

MVDiff: Scalable and Flexible Multi-view Diffusion for 3D Object Reconstruction from Single-View

Emmanuelle Bourigault

Visual Geometry Group, Department of Engineering Science, University of Oxford

emmanuelle@robots.ox.ac.uk

Pauline Bourigault

Department of Computing, Imperial College London

p.bourigault22@imperial.ac.uk

Abstract

Generating consistent multiple views for 3D reconstruction tasks is still a challenge to existing image-to-3D diffusion models. Generally, incorporating 3D representations into diffusion model decrease the model's speed as well as generalizability and quality. This paper proposes a general framework to generate consistent multi-view images from single image or leveraging scene representation transformer and view-conditioned diffusion model. In the model, we introduce epipolar geometry constraints and multi-view attention to enforce 3D consistency. From as few as one image input, our model is able to generate 3D meshes surpassing baselines methods in evaluation metrics, including PSNR, SSIM and LPIPS.

1. Introduction

Consistent and high-quality novel view synthesis of real-world objects from a single input image is a remaining challenge in computer vision. There is a myriad of applications in virtual reality, augmented reality, robotic navigation, content creation, and filmmaking. Recent advances in the field of deep learning such as diffusion-based models [2, 13, 22, 36, 37] significantly improved mesh generation by denoising process from Gaussian noise. Text-to-image generation has shown great progress with the development of efficient approaches as generative adversarial networks [3, 11, 16], autoregressive transformers [9, 28, 39], and more recently, diffusion models [12, 14, 27, 32]. DALL-E 2 [27] and Imagen [32] are such models capable of generating of photo-realistic images with large-scale diffusion models. Latent diffusion models [31] apply the diffusion process in the latent space, enabling for faster image synthesis.

Although, image-to-3D generation has shown impres-

sive results, there is still room for improvement in terms of consistency, rendering and efficiency. Generating 3D representations from single view is a difficult task. It requires extensive knowledge of the 3D world. Although diffusion models have achieved impressive performance, they require expensive per-scene optimization.

Zero123[18] proposes a diffusion model conditioned on view features and camera parameters trained on perspective images [6]. However, the main drawback is the lack of multiview consistency in the generation process impeding high-quality 3D shape reconstruction with good camera control. SyncDreamer[19] proposes a 3D feature volume into the Zero123[18] backbone to improve the multi-view consistency. However, the volume conditioning significantly reduces the speed of generation and it overfits to some viewpoints, with 3D shapes displaying distortions.

In this paper, we present MVDiff, a multiview diffusion model using epipolar geometry and transformers to generate consistent target views. The main idea is to incorporate epipolar geometry constraints in the model via self-attention and multi-view attention in the UNet to learn the geometry correspondence. We first need to define a scene transformation transformer (SRT) to learn an implicit 3D representation given a set of input views. Then, given an input view and its relative camera pose, we use a view-conditioned diffusion model to estimate the conditional distribution of the target view.

We show that this framework presents dual improvements compared to existing baselines in improving the 3D reconstruction from generated multi-view images and in terms of generalization capability.

In summary, the paper presents a multi-view generation framework from single image that is transferable to various datasets requiring little amount of changes. We show high performance on the GSO dataset for 3D mesh generation. The model is able to extrapolate one view image of a 3D

object to 360-view with high fidelity. Despite being trained on one dataset of natural objects, it can create diverse and realistic meshes. We summarise our contributions as follows:

- Implicit 3D representation learning with geometrical guidance
- Multi-view self-attention to reinforce view consistency
- Scalable and flexible framework

2. Related Work

Diffusion for 3D Generation.

Recently, the field of 3D generation has demonstrated rapid progress with the use of diffusion models. Several studies showed remarkable performance by training models from scratch on large datasets to generate point clouds [21, 24], meshes [10, 20] or neural radiance fields (NeRFs) at inference. Nevertheless, these models lack generalizability as they are trained on specific categories of natural objects. DreamFusion [26] explored leveraging 2D priors to guide 3D generation. Inspired by DreamFusion, several studies adopted a similar pipeline using distillation of a pre-trained 2D text-to-image generation model for generating 3D shapes [1, 4, 5, 23, 43]. The per-scene optimisation process typically lacks in efficiency with times ranging from minutes to hours to generate single scenes.

