k) \times F},$$

$$Cross-Attn(Q, K, V) = \operatorname{softmax}\left(\frac{QK^{T}}{\sqrt{d_{k}}}\right)V,$$
(7)

where B is the batch size, N is the total number of views, k is the number of samples along the epipolar "line segment", R is the number of neighboring views and F is the number of feature channels. At denoising step t, for any pixel v_{ij} at a source depth map \mathbf{x}_t^v , we first back project its depth value $\mathbf{x}_t^{v_{ij}}$ into the 3D space to obtain a 3D point $\rho^{v_{ij}}$, and project it to a coordinate r_{mn} on neighboring view r:

$$\rho^{v_{ij}} = \mathbf{x}_t^{v_{ij}} A^{-1} v_{ij}, \text{ where } v_{ij} := [i, j, 1]^T,$$
(8)

$$r_{mn} = A \ \pi_{v \to r} \rho^{v_{ij}},\tag{9}$$

where $\rho^{v_{ij}}$ is in the camera coordinate of view v, $\pi_{v \to r}$ is the relative pose transformation matrix and A is the intrinsic matrix. Since $\mathbf{x}_t^{v_{ij}}$ is noisy, we select another k-1 evenly spaced points around $\rho^{v_{ij}}$ along the ray and project these points, $\{\rho_1^{v_{ij}}, ..., \rho_k^{v_{ij}}\}$, into each neighboring view, as shown in Fig. 2 (right). The k projected pixels lay on a epipolar "line segment" on view r and provides features for K, V in Eq. (7).

Cross attention thresholding. To ensure that depth features from a neighboring view r are relevant to $\mathbf{x}_t^{v_{ij}}$, we need to cull only the r_{mn} that are also visible from source view v. Let $z(\cdot)$ denote the operator to extract the z value from a vector [x, y, z], we create the visibility mask by thresholding the Euclidean distance between the depth value of the 3D point in r's camera coordinate, $\rho^{r_{mn}} = \pi_{v \to r} \rho^{v_{ij}}$, and the predicted depth value on the pixel r_{mn} that ρ^r projects onto:

$$M(r_{mn}) = \|z(\pi_{v \to r} \rho^{v_{ij}}) - \mathbf{x}_t^{r_{mn}}\| < \tau.$$
(10)

For projected pixels that do not satisfy the above requirement, in Eq. (7), we manually overwrite their attention weights as a very small value to minimize its effect.

Depth concatenation. For pixel v_{ij} , since the sampled points $\{\rho_1^{v_{ij}}, ..., \rho_k^{v_{ij}}\}$ query geometric features K, V from neighboring views, the attention mechanism conditions the denoising step of $\mathbf{x}_t^{v_{ij}}$ with the features V weighted by the similarity between Q and K. To enhance awareness of the locations of these points, we propose concatenating the depth values $\{z(\rho_1^{v_{ij}}), ..., z(\rho_k^{v_{ij}})\}$ to the feature dimension of V, resulting in the last dimension of V being F + 1.

The intuition behind this is if the geometric features of v_{ij} are very similar to features queried by $\rho_1^{v_{ij}}$, the depth value $\mathbf{x}_{t-1}^{v_{ij}}$ should move toward $z(\rho_1^{v_{ij}})$. We empirically verify the effectiveness of the depth concatenation in Tab. 3.

Denoising Depth Fusion To further enforce alignment across multi-view depth maps, MVDD incorporates depth fusion in diffusion steps during ancestral sampling.

Assuming we have multi-view depth maps $\{\mathbf{x}_1, ..., \mathbf{x}_N\}$, following multi-view stereo methods [28, 55], a pixel v_{ij} will be projected to another view r at r_{mn} as described in Eq. (8). Subsequently, we reproject r_{mn} with its depth value $\mathbf{x}^{r_{mn}}$ towards view v:

$$\rho^{v_{\tilde{i}\tilde{j}}} = \pi_{r \to v} \mathbf{x}^{r_{mn}} A^{-1} r_{mn}, \tag{11}$$

$$v_{\tilde{i}\tilde{j}} = A\rho^{v_{\tilde{i}\tilde{j}}},\tag{12}$$

where $\rho^{v_{ij}}$ is the reprojected 3D point in view v's camera coordinate. To determine the visibility of pixel v_{ij} from view r, we set two thresholds:

