

# EpiDiff: Enhancing Multi-View Synthesis via Localized Epipolar-Constrained Diffusion

Zehuan Huang<sup>1\*</sup> Hao Wen<sup>1\*</sup> Junting Dong<sup>2\*</sup> Yaohui Wang<sup>2</sup> Yangguang Li<sup>3</sup> Xinyuan Chen<sup>2</sup>  
 Yan-Pei Cao<sup>3</sup> Ding Liang<sup>3</sup> Yu Qiao<sup>2</sup> Bo Dai<sup>2†</sup> Lu Sheng<sup>1†</sup>  
<sup>1</sup>Beihang University <sup>2</sup>Shanghai AI Laboratory <sup>3</sup>VAST




Figure 1. **EpiDiff** is able to efficiently generate multi-view consistent and high-quality images from a single view input. Instead of limited to fixed viewpoints, EpiDiff can generate relatively arbitrary multi-views. EpiDiff is lightweight and only takes 12 seconds to generate 16 multi-view images. The generated multiview images can be used to recover 3D shapes by neural reconstruction methods [17, 57].

## Abstract

Generating multiview images from a single view facilitates the rapid generation of a 3D mesh conditioned on a single image. Recent methods [31] that introduce 3D global representation into diffusion models have shown the potential to generate consistent multiviews, but they have reduced generation speed and face challenges in maintaining generalizability and quality. To address this issue, we propose EpiDiff, a localized interactive multiview diffusion model. At the core of the proposed approach is to insert a lightweight epipolar attention block into the frozen diffusion model, leveraging epipolar constraints to enable cross-view interaction among feature maps of neighboring views. The newly initialized 3D modeling module preserves the

original feature distribution of the diffusion model, exhibiting compatibility with a variety of base diffusion models. Experiments show that EpiDiff generates 16 multiview images in just 12 seconds, and it surpasses previous methods in quality evaluation metrics, including PSNR, SSIM and LPIPS. Additionally, EpiDiff can generate a more diverse distribution of views, improving the reconstruction quality from generated multiviews. Please see the project page at [huanngzh.github.io/EpiDiff/](https://huanngzh.github.io/EpiDiff/).

## 1. Introduction

3D reconstruction from a single image has a wide range of applications in augmented/virtual reality, robotic simulation, and gaming. However, this task presents significant challenges due to the complexity of accurately interpreting

\*Equal Contribution

†Corresponding author

spatial structure and details from a 2D perspective observation. It requires extensive knowledge of the 3D world.

Recently, with the introduction of diffusion models [21, 45], the field of 3D generation has experienced rapid and vigorous development. A series of research [1, 3, 6, 9, 14, 15, 19, 24, 25, 32, 34, 37–39, 59, 67, 70] attempts to train 3D generative diffusion models from scratch on 3D assets. However, due to the limited size of publicly available 3D datasets, these methods demonstrate poor generalizability. Therefore, another group of works [7, 22, 35, 40–42, 51, 52, 56, 61, 69] propose to distill prior knowledge of 2D diffusion models to create 3D models from text or images via Score Distillation Sampling (SDS) [40]. Although these methods have achieved impressive results, they require expensive per-scene optimization that usually takes tens of minutes or even hours.

More recently, several works [29, 31] have emerged that finetune pretrained 2D image diffusion models [45] on 3D datasets to generate novel view images, which can be utilized for 3D shape and appearance recovery with reconstruction methods [36, 57]. This feed-forward pipeline not only retains generalization capacity of image diffusion models, but also allows for faster synthesis of multiview images. Zero123 [29] proposes a view-conditioned diffusion model by embedding view features and camera information. However, the multiview consistency of generated images is limited, preventing high-quality 3D shape reconstruction. To improve the multiview consistency, SyncDreamer [31] introduces a 3D feature volume constructed with global multiview images into the Zero123 backbone. However, the heavy global feature volume modeling not only slows down the generation process but also easily overfits to specific viewpoints and tends to produce low-quality images with distorted appearances and blurriness (see Fig. 4).

In this paper, we present EpiDiff, a novel localized interactive multiview diffusion model. The key idea is to utilize the feature maps of neighboring views that suffer fewer occlusions for the generation of target view images, which improves the generalization capability and efficiency. To fuse the neighboring feature maps, during the generation of  $N$  target views, we insert a lightweight epipolar attention block into the UNet, which leverages epipolar constraints to enable interaction among feature maps of  $F$  neighboring views within a smaller range. The 3D modeling module preserves the original feature distribution of the diffusion model, exhibiting compatibility with various base diffusion models. Compared to global modeling methods, our localized modeling approach not only speeds up the generation process, but also facilitating the generation of more diversely distributed views, contributing to improved reconstruction from generated multiviews.

We conduct extensive experiments on the Google Scanned Object dataset [12] and various styles of 2D im-

ages. The experiments validate that, compared to baseline methods, EpiDiff can generate higher quality multiview images more rapidly. The contributions of EpiDiff are summarized as follows:

- EpiDiff employs a 3D modeling module into the frozen diffusion model. The module preserves the original feature distribution of the diffusion model, exhibiting compatibility with various diffusion models.
- We propose an epipolar attention block to learn the inter-correlations among neighboring views based on epipolar geometry relationships. This localized interactive and lightweight module models consistency effectively.
- Experiments show that EpiDiff generates 16 multiviews in just 12 seconds, and it surpasses previous methods in metrics including PSNR, SSIM, and LPIPS. Additionally, it can generate more freely distributed views, improving the reconstruction quality from generated multiviews.

## 2. Related work

**3D diffusion models** Recently, the field of 3D generation has witnessed rapid and significant advancements with the introduction of diffusion models. Motivated by the impressive generative performance of diffusion models trained on extensive image datasets, several studies have explored training 3D generative diffusion models from scratch on 3D assets. These 3D diffusion models directly generate point clouds [34, 38, 70], meshes [14, 32], and neural radiance fields [1, 3, 6, 9, 15, 19, 24, 25, 37, 39, 59, 67] through network inference. However, owing to the considerably smaller size of publicly accessible 3D datasets compared to image datasets, these methods exhibit limited generalizability. They often generate shapes restricted to specific categories and face challenges in adapting to more intricate input conditions.

