

# ZeroRF: Fast Sparse View 360° Reconstruction with Zero Pretraining

Ruoxi Shi\* Xinyue Wei\* Cheng Wang Hao Su  
 UC San Diego




Figure 1. We demonstrate fast 360° reconstruction from sparse training views via ZeroRF. ZeroRF is able to perform novel view synthesis from few views (6 as shown in the figure) with exceptional quality, while also being fast, obtaining competitive results within 2 minutes and finishing in around 25 minutes at the full 800<sup>2</sup> resolution. For common resolutions like 256<sup>2</sup> or 320<sup>2</sup> in 3D generation applications, ZeroRF reconstructs an object from sparse-view generations in only 30 seconds.

## Abstract

We present *ZeroRF*, a novel per-scene optimization method addressing the challenge of sparse view 360° reconstruction in neural field representations. Current breakthroughs like *Neural Radiance Fields (NeRF)* have demonstrated high-fidelity image synthesis but struggle with sparse input views. Existing methods, such as *Generalizable NeRFs* and *per-scene optimization approaches*, face limitations in data dependency, computational cost, and generalization across diverse scenarios. To overcome these challenges, we propose *ZeroRF*, whose key idea is to integrate a tailored *Deep Image Prior* into a factorized *NeRF* representation. Unlike traditional methods, *ZeroRF* parametrizes feature grids with a neural network generator, enabling efficient sparse view 360° reconstruction without any pretraining or additional regularization. Extensive experiments showcase *ZeroRF*'s versatility and superiority in terms of both quality and speed, achieving state-of-the-art results on benchmark datasets. *ZeroRF*'s significance extends to applications in 3D content generation and editing. Project page: <https://sarahweiii.github.io/zerorf/>.

<sup>1</sup>\* Equal contribution.

## 1. Introduction

Breakthroughs in neural field representations, like *Neural Radiance Fields (NeRF)* [38] and its subsequent developments [3, 8–10, 12, 16, 39, 58, 59, 64, 69, 78], have paved the way for high-fidelity image synthesis, expedited optimization processes, and various downstream applications. Nevertheless, these approaches hinge on having a rich set of input views, and they exhibit a marked degradation in performance when confronted with sparse input views. In practical scenarios, it is not always feasible to obtain a comprehensive set of high-resolution images along with precise camera data, especially when it comes to 3D content generation [31, 34, 56]. Therefore, addressing the reconstruction from sparse views presents a notable challenge, yet it remains a critical and pivotal area of interest.

In recent years, there has been a growing focus on methods tailored for sparse view reconstruction [7, 23, 26, 33, 41, 61, 65, 67, 77]. One line of approaches [7, 28, 33, 77], commonly referred to as *Generalizable NeRFs*, rely on extensive pretraining with substantial time and data requirements to directly reconstruct the scenes of interest. Performances of these models are thus closely related to the quality of the training data, and their resolutions are limited due to the heavy computation cost of large neural networks. Moreover, it is also hard for these models to generalize ef-




Figure 2. **Visualization of features obtained by fitting a vanilla TensorRF on sparse and dense views.** With dense views the features obtained are clean, while with sparse views the features are distorted with lots of noise and unwanted artifacts.

fectively across diverse scenarios. Other approaches that follow the per-scene optimization paradigm incorporate extra modules, like vision language models [23] and depth estimators [65] to help with the reconstruction. While these methods prove effective in managing narrow baselines, they fall short in achieving optimal performance in 360° reconstruction. Additionally, their applicability to real-world data is limited due to their dependence on additional supervision, which may not always be available or accurate. People have also manually designed priors spanning continuity [41], information theory [26], symmetry [51] and frequency [73] regularizations for the task. However, the extra regularizations may prevent the NeRFs from reconstructing the scenes faithfully [73]. Furthermore, handcrafted priors are often not robust to even quite subtle setting changes.

We also observe that existing per-scene optimization approaches for 360° reconstruction typically demand hours of training even on the most powerful GPUs today, which hinders their use in real applications. All of them are based on the original NeRF representation, which converges much slower compared to factorized NeRF representations like iNGP [39] or TensorRF [8]. The reason is that those handcrafted priors can hardly be applied to new representations. FreeNeRF [73], for example, uses a regularization technique upon positional encodings that are specific to NeRF.

We fit a TensorRF [8] with different number of training views (4 and 100) on the Lego scene from the NeRF-Synthetic [38] dataset and visualize one channel from the plane features after the training converges. From Fig. 2 we can clearly see ray artifacts that result in noisy and distorted features under the sparse (4) view setting, while with dense (100) views the feature plane looks exactly like an orthogonal projection image. We carried out similar experiments on the triplane [6, 17] and DiF [9] representations and find that this is not specific to TensorRF but is a general phenomenon for these grid-based factorized representations. Thus, we hypothesize that *fast sparse view reconstruction with optimization can be achieved if the factorization features remain clean under sparse view supervision.*

To verify and achieve this, we propose to integrate a tailored version of the Deep Image Prior [62] into a fac-

torized NeRF representation (See Fig. 3). Instead of directly optimizing feature grids as in TensorRF, K-planes or DiF [8, 9, 17], we parametrize the feature grids with a randomly-initialized deep neural network (*generator*). The intuition behind this is that with under-determined supervision, neural networks generalize much better than look-up grids for the vast majority of cases, if not always. More theoretically speaking, neural networks have much higher impedance on noise and artifacts compared to data easy to perceive and remember [21, 22, 62]. The design works without any extra regularizations or pretraining, and can uniformly apply to multiple representations. The parametrization is also “lossless” as there exist a set of deep network parameters such that any given target feature grid could be achieved [62].