Recently, 2D diffusion models for multi-view synthesis from single view have raised interest for their fast 3D shape generation with appealing visuals [17, 18, 34]. However, they generally do not consider consistency of multi-view in the network design. Zero123 proposes relative viewpoint as conditioning in 2D diffusion models, in order to generate novel views from a single image [18]. However, this work does not consider other views in the learning process and this causes inconsistencies for complex shapes. One-2-3-45 [17] decodes signed distance functions (SDF) [25] for 3D shape generation given multi-view images from Zero123 [18], but the 3D reconstruction is not smooth and artifacts are present.

More recently, SyncDreamer [19] suggests a 3D global feature volume, in order to tackle inconsistencies in multi-view generation. 3D volumes are used with depth-wise attention for maintaining multi-view consistency. The heavy 3D global modeling tend to reduce the speed of the generation and quality of the generated meshes. MVDream [35] on the other hand incorporates 3D self-attention with improved generalisability to unseen datasets.

Sparse-View Reconstruction. Sparse-view image reconstruction [15, 45] is a challenging task where only a limited number of images, generally less than 10, are given. Traditional 3D reconstruction methods start by estimating camera poses, then as a second step perform dense reconstruction with multi-view stereo [38, 46] or NeRF [40]. Estimating camera poses in the context of sparse-view recon-

struction is a challenging task as there is little or no overlap between views. [45] aimed at addressing this challenge by optimising camera poses and 3D shapes simultaneously. In the same line of research, PF-LRM [42] suggests a pose-free approach to tackle the uncertainty in camera poses. In our work, we learn the relative camera poses of the 3D representation implicitly via a transformer encoder-decoder network and a view-conditioned diffusion model capable of generating consistent multi-view images directly. We then employ a reconstruction system Neus [41] to recover a mesh.

3. Methodology

3.1. Multi-view Conditional Diffusion Model

The rationale behind multi-view conditioning in diffusion models is to infer precisely the 3D shape of an object with the constraint that regions of the 3D object are unobserved. Direct 3D predictions for sequential targets as in Zero123 [18] might lead to implausible novel views. To control the uncertainty in novel view synthesis, we choose to enforce multi-view consistency during training.

Given an input image or sparse-view input images of a 3D object, denoted as x_I , with known camera parameters π_I , and target camera parameters π_T , our aim is to synthesize novel views that recover the geometry of the object.

Our framework can be broken down into two parts: (i) first a scene representation transformer (SRT) [33] that learns the latent 3D representation given a single or few input views, and (ii) second a view-conditioned diffusion model to generate novel views.

3.2. Novel View Synthesis via Epipolar Geometry

To perform novel view synthesis, we employ a scene representation transformer (SRT) [33]. In the work of (SRT) [33], a transformer encoder-decoder architecture learns an implicit 3D latent representation given a set of images with camera poses (x_I, π_I). First, a CNN extracts features from x_I and feeds them as tokens to the transformer encoder f_E . The transformer encoder then outputs a set-latent scene representation z via self-attention.

For novel view rendering, the decoder transformer of SRT queries the pixel color via cross-attention between the ray associated to that pixel r and the set-latent scene representation z .

The aim is to minimize the pixel-level reconstruction loss in Equation 1:

$$\mathcal{L}_{\text{recon}} = \sum_{r \in \mathcal{R}} \left\| C(r) - \hat{C}(r) \right\|_2^2, \quad (1)$$

where $C(r)$ is the ground truth color of the ray and \mathcal{R} is the set of rays sampled from target views.

We aim to leverage cross-interaction between images through relative camera poses using epipolar geometrical