**3D generation guided by 2D prior** Given the robust generative capabilities of 2D diffusion models [43, 45, 46], some research efforts directly leverage pretrained image diffusion models. With DreamFusion [40] and SJC [56] proposing to distill a pretrained 2D text-to-image generation model for generating 3D shapes, a series of subsequent studies [2, 7, 18, 27, 35, 41, 42, 51, 52, 61, 68, 69] adopt such a pre-scene optimization pipeline. These methods optimize a 3D representation (e.g., NeRF, mesh, or SDF) conditioned on text or images through Score Distillation Sampling (SDS) [40]. However, the per-scene optimization approach has issues with efficiency. Generating a single scene may take tens of minutes or even hours.

**Using 2D diffusion for multiview synthesis** 2D diffusion models have introduced significant progress in generating multiview images from a single view. While some approaches [4, 11, 13, 16, 53, 55, 62, 63, 66] integrate 3D representations or utilize shared attention for multiview or

texture synthesis, they lack design for 3D reconstruction. Zero123 [29] introduces relative view condition to diffusion models, enabling novel view synthesis (NVS) from a single image. However, due to its limited consistency, it utilizes SDS loss [40] to achieve 3D reconstruction. One-2-3-45 [28] leverages a generalizable neural reconstruction method [33] for producing 3D shapes from multiview images generated by Zero123 [29], yet the resultant quality is suboptimal. More recently, SyncDreamer [31] has introduced a 3D global feature volume during multiview generation to ensure consistency. However, the heavy global 3D representation modeling not only slows down the generation but also tends to produce low-quality images. In contrast, our localized modeling method facilitates the generation of multiview-consistent and high-quality images.

**Generalizable radiance field** Since NeRF [36] achieved advancements in fitting single scene for novel view synthesis, numerous studies have explored generalizable neural radiance field representation for various scenes. These methods take a few input view images and generate a 3D implicit field in a single feed-forward pass, eliminating the need for per-scene optimization. Among the methods, some explicitly construct the 3D feature or cost volume [5, 23, 33, 72] and utilize the voxel feature for decoding density and color, inspiring SyncDreamer[31] to model 3D global representation, while others [8, 20, 26, 30, 44, 47, 48, 50, 54, 58, 64, 65] directly aggregate 2D features using MLPs or transformers, guided by spatial projection relationships. These methods, either explicitly or implicitly modeling generalized scenes, provide valuable insights for consistency modeling in the feed-forward multiview generation process.

### 3. Method

Given an input view  $y$  of an object, our goal is to generate multiview images of the object. Let  $R \in \mathbb{R}^{3 \times 3}$  and  $T \in \mathbb{R}^3$  represent the rotation and translation of the camera relative to the input view for the target viewpoints, respectively. We denote the target viewpoints for the  $N$  views as  $[R, T]^{1:N}$ . Therefore, the aim is to learn a model  $f$  that accepts the input view and synthesizes multiview images under specified camera transformations:

$$\hat{x}^{1:N} = f(y, [R, T]^{1:N}) \quad (1)$$

where  $\hat{x}^{1:N}$  is denoted as the generated multiview images.

To enhance the consistency, quality and efficiency of the multiview generation, we introduce a localized interactive multiview diffusion scheme based on the novel view synthesis base model. We design a lightweight attention block in the original feature space of the diffusion model, with the aim of establishing nearby view associations to model 3D representation. As illustrated in Fig. 2, the proposed approach involves an interaction process among the intermediate features of multiview images (refer to Sec. 3.1),

thereby extending single-view synthesis to multiview synthesis. Specifically, we design a 3D modeling module, named Epipolar-constrained Attention Block (ECA Block) based on the attention mechanism and insert it into the base diffusion model (refer to Sec. 3.2). The ECA Block facilitates information exchange among latent features of neighboring views, leveraging epipolar constraints in stereo vision to enhance consistency and quality in multiview images. The ECA Block, with its lightweight and adaptable characteristics, is suitable for various base models.

### 3.1. Multiview diffusion model

**Base diffusion model** The proposed multiview diffusion model is rooted in image diffusion models, typically constructed on the Latent Diffusion Models (LDM) [45]. The LDM operates the denoising process in the latent space of an autoencoder, namely  $\mathcal{E}(\cdot)$  and  $\mathcal{D}(\cdot)$ . During the training phase, an input image  $x_0$  is initially mapped to the latent space using a frozen encoder, yielding  $z_0 = \mathcal{E}(x_0)$ , then perturbed by a pre-defined Markov process:

$$q(z_t|z_{t-1}) = \mathcal{N}(z_t; \sqrt{1 - \beta_t}z_{t-1}, \beta_t I) \quad (2)$$

for  $t = 1, \dots, T$ , with  $T$  represents the number of steps in the forward diffusion process. The sequence of hyperparameters  $\beta_t$  determines the noise strength at each step. The denoising UNet  $\epsilon_\theta$ , featuring four downsample blocks, one middle block and four upsample blocks to attain four resolution levels, is trained to approximate the reverse process  $q(z_{t-1}|z_t)$ . The vanilla training objective is expressed as:

$$\mathcal{L} = \mathbb{E}_{\mathcal{E}(x_0), \epsilon \sim \mathcal{N}(0, I), t} [\|\epsilon - \epsilon_\theta(z_t, t)\|_2^2] \quad (3)$$

Zero123 [29], a base model for novel view synthesis (NVS), employs an input view  $y$  and camera parameters  $[R, T]$  as conditions to extend from image to 3D domain. This model, built on a pretrained image diffusion model and augmented with camera control, aims to optimize the following objective:

$$\mathcal{L} = \mathbb{E}_{\mathcal{E}(x_0), \epsilon \sim \mathcal{N}(0, I), t} [\|\epsilon - \epsilon_\theta(z_t, t, c(y, R, T))\|_2^2] \quad (4)$$

where  $c(x, R, T)$  is the embedding of the input view and relative camera extrinsics.