We carried out extensive experiments on different generator networks for parametrization and different factorized representations to find the most suitable combinations for sparse view 360° reconstruction (Sec. 5.3), and come up with ZeroRF, a novel per-scene optimization method for this challenging task. ZeroRF 1) does not require any sort of model pretraining, avoiding any potential bias towards training data and any limits on settings like resolution or camera distribution; 2) is fast in training and inference, as it is built upon factorized NeRF representations, running in as low as 30 seconds; 3) has the same theoretical expressiveness as the underlying factorized representations; 4) achieves state-of-the-art quality for novel view synthesis with sparse-view input on NeRF-Synthetic [38] and OpenIllumination [30] benchmarks (Sec. 5.2).

Given the high-quality 360° reconstruction capabilities of ZeroRF, our method finds applications in various domains, including 3D content generation and editing. The potential of our approach in addressing these tasks is demonstrated in Sec. 6.

## 2. Related Work

### 2.1. Novel View Synthesis

Neural rendering techniques have paved the way for achieving photo-realistic rendering quality in novel view synthesis. It all started with the inception of NeRF [38], which was the first to introduce a MLP for storing the radiance field and achieving remarkable rendering quality through volume rendering. Subsequently, many follow-up studies have presented various representations aimed at further enhancing performance. For instance, approaches like Plenoxels [16] and DVGO [58] employed voxel-based representations, while TensorRF [8], iNGP [39], and DiF [9] put forward decomposition strategies to expedite training. MipNeRF [3] and RefNeRF [64] are founded on coordinate-based MLPs, and Point-NeRF [72] relies on a point-cloud-based representation.

Some methods replace the density field with the SDF [29, 42, 49, 66, 68, 75] or turn density fields into mesh representation [12, 40, 59, 69, 76] to improve surface reconstruction. These methods can extract superior-quality meshes without a substantial compromise in their rendering quality. Additionally, recent works [25, 70, 74] have used Gaussian splatting to achieve real-time radiance field rendering.

## 2.2. Deep Network Priors

While people commonly believe that the success of deep neural networks is due to their capability to learn from large-scale datasets, the architecture of deep networks actually capture a great amount of features prior to any learning. Training a linear classifier on features from a random convolutional network can yield performance much higher than random guess [20]. Features from randomly initialized networks are also good for few-shot learners [1, 18, 50]. Via distillation upon this random features, the prior can be pushed further, with a line of self-supervised methods including BYOL [20], DeepCluster [5] and Selective Pseudo-labeling [37] starting from this inductive bias and use different methods to boost this prior for representation learning for images.

In contrast to these works, Deep Image Prior [62] directly exploits this deep prior without further distillation. It shows that a GAN generator architecture can act as a parametrization with high noise impedance, and thus can be applied to image restoration tasks such as denoising, super-resolution and inpainting. This is further applied to various imaging and microscopy applications [36, 43, 54, 55, 63], and extended with theoretical and practical improvements in Deep Decoders [21, 22]. ZeroRF follows a similar paradigm to embed the deep prior into the parametrization of radiance fields.

## 2.3. Sparse View Reconstruction

Despite of the exceptional performance, NeRF models exhibit limitations in producing accurate solutions when trained with sparse observations due to insufficient information. To address this challenge, some methods opt for *pretraining* [7, 13, 47, 60, 77] on extensive datasets to impart prior knowledge and fine-tune the model on the target scene. Conversely, an alternative line of research focuses on per-scene optimization through *manually designed regularizations* [23, 48, 51–53, 57, 61, 73]. For example, to increase semantic consistency, DietNeRF [23] extracts high-level features with the CLIP Vision Transformer [45]. Many of them design loss functions to alleviate cross-view inconsistency, either based on information theory [26, 41]. SPARF [14, 61, 65] leverages pretrained networks for correspondence or depth estimation to compensate for the lack of 3D information. Different from these existing arts, ZeroRF

demonstrates a remarkable ability to synthesize novel views without relying on pretraining or explicit regularizations.

## 3. Preliminaries

**Neural Radiance Field (NeRF)** represents a 3D scene radiance field by an MLP, where given an input 3D location  $x$  and the view direction  $d$ , it outputs the volume density  $\sigma_x$  and view-dependent color  $c_x$ :

$$\sigma_x, c_x = F(x, d) \quad (1)$$

Then the density  $\sigma$  and color  $c$  are used in the differentiable volume rendering:

$$\hat{C}(r) = \sum_{i=1}^N T_i (1 - \exp(-\sigma_i \delta_i)) c_i, T_i = \exp\left(-\sum_{j=1}^{i-1} \sigma_j \delta_j\right) \quad (2)$$

where  $\hat{C}(r)$  is the volume rendering predicted RGB colors for ray  $r$ ,  $T$  is the volume transmittance and  $\delta$  is the ray marching step size. The whole rendering process is differentiable, which allows the neural network to be optimized by rendering loss:

$$\mathcal{L} = \sum_{r \in R} \|\hat{C}(r) - C(r)\|_2^2 \quad (3)$$

where  $C(r)$  is the ground truth RGB colors.

**TensorRF** swapped out the initial MLP utilized in NeRF, opting for a feature volume to expedite training. It further breaks down this feature volume into factors using CANDECOMP/PARAFAC decomposition or Vector-Matrix (VM) decomposition. In our work, we mainly focus on the VM decomposition, where given a 3D tensor  $\mathcal{T} \in \mathbb{R}^{I,J,K}$ , it decomposes a tensor into multiple vectors and matrices:

$$\mathcal{T} = \sum_{r=1}^{R_1} v_r^1 \circ M_r^{2,3} + \sum_{r=1}^{R_2} v_r^2 \circ M_r^{1,3} + \sum_{r=1}^{R_3} v_r^3 \circ M_r^{1,2} \quad (4)$$

where  $v_r^a$  are vector factors and  $M_r^{b,c}$  are matrix factors.