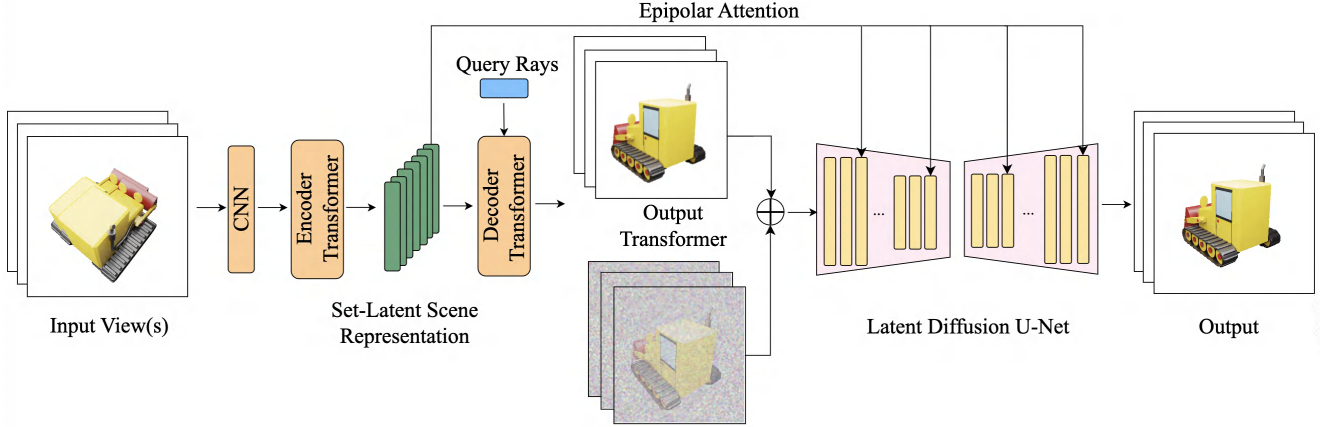


Figure 1. **Pipeline of MVDiff.** From a single input or few input images, the transformer encoder translates the image(s) into latent scene representations, implicitly capturing 3D information. The intermediate outputs from the scene representation transformer are used as input by the view-conditioned latent diffusion UNet, generating multi-view consistent images from varying viewpoints.

constraints. For each pixel in a given view i , we compute the epipolar line and the epipolar distance for all pixels in view j to build a weighted affinity matrix $A'_{i,j} = A_{i,j} + W_{i,j}$ where $W_{i,j}$ is the weighted map obtained from the inverse epipolar distance. **View-Conditioned Latent Diffusion.** The outputs from SRT do not recover fine details with simple pixel-level reconstruction loss. We employ a view-conditioned diffusion model LDM from [29] to estimate the conditional distribution of the target view given the source view and the relative camera pose: $p(\mathbf{x}_T | \pi_T, \mathbf{x}_I, \pi_I)$.

First, the SRT predicts a low-resolution 32×32 latent image $\tilde{\mathbf{x}}_T$ based on the target view π_T for computationally efficiency. The latent image from SRT is concatenated with the noisy image \mathbf{y} and fed into the latent diffusion UNet \mathcal{E}_θ . In addition, we condition \mathcal{E}_θ on the latent scene representation \mathbf{z} via cross-attention layers (see Figure 1).

The generated images $\hat{\epsilon}_t$ can be denoted as:

$$\hat{\epsilon}_t = \mathcal{E}_\theta(\mathbf{y}, \tilde{\mathbf{x}}_I, \mathbf{z}, t), \quad (2)$$

where t is the timestep.

We optimize a simplified variational lower bound:

$$\mathcal{L}_{\text{VLDM}} = \mathbb{E} \left[\|\mathcal{E}_t - \mathcal{E}_\theta(\mathbf{y}, \tilde{\mathbf{x}}_T, \mathbf{z}, t)\|^2 \right] \quad (3)$$

Multi-View Attention. As previously stated, in Zero123[18], multiple images are generated in sequence from a given input view based on camera parameters. This approach can introduce inconsistencies between generated views. To address this issue, we apply modifications to the UNet in order to feed multi-view images. This way, we can predict simultaneously multiple novel views. We employ self-attention block to ensure consistency for different viewpoints.

4. Experiments

This section presents the novel view synthesis experiments in 4.1, and the 3D generation experiments in 4.2. We present ablation experiments in 4.3 and ethical considerations in 4.4.

Training Data. For training our model for novel view synthesis, we use 800k 3D object models from Objaverse[6]. For a fair comparison with other 3D diffusion baselines, we use the same training dataset.

Input condition views are chosen in a similar way as Zero123[18]. An azimuth angle is randomly chosen from one of the eight discrete angles of the output cameras. The elevation angle is randomly selected in the range $[-10^\circ, 45^\circ]$. For data quality purposes, we discard empty rendered images. This represents about one per cent of the training data. 3D objects are centered and we apply uniform scaling in the range $[-1, 1]$ so that dimensions matches. Input images to our pipeline are RGB images 256×256 .