$$\left\| v_{ij} - v_{\tilde{i}\tilde{j}} \right\| < \psi_{\max}, \ \frac{\left| \mathbf{x}^{v_{ij}} - z(\rho^{v_{\tilde{i}\tilde{j}}}) \right|}{\mathbf{x}^{v_{ij}}} < \epsilon_{\theta}, \tag{13}$$

where $z(\rho^{v_{ij}})$ represents the reprojected depth, ψ_{max} and ϵ_{θ} are the thresholds for discrepancies between reprojected pixel coordinates and depth compared to the original ones.

Integration with denosing steps. For a diffusion step t described in Eq. (6), after obtaining $\mu_{\theta}(\mathbf{x}_t, t)$, we apply depth averaging. For each pixel, we average the reprojected depths from other visible views to refine this predicted value. Subsequently, we add $\mathcal{N}(\mathbf{0}, \beta_t \mathbf{I})$ on top to obtain $\{\mathbf{x}_{t-1}^v | v = 1, 2, ..., N\}$. Only at the last step, we also apply depth filtering to X_0 to filter out the back-projected 3D points that are not visible from neighboring views.

3.2 Training Objectives

Aiming to maximize $p_{\theta}(\mathbf{x}_{0:T})$, we can minimize the objective, following DDPM [11]:

$$L_{t} = \mathbb{E}_{t \sim [1,T], \mathbf{x}_{0}, \boldsymbol{\epsilon}_{t}} \left[\left\| \boldsymbol{\epsilon}_{t} - \boldsymbol{\epsilon}_{\theta} \left(\mathbf{x}_{t}, t \right) \right\|^{2} \right]$$

$$= \mathbb{E}_{t \sim [1,T], \mathbf{x}_{0}, \boldsymbol{\epsilon}_{t}} \left[\left\| \boldsymbol{\epsilon}_{t} - \boldsymbol{\epsilon}_{\theta} \left(\sqrt{\bar{\alpha}_{t}} \mathbf{x}_{0} + \sqrt{1 - \bar{\alpha}_{t}} \boldsymbol{\epsilon}_{t}, t \right) \right\|^{2} \right],$$
(14)

where \mathbf{x}_0 is the ground-truth multiview depth maps, β_t and $\bar{\alpha}_t := \prod_{s=1}^t (1 - \beta_s)$ are predefined coefficients of noise scheduling at step t.

3.3 Implementation Details

Our model is implemented in PyTorch [34] and employs the Adam optimizer [16] with the first and the second momentum set to 0.9 and 0.999, respectively, and a learning rate of $2e^{-4}$ to train all our models. Unless otherwise noted, we set the height H and width W of depth map to both be 128 and number of views of depth map N to be 8. The first camera is free to be placed anywhere on a sphere, facing the object center, and form a cube with the other 7 cameras. The number of sampled points along the epipolar line segment k is 10. The threshold τ for cross attention thresholding is 0.15. We apply denoising depth fusion only in the last 20 steps. For training, we uniformly sample time steps t = 1, ..., T = 1000 for all experiments with cosine scheduling [32]. We train our model on 8 Nvidia A100-80GB and the model usually converges within 3000 epochs. Please refer to supplemental material for more details on network architecture, camera setting, and other aspects.

9



Fig. 3: Our generated meshes exhibit superior quality compared to point cloud diffusion model [56] and implicit diffusion model [31].

Table 1: Unconditional generation on ShapeNet categories. MMD (EMD) is multiplied by 10^2 . • represents the best result.