## 4. Method

### 4.1. Overview

The ZeroRF pipeline is illustrated in Fig. 3. We use deep generator networks with a frozen standard Gaussian noise sample as input to generate the planes and vectors in the TensorRF-VM style, forming a decomposed tensorial feature volume. The feature volume is then sampled among render rays and decoded by a multi-layer perceptron (MLP). We employ the standard volume rendering process and a plain MSE loss.




Figure 3. **Architecture of ZeroRF.** It parametrizes TensorRF-VM tensors with randomly-initialized deep generator networks (Sec. 4.3), with the input to the networks set to a frozen Gaussian noise on start of training. The system performs per-scene optimization using the standard volume rendering procedure with a plain rendering loss.

The main idea of ZeroRF is to apply untrained deep generator networks as a parametrization of spatial feature grids. The network can learn patterns of different scales from the sparse view observations and naturally generalize to unseen views, without the need of further progressive upsampling tricks or explicit regularizations that typically require a lot of manual labor to tune, as opposed to prior works for sparse view reconstruction. Nevertheless, there are still several points of design left in the pipeline: the spatial organization, or the representation of the feature volume; the architecture of the representation generator; and the architecture of the feature decoder. We will detail these designs in the following sections.

## 4.2. Factorizing the Feature Volume

The principle of applying deep generator networks for parametrization is universal to any grid-based representation. The most straightforward solution is to parameterize a feature volume directly. However, this is memory and compute inefficient as we would need a very large feature volume if we want a decent volume rendering quality. This is not peculiar to ZeroRF; many prior arts for dense-view reconstruction actually work on this factorization. TensorRF [8] uses tensorial decompositions to exploit the low-rankness of feature volumes. The triplane representation used in EG3D [6] and K-planes [17] can be seen as a special case of TensorRF-VM representation when the vectors are constants. Dictionary Fields (DiF) [9] factorizes the feature volume into multiple smaller volumes encoding different frequencies. Instant-NGP [39] employ a multi-resolution hashmap as information in the feature volume is sparse in nature.

Among these factorizations, hashing breaks the spatial correlation between adjacent cells, so deep priors cannot be applied. Deep generator networks can be used to parameterize all the rest three representations (TensorRF, triplane and DiF). We built generator architectures for generating 1D vectors, 2D matrices, and 3D volumes, upon which we experimented with all three factorizations. All of them work

similarly and achieve better performance than previous arts, but the TensorRF-VM representation achieves slightly better performance on our test benchmarks overall. Thus, we employ the TensorRF-VM representation as our final choice of factorization, as is shown in Fig. 3.

We include the comparison between different factorizations in Sec. 5.3.

## 4.3. Generator Architecture

The quality of deep parametrization highly depends on the architecture. Most generator networks to date are convolutional and attention architectures. When designing ZeroRF, we investigated various structures including Deep Decoders (DD) [21], the variational autoencoder (VAE) used in Stable Diffusion (SD) [44], the decoder used in Kadinsky models [46], and the SimMIM generator based on a ViT decoder [71]. We change the 2D convolution, pooling and upsampling layers into 1D and 3D to obtain the corresponding 1D and 3D generators required in different factorizations.

These generators are originally quite large in size as they were designed to be fit on a very large dataset for generation of high-quality contents. This results in both unnecessarily long run-times and slower convergence when it comes to fitting to a single NeRF scene. Fortunately, we find that the performance of ZeroRF after convergence remains intact when we shrink the models in width and depth. Thus, we keep the block composition but modify these architectures in size to boost the training speed. Note that we only need to store the radiance field representation and not the generator during inference, so ZeroRF has zero overhead compared to its underlying representation during rendering.

We found that the SD VAE and its decoder part, as well as Kadinsky decoder work similarly well for novel view synthesis, followed by Deep Decoders, while the SimMIM architecture proves to be invalid as a deep prior for radiance fields. SD/Kadinsky coders are mostly convolutional architectures, with Kadinsky adding self-attention to the first two blocks. We took the (modified) SD decoder as our final choice of generator architecture as it has the least computa-

tion. We carry out a more complete analysis on the results of using different generators in Sec. 5.3.

#### 4.4. Decoder Architecture

Our decoder architecture follows that of SSDNeRF [11]. We sample with linear interpolation (or bilinear, trilinear according to the dimension) from the feature grid at the point to decode, and project it with a first linear layer to get a base feature code that is shared between density and appearance decodings. We find that sharing the feature code can help reduce floaters by coupling geometry and appearance closely. We apply SiLU activation and invoke another linear layer for density prediction. For color prediction, we encode the view direction with Sphere Harmonics (SH) and add its projection by a linear layer to the base feature to involve view dependence. We then apply SiLU activation and use another linear layer, similar to the density prediction, to predict RGB values. Formally, we have

$$\sigma_x = \exp(\Theta_\sigma(\text{SiLU}(\Theta_b(F_x)))) , \quad (5)$$

$$c_x = \sigma(\Theta_c(\text{SiLU}(\Theta_b(F_x) + \Theta_d(\text{SH}(d)))) , \quad (6)$$

where  $F_x$  is the feature field,  $\sigma(\cdot)$  is the sigmoid function, and  $\Theta_\bullet$  denotes a linear weight layer.

Note that different from the decoders used in TensorRF and DiF, this decoder does not consume any positional encodings, as there would otherwise be a chance to break or degrade ZeroRF by leaking the position information outside the deep prior.

#### 4.5. Implementation Details

In our experiments, we use the AdamW optimizer [27, 35] with  $\beta_1 = 0.9, \beta_2 = 0.98$  and a weight decay of 0.2. The learning rate starts at 0.002 and decays to 0.001 with a cosine schedule. We train ZeroRF for  $10k$  iterations. We uniformly sample 1024 points per ray during volume rendering, and employ occupancy pruning and occlusion culling to accelerate the process. We include figures for detailed architecture of our generator and decoder in Appendix C.