**Interactive multiview diffusion** Drawing upon the foundational image or NVS diffusion models, our multiview diffusion model promotes spatial interaction among images at the latent feature level. It incorporates 3D geometric significance during the step-by-step denoising process, thus generating 3D-consistent multiview images.

Given a dataset of multiview images and their associated camera extrinsics  $\{(x^{1:N}, [R, T]^{1:N})\}$ , we finetune a foundational diffusion model to learn 3D consistency.  $x^1$  is selected as the input view image  $y$  in the training stage. Multiview images  $x^{1:N}$  are mapped to latent codes  $z_0^{1:N}$  through the VAE encoder  $\mathcal{E}$ .

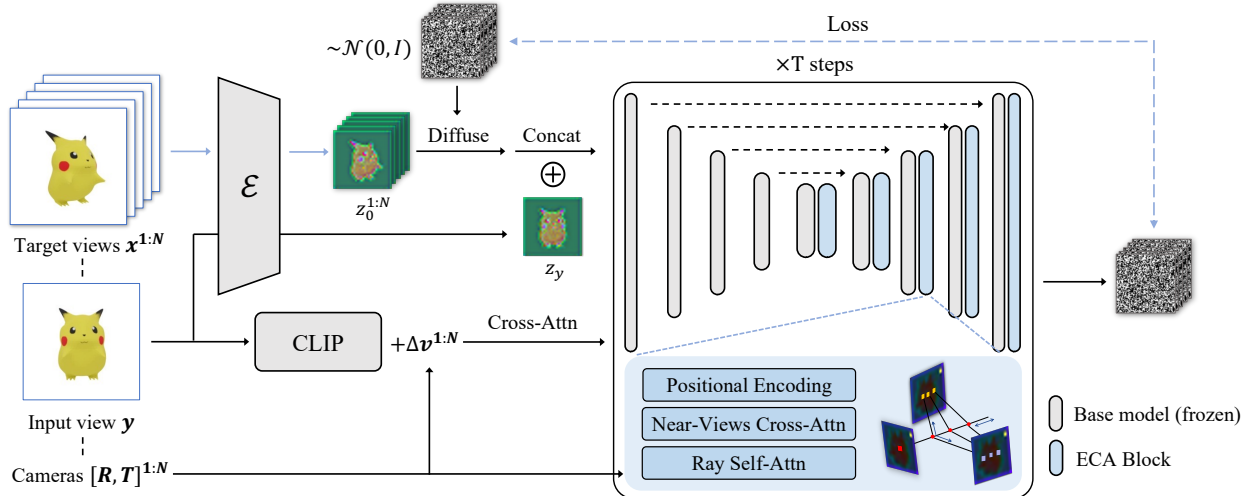


Figure 2. **Pipeline of EpiDiff.** Based on a base NVS model (e.g., Zero123 [29]), our method designs a module for modeling 3D consistency, which is inserted into the mid-sample and up-sample stages of the UNet. We use attention networks to construct the module, aimed at learning generalized epipolar geometry, termed Epipolar-constrained Attention Block (ECA Block). During training, only parameters of the ECA Block are updated, thereby preserving the feature space of the base model and encouraging the module to extract 3D priors.

We introduce a generalized 3D consistency module into the UNet of the base model. This module, inserted into the middle and each upsample block of the UNet, allows multiview image feature maps to interact in 3D space, creating consistent interconnections. The integration of the module transforms the generation of multiple images from independent processes into a synchronous interactive process. The final denoising network  $\epsilon_\theta$  is trained on a dataset comprising multiview images and camera extrinsics. The optimization objective is defined as:

$$\mathcal{L} = \mathbb{E}_{\mathcal{E}(x_0^{1:N}), \epsilon \sim \mathcal{N}(0, I), t} [\|\epsilon - \epsilon_\theta(z_t^{1:N}, t, y, [R, T]^{1:N})\|_2^2] \quad (5)$$

By freezing the base model and exclusively training the 3D modeling module, we maintain the original generative capabilities of the base model. This add-on module, modeling multiview associations, enhances the base diffusion models’ ability to generate multiview-consistent images. Our design of the multiview diffusion framework is adaptable to various base diffusion models.

### 3.2. Epipolar-constrained attention

We introduce a module that uses a series of attention networks for cross-view interaction, in accordance with the principles of 3D projection. This module facilitates the learning of a generalizable epipolar geometry, which we refer to as the Epipolar-constrained Attention Block (ECA Block). The ECA Block identifies correlations between latent features of the current view and corresponding features along the epipolar line in other views.

As illustrated in Fig. 3, the ECA Block comprises two primary attention blocks, supplemented with linear layers:

(i) the Near-Views Cross-Attention module, tasked with aggregating epipolar line features from neighboring views for current view ray features, and (ii) the Ray Self-Attention module, designed to model the depth weights of the sampled rays, fuse spatial points into pixel features.

**Near-Views Cross-Attention** Given a target view, our method relies on the features of the  $F$  closest views among the total  $N$  views to produce the output feature map. We start by emitting rays from patches on the target view, and sample  $S$  points uniformly along each ray. We then project them onto  $F - 1$  neighboring views. For each point  $s$ , we get the feature  $v^s$  by positional encoding and projected feature  $\{f_i^s \mid 1 \leq i \leq F - 1\}$  on reference views. We enable cross-view interaction by a cross-attention layer. Specifically, using point feature  $v^s$  as query, we implement  $\text{Attention}(Q, K, V) = \text{Softmax}(\frac{QK^T}{\sqrt{d}}) \cdot V$  with

$$Q = W^Q v^s, K = W^K [f_1^s, \dots, f_{F-1}^s], V = W^V [f_1^s, \dots, f_{F-1}^s], \quad (6)$$

where  $[\cdot]$  denotes concatenation operation, and  $W^Q, W^K, W^V$  are learnable matrices that project the inputs to query, key, and value, respectively.

During the synthesis of the target view feature map, features from nearby views with a 3D geometric relationship to the target view are integrated, as described in Eq. (6). This approach effectively captures inter-view feature correlations, thereby enhancing multiview consistency efficiently.