		Vox-diff [60]	DPM [26]	3D-LDM [31]	IM-GAN [5]	PVD [60]	LION [56]	MVDD (Ours)
Airplane	MMD (EMD)	1.561	0.990	3.520	0.980	1.000	0.920 😐	0.920 😐
	COV (EMD)	25.43	40.40	42.60	52.07	49.33	48.27	53.00 $-$
	1-NNA (EMD)	98.13	73.47	80.10	64.04	64.89	63.49	62.50 $-$
Car	MMD (EMD)	1.551	0.710	-	0.640	0.820	0.900	0.620 😐
	COV (EMD)	22.15	36.05	-	47.27	39.51	42.59	49.53 $-$
	1-NNA (EMD)	96.83	80.33	-	57.04	71.29	65.70	56.80 😐
Chair	MMD (EMD)	2.930	2.140	8.200	2.200	2.330	1.720 -	2.110
	COV (EMD)	21.75	46.17	42.20	49.51	46.47	50.52	51.55 😐
	1-NNA (EMD)	96.74	65.73	65.30	55.54	56.14	57.31	54.51 😐

4 Application

4.1 3D Shape Generation

Inference strategy. Initialized as 2D maps of pure Gaussian noise, the multi-view depth maps can be generated by MVDD following ancestral sampling [11]:

$$\mathbf{x}_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left(\mathbf{x}_t - \frac{1 - \alpha_t}{\sqrt{1 - \tilde{\alpha}_t}} \boldsymbol{\epsilon}_{\theta} \left(\mathbf{x}_t, t \right) \right) + \sqrt{\beta_t} \boldsymbol{\epsilon}, \tag{15}$$

where ϵ follows a isotropic multivariate normal distribution. We iterate the above process for T = 1000 steps, utilizing the effective epipolar "line segment" attention and denoising depth fusion . Finally, we back-project the multi-view depth maps to form a dense (> 20K) 3D point cloud with fine-grained details. Optionally, high-quality meshes can be created with SAP [35] from the dense point cloud.

Datasets and comparison methods. To assess the performance of our method compared to established approaches, we employ the ShapeNet dataset [4], which is the commonly adopted benchmark for evaluating 3D shape generative models. In line with previous studies of 3D shape generation [5, 54, 56, 60], we evaluate our model on standard shape generation bench mark categories: airplanes, chairs,



Fig. 4: Unconditional generation on ShapeNet car, airplane and chair category.

and cars, with the same train/test split. We compare MVDD with state-of-theart point cloud generation methods such as DPM [26], PVD [60] and LION [56], implicit functions-based methods such as IM-GAN [5] and 3D-LDM [31], as well as a voxel diffusion model Vox-diff [60]. As our method generates varying number of points and point cloud backprojected from depth maps is not uniform, we sample 2048 points from meshes using SAP [35] and measure against groundtruth points with inner surface removed. For those implicit methods that are not impacted by inner surface, we directly use the number reported for comparison.

Metrics. We follow [5,56,60] and primarily employ: 1) Minimum Matching Distance (MMD), which calculate the average distance between the point clouds in the reference set and their closest neighbors in the generated set; 2) Coverage (COV), which measures the number of reference point clouds that are matched to at least one generated shape; 3) 1-Nearest Neighbor Alignment (1-NNA), which measures the distributional similarity between the generated shapes and the validation set. MMD focus on the shape fidelity and quality, while COV focus on the shape diversity. 1-NNA can assess both quality and diversity of the generation results. For methods generate mesh or voxel, we transform it to point cloud and apply these metrics. Please refer to supplemental materials for more details.

Evaluation. We report the quantitative results of all methods in Tab. 1. Due to space constraints, we defer the performance in metric CD to the supplemental material. Our method MVDD exhibits strong competitiveness across all categories and surpassed comparison methods, particularly excelling in the 1-NNA (EMD) metric. This metric holds significant importance as it addresses the limitations of MMD and COV [56]. We augmented our generated point cloud and visualize the mesh quality in Fig. 3 together with LION [56] and 3D-LDM [31]. Our method generates more diverse and plausible 3D shapes compared with



Fig. 5: We report the performance of our method and LION with varying number of point clouds measured by 1-NNA with CD and EMD, respectively, in the ShapeNet [4] chair category.



Fig. 6: Without proper shape regularization, 3D GAN inversion [3] using PTI [39] fails to reconstruct input image under extreme pose. Our model can serve as a shape prior for 3D GAN inversion and yield better reconstruction performance in novel frontal view.

all baselines. The visualization of our meshes shows that our method excels in synthesizing intricate details, e.g. slats of the chair and thin structure in chair base. We also visualize point clouds in Fig. 1(a) and Fig. 4. The clean point cloud back-projected from our generated depth maps demonstrates 3D consistency and also validates the effectiveness of the proposed epipolar "line segment" attention and denoising depth fusion. In contrast, the number of points (2048) that can be generated by point cloud-based diffusion models [26, 56, 60] limits their capabilities to capture fine-grained details of the 3D shapes. To show the generalizability of our method, we also train our model on ShapeNet 13 classes simultaneously and compare the performance with the model trained separately in the supplemental material.