### 5. Experiments

#### 5.1. Experiment Setups

**Datasets and Metrics.** We evaluate our method on sparse view reconstruction using NeRF-Synthetic [38], OpenIllumination [30] and DTU [24] datasets. We use the standard PSNR, SSIM and LPIPS [79] metrics for evaluation.

**NeRF-Synthetic** is a synthetic dataset rendered by Blender, which contains 8 objects with various materials and geometric structures. We use 4 or 6 views as the input and evaluate the model on 200 testing views.

**OpenIllumination** is a real-world dataset featuring 360° reconstruction. We narrowed our focus to 8 objects display-

ing intricate geometry under a single illumination setup, extracting 4 or 6 views from the available pool of 38 training views and evaluating on 10 testing views.

**DTU** mainly focuses on forward-facing objects instead of 360° reconstruction, but for the sake of completeness, we include our results on DTU in Fig. 6. We use 3 views as the input and test the model on the rest of the views. We include more comparisons and quantitative results in Appendix B.

All the input views are selected by running KMeans [2, 32] on the camera translation vector and picking the views closest to cluster centroids.

**Baselines.** We compare our ZeroRF against a few state-of-the-art few-shot NeRF methods: RegNeRF based on continuity and pretrained RealNVP [15] regularization [41], DietNeRF [23] that uses a pretrained CLIP [45] prior, InfoNeRF [26] using entropy as regularizer, FreeNeRF [73] based on frequency regularization, and FlipNeRF [51] using a spatial symmetry prior.

#### 5.2. Results

The quantitative results for NeRF-Synthetic and OpenIllumination are presented in Tab. 1 and Tab. 2, respectively. Across both the 4-view and 6-view experiments, our approach consistently outperforms all other methods, as evidenced by superior PSNR, SSIM, and LPIPS scores. Moreover, our method achieves these results in significantly less time. Even with only 2 minutes of training, ZeroRF remains superior or competitive to the best baselines.

Visual comparisons between ZeroRF and baseline methods are illustrated in Fig. 4 and Fig. 5. Most of the baseline models exhibit noticeable flaws of varying degrees, including floaters and apparent color shifts in synthesis results (highlighted within red boxes in the figure). For pretrained priors, the RegNeRF prior model was not trained on wide-baseline images, and fails to reconstruct objects under 360° settings; DietNeRF using CLIP as a prior model interestingly works better on real images than on synthetic images, which is consistent with CLIP’s pretraining data distribution. For non-pretrained models, InfoNeRF and FreeNeRF applying information-theoretical and frequency regularizers fail to represent intricate structures like Ficus leaves. Notably, FreeNeRF and FlipNeRF perform relatively well on NeRF-Synthetic, but fail catastrophically on OpenIllumination. This shows that handcrafted priors are not robust to setting changes. FlipNeRF fails with numerical instabilities during training on OpenIllumination, which is also observed by Wang on their own data <sup>1</sup>. ZeroRF shows the best visual quality and robustness across diverse datasets and is much less prone to floaters in comparison or unrealistic color shifts on all scenes.

<sup>1</sup><https://github.com/shawn615/FlipNeRF/issues/3>

Table 1. Comparison with the state-of-the-art sparse view reconstruction methods on NeRF-Synthetic dataset.

Method	Prior	Train Time	4 views			6 views		
			PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
RegNeRF [41]	RealNVP	$\sim 6$ h	9.93	0.419	0.572	9.82	0.685	0.580
DietNeRF [23]	CLIP	$\sim 6$ h	10.92	0.557	0.446	16.92	0.727	0.267
InfoNeRF [26]	Info. Theory	$\sim 3.5$ h	17.83	0.802	0.212	21.44	0.854	0.159
FreeNeRF [73]	Frequency	$\sim 3$ h	18.81	0.808	0.188	22.77	0.865	0.149
FlipNeRF [51]	Symmetry	$\sim 3$ h	19.78	0.822	0.212	21.73	0.845	0.202
Ours (1k iters)	Deep Parametrization	2 min	21.42	0.843	0.192	23.47	0.860	0.177
Ours (2k iters)		5 min	21.85	0.852	0.170	24.29	0.875	0.153
Ours (Full)		$\sim 25$ min	<b>21.94</b>	<b>0.856</b>	<b>0.139</b>	<b>24.73</b>	<b>0.889</b>	<b>0.113</b>




Figure 4. **Qualitative comparison between ZeroRF and previous works on NeRF-Synthetic dataset.** The top two rows (Hotdog and Mic) are reconstruction results from 4 views and the bottom two rows (Ficus and Ship) are reconstruction results from 6 views. ZeroRF results have the best visual quality, and is free of walls or floaters.

### 5.3. Analysis

**Effect of the Number of Training Views.** We designed experiments to show the benefit of our proposed method with the number of input views. We plot the results in Fig. 7. ZeroRF has a significant advantage to the base TensorRF representation on sparse views (3 to 8). When the views become denser, ZeroRF remains competitive, though by a smaller margin.

**Feature Volume Factorization Choices.** We apply ZeroRF generators to Triplane, TensorRF and DiF and compare the performance of resulting parametrizations on the NeRF-Synthetic dataset (6-view setting). The results are shown in

Tab. 3. The inclusion of generators consistently improves upon base representations, and they all achieve state-of-the-art performance. This shows that the principles of using a deep parametrization is generally applicable to grid-based representations (also see Sec. 4.2). Among the factorization methods, TensorRF with our prior performs the best, so we chose TensorRF as our final feature volume factorization.