**Ray Self-Attention** To model spatial points’ contributions to pixel-level features, the ray aggregation module is implemented using self-attention. Upon acquiring the feature of each sampled point  $\tilde{v}^s (1 \leq s \leq S)$ , we apply self-

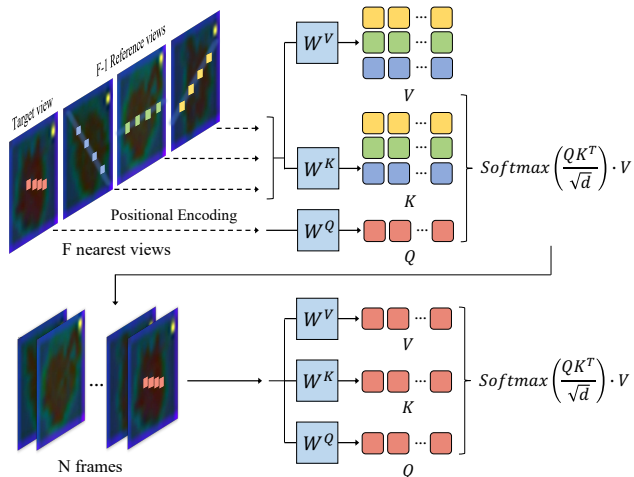


Figure 3. **Illustration of our Epipolar-constrained Attention Block (ECA Block).** Latent features of multiview images are associated in 3D space through two key attention blocks. The Near-Views Cross-Attn initially aggregates features from nearby views onto the target view’s ray points, guided by epipolar geometry. Subsequently, Ray Self-Attn models depth information to fuse ray features into 2D feature maps. The ECA Block facilitates multiview interaction by effectively harnessing spatial geometry.

attention along the ray with dimension of  $S$ , deriving all of query, key and value features from  $[v^1, \dots, v^S]$ . This step is crucial for revealing depth relationships among sampled points along rays and capturing their spatial correlations. Finally, we perform a weighted summation to fuse 3D features into 2D feature of target patch  $\tilde{f}$ .

Our method aligns features precisely within four resolution levels of the UNet latent space by leveraging epipolar projection, which preserves the feature integrity of the base model. The subdivision of the 3D attention process into two modules further allows for efficient multiview consistency modeling.

**Positional Encoding** To enhance the module’s understanding of perspective relationships, we have adopted a positional encoding technique inspired by Light Field Networks (LFN) [49]. This approach is distinct in that each pixel is associated with a ray, utilizing Plücker coordinates for representation, denoted as  $r = (o \times d, d)$ , where  $o$  signifies the ray’s origin, and  $d$  its direction. Unlike traditional position encoding methods where all pixels of an image share the same positional encoding, the LFN [49] method employs MVP transformation to back-project each pixel of the image into the world coordinate system, precisely locating the corresponding ray for each pixel. It introduces strict geometric constraints and sufficiently incorporates camera parameters, effectively applying encoding at the pixel level. By integrating these encoded coordinates obtained by LFN with latent feature maps via a harmonic transformation, our

method significantly enriches the spatial information within the feature representation.

Furthermore, to address the potential biases caused by using absolute positions, we also adopt the strategy from GPNR [50], establishing a ray-relative positioning system through its corresponding camera extrinsics  $[R, T]$ . We construct canonicalizing transformation  $\mathcal{T}$  using the Y-axis of  $R(y)$  and the ray direction ( $v$ ) as the Z-axis. Let  $v' = \frac{v}{\|v\|}$  and  $y' = y - (y \cdot v')v'$ , the transformation  $\mathcal{T}$  is formed as:

$$R_c = \begin{bmatrix} y' & y' \times v' & y' \\ \frac{y'}{\|y'\|} & \frac{y'}{\|y'\|} & v' \end{bmatrix} \quad (7)$$

$$\mathcal{T} = [R_c^T \mid -R_c^T T] \quad (8)$$

By applying  $\mathcal{T}$  to each ray, we adjust them to originate from  $(0, 0, 0)$  and align along the direction  $(0, 0, 1)$ . Finally, we transform the rays from all neighboring viewpoints into the relative coordinate system. It ensures that the rays are consistently represented relative to a common reference, mitigating biases from absolute positioning.

## 4. Experiments

### 4.1. Experiment settings

**Dataset** Our model is trained on the LVIS subset of the Objaverse [10], comprising over 40,000 3D models. For each object, we setup a camera layout to uniformly sample 96 views with a size of  $256 \times 256$ . The layout comprises six concentric circles, each as a layer, with elevation angles set at  $\{-10^\circ, 0^\circ, 10^\circ, 20^\circ, 30^\circ, 40^\circ\}$ . Each circle contains 16 cameras spaced evenly from  $0^\circ$  to  $360^\circ$  in azimuth angle. During training, 16 views from each model are randomly sampled.

We evaluate EpiDiff using the Google Scanned Object dataset [12], which consists of over one thousand scanned models. To assess multiview synthesis, we employ two rendering settings. Firstly, aligning with SyncDreamer’s training setting [31], we fix elevation at  $30^\circ$  and sample 16 cameras with azimuths evenly spread from  $0^\circ$  to  $360^\circ$ . Secondly, we render 16 views with azimuths also uniformly distributed in  $[0^\circ, 360^\circ]$ , but with varying elevations from  $-10^\circ$  to  $40^\circ$ . This tests the model’s ability to generate multiview images across a wider range of viewpoints.

**Implementation details** We use Zero123 [29] as our backbone. EpiDiff is trained for 30,000 steps (around 2 days) using 8 80G A800 GPUs, with a learning rate of  $1e-5$  and a total batch size of 512. In each iteration, we generate 16 multiview images, i.e.,  $N = 16$ . Within the ECA Block, the number of sampled rays is aligned with the resolution of feature maps and the number of sampled points per ray  $S$  is set to 16. For the Near-Views Cross Attention module, the number of nearby views  $F$ , is set to 4. For reconstruction,

Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	Runtime $\downarrow$
Zero123 [29]	17.79	0.796	0.201	7s
SyncDreamer [31]	20.11	0.829	0.159	60s
Ours	<b>20.49</b>	<b>0.855</b>	<b>0.128</b>	12s

Table 1. Quantitative comparison in multiview synthesis under elevation  $30^\circ$  setting. We report PSNR, SSIM [60], LPIPS [71], Runtime on the GSO [12] dataset.

the foreground masks of the generated images are initially predicted using CarveKit. Then, we use the instant-NGP based SDF reconstruction method [17] to reconstruct the shape. Each object undergoes training for 3k steps, which costs about 2 minutes.