Generated dense point cloud vs up-sampled sparse point cloud. Since our method can directly generate 20K points, while LION [56] is limited to producing sparse point cloud with 2048 points, we up-sample varying number of points from LION's meshes. We then compare the performance of our method with LION. As shown in Fig. 5, the performance of LION deteriorates significantly as the number of points increases. It is because LION struggles to faithfully capture necessary 3D shape details with its sparse point cloud. In contrast, the performance of our method is robust with the increased number of 3D points and outperforms LION by larger margins as the point cloud density increases. We also utilize a SOTA point cloud upsampler and please see supplemental materials for the comparison.

4.2 Depth Completion

Inference strategy. We reuse an unconditional generative model to perform shape completion task, where depth maps from other views $\mathbf{x}^{\text{other}}$ can be generated

Table 2: Depth completion comparison against baselines. EMD is multiplied by 10^2 . • represents the best result.

	SoftFlow [14]	PointFlow [54]] DPF-Net [18]	PVD [60]	MVDD (Ours)
Airplane	1.198	1.180	1.105	1.030	0.900 😐
Chair	3.295	3.649	3.320	2.939	2.400 $-$
Car	2.789	2.851	2.318	2.146	1.460 -

conditioned on the single input view of depth map \mathbf{x}^{in} . In each reverse step of the diffusion model, we have:

$$\mathbf{x}_{t-1}^{\text{in}} \sim \mathcal{N}\left(\sqrt{\bar{\alpha}_t}\mathbf{x}_0^{\text{in}}, (1-\bar{\alpha}_t)\,\mathbf{I}\right),$$
1st pass: $\hat{\mathbf{x}}_{t-1}^{\text{other}} \sim \mathcal{N}(\sqrt{1-\beta_t}\ \mu_{\theta}(\mathbf{x}_t^{r_1:r_R}, t), \beta_t\mathbf{I}),$
2nd pass: $\mathbf{x}_{t-1}^{\text{other}} \sim \mathcal{N}(\mu_{\theta}(\hat{\mathbf{x}}_{t-1}^{\text{other}}, \mathbf{x}_0^{\text{in}}, t), \beta_t\mathbf{I}),$
(16)

where $\mathbf{x}_{t-1}^{\text{in}}$ is sampled using the given depth map \mathbf{x}^{in} , while $\mathbf{x}_{t-1}^{\text{other}}$ is sampled from the model, given the previous iteration \mathbf{x}_t . Different from unconditional generation, to enhance the consistency with the input view, we do two passes to denoise the other views. In the first pass each view attends to every other views and in the second pass each view only attends to the input view \mathbf{x}^{in} . We scale back noise at first pass, following the Langevin dynamics steps [46, 47].

Datasets and comparison methods. Following the experimental setup of PVD [60], we use the benchmark provided by GenRe [58], which contains renderings of shapes in ShapeNet from 20 random views. For shape completion [7, 38], as the ground-truth data are involved, Chamfer Distance and Earth Mover's Distance suffice to evaluate the reconstruction results. We select models PointFlow [54], DPF-Net [18], SoftFlow [14], and PVD [60] for comparison.

Evaluation. We show the quantitative results of our method and baselines in Tab. 2. Our method consistently outperforms all the baselines with EMD metric on all categories. The qualitative results in Fig. 1(b) also showcases that our inference strategy for depth completion can effectively "pull" the learned depth map of other views to be geometrically consistent with the input depth map.

