

Multiplane Prior Guided Few-Shot Aerial Scene Rendering

Zihan Gao Licheng Jiao* Lingling Li Xu Liu Fang Liu
 Puhua Chen Yuwei Guo
 School of Artificial Intelligence, Xidian University

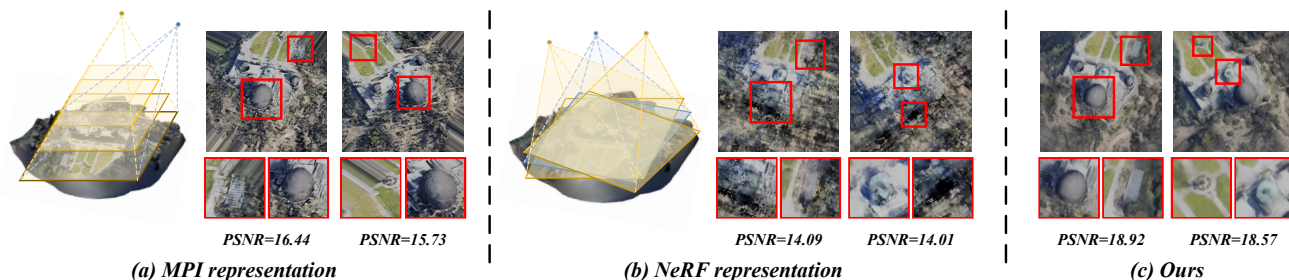


Figure 1. **Visualization of failure modes in NeRF and MPI.** (a) MPI models scenes only in each single camera frustums and performs homography warping to render novel views. Insufficient sampling leads to incorrect depth and thus results in an overlapping ghosting effect. Large camera movement leads to cropped corners. However, high-frequency details seem to be successfully preserved. (b) NeRF models scenes in a continuous volumetric manner. If only sparse views with large camera movement are provided, some parts of the scene may be sampled very little or even never. Insufficient sampling leads to collapsed details and unexpected floaters. (c) Our approach combines the capabilities of NeRF with the perspective-friendly nature of MPI in aerial scenes to achieve photorealistic novel view renderings.

Abstract

Neural Radiance Fields (NeRF) have been successfully applied in various aerial scenes, yet they face challenges with sparse views due to limited supervision. The acquisition of dense aerial views is often prohibitive, as unmanned aerial vehicles (UAVs) may encounter constraints in perspective range and energy constraints. In this work, we introduce Multiplane Prior guided NeRF (MPNeRF), a novel approach tailored for few-shot aerial scene rendering—marking a pioneering effort in this domain. Our key insight is that the intrinsic geometric regularities specific to aerial imagery could be leveraged to enhance NeRF in sparse aerial scenes. By investigating NeRF’s and Multiplane Image (MPI)’s behavior, we propose to guide the training process of NeRF with a Multiplane Prior. The proposed Multiplane Prior draws upon MPI’s benefits and incorporates advanced image comprehension through a SwinV2 Transformer, pre-trained via

*Corresponding author. This work was supported in part by the Key Scientific Technological Innovation Research Project by Ministry of Education, the Joint Funds of the National Natural Science Foundation of China (U22B2054), the National Natural Science Foundation of China (62076192, 61902298, 61573267, 61906150 and 62276199), the 111 Project, the Program for Cheung Kong Scholars and Innovative Research Team in University (IRT 15R53), the ST Innovation Project from the Chinese Ministry of Education, the Key Research and Development Program in Shaanxi Province of China(2019ZDLGY03-06), the China Postdoctoral fund(2022T150506).

SimMIM. Our extensive experiments demonstrate that MPNeRF outperforms existing state-of-the-art methods applied in non-aerial contexts, by tripling the performance in SSIM and LPIPS even with three views available. We hope our work offers insights into the development of NeRF-based applications in aerial scenes with limited data.

1. Introduction

Neural Radiance Field (NeRF) [33] has succeeded in rendering high-fidelity novel views and many 3D applications by modeling 3D scenes as a continuous implicit function. In contrast to indoor or synthetic scenes capturing simple objects using cell phones, aerial images provide a unique bird’s-eye view and overview of a landscape. Based on NeRF, many applications in aerial scenes have been developed, such as navigation, urban planning, data augmentation, autonomous vehicles, and environmental mapping [1, 7, 21, 29, 32, 40, 59].