Test Data. We use the Google Scanned Object (GSO)[8] as our testing dataset, and use the same 30 objects as SyncDreamer[19]. There are 16 images per 3D object, with a fixed elevation of 30° and every 22.5° for azimuth.

Implementation Details. Our model is trained using the AdamW optimiser [24] with a learning rate of 10^{-4} and weight decay of 0.01. We reduce the learning rate to 10^{-5} for a total of 100k training steps. For our training batches, we use 3 input views and 3 target views randomly sampled with replacement from 12 views for each object, with a batch size of 356. We train our model for 6 days on 4 A6000 (48GB) GPUs.

Evaluation Metrics. For novel view synthesis, we report the PSNR, SSIM[44], and LPIPS[47]. For 3D reconstruction from single-view or few views, we use the Chamfer Distances (CD) and 3D IoU between the ground-truth

and reconstructed volumes.

4.1. Novel View Synthesis

We show in Table 1 the performance of MVDiff compared to baselines for novel view synthesis on an unseen dataset [8]. Qualitative results are shown in Figure 2. Our model surpasses baseline Zero-123XL by a margin and benefits from additional views. Given the probabilistic nature of the model it is able to generate diverse and realistic shapes given a single view (see Figure 3).

	Training # Sample	Ref. Views	GSO				NeRF Synthetic			
			PSNR↑	SSIM↑	LPIPS↓	Runtime↓	PSNR↑	SSIM↑	LPIPS↓	Runtime↓
Zero123	800K	1	18.51	0.856	0.127	7s	12.13	0.601	0.421	7s
Zero123-XL	10M	1	18.93	0.856	0.124	8s	12.61	0.620	0.381	8s
MVDiff	800k	1	20.24	0.884	0.095	9s	12.66	0.638	0.342	9s
MVDiff	800k	2	22.91	0.908	0.064	9s	13.42	0.685	0.321	10s
MVDiff	800k	3	24.09	0.918	0.052	10s	13.58	0.741	0.301	11s
MVDiff	800k	5	25.09	0.927	0.043	11s	14.55	0.833	0.288	12s
MVDiff	800k	10	25.90	0.935	0.036	12s	14.51	0.657	0.215	13s

Table 1. **Novel view synthesis performance on GSO and NeRF Synthetic datasets.** MVDiff outperforms Zero-123XL with significantly less training data. Additionally, MVDiff performance exhibits further improvement with the inclusion of more reference views.

4.2. 3D Generation

We showed in section 4.1 that our model can generate multiple consistent novel views. In this section, we perform single and few-images 3D generation on the GSO dataset. We generate 16 views with azimuths uniformly distributed in the range 0° to 360°. For a fixed elevation angle of 30°, SyncDreamer may fail to recover the shape of 3D objects at the top and bottom since the camera angle does not cover those regions. Therefore, we also use different elevation angles from -10° to 40°. Then, we adopt NeuS[40] for 3D reconstruction. The foreground masks of the generated images are initially predicted using CarveKit. It takes around 3 minutes to reconstruct a textured mesh.

We compare our 3D reconstructions with SoTA 3D generation models, including One-2-3-45[17] for decoding an SDF using multiple views predicted from Zero123, and SyncDreamer[19] for fitting an SDF using NeuS[40] from 16 consistent fixed generated views. Given two or more reference views, MVDiff outperforms all other baselines (see Table 2). MVDiff generates meshes that are visually consistent and resembles the ground-truth (see Figure 4).

4.3. Ablation Study

Multi-View Consistency. The generated images may not always plausible and we need to generate multiple instances with different seeds and select a desirable instance for 3D reconstruction based on higher overall PSNR, SSIM and LPIPS for the view generated. Experiments show that we need 5 generations to obtain optimal reconstruction.

	# Input Views	Chamfer Dist. ↓	Volume IoU ↑
Point-E	1	0.0561	0.2034
Shape-E	1	0.0681	0.2467
One2345	1	0.0759	0.2969
LGM	1	0.0524	0.3851
SyncDreamer	1	0.0493	0.4581
MVDiff	1	0.0411	0.4357
MVDiff	2	0.0341	0.5562
MVDiff	3	0.0264	0.5894
MVDiff	5	0.0252	0.6635
MVDiff	10	0.0254	0.6721

Table 2. **3D reconstruction performance on GSO dataset.** MVDiff outperforms other image-to-3D baselines in generating high-quality 3D objects, with improved performance for multiple input views.