4.3 3D Prior for GAN Inversion

We illustrate how our trained multi-view depth diffusion model can be plugged into downstream tasks, such as 3D GAN inversion [3]. As in the case of 2D GAN inversion, the goal of 3D GAN inversion is to map an input image I onto the space represented by a pre-trained unconditional 3D GAN model, denoted as $G_{3D}(\cdot; \theta)$, which is defined by a set of parameters θ . Upon successful inversion, G_{3D} has the capability to accurately recreate the input image when provided with the corresponding camera pose. One specific formulation of the 3D GAN inversion problem [39] can be defined as follows:

$$w^*, \theta^* = \underset{w,\theta}{\operatorname{arg\,max}} = \mathcal{L}\left(G_{3D}(w, \pi; \theta), I\right), \tag{17}$$

where w is the latent representation in \mathcal{W}^+ space and π is the corresponding camera matrix of input image. w and θ are optimized alternatively, i.e., w is optimized first and then θ is also optimized together with the photometric loss:

$$\mathcal{L}_{\text{photo}} = \mathcal{L}_2 \left(G_{3\text{D}} \left(w, \pi_s; \theta \right), I_s \right) + \mathcal{L}_{\text{LPIPS}} \left(G_{3\text{D}} \left(w, \pi_s; \theta \right), I_s \right),$$
(18)

where $\mathcal{L}_{\text{LPIPS}}$ is the perceptual similarity loss [57]. However, with merely supervision from single or sparse views, this 3D inversion problem is ill-posed without proper regularization, so that the geometry could collapse (shown in Fig. 1(c) and Fig. 6 2nd row). To make the inversion look plausible from other views, a 3D geometric prior is needed, as well as a pairing regularization method which can preserve diversity. Score distillation sampling has been proposed in DreamFusion [36] to utilize a 2D diffusion model as a 2D prior to optimize the parameters of a radiance field. In our case, we use our well-trained MVDD model as a 3D prior to regularize on the multi-view depth maps extracted from the tri-plane radiance fields. As a result, the following gradient direction would not lead to collapsed geometry after inversion:

$$\nabla \mathcal{L} = \nabla \mathcal{L}_{\text{photo}} + \nabla \lambda_{\text{SDS}} \mathcal{L}_{\text{SDS}}, \qquad (19)$$

where λ_{SDS} is the weighting factor of \mathcal{L}_{SDS} [36].

To learn the shape prior for this 3D GAN inversion task, we render multi-view depth maps from the randomly generated radiance fields of EG3D [3] trained with FFHQ [12] dataset. We then use them as training data and train our multi-view depth diffusion model. Using Eq. (19), we perform test-time optimization for each input image to obtain the optimized radiance field. In Fig. 1(c) and Fig. 6, we show the rendering and geometry results of 3D GAN inversion with and without regularization by MVDD. With the regularization of our model, the "wall" artifact is effectively removed and it results in better visual quality in the rendered image from novel frontal view.

4.4 Ablation study

We perform ablation study to further examine the effectiveness of each module described in the method section. Specifically, in Tab. 3 we report the ablated results of epipolar "line segment" attention, depth concatenation, and cross attention thresholding (Sec. 3.1) and depth fusion (Sec. 3.1) in ShapeNet chair category for the unconditional generation task as we describe in Sec. 4.1. Without the designed cross attention, the model could barely generate plausible 3D shapes as measured by 1NN-A metric. With designs such as depth concatenation and cross attention thresholding being added, the 3D consistency along with the



(a) Input depth map (b) Completed depth maps
 (c) W/o denoising depth fusion
 (d) W/ denoising depth fusion

Fig. 7: Depth completion results prove the effectiveness of the proposed denoising depth fusion strategy (Sec. 3.1).

Table 3: Ablation study on the chair category. • is the top result.

Cross attn. Depth concat. Cross attn. thresholding Depth fusion					1NN-A	
(Sec. 3.1)	(Sec. 3.1)	(Sec. 3.1)	(Sec. 3.1)	CD	EMD	
×	×	×	×	92.00	90.00	
\checkmark	×	×	×	61.78	59.65	
\checkmark	\checkmark	×	×	60.72	59.00	
\checkmark	\checkmark	\checkmark	×	59.82	57.75	
\checkmark	\checkmark	\checkmark	\checkmark	57.90 😐	54.51 😐	

performance of our model is progressively improving. Last but not least, denoising depth fusion align the depth maps and further boost the performance. Qualitatively, Fig. 7 illustrates how the denoising depth fusion help eliminate double layers in depth completion task.