In many real-world scenarios, unmanned aerial vehicles (UAVs) encounter constraints such as limited perspectives, energy limitations, or adverse weather conditions, which restrict their ability to acquire dense observational data. While NeRF forms a foundational technology for many aerial applications, it is prone to overfitting on sparse training views [13, 17, 34]. This limitation of NeRF becomes particularly salient in the context of aerial imagery. Alleviating this prob-

lem could save resources and may benefit numerous applications.

As the first to explore few-shot NeRF for aerial imagery, we stand at the forefront of this emerging field. The landscape of few-shot neural rendering to date has been predominantly shaped by its application to indoor and synthetic scenes. Transfer learning based methods [3, 46, 47, 56] aim to pre-train the model on a large number of scenes. Yet, these approaches necessitate extensive datasets for pre-training. This is not only resource-heavy but also impractical to fulfill for varied aerial scenarios. Another line of works [10, 13, 15, 17, 34, 38, 54, 55] seeks to impose regularization on NeRF by exploring the universal attributes of 3D scenes like local continuity and semantic consistency. Yet, in situations where available data is significantly sparse relative to the complexity of the scene, these methods might struggle to maintain their effectiveness. And the last is depth-prior-based methods [8, 45] gain additional supervision from the scene’s depth. These methods can be problematic in aerial images due to the frequent occurrence of ambiguous depth cues and the high cost of obtaining accurate depth maps. Despite their efficacy in controlled environments, these methods fall short of addressing the unique complexities of aerial scenes, leaving a gap that our work aims to fill. We therefore ask: **Can we harness the intrinsic geometric regularities specific to aerial imagery to broaden the capabilities of NeRF under sparse data conditions, thereby easing the data collection constraints?**

To answer this question, we turn to earlier works on 2.5D representations such as Multiplane Image (MPI) [23, 43, 50, 58]. MPI typically operates by extracting multiple RGB and density planes from a single image input by an encoder-decoder style MPI generator to represent scene geometry within the camera’s frustum. Although NeRF provides a continuous representation of a scene, MPI offers discrete, frustum-confined layers that can be particularly advantageous in the context of aerial imagery. This is due to UAVs frequently capturing images from overhead perspectives that align well with MPI’s planar representation. Additionally, the encoder-decoder architecture of the MPI generator can exploit the inductive biases inherent in advanced convolutional and self-attention-based image processing compared to the simple multi-layer perceptron (MLP) of NeRF, thus enhancing the rendering of local and global scene details. However, while MPIs present certain benefits in terms of their adaptability to aerial perspectives, their partial scene recovery and limitation to individual frustums pose challenges in creating a comprehensive 3D understanding.

In this work, we present Multiplane Prior guided NeRF (MPNeRF), a novel method for enhancing NeRF models in few-shot aerial scene rendering. We guide NeRF’s learning process by using a multiplane prior—a concept drawn from the strengths of MPI and refined with cutting-edge im-

age understanding from a Swin TransformerV2 pre-trained with SimMIM [28]. This approach unites the capabilities of NeRF with the perspective-friendly nature of MPI, tailored for the unique vantage points of aerial scenes. Concretely, our approach updates the NeRF branch using pseudo labels generated from the MPI branch. As training proceeds, NeRF can effectively pick up finer details from the MPI branch and the advantage of the MPI branch is implicitly distilled into NeRF. This strategy implicitly folds a multiplane prior to NeRF, boosting its performance in handling sparse aerial imagery data. Our contributions can be summarized as follows:

1. We introduce Multiplane Prior guided NeRF (MPNeRF), a novel framework that synergistically combines NeRF and MPIs for enhanced few-shot neural rendering in aerial scenes. To the best of our knowledge, this is the first method specially designed for this task.
2. Through an investigation, we pinpoint and analyze the typical failure modes of NeRF and MPI in aerial scenes. We devise a simple yet effective learning strategy that guides the training process of NeRF by learning a multiplane prior, effectively circumventing NeRF’s typical pitfalls in sparse aerial scenes.
3. We compare MPNeRF against a suite of state-of-the-art non-aerial scene methods, rigorously testing its adaptability and performance in aerial scenarios. Our experiments demonstrate MPNeRF’s superior performance, showcasing its significant leap over methods previously confined to non-aerial contexts.