**Generator Architecture.** We applied different generator architectures introduced in Sec. 4.3 upon the TensorRF factorization and compared their performance in Tab. 4. To further investigate the effect of different priors, we visualize one channel of the plane features with different genera-

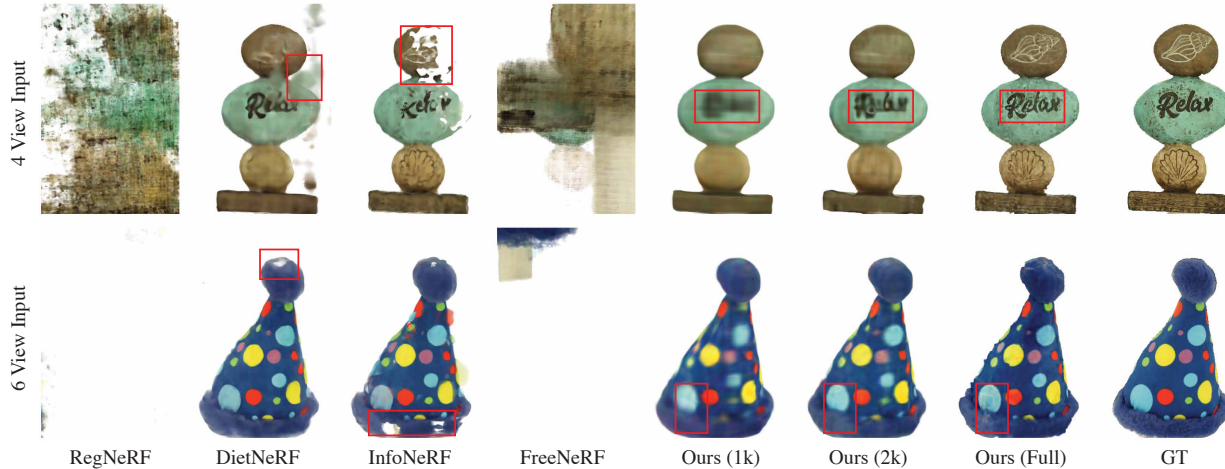


Figure 5. Qualitative comparison between ZeroRF and previous works on OpenIllumination dataset.



Figure 6. Results of our method on DTU with only 3 views as input. See Appendix B for more details and comparisons.

Table 2. Comparison with the state-of-the-art sparse view reconstruction methods on OpenIllumination.

	Train Time	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
RegNeRF [41]	~ 6 h	14.08	0.859	0.303
DietNeRF [23]	~ 6 h	24.20	0.940	0.095
InfoNeRF [26]	~ 3.5 h	22.28	0.935	0.090
FreeNeRF [73]	~ 3 h	11.47	0.813	0.351
Ours (1k iters)	2 min	26.26	0.939	0.097
Ours (2k iters)	5 min	26.49	0.941	0.090
Ours (Full)	~ 25 min	26.15	0.939	0.080

tors in Fig. 8. Without any prior and directly optimizing the planes, the features are noisy with high-frequency glitches and visible view boundary lines all across. In contrast, the SD Decoder and Kadinsky models produce clean and well-post features. The fully attentional ViT decoder of SimMIM works with patches, and we can see visible blocky artifacts. MLP assumes a very smooth transition over the grid and thus is unable to represent scene content faithfully. Overall, the convolutional architectures produce features that align the best with the scene.

ZeroRF is robust to over-parametrization of the net-




Figure 7. PSNR of ZeroRF versus vanilla TensorRF on NeRF-Synthetic dataset.

Table 3. Applying our prior to different grid-based representations. ZeroRF parametrization consistently enhances the models to better generalize to unseen views. Results are from the 6-view NeRF-Synthetic setting.

Grid Repr.	Our Prior	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
Triplane [6]	$\times$	20.55	0.822	0.201
	$\checkmark$	<b>24.38</b>	<b>0.887</b>	<b>0.117</b>
DiF [9]	$\times$	17.45	0.775	0.258
	$\checkmark$	<b>23.56</b>	<b>0.878</b>	<b>0.125</b>
TensorRF [8]	$\times$	21.73	0.844	0.169
	$\checkmark$	<b>24.73</b>	<b>0.889</b>	<b>0.113</b>

works. The results are similar if we scale the decoder with 2x more layers (the second last row in Tab. 4).

**Importance of the Noise.** The input noise is the key to our prior. Swapping it with a trainable feature initialized with zeros breaks the system completely (the last row in Tab. 4). We do not observe performance improvements if we unfreeze the noise – as the learning rate would be small compared to the scale of the noise, the structure of the noise is kept unchanged throughout training. But it introduces extra overhead and slows down the convergence. Thus, we

Table 4. **Ablation study on VM generator architecture.** Results are from the 6-view NeRF-Synthetic setting. ‘Up’ in the table refers to bilinear upsampling.

Generator	Architecture	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
No Generator	-	21.73	0.844	0.169
MLP on Grid	MLP	20.69	0.827	0.214
SimMIM [71]	Attn. (ViT)	21.40	0.839	0.156
SD VAE [44]	Conv.	24.68	0.890	0.116
Kadinsky [46]	Attn. + Conv.	24.69	0.889	0.112
DD [21]	Lin. + Up	23.93	0.876	0.148
SD Decoder	Conv.	24.73	0.889	0.113
2 $\times$ layers	Conv.	24.69	0.889	0.122
No Noise	Conv.	12.21	0.785	0.375




Figure 8. **Visualization of plane features from different generators.** Different architecture impose different priors on features.

keep the noise frozen during training. We also found that noise scales of 0.5 and 2.0 both lead to a slight performance degradation while using a uniform noise slightly improves. We believe this is due to distribution assumptions of the initialization schemes in PyTorch.




Figure 9. **Text-to-3D and Image-to-3D generation results with ZeroRF.** ZeroRF can naturally handle model-generated multi-view images, and reconstruct 360 $^\circ$  views from the sparse view generations with high quality in 30 seconds.