	PSNR↑	SSIM↑	LPIPS↓
MVDiff	20.24	0.884	0.095
w/o epipolar att.	19.14	0.864	0.118
w/o multi-view att.	19.92	0.871	0.113

Table 3. **Effect of Self-Attention Mechanisms.** We report PSNR, SSIM [44], and LPIPS [47] for novel view synthesis from single view on GSO dataset. Results show that epipolar attention and multi-view attention lead to superior performance.

Effect of Epipolar and Multi-View Attention. We evaluate the benefits of epipolar attention and multi-view attention on novel view synthesis performing ablation experiments on those components. In particular, we observe a significant drop in performance metrics when removing epipolar attention suggesting that the model is effectively able to implicitly learn 3D object geometry by enforcing geometrical guidance (see Table 3).

Weight Initialisation. An alternative to initializing weights trained from Zero123 on view-dependent objects[7] is to use weights from Stable Diffusion[30]. We compare the performance of our model initializing weights from Stable Diffusion v2[30] with a drop in performance of -2.58 PSNR compared to Zero123[18] weight initialisation. This shows that initializing from Stable Diffusion v2 leads to poorer performance on the novel view task and worse generalisability.

4.4. Risks and Ethical Considerations

There are several promising applications of synthetic data, notably in medicine. Synthetic data could make significant improvement in surgery planning and tailored patient diagnosis leveraging 3D information and its assets of quantitative parameters. Nevertheless, there are ethical considerations associated with the use of synthetic data in medicine. We should ensure the synthetic data is anonymised such that no particular features of the synthetic meshes could link back to a specific patient. In that light, there are transformations that can be applied to the meshes. We should also




Figure 2. **Zero-Shot Novel View Synthesis on GSO.** MVDiff outperforms Zero123-XL for single view generation with greater camera control and generation quality. As more views are added, MVDiff resembles the ground-truth with fine details being captured such as elephant tail and turtle shell design.

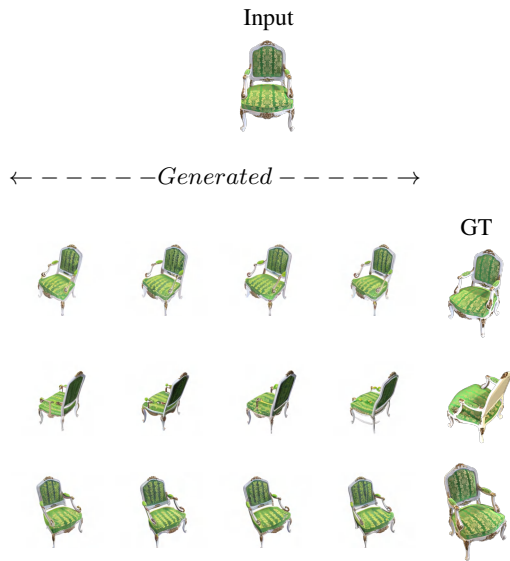


Figure 3. **Diversity of Novel View Diffusion with MVDiff on NeRF-Synthetic Dataset.** We show nearby views (*top and bottom row*) displaying good consistency, while more distant views (*middle*) are more diverse but still realistic.

make sure that the synthetic data is not used in a way it could harm or be detrimental. Further validation on different cohorts of people is required before using these synthetic data in clinical settings.

Despite important ethical considerations we shed light

on, we believe these 3D representations of organs could be of great use, on hand for research purposes to run large-scale statistical analysis on different cohorts and highlight associations with patient metadata. These cost effective synthetic data could be beneficial to improve the visualizations of bones and organs and be deployed widely.

4.5. Limitations

A limitation of this work lies in its computational time and resource requirements. Despite advances in sampling approaches, our model still requires more than 50 steps to generate high-quality images. This is a limit of all diffusion based generation models. Moreover, the reconstructed meshes may not always be plausible. To increase the quality, we may need to use a larger object dataset like Objaverse-XL[7] and manually curate the dataset to filter out uncommon shapes such as point clouds, textureless 3D models and more complex scene representation.