We report the shape generation results with full self-attention (as in MV-Dream [41]) and epipolar attention in the supplemental material. This ablation study shows that our model excels in both efficiency and effectiveness. The proposed attention obtains more useful cross-view features by avoiding the irrelevant pixels in calculating the keys and values.

5 Conclusion

We leveraged multi-view depth representation in 3D shape generation and proposed a novel denoising diffusion model MVDD. To enforce 3D consistency among different view of depth maps, we proposed an epipolar "line segment" attention and denoising depth fusion technique. Through extensive experiments in various tasks such as shape generation, shape completion and shape regularization, we demonstrated the scalability, faithfulness and versatility of our multi-view depth diffusion model.

Acknowledgement: A.K. was supported by a National Science Foundation (NSF) CAREER award IIS-2046737, Army Young Investigator Program Award, and Defense Advanced Research Projects Agency (DARPA) Young Faculty Award.

15

References

- Anvekar, T., Tabib, R.A., Hegde, D., Mudengudi, U.: Vg-vae: A venatus geometry point-cloud variational auto-encoder. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2978–2985 (2022) 4
- Cai, R., Yang, G., Averbuch-Elor, H., Hao, Z., Belongie, S., Snavely, N., Hariharan, B.: Learning gradient fields for shape generation. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16. pp. 364–381. Springer (2020) 4
- Chan, E.R., Lin, C.Z., Chan, M.A., Nagano, K., Pan, B., De Mello, S., Gallo, O., Guibas, L.J., Tremblay, J., Khamis, S., et al.: Efficient geometry-aware 3d generative adversarial networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16123–16133 (2022) 2, 11, 12, 13
- 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) 2, 9, 11
- Chen, Z., Zhang, H.: Learning implicit fields for generative shape modeling. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5939–5948 (2019) 2, 3, 9, 10
- Chou, G., Bahat, Y., Heide, F.: Diffusion-sdf: Conditional generative modeling of signed distance functions. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2262–2272 (2023) 2, 3
- Chu, R., Xie, E., Mo, S., Li, Z., Nießner, M., Fu, C.W., Jia, J.: Diffcomplete: Diffusion-based generative 3d shape completion. Advances in Neural Information Processing Systems 36 (2024) 12
- Galliani, S., Lasinger, K., Schindler, K.: Massively parallel multiview stereopsis by surface normal diffusion. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 873–881 (2015) 6
- Gao, J., Shen, T., Wang, Z., Chen, W., Yin, K., Li, D., Litany, O., Gojcic, Z., Fidler, S.: Get3d: A generative model of high quality 3d textured shapes learned from images. Advances In Neural Information Processing Systems 35, 31841–31854 (2022) 4
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial networks. Communications of the ACM 63(11), 139–144 (2020) 1
- 11. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. Advances in neural information processing systems **33**, 6840–6851 (2020) **1**, **4**, **5**, **8**, **9**
- Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4401–4410 (2019) 13
- Kazhdan, M., Bolitho, M., Hoppe, H.: Poisson surface reconstruction. In: Proceedings of the fourth Eurographics symposium on Geometry processing. vol. 7, p. 0 (2006) 5
- Kim, H., Lee, H., Kang, W.H., Lee, J.Y., Kim, N.S.: Softflow: Probabilistic framework for normalizing flow on manifolds. Advances in Neural Information Processing Systems 33, 16388–16397 (2020) 12
- Kim, J., Yoo, J., Lee, J., Hong, S.: Setvae: Learning hierarchical composition for generative modeling of set-structured data. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15059–15068 (2021) 4