Figure 10. **Texture generation with ZeroRF.** ZeroRF can be used to apply new appearance to a given geometry, with the assistance of language-based image editing models.

## 6. Applications

**Text to 3D and Image to 3D.** Given the powerful sparse-view reconstruction capability of ZeroRF, a straightforward idea is to use an existing model to perform consistent multi-view generation, and apply ZeroRF to lift the sparse view into 3D. In this example, for image-to-3D we employ Zero123++ [56] to lift single image input into 6-view images and directly fit a ZeroRF on the generated images. For text-to-3D, we first invoke SDXL [44] to generate an image from the text, and apply the image-to-3D procedure described before. As shown in Fig. 9, ZeroRF is able to produce faithful high-quality reconstructions from generated multi-view images. Fitting the ZeroRF only costs 30 seconds on a single A100 GPU.

**Mesh Texturing and Texture Editing.** ZeroRF can also be utilized to reconstruct the appearance with a frozen provided geometry. To do this, we render 4 images of the mesh from random views, tile them into one large image and apply Instruct-Pix2Pix [4] to edit the images according to a text prompt. We then fit a ZeroRF on the four images, and bake the color values back to the mesh surface. In this case, fitting the ZeroRF only requires 20 seconds. Fig. 10 (and also the rightmost column in Fig. 1) demonstrate results of texture editing on the Bob mesh (Fig. 10 Left) and mesh texturing on Stanford Bunny (Fig. 10 Right).

## 7. Conclusion and Future Work

In this work we present ZeroRF, a novel method for fast and high quality sparse view 360 $^\circ$  reconstruction. Based on a deep parametrization technique, it can be applied on various factorized grid-based radiance fields, achieving state-of-the-art performance for sparse view 360 $^\circ$  reconstruction without the need of designing any specific regularizations or incorporating any pretraining priors. One possible future work would be extending ZeroRF to unbounded scenes; we discuss more limitations and future work in Appendix E.

## Acknowledgement

This work is supported in part by gifts from Qualcomm.



## References

- [1] Ehsan Amid, Rohan Anil, Wojciech Kotłowski, and Manfred K Warmuth. Learning from randomly initialized neural network features. *arXiv preprint arXiv:2202.06438*, 2022. [3](#)
- [2] David Arthur and Sergei Vassilvitskii. K-means++ the advantages of careful seeding. In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 1027–1035, 2007. [5](#)
- [3] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5470–5479, 2022. [1, 2](#)
- [4] Tim Brooks, Aleksander Holynski, and Alexei A Efros. Instructpix2pix: Learning to follow image editing instructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18392–18402, 2023. [8](#)
- [5] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In *Proceedings of the European conference on computer vision (ECCV)*, pages 132–149, 2018. [3](#)
- [6] Eric R. Chan, Connor Z. Lin, Matthew A. Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas Guibas, Jonathan Tremblay, Sameh Khamis, Tero Karras, and Gordon Wetzstein. Efficient geometry-aware 3D generative adversarial networks. In *arXiv*, 2021. [2, 4, 7](#)
- [7] Anpei Chen, Zexiang Xu, Fuqiang Zhao, Xiaoshuai Zhang, Fanbo Xiang, Jingyi Yu, and Hao Su. Mvsnerf: Fast generalizable radiance field reconstruction from multi-view stereo. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14124–14133, 2021. [1, 3](#)
- [8] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXII*, pages 333–350. Springer, 2022. [1, 2, 4, 7](#)
- [9] Anpei Chen, Zexiang Xu, Xinyue Wei, Siyu Tang, Hao Su, and Andreas Geiger. Dictionary fields: Learning a neural basis decomposition. *ACM Transactions on Graphics (TOG)*, 42(4):1–12, 2023. [2, 4, 7](#)
- [10] Anpei Chen, Zexiang Xu, Xinyue Wei, Siyu Tang, Hao Su, and Andreas Geiger. Factor fields: A unified framework for neural fields and beyond. *arXiv preprint arXiv:2302.01226*, 2023. [1](#)
- [11] 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. [5](#)
- [12] Zhiqin Chen, Thomas Funkhouser, Peter Hedman, and Andrea Tagliasacchi. Mobilenerf: Exploiting the polygon rasterization pipeline for efficient neural field rendering on mobile architectures. *arXiv preprint arXiv:2208.00277*, 2022. [1, 3](#)
- [13] Julian Chibane, Aayush Bansal, Verica Lazova, and Gerard Pons-Moll. Stereo radiance fields (srf): Learning view synthesis for sparse views of novel scenes, 2021. [3](#)
- [14] Kangle Deng, Andrew Liu, Jun-Yan Zhu, and Deva Ramanan. Depth-supervised nerf: Fewer views and faster training for free, 2022. [3](#)
- [15] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*, 2016. [5](#)
- [16] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5501–5510, 2022. [1, 2](#)
- [17] Sara Fridovich-Keil, Giacomo Meanti, Frederik Rahbæk Warburg, Benjamin Recht, and Angjoo Kanazawa. K-planes: Explicit radiance fields in space, time, and appearance. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12479–12488, 2023. [2, 4](#)
- [18] Adam Gaier and David Ha. Weight agnostic neural networks. *Advances in neural information processing systems*, 32, 2019. [3](#)
- [19] Quankai Gao, Qiangeng Xu, Hao Su, Ulrich Neumann, and Zexiang Xu. Strivec: Sparse tri-vector radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. [12](#)
- [20] Jean-Bastien Grill, Florian Strub, Florent Alché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent—a new approach to self-supervised learning. *Advances in neural information processing systems*, 33:21271–21284, 2020. [3](#)
- [21] Reinhard Heckel and Paul Hand. Deep decoder: Concise image representations from untrained non-convolutional networks. *arXiv preprint arXiv:1810.03982*, 2018. [2, 3, 4, 8](#)
- [22] Reinhard Heckel and Mahdi Soltanolkotabi. Denoising and regularization via exploiting the structural bias of convolutional generators. *arXiv preprint arXiv:1910.14634*, 2019. [2, 3](#)
- [23] Ajay Jain, Matthew Tancik, and Pieter Abbeel. Putting nerf on a diet: Semantically consistent few-shot view synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5885–5894, 2021. [1, 2, 3, 5, 6, 7](#)
- [24] Rasmus Jensen, Anders Dahl, George Vogiatzis, Engin Tola, and Henrik Aanæs. Large scale multi-view stereopsis evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 406–413, 2014. [5](#)
- [25] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics (ToG)*, 42(4):1–14, 2023. [3](#)
- [26] Mijeong Kim, Seonguk Seo, and Bohyung Han. Infonerf: Ray entropy minimization for few-shot neural volume rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12912–12921, 2022. [1, 2, 3, 5, 6, 7](#)
- [27] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [5](#)