**Baselines** We adopt Zero123 [29], One-2-3-45 [28], Point-E [38], Shap-E [24], SyncDreamer [31] as baseline methods. We adopt the official implementation of Zero123 XL[29]. One-2-3-45 [28] uses Zero123 to synthesize multiview images and regresses SDFs from them. Point-E [38] and Shap-E [24] are 3D generative models trained on a large internal OpenAI 3D dataset. We adopt the best released version of Point-E, *base1B*, which is a 1 billion parameter image to point cloud diffusion model. We use it to generate point clouds, then convert the generated point clouds to SDFs using the official models. For Shap-E [24], We use its STF rendering mode to predict SDF values and texture colors, enabling direct generation of colored vertex meshes. SyncDreamer [31] builds a multiview shared 3D feature volume based on Zero123, aiming to generate multiview consistent images from a single input image. We utilize the official source code in our evaluation.

**Metrics** We primarily focus on two tasks: multiview synthesis and single-view 3D reconstruction. For multiview synthesis, we employ the commonly used evaluation metrics, i.e., PSNR, SSIM [60], and LPIPS [71] to assess our method’s performance. In single-view 3D reconstruction task, we measure the quality of reconstructed shapes using Chamfer Distance and Volumetric Intersection over Union (Volume IoU) compared to the ground truth. For comparing reconstruction performance between Zero123 [29], SyncDreamer [31], and our method EpiDiff, we use multiview images generated under the fixed elevation  $30^\circ$  setting.

## 4.2. Consistent multiview synthesis

For this task, we use two viewpoint sampling settings. Under the setting with a fixed elevation of  $30^\circ$ , the quantitative results are shown in Tab. 1 and the qualitative results are shown in Fig. 4. Zero123 [29] generates visually plausible images, but the generated images are inconsistent. SyncDreamer [31] generates more consistent multiview images, but the image quality is relatively poor, exhibiting blurry and unrealistic appearances. Our method efficiently gener-

Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	Runtime $\downarrow$
Zero123 [29]	15.91	0.772	0.231	7s
SyncDreamer [31]	15.90	0.773	0.246	60s
Ours	<b>18.83</b>	<b>0.821</b>	<b>0.163</b>	12s

Table 2. Quantitative comparison in multiview synthesis under uniform elevation setting. We report PSNR, SSIM [60], LPIPS [71], Runtime on the GSO [12] dataset.

Method	Chamfer Dist. $\downarrow$	Volume IoU $\uparrow$
One-2-3-45 [28]	0.0768	0.2936
Point-E [38]	0.0570	0.2027
Shap-E [24]	0.0689	0.2473
Zero123 [29]	0.0543	0.3358
SyncDreamer [31]	0.0496	0.4149
Ours	<b>0.0429</b>	<b>0.4518</b>

Table 3. Quantitative comparison of surface reconstruction. We report Chamfer Distance and Volume IoU on the GSO dataset [12].

ates multiview images that not only maintain semantic consistency with the input image, but also demonstrate higher performance in terms of structure and realism.

Under the uniform elevation setting, the quantitative results are shown in Tab. 2, and the qualitative results are shown in Fig. 5. The uniform view setting offers a more flexible viewpoint distribution. SyncDreamer [31] lacks precise cross-view geometric constraints, leading to overfitting training viewpoints and a bias towards top-down perspectives in multiview generation, as shown in Fig. 5. Our method can still generate relatively realistic and multiview-consistent images, benefiting from (a) the introduction of localized epipolar geometric constraints along with effective camera position encoding, and (b) the random viewpoint sampling training method. Meanwhile, for the same input image, Our method can generate different plausible instances using different random seeds as shown in Fig. 6.

## 4.3. Single view reconstruction

We present the quantitative results in Tab. 3 and the qualitative results in Fig. 7. One-2-3-45 [28] generates smooth meshes using a generalizable neural reconstruction approach [33], but lacks detail. Point-E [38] and Shap-E [24] tend to generate incomplete shapes and exhibit differences from the objects shown in the input images. Zero123 [29], due to its poor multiview consistency, results in missing or rough geometry in the reconstruction. SyncDreamer [31] generates a roughly consistent geometric framework but struggles to handle details. In comparison, our method achieves the best reconstruction quality in terms of smooth surfaces and detailed geometric structures.

We also observe, as shown in Fig. 8, that reconstruct-




Figure 4. Qualitative comparison with Zero123 [29] and SyncDreamer [31] under the elevation  $30^\circ$  setting.




Figure 5. Qualitative comparison with Zero123 [29] and SyncDreamer [31] under the uniform elevation setting. We present the generated results of each method when the elevation is  $-10^\circ$ ,  $10^\circ$  and  $30^\circ$ .

	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	Runtime $\downarrow$
$N = 16, F = 4$	<b>20.49</b>	<b>0.855</b>	<b>0.128</b>	<b>12s</b>
$N = 16, F = 8$	20.27	0.847	0.135	18s
$N = 16, F = 16$	20.32	0.833	0.152	30s
w/o light field PE	19.83	0.845	0.139	–
w/o ray-relative PE	20.01	0.841	0.141	–

Table 4. Ablation study under the elevation  $30^\circ$  setting. We report PSNR, SSIM [60], LPIPS [71] and Runtime for each ablation. Results show that our nearby views aggregation and the positional encoding lead to superior generalization performance.

ing using 16 generated views with a fixed elevation leads to unrealistic shapes at the bottom of objects, which can be attributed to the unseen area underneath. Conversely, employing a flexible viewpoint distribution with a wider coverage angle enhances the quality of reconstruction.