- 16 Z. Wang et al.
- Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014) 8
- 17. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114 (2013) 1
- Klokov, R., Boyer, E., Verbeek, J.: Discrete point flow networks for efficient point cloud generation. In: European Conference on Computer Vision. pp. 694–710. Springer (2020) 4, 12
- Lan, Y., Hong, F., Yang, S., Zhou, S., Meng, X., Dai, B., Pan, X., Loy, C.C.: Ln3diff: Scalable latent neural fields diffusion for speedy 3d generation. arXiv preprint arXiv:2403.12019 (2024) 2
- Lan, Y., Meng, X., Yang, S., Loy, C.C., Dai, B.: E3dge: Self-supervised geometryaware encoder for style-based 3d gan inversion (2023) 1
- Lan, Y., Tan, F., Qiu, D., Xu, Q., Genova, K., Huang, Z., Fanello, S., Pandey, R., Funkhouser, T., Loy, C.C., et al.: Gaussian3diff: 3d gaussian diffusion for 3d full head synthesis and editing. arXiv preprint arXiv:2312.03763 (2023) 2
- 22. Li, W., Chen, R., Chen, X., Tan, P.: Sweetdreamer: Aligning geometric priors in 2d diffusion for consistent text-to-3d. arXiv preprint arXiv:2310.02596 (2023) 4
- Liu, Y., Lin, C., Zeng, Z., Long, X., Liu, L., Komura, T., Wang, W.: Syncdreamer: Generating multiview-consistent images from a single-view image. arXiv preprint arXiv:2309.03453 (2023) 3, 4
- Liu, Z., Feng, Y., Black, M.J., Nowrouzezahrai, D., Paull, L., Liu, W.: Meshdiffusion: Score-based generative 3d mesh modeling. arXiv preprint arXiv:2303.08133 (2023) 4
- Long, X., Guo, Y.C., Lin, C., Liu, Y., Dou, Z., Liu, L., Ma, Y., Zhang, S.H., Habermann, M., Theobalt, C., et al.: Wonder3d: Single image to 3d using crossdomain diffusion. arXiv preprint arXiv:2310.15008 (2023) 4
- Luo, S., Hu, W.: Diffusion probabilistic models for 3d point cloud generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2837–2845 (2021) 2, 3, 4, 9, 10, 11
- Melas-Kyriazi, L., Laina, I., Rupprecht, C., Vedaldi, A.: Realfusion: 360deg reconstruction of any object from a single image. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8446–8455 (2023) 4
- Merrell, P., Akbarzadeh, A., Wang, L., Mordohai, P., Frahm, J.M., Yang, R., Nistér, D., Pollefeys, M.: Real-time visibility-based fusion of depth maps. In: 2007 IEEE 11th International Conference on Computer Vision. pp. 1–8. Ieee (2007) 7
- Mittal, P., Cheng, Y.C., Singh, M., Tulsiani, S.: Autosdf: Shape priors for 3d completion, reconstruction and generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 306–315 (2022) 2, 3
- Mo, S., Xie, E., Chu, R., Hong, L., Niessner, M., Li, Z.: Dit-3d: Exploring plain diffusion transformers for 3d shape generation. Advances in Neural Information Processing Systems 36 (2024) 3
- Nam, G., Khlifi, M., Rodriguez, A., Tono, A., Zhou, L., Guerrero, P.: 3d-ldm: Neural implicit 3d shape generation with latent diffusion models. arXiv preprint arXiv:2212.00842 (2022) 2, 3, 9, 10
- Nichol, A.Q., Dhariwal, P.: Improved denoising diffusion probabilistic models. In: International Conference on Machine Learning. pp. 8162–8171. PMLR (2021) 8
- Papamakarios, G., Nalisnick, E., Rezende, D.J., Mohamed, S., Lakshminarayanan, B.: Normalizing flows for probabilistic modeling and inference. The Journal of Machine Learning Research 22(1), 2617–2680 (2021) 1