- [28] Sixu Li, Chaojian Li, Wenbo Zhu, Boyang Yu, Yang Zhao, Cheng Wan, Haoran You, Huihong Shi, and Yingyan Lin. Instant-3d: Instant neural radiance field training towards on-device ar/vr 3d reconstruction. In *Proceedings of the 50th Annual International Symposium on Computer Architecture*, pages 1–13, 2023. 1
- [29] Zhaoshuo Li, Thomas Müller, Alex Evans, Russell H Taylor, Mathias Unberath, Ming-Yu Liu, and Chen-Hsuan Lin. Neuralangelo: High-fidelity neural surface reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8456–8465, 2023. 3
- [30] Isabella Liu, Linghao Chen, Ziyang Fu, Liwen Wu, Haian Jin, Zhong Li, Chin Ming Ryan Wong, Yi Xu, Ravi Ramamoorthi, Zexiang Xu, et al. Openillumination: A multi-illumination dataset for inverse rendering evaluation on real objects. *arXiv preprint arXiv:2309.07921*, 2023. 2, 5
- [31] Minghua Liu, Chao Xu, Haian Jin, Linghao Chen, Zexiang Xu, Hao Su, et al. One-2-3-45: Any single image to 3d mesh in 45 seconds without per-shape optimization. *arXiv preprint arXiv:2306.16928*, 2023. 1
- [32] Stuart Lloyd. Least squares quantization in pcm. *IEEE transactions on information theory*, 28(2):129–137, 1982. 5
- [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. 1
- [34] Xiaoxiao Long, Yuan-Chen Guo, Cheng Lin, Yuan Liu, Zhiyang Dou, Lingjie Liu, Yuexin Ma, Song-Hai Zhang, Marc Habermann, Christian Theobalt, et al. Wonder3d: Single image to 3d using cross-domain diffusion. *arXiv preprint arXiv:2310.15008*, 2023. 1
- [35] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 5
- [36] Michael Lustig, David Donoho, and John M Pauly. Sparse mri: The application of compressed sensing for rapid mr imaging. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 58(6):1182–1195, 2007. 3
- [37] Louis Mahon and Thomas Lukasiewicz. Selective pseudo-label clustering. In *KI 2021: Advances in Artificial Intelligence: 44th German Conference on AI, Virtual Event, September 27–October 1, 2021, Proceedings 44*, pages 158–178. Springer, 2021. 3
- [38] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. 1, 2, 5
- [39] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *arXiv preprint arXiv:2201.05989*, 2022. 1, 2, 4
- [40] Jacob Munkberg, Jon Hasselgren, Tianchang Shen, Jun Gao, Wenzheng Chen, Alex Evans, Thomas Müller, and Sanja Fidler. Extracting triangular 3d models, materials, and lighting from images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8280–8290, 2022. 3
- [41] Michael Niemeyer, Jonathan T Barron, Ben Mildenhall, Mehdi SM Sajjadi, Andreas Geiger, and Noha Radwan. Regnerf: Regularizing neural radiance fields for view synthesis from sparse inputs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5480–5490, 2022. 1, 2, 3, 5, 6, 7
- [42] Michael Oechsle, Songyou Peng, and Andreas Geiger. Unisurf: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. In *International Conference on Computer Vision (ICCV)*, 2021. 3
- [43] Gregory Ongie, Ajil Jalal, Christopher A Metzler, Richard G Baraniuk, Alexandros G Dimakis, and Rebecca Willett. Deep learning techniques for inverse problems in imaging. *IEEE Journal on Selected Areas in Information Theory*, 1(1):39–56, 2020. 3
- [44] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952*, 2023. 4, 8
- [45] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 3, 5
- [46] Anton Razzhigaev, Arseniy Shakhmatov, Anastasia Maltseva, Vladimir Arkhipkin, Igor Pavlov, Ilya Ryabov, Angelina Kuts, Alexander Panchenko, Andrey Kuznetsov, and Denis Dimitrov. Kandinsky: an improved text-to-image synthesis with image prior and latent diffusion. *arXiv preprint arXiv:2310.03502*, 2023. 4, 8
- [47] Konstantinos Rematas, Ricardo Martin-Brualla, and Vittorio Ferrari. Sharf: Shape-conditioned radiance fields from a single view, 2021. 3
- [48] Barbara Roessle, Jonathan T. Barron, Ben Mildenhall, Pratul P. Srinivasan, and Matthias Nießner. Dense depth priors for neural radiance fields from sparse input views, 2022. 3
- [49] Radu Alexandru Rosu and Sven Behnke. Permutosdf: Fast multi-view reconstruction with implicit surfaces using permutohedral lattices. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8466–8475, 2023. 3
- [50] Aditya Sanghi and Pradeep Kumar Jayaraman. How powerful are randomly initialized pointcloud set functions? *arXiv preprint arXiv:2003.05410*, 2020. 3
- [51] Seunghyeon Seo, Yeonjin Chang, and Nojun Kwak. Flipnerf: Flipped reflection rays for few-shot novel view synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22883–22893, 2023. 2, 3, 5, 6
- [52] Seunghyeon Seo, Donghoon Han, Yeonjin Chang, and Nojun Kwak. Mixnerf: Modeling a ray with mixture density for novel view synthesis from sparse inputs, 2023.