#### 4.4. Ablation study

To assess the effectiveness of our design for ECA Block (Sec. 3.2), we conducted ablation experiments on multiview

synthesis at elevation  $30^\circ$  setting (Tab. 4). We investigated the impact of the number of neighboring views used for aggregating features in Near-Views Cross Attention module, setting  $F$  to  $\{4, 8, 16\}$ , where  $F = 16$  indicates the use of all perspective features aggregated to the target view. Results show that increasing  $F$  consumes more time and reduces the quality of the generated multiview structure and appearance. Additionally, we examined the positional encoding (PE) strategies and found that removing either the light field PE or the ray-relative PE leads to a decrease in the consistency and quality of the generated images.

## 5. Limitations and Conclusion

**Limitations and future works** The proposed EpiDiff still has the following limitations. Firstly, EpiDiff’s efficacy in generating multiview images from arbitrary viewpoints is less pronounced compared to closer views. This is primarily due to the instability of the base NVS model when generating novel views significantly distant from the input view. We anticipate the emergence of more robust base models and enhanced stability in multiview synthesis upon integra-




Figure 6. Different plausible instances generated by EpiDiff from the same input image.

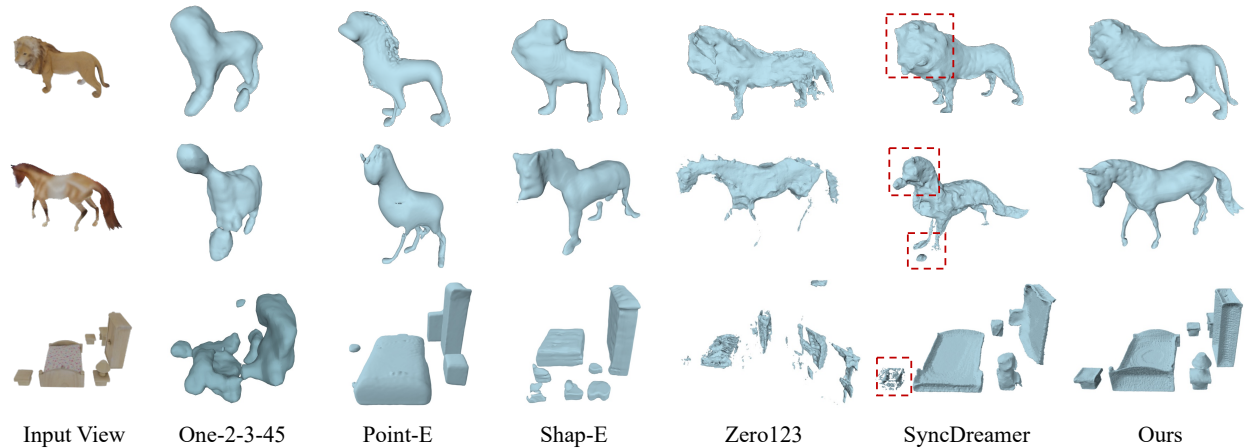


Figure 7. Qualitative comparison of surface reconstruction from single image. For Zero123 [29], SyncDreamer [31] and EpiDiff, we use multiview images generated under elevation  $30^\circ$  setting.




Figure 8. The common failure mode when reconstructing using multiviews at a fixed elevation of  $30^\circ$ . When the bottom of the object is round-shaped, the reconstructed bottom appears protruded because the camera’s visible frustum can not cover the bottom. Using a uniform distribution of elevations can solve it.

tion with EpiDiff. Secondly, the process of generating a 3D model from a single image is divided into two separate steps: multiview synthesis and 3D reconstruction. Notably,

the acquisition of 3D knowledge is an integral component of both these stages. Therefore, a unified pipeline, trained on extensive 2D image datasets and 3D datasets, would offer a marked improvement in efficiency and generalizability.

**Conclusion** In this paper, we present EpiDiff, a localized interactive multiview diffusion model designed for generating multiview-consistent and high-quality images from a single-view image. With the Zero123 [29] as the backbone, EpiDiff adopts a lightweight epipolar attention block in the UNet, which utilizes epipolar constraints to enable cross-view interaction among feature maps of neighboring views. This module models 3D consistency within the original feature space of the diffusion model, exhibiting adaptability to various base diffusion models. Extensive experiments demonstrate that EpiDiff not only efficiently generates consistent multiview images, but is also capable of generating more freely distributed views, help to better reconstruction from generated multiviews.

### Acknowledgment

This work was supported by National Natural Science Foundation of China (62132001).