17

- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: Pytorch: An imperative style, highperformance deep learning library. Advances in neural information processing systems 32 (2019) 8
- Peng, S., Jiang, C., Liao, Y., Niemeyer, M., Pollefeys, M., Geiger, A.: Shape as points: A differentiable poisson solver. Advances in Neural Information Processing Systems 34, 13032–13044 (2021) 5, 9, 10
- Poole, B., Jain, A., Barron, J.T., Mildenhall, B.: Dreamfusion: Text-to-3d using 2d diffusion. arXiv preprint arXiv:2209.14988 (2022) 4, 13
- Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PMLR (2021) 2
- Rao, Y., Nie, Y., Dai, A.: Patchcomplete: Learning multi-resolution patch priors for 3d shape completion on unseen categories. Advances in Neural Information Processing Systems 35, 34436–34450 (2022) 12
- Roich, D., Mokady, R., Bermano, A.H., Cohen-Or, D.: Pivotal tuning for latentbased editing of real images. ACM Transactions on graphics (TOG) 42(1), 1–13 (2022) 2, 11, 13
- Schönberger, J.L., Zheng, E., Frahm, J.M., Pollefeys, M.: Pixelwise view selection for unstructured multi-view stereo. In: Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part III 14. pp. 501–518. Springer (2016) 6
- Shi, Y., Wang, P., Ye, J., Long, M., Li, K., Yang, X.: Mvdream: Multi-view diffusion for 3d generation. arXiv preprint arXiv:2308.16512 (2023) 3, 4, 6, 14
- 42. Shonenkov, A., Konstantinov, M., Bakshandaeva, D., Schuhmann, C., Ivanova, K., Klokova, N.: (2023), https://github.com/deep-floyd/IF/tree/develop 2
- 43. Shu, D.W., Park, S.W., Kwon, J.: 3d point cloud generative adversarial network based on tree structured graph convolutions. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 3859–3868 (2019) 4
- 44. Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S.: Deep unsupervised learning using nonequilibrium thermodynamics. In: International conference on machine learning. pp. 2256–2265. PMLR (2015) 5
- 45. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. arXiv preprint arXiv:2010.02502 (2020) 1, 4
- 46. Song, Y., Ermon, S.: Generative modeling by estimating gradients of the data distribution. Advances in neural information processing systems 32 (2019) 12
- Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Scorebased generative modeling through stochastic differential equations. arXiv preprint arXiv:2011.13456 (2020) 12
- Tsalicoglou, C., Manhardt, F., Tonioni, A., Niemeyer, M., Tombari, F.: Textmesh: Generation of realistic 3d meshes from text prompts. arXiv preprint arXiv:2304.12439 (2023) 4
- Tseng, H.Y., Li, Q., Kim, C., Alsisan, S., Huang, J.B., Kopf, J.: Consistent view synthesis with pose-guided diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16773–16783 (2023) 6
- Valsesia, D., Fracastoro, G., Magli, E.: Learning localized generative models for 3d point clouds via graph convolution. In: International conference on learning representations (2018) 4

- 18 Z. Wang et al.
- Wang, Z., Zhou, S., Park, J.J., Paschalidou, D., You, S., Wetzstein, G., Guibas, L., Kadambi, A.: Alto: Alternating latent topologies for implicit 3d reconstruction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 259–270 (2023) 2
- 52. Xu, Q., Wang, W., Ceylan, D., Mech, R., Neumann, U.: Disn: Deep implicit surface network for high-quality single-view 3d reconstruction. Advances in neural information processing systems **32** (2019) 2
- 53. Xu, Q., Xu, Z., Philip, J., Bi, S., Shu, Z., Sunkavalli, K., Neumann, U.: Point-nerf: Point-based neural radiance fields. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5438–5448 (2022) 6
- Yang, G., Huang, X., Hao, Z., Liu, M.Y., Belongie, S., Hariharan, B.: Pointflow: 3d point cloud generation with continuous normalizing flows. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 4541–4550 (2019) 4, 9, 12
- Yao, Y., Luo, Z., Li, S., Fang, T., Quan, L.: Mvsnet: Depth inference for unstructured multi-view stereo. In: Proceedings of the European conference on computer vision (ECCV). pp. 767–783 (2018) 7
- Zeng, X., Vahdat, A., Williams, F., Gojcic, Z., Litany, O., Fidler, S., Kreis, K.: Lion: Latent point diffusion models for 3d shape generation. arXiv preprint arXiv:2210.06978 (2022) 2, 4, 9, 10, 11
- 57. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 586–595 (2018) 13
- Zhang, X., Zhang, Z., Zhang, C., Tenenbaum, J., Freeman, B., Wu, J.: Learning to reconstruct shapes from unseen classes. Advances in neural information processing systems **31** (2018) 12
- Zhao, M., Zhao, C., Liang, X., Li, L., Zhao, Z., Hu, Z., Fan, C., Yu, X.: Efficientdreamer: High-fidelity and robust 3d creation via orthogonal-view diffusion prior. arXiv preprint arXiv:2308.13223 (2023) 4
- Zhou, L., Du, Y., Wu, J.: 3d shape generation and completion through point-voxel diffusion. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5826–5835 (2021) 2, 3, 4, 9, 10, 11, 12