- [53] Min seop Kwak, Jiuhn Song, and Seungryong Kim. Geconerf: Few-shot neural radiance fields via geometric consistency, 2023. 3
- [54] Fahad Shamshad, Salman Khan, Syed Waqas Zamir, Muhammad Haris Khan, Munawar Hayat, Fahad Shahbaz Khan, and Huazhu Fu. Transformers in medical imaging: A survey. *Medical Image Analysis*, page 102802, 2023. 3
- [55] Liyue Shen, John Pauly, and Lei Xing. Nerp: implicit neural representation learning with prior embedding for sparsely sampled image reconstruction. *IEEE Transactions on Neural Networks and Learning Systems*, 2022. 3
- [56] 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. 1, 8
- [57] Nagabhushan Somraj and Rajiv Soundararajan. Vip-nerf: Visibility prior for sparse input neural radiance fields. In *ACM SIGGRAPH 2023 Conference Proceedings*, pages 1–11, 2023. 3
- [58] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In *CVPR*, 2022. 1, 2
- [59] Jiaxiang Tang, Hang Zhou, Xiaokang Chen, Tianshu Hu, Er-rui Ding, Jingdong Wang, and Gang Zeng. Delicate textured mesh recovery from nerf via adaptive surface refinement. *arXiv preprint arXiv:2303.02091*, 2022. 1, 3
- [60] Alex Trevithick and Bo Yang. Grf: Learning a general radiance field for 3d representation and rendering, 2021. 3
- [61] Prune Truong, Marie-Julie Rakotosaona, Fabian Manhardt, and Federico Tombari. Sparf: Neural radiance fields from sparse and noisy poses. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4190–4200, 2023. 1, 3
- [62] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Deep image prior. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9446–9454, 2018. 2, 3
- [63] Dave Van Veen, Ajil Jalal, Mahdi Soltanolkotabi, Eric Price, Sriram Vishwanath, and Alexandros G Dimakis. Compressed sensing with deep image prior and learned regularization. *arXiv preprint arXiv:1806.06438*, 2018. 3
- [64] Dor Verbin, Peter Hedman, Ben Mildenhall, Todd Zickler, Jonathan T Barron, and Pratul P Srinivasan. Ref-nerf: Structured view-dependent appearance for neural radiance fields. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5481–5490. IEEE, 2022. 1, 2
- [65] Guangcong Wang, Zhaoxi Chen, Chen Change Loy, and Ziwei Liu. Sparsenerf: Distilling depth ranking for few-shot novel view synthesis. *arXiv preprint arXiv:2303.16196*, 2023. 1, 2, 3
- [66] 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. *NeurIPS*, 2021. 3
- [67] Qianqian Wang, Zhicheng Wang, Kyle Genova, Pratul Srinivasan, Howard Zhou, Jonathan T. Barron, Ricardo Martin-Brualla, Noah Snavely, and Thomas Funkhouser. Ibrnet: Learning multi-view image-based rendering. In *CVPR*, 2021. 1
- [68] Yiming Wang, Qin Han, Marc Habermann, Kostas Daniilidis, Christian Theobalt, and Lingjie Liu. Neus2: Fast learning of neural implicit surfaces for multi-view reconstruction. *arXiv preprint arXiv:2212.05231*, 2022. 3
- [69] Xinyue Wei, Fanbo Xiang, Sai Bi, Anpei Chen, Kalyan Sunkavalli, Zexiang Xu, and Hao Su. Neumanifold: Neural watertight manifold reconstruction with efficient and high-quality rendering support. *arXiv preprint arXiv:2305.17134*, 2023. 1, 3
- [70] Guanjun Wu, Taoran Yi, Jiemin Fang, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu, Qi Tian, and Xinggang Wang. 4d gaussian splatting for real-time dynamic scene rendering. *arXiv preprint arXiv:2310.08528*, 2023. 3
- [71] Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Jianmin Bao, Zhuliang Yao, Qi Dai, and Han Hu. Simmim: A simple framework for masked image modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9653–9663, 2022. 4, 8
- [72] Qiangeng Xu, Zexiang Xu, Julien Philip, Sai Bi, Zhixin Shu, Kalyan Sunkavalli, and Ulrich Neumann. Point-nerf: Point-based neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5438–5448, 2022. 2
- [73] Jiawei Yang, Marco Pavone, and Yue Wang. Freenerf: Improving few-shot neural rendering with free frequency regularization. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2, 3, 5, 6, 7
- [74] Ziyi Yang, Xinyu Gao, Wen Zhou, Shaohui Jiao, Yuqing Zhang, and Xiaogang Jin. Deformable 3d gaussians for high-fidelity monocular dynamic scene reconstruction. *arXiv preprint arXiv:2309.13101*, 2023. 3
- [75] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. *Advances in Neural Information Processing Systems*, 34:4805–4815, 2021. 3
- [76] Lior Yariv, Peter Hedman, Christian Reiser, Dor Verbin, Pratul P. Srinivasan, Richard Szeliski, Jonathan T. Barron, and Ben Mildenhall. Baked sdf: Meshing neural sdf for real-time view synthesis. *arXiv*, 2023. 3
- [77] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from one or few images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4578–4587, 2021. 1, 3
- [78] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. Nerf++: Analyzing and improving neural radiance fields. *arXiv preprint arXiv:2010.07492*, 2020. 1
- [79] 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. 5