## 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. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12608–12618, 2023. [2](#)
- [2] Yukang Cao, Yan-Pei Cao, Kai Han, Ying Shan, and Kwan-Yee K Wong. Dreamavatar: Text-and-shape guided 3d human avatar generation via diffusion models. *arXiv preprint arXiv:2304.00916*, 2023. [2](#)
- [3] Ziang Cao, Fangzhou Hong, Tong Wu, Liang Pan, and Ziwei Liu. Large-vocabulary 3d diffusion model with transformer. *arXiv preprint arXiv:2309.07920*, 2023. [2](#)
- [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. [2](#)
- [5] Anpei Chen, Zexiang Xu, Fuqiang Zhao, Xiaoshuai Zhang, Fanbo Xiang, Jingyi Yu, and Hao Su. Mvnerf: Fast generalizable radiance field reconstruction from multi-view stereo. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14124–14133, 2021. [3](#)
- [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](#)
- [7] Rui Chen, Yongwei Chen, Ningxin Jiao, and Kui Jia. Fantasia3d: Disentangling geometry and appearance for high-quality text-to-3d content creation. *arXiv preprint arXiv:2303.13873*, 2023. [2](#)
- [8] Weihao Cheng, Yan-Pei Cao, and Ying Shan. Sparsegnv: Generating novel views of indoor scenes with sparse input views. *arXiv preprint arXiv:2305.07024*, 2023. [3](#)
- [9] Yen-Chi Cheng, Hsin-Ying Lee, Sergey Tulyakov, Alexander G Schwing, and Liang-Yan Gui. Sdfusion: Multimodal 3d shape completion, reconstruction, and generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4456–4465, 2023. [2](#)
- [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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13142–13153, 2023. [5](#)
- [11] Congyue Deng, Chiyu Jiang, Charles R Qi, Xinchun Yan, Yin Zhou, Leonidas Guibas, Dragomir Anguelov, et al. Nerdi: Single-view nerf synthesis with language-guided diffusion as general image priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20637–20647, 2023. [2](#)
- [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. [2](#), [5](#), [6](#)
- [13] Chenjian Gao, Boyan Jiang, Xinghui Li, Yingpeng Zhang, and Qian Yu. Genesistex: Adapting image denoising diffusion to texture space, 2024. [2](#)
- [14] 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. *Advances In Neural Information Processing Systems*, 35:31841–31854, 2022. [2](#)
- [15] Jiatao Gu, Qingzhe Gao, Shuangfei Zhai, Baoquan Chen, Lingjie Liu, and Josh Susskind. Learning controllable 3d diffusion models from single-view images. *arXiv preprint arXiv:2304.06700*, 2023. [2](#)
- [16] 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. [2](#)
- [17] Yuan-Chen Guo. Instant neural surface reconstruction, 2022. <https://github.com/bennyguo/instant-nsr-pl>. [1](#), [6](#)
- [18] Yuan-Chen Guo, Ying-Tian Liu, Ruizhi Shao, Christian Laforte, Vikram Voleti, Guan Luo, Chia-Hao Chen, Zi-Xin Zou, Chen Wang, Yan-Pei Cao, and Song-Hai Zhang. threestudio: A unified framework for 3d content generation. <https://github.com/threestudio-project/threestudio>, 2023. [2](#)
- [19] Anchit Gupta, Wenhan Xiong, Yixin Nie, Ian Jones, and Barlas Oğuz. 3dgen: Triplane latent diffusion for textured mesh generation. *arXiv preprint arXiv:2303.05371*, 2023. [2](#)
- [20] Philipp Henzler, Jeremy Reizenstein, Patrick Labatut, Roman Shapovalov, Tobias Ritschel, Andrea Vedaldi, and David Novotny. Unsupervised learning of 3d object categories from videos in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4700–4709, 2021. [3](#)
- [21] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. [2](#)
- [22] Ajay Jain, Ben Mildenhall, Jonathan T Barron, Pieter Abbeel, and Ben Poole. Zero-shot text-guided object generation with dream fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 867–876, 2022. [2](#)
- [23] Mohammad Mahdi Johari, Yann Lepoittevin, and François Fleuret. Geonerf: Generalizing nerf with geometry priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18365–18375, 2022. [3](#)
- [24] Heewoo Jun and Alex Nichol. Shap-e: Generating conditional 3d implicit functions. *arXiv preprint arXiv:2305.02463*, 2023. [2](#), [6](#)
- [25] Seung Wook Kim, Bradley Brown, Kangxue Yin, Karsten Kreis, Katja Schwarz, Daiqing Li, Robin Rombach, Antonio Torralba, and Sanja Fidler. Neuralfield-ldm: Scene generation with hierarchical latent diffusion models. In *Proceed-*

- ings of the *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8496–8506, 2023. 2
- [26] Jonáš Kulháněk, Erik Derner, Torsten Sattler, and Robert Babuška. Viewformer: Nerf-free neural rendering from few images using transformers. In *European Conference on Computer Vision*, pages 198–216. Springer, 2022. 3
- [27] Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiaohui Zeng, Xun Huang, Karsten Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. Magic3d: High-resolution text-to-3d content creation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2
- [28] Minghua Liu, Chao Xu, Haian Jin, Linghao Chen, Mukund Varma T, Zexiang Xu, and Hao Su. One-2-3-45: Any single image to 3d mesh in 45 seconds without per-shape optimization, 2023. 3, 6
- [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, 2023. 2, 3, 4, 5, 6, 7, 8
- [30] Yuan Liu, Sida Peng, Lingjie Liu, Qianqian Wang, Peng Wang, Christian Theobalt, Xiaowei Zhou, and Wenping Wang. Neural rays for occlusion-aware image-based rendering. arxiv cs. *CV*, 2107:1, 2021. 3
- [31] Yuan Liu, Cheng Lin, Zijiao Zeng, Xiaoxiao Long, Lingjie Liu, Taku Komura, and Wenping Wang. Syncdreamer: Learning to generate multiview-consistent images from a single-view image. *arXiv preprint arXiv:2309.03453*, 2023. 1, 2, 3, 5, 6, 7, 8
- [32] Zhen Liu, Yao Feng, Michael J Black, Derek Nowrouzezahrai, Liam Paull, and Weiyang Liu. Meshdiffusion: Score-based generative 3d mesh modeling. *arXiv preprint arXiv:2303.08133*, 2023. 2
- [33] Xiaoxiao Long, Cheng Lin, Peng Wang, Taku Komura, and Wenping Wang. Sparseneus: Fast generalizable neural surface reconstruction from sparse views. In *European Conference on Computer Vision*, pages 210–227. Springer, 2022. 3, 6
- [34] Shitong Luo and Wei Hu. Diffusion probabilistic models for 3d point cloud generation. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [35] Luke Melas-Kyriazi, Iro Laina, Christian Rupprecht, and Andrea Vedaldi. Realfusion: 360deg reconstruction of any object from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8446–8455, 2023. 2
- [36] 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. In *ECCV*, 2020. 2, 3
- [37] Norman Müller, Yawar Siddiqui, Lorenzo Porzi, Samuel Rota Buló, Peter Kotschieder, and Matthias Nießner. Diffri: Rendering-guided 3d radiance field diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4328–4338, 2023. 2
- [38] 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. 2, 6
- [39] Evangelos Ntavelis, Aliaksandr Siarohin, Kyle Olszewski, Chaoyang Wang, Luc Van Gool, and Sergey Tulyakov. Autodecoding latent 3d diffusion models. *arXiv preprint arXiv:2307.05445*, 2023. 2
- [40] Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv*, 2022. 2, 3
- [41] Guocheng Qian, Jinjie Mai, Abdullah Hamdi, Jian Ren, Aliaksandr Siarohin, Bing Li, Hsin-Ying Lee, Ivan Skokhodov, Peter Wonka, Sergey Tulyakov, and Bernard Ghanem. Magic123: One image to high-quality 3d object generation using both 2d and 3d diffusion priors. *arXiv preprint arXiv:2306.17843*, 2023. 2
- [42] Amit Raj, Srinivas Kaza, Ben Poole, Michael Niemeyer, Nataniel Ruiz, Ben Mildenhall, Shiran Zada, Kfir Aberman, Michael Rubinstein, Jonathan Barron, et al. Dreambooth3d: Subject-driven text-to-3d generation. *arXiv preprint arXiv:2303.13508*, 2023. 2
- [43] 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. 2
- [44] Jeremy Reizenstein, Roman Shapovalov, Philipp Henzler, Luca Sbordone, Patrick Labatut, and David Novotny. Common objects in 3d: Large-scale learning and evaluation of real-life 3d category reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10901–10911, 2021. 3
- [45] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models, 2021. 2, 3
- [46] 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*, 35:36479–36494, 2022. 2
- [47] Mehdi SM Sajjadi, Daniel Duckworth, Aravindh Mahendran, Sjoerd van Steenkiste, Filip Pavetic, Mario Lucic, Leonidas J Guibas, Klaus Greff, and Thomas Kipf. Object scene representation transformer. *Advances in Neural Information Processing Systems*, 35:9512–9524, 2022. 3
- [48] Mehdi SM Sajjadi, Henning Meyer, Etienne Pot, Urs Bergmann, Klaus Greff, Noha Radwan, Suhani Vora, Mario Lučić, Daniel Duckworth, Alexey Dosovitskiy, et al. Scene representation transformer: Geometry-free novel view synthesis through set-latent scene representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6229–6238, 2022. 3
- [49] Vincent Sitzmann, Semon Rezkikov, Bill Freeman, Josh Tenenbaum, and Frédo Durand. Light field networks: Neural scene representations with single-evaluation rendering. In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 19313–19325, 2021. 5