5. Conclusion

In our work, we aimed to address the problem of inconsistencies in multi-view synthesis from single view. We specifically apply epipolar attention mechanisms as well as multi-view attention to aggregate features from multiple views. We propose a simple and flexible framework capable of generating high-quality multi-view images conditioned on an arbitrary number of images.




Figure 4. **3D reconstruction from single-view on GSO dataset.** MVDiff produces consistent novel views and improves the 3D geometry compared to baselines. One-2-3-45 and SyncDreamer tend to generate overly-smoothed and incomplete 3D objects, in particular the sofa.

5.1. Future Work

Combining with graphics. In this study, we show that we can generate view consistent 3D objects by learning geometrical correspondences between views during training. We modified the latent diffusion U-Net model to feed multi view in order to generate consistent multi view for 3D reconstruction. Future work can explore utilising knowledge about lighting, and texture to generate more diverse range of 3D shapes with varying lighting and texture.

References

- [1] Titas Anciukevičius, Zexiang Xu, Matthew Fisher, Paul Henderson, Hakan Bilen, Niloy J. Mitra, and Paul Guerrero. Renderdiffusion: Image diffusion for 3d reconstruction, inpainting and generation, 2024. [2](#)
- [2] Andreas Blattmann, Robin Rombach, Huan Ling, Tim Dockhorn, Seung Wook Kim, Sanja Fidler, and Karsten Kreis. Align your latents: High-resolution video synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22563–22575, 2023. [1](#)
- [3] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018. [1](#)
- [4] Ziang Cao, Fangzhou Hong, Tong Wu, Liang Pan, and Ziwei Liu. Large-vocabulary 3d diffusion model with transformer, 2023. [2](#)
- [5] Yen-Chi Cheng, Hsin-Ying Lee, Sergey Tulyakov, Alexander Schwing, and Liangyan Gui. Sdfusion: Multimodal 3d shape completion, reconstruction, and generation, 2023. [2](#)
- [6] 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. *arXiv preprint arXiv:2212.08051*, 2022. [1](#), [3](#)
- [7] Matt Deitke, Ruoshi Liu, Matthew Wallingford, Huong Ngo, Oscar Michel, Aditya Kusupati, Alan Fan, Christian Laforte, Vikram Voleti, Samir Yitzhak Gadre, Eli VanderBilt, Aniruddha Kembhavi, Carl Vondrick, Georgia Gkioxari, Kiana Ehsani, Ludwig Schmidt, and Ali Farhadi. Objaverse-xl: A universe of 10m+ 3d objects. *arXiv preprint arXiv:2307.05663*, 2023. [4](#), [5](#)
- [8] Laura Downs, Anthony Francis, Nate Koenig, Brandon Kinman, Ryan Michael Hickman, Krista Reymann, Thomas Barlow McHugh, and Vincent Vanhoucke. Google scanned objects: A high-quality dataset of 3d scanned household items. *2022 International Conference on Robotics and Automation (ICRA)*, pages 2553–2560, 2022. [3](#), [4](#)
- [9] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12873–12883, 2021. [1](#)
- [10] Jun Gao, Tianchang Shen, Zian Wang, Wenzheng Chen, Kangxue Yin, Daiqing Li, Or Litany, Zan Gojcic, and Sanja Fidler. Get3d: A generative model of high quality 3d textured shapes learned from images, 2022. [2](#)