- [50] Mohammed Suhail, Carlos Esteves, Leonid Sigal, and Ameesh Makadia. Generalizable patch-based neural rendering. In *European Conference on Computer Vision*, pages 156–174. Springer, 2022. 3, 5
- [51] Jingxiang Sun, Bo Zhang, Ruizhi Shao, Lizhen Wang, Wen Liu, Zhenda Xie, and Yebin Liu. Dreamcraft3d: Hierarchical 3d generation with bootstrapped diffusion prior. *arXiv preprint arXiv:2310.16818*, 2023. 2
- [52] 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. 2
- [53] 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. 2
- [54] Alex Trevithick and Bo Yang. GRF: learning a general radiance field for 3d scene representation and rendering. *CoRR*, abs/2010.04595, 2020. 3
- [55] Hung-Yu Tseng, Qinbo Li, Changil Kim, Suhub Alsisan, Jia-Bin Huang, and Johannes Kopf. Consistent view synthesis with pose-guided diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16773–16783, 2023. 2
- [56] Haochen Wang, Xiaodan Du, Jiahao Li, Raymond A Yeh, and Greg Shakhnarovich. Score jacobian chaining: Lifting pretrained 2d diffusion models for 3d generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12619–12629, 2023. 2
- [57] 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. 1, 2
- [58] Qianqian Wang, Zhicheng Wang, Kyle Genova, Pratul Sriniwasan, Howard Zhou, Jonathan T. Barron, Ricardo Martin-Brualla, Noah Snavely, and Thomas Funkhouser. Ibrnet: Learning multi-view image-based rendering. In *CVPR*, 2021. 3
- [59] Tengfei Wang, Bo Zhang, Ting Zhang, Shuyang Gu, Jianmin Bao, Tadas Baltrusaitis, Jingjing Shen, Dong Chen, Fang Wen, Qifeng Chen, et al. Rodin: A generative model for sculpting 3d digital avatars using diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4563–4573, 2023. 2
- [60] 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, 7
- [61] Zhengyi Wang, Cheng Lu, Yikai Wang, Fan Bao, Chongxuan Li, Hang Su, and Jun Zhu. Prolificdreamer: High-fidelity and diverse text-to-3d generation with variational score distillation. *arXiv preprint arXiv:2305.16213*, 2023. 2
- [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. 2
- [63] Jianfeng Xiang, Jiaolong Yang, Binbin Huang, and Xin Tong. 3d-aware image generation using 2d diffusion models. *arXiv preprint arXiv:2303.17905*, 2023. 2
- [64] Hao Yang, Lanqing Hong, Aoxue Li, Tianyang Hu, Zhenguo Li, Gim Hee Lee, and Liwei Wang. Contranerf: Generalizable neural radiance fields for synthetic-to-real novel view synthesis via contrastive learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16508–16517, 2023. 3
- [65] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from one or few images. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 4578–4587. Computer Vision Foundation / IEEE, 2021. 3
- [66] Jason J Yu, Fereshteh Forghani, Konstantinos G Derpanis, and Marcus A Brubaker. Long-term photometric consistent novel view synthesis with diffusion models. *arXiv preprint arXiv:2304.10700*, 2023. 2
- [67] Wang Yu, Xuelin Qian, Jingyang Huo, Tiejun Huang, Bo Zhao, and Yanwei Fu. Pushing the limits of 3d shape generation at scale. *arXiv preprint arXiv:2306.11510*, 2023. 2
- [68] Wangbo Yu, Li Yuan, Yan-Pei Cao, Xiangjun Gao, Xiaoyu Li, Long Quan, Ying Shan, and Yonghong Tian. Hifi-123: Towards high-fidelity one image to 3d content generation. *arXiv preprint arXiv:2310.06744*, 2023. 2
- [69] Xin Yu, Yuan-Chen Guo, Yangguang Li, Ding Liang, Song-Hai Zhang, and Xiaojuan Qi. Text-to-3d with classifier score distillation, 2023. 2
- [70] Xiaohui Zeng, Arash Vahdat, Francis Williams, Zan Gojcic, Or Litany, Sanja Fidler, and Karsten Kreis. LION: latent point diffusion models for 3d shape generation. In *NeurIPS*, 2022. 2
- [71] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. 6, 7
- [72] Xiaoshuai Zhang, Sai Bi, Kalyan Sunkavalli, Hao Su, and Zexiang Xu. Nerfusion: Fusing radiance fields for large-scale scene reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5449–5458, 2022. 3