- [11] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020. 1
- [12] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. 1
- [13] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 1
- [14] Jonathan Ho, Chitwan Saharia, William Chan, David J Fleet, Mohammad Norouzi, and Tim Salimans. Cascaded diffusion models for high fidelity image generation. *J. Mach. Learn. Res.*, 23(47):1–33, 2022. 1
- [15] Hanwen Jiang, Zhenyu Jiang, Yue Zhao, and Qixing Huang. Leap: Liberate sparse-view 3d modeling from camera poses. *arXiv preprint arXiv:2310.01410*, 2023. 2
- [16] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. *Advances in Neural Information Processing Systems*, 34:852–863, 2021. 1
- [17] 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. 2, 4
- [18] 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. In *ICCV*, 2023. 1, 2, 3, 4
- [19] 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. *arXiv preprint arXiv:2309.03453*, 2023. 1, 2, 3, 4
- [20] Zhen Liu, Yao Feng, Michael J. Black, Derek Nowrouzezahrai, Liam Paull, and Weiyang Liu. Meshdiffusion: Score-based generative 3d mesh modeling, 2023. 2
- [21] Shitong Luo and Wei Hu. Diffusion probabilistic models for 3d point cloud generation. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2836–2844, 2021. 2
- [22] Oscar Michel, Roi Bar-On, Richard Liu, Sagie Benaim, and Rana Hanocka. Text2mesh: Text-driven neural stylization for meshes, 2021. 1
- [23] Norman Müller, Yawar Siddiqui, Lorenzo Porzi, Samuel Rota Bulò, Peter Kontschieder, and Matthias Nießner. Diffrr: Rendering-guided 3d radiance field diffusion, 2023. 2
- [24] Alex Nichol, Heewoo Jun, Prafulla Dhariwal, Pamela Mishkin, and Mark Chen. Point-e: A system for generating 3d point clouds from complex prompts, 2022. 2
- [25] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *CVPR*, 2019. 2
- [26] Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion, 2022. 2
- [27] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022. 1
- [28] Ali Razavi, Aaron Van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with vq-vae-2. *Advances in neural information processing systems*, 32, 2019. 1
- [29] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models, 2021. 3
- [30] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models, 2021. 4
- [31] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022. 1
- [32] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. 1
- [33] Mehdi S. M. Sajjadi, Henning Meyer, Etienne Pot, Urs Bergmann, Klaus Greff, Noha Radwan, Suhani Vora, Mario Lucic, Daniel Duckworth, Alexey Dosovitskiy, Jakob Uszkoreit, Thomas Funkhouser, and Andrea Tagliasacchi. Scene representation transformer: Geometry-free novel view synthesis through set-latent scene representations, 2022. 2
- [34] Ruoxi Shi, Hansheng Chen, Zhuoyang Zhang, Minghua Liu, Chao Xu, Xinyue Wei, Linghao Chen, Chong Zeng, and Hao Su. Zero123++: a single image to consistent multi-view diffusion base model. *arXiv preprint arXiv:2310.15110*, 2023. 2
- [35] Yichun Shi, Peng Wang, Jianguo Ye, Mai Long, Kejie Li, and Xiao Yang. Mvdream: Multi-view diffusion for 3d generation. *arXiv preprint arXiv:2308.16512*, 2023. 2
- [36] Yang Song and Stefano Ermon. Improved techniques for training score-based generative models. *Advances in neural information processing systems*, 33:12438–12448, 2020. 1
- [37] Yang Song, Jascha Narain Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *ICLR*, 2021. 1
- [38] Robust Multiview Stereopsis. Accurate, dense, and robust multiview stereopsis. *IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE*, 32(8), 2010. 2
- [39] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017. 1
- [40] 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. In *Advances in Neural Information Processing Systems*, pages 27171–27183. Curran Associates, Inc., 2021. 2, 4
- [41] 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. In *NeurIPS*, 2021. 2
- [42] Peng Wang, Hao Tan, Sai Bi, Yinghao Xu, Fujun Luan, Kalyan Sunkavalli, Wenping Wang, Zexiang Xu, and Kai Zhang. Pflrm: Pose-free large reconstruction model for joint pose and shape prediction. *arXiv preprint arXiv:2311.12024*, 2023. 2
- [43] Tengfei Wang, Bo Zhang, Ting Zhang, Shuyang Gu, Jianmin Bao, Tadas Baltrusaitis, Jingjing Shen, Dong Chen, Fang Wen, Qifeng Chen, and Baining Guo. Rodin: A generative model for sculpting 3d digital avatars using diffusion, 2022. 2
- [44] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE TIP*, 2004. 3, 4
- [45] Zhenpei Yang, Zhile Ren, Miguel Angel Bautista, Zaiwei Zhang, Qi Shan, and Qixing Huang. Fvor: Robust joint shape and pose optimization for few-view object reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2497–2507, 2022. 2
- [46] Yao Yao, Zixin Luo, Shiwei Li, Tian Fang, and Long Quan. Mvsnet: Depth inference for unstructured multi-view stereo. In *Proceedings of the European conference on computer vision (ECCV)*, pages 767–783, 2018. 2
- [47] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 3, 4