2. Related Work

2.1. Scene Representations for View Synthesis.

Earlier works on light fields [6, 11, 22] achieve view synthesis by interpolating nearby views given a dense set of input images. Later works utilize explicit mesh [4, 24, 26, 44], or volumetric [12, 16, 19, 37, 39, 49] representation to represent the scene. More recently, layered representations have gained attention due to their efficiency in modeling occluded content. One such layered representation is the MPI [23, 43, 50, 58]. An MPI consists of multiple planes of RGB and α values at fixed depths. Given an input image, an encoder-decoder network typically generates the MPI within the camera frustum. This MPI is then homography warped to the target camera position and integrated over the planes to produce novel views. It’s important to note that the generated MPI only models the geometry within each camera frustum at given depths, and the complete 3D scene is not fully recovered.

Recently, NeRF [33] has shown significant potential in novel view synthesis. NeRF works by modeling the scene with a continuous function of 3D coordinates and viewing directions to output the corresponding RGB and volume density values. Following NeRF, many methods have been proposed. mip-NeRF [2] introduces a more robust representa-

tion that uses a cone tracing technique and samples the cone with multivariate Gaussian. NeRF-W [31] and Ha-NeRF[5] have extended the applicability of NeRF to in-the-wild photo collections through object decomposition and hallucination techniques. For large-scale scenes, BungeeNeRF [51] proposes a multiscale progressive learning method to reconstruct cities from satellite imagery, while Block-NeRF [41] leverages individual NeRFs for each component of the scene to achieve large-scale scene rendering. ShadowNeRF [9] and Sat-NeRF [30] address the issue of strong noncorrelation between satellite images taken at different times by modeling solar light and transient objects. However, these NeRF-based methods are still limited by the need for densely sampled views of the scene.

2.2. NeRF with Sparse Input

Many approaches have been developed to train a NeRF from sparse input in different directions. One straightforward direction [3, 27, 42, 46, 47, 56] is to learn the general appearance of a scene or object from a large number of data. PixelNeRF [56] adopts a CNN feature extractor to condition each input coordinate with image features. MVSNerF [3] uses 3D CNN to process cost volume acquired by image warping. These methods often require a large number of multi-view images to be pre-trained on, which is sometimes hard to acquire in aerial imagery. Some other techniques [10, 17, 20, 34, 38] find it is more data-efficient to regularize NeRF with common properties of the 3D geometry. InfoNeRF [17] regularizes NeRF by putting a sparsity constraint on the density of each ray. RegNeRF [34] regularizes NeRF by local smoothness. Other works [8, 13, 45, 48] aim to take advantage of supervision from other sources, such as depth or appearance. DS-NeRF [8] supervises the geometry with sparse point cloud generated with structure from motion. DietNeRF [13] regularize NeRF by ensuring perceptual consistency within different views. ManifoldNeRF [15] builds upon DietNeRF and takes into account viewpoint-dependent perceptual consistency to refine supervision in unknown viewpoints. However, we noticed none of these methods is designed for aerial scenes and thus left a gap our work aims to fill.

3. Method

Our objective is to train a standard NeRF model to create highly realistic novel views of an aerial scene from a limited number of captured perspectives. To address the challenges of training NeRF with sparse aerial views, we introduce a novel training approach that leverages a Multiplane Prior. The proposed Multiplane Prior harnesses the strengths of MPI and is enriched by advanced image understanding capabilities derived from a SwinV2 Transformer pre-trained using SimMIM[52]. An overview of our approach is presented in Figure 2.

In Sec. 3.1, we briefly review the background related to our method. Sec. 3.2 provides an investigation into NeRF and MPI’s behavior in sparse aerial scenes. Sec. 3.3 describe our overall framework.

3.1. Preliminaries

Neural Radiance Field. Given a 3D coordinate $\mathbf{x} = (x, y, z)$ and a 2D viewing direction $\mathbf{d} = (\theta, \varphi)$, NeRF aims to model the scene by solving a continuous function $f(\mathbf{x}, \mathbf{d}) = (\mathbf{c}, \sigma)$ using multi-layer perceptron (MLP) network, where \mathbf{c} and σ represent the emissive color and volume density at the given coordinate. NeRF cast rays $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$ from the camera origin \mathbf{o} along the direction \mathbf{d} to pass through a pixel. NeRF then samples M points along this ray and computes its color by volume rendering:

$$\hat{\mathbf{C}}(\mathbf{r}) = \sum_{i=1}^M T_i (1 - \exp(-\sigma_i \delta_i)) \mathbf{c}_i, \quad (1)$$

$$T_i = \exp\left(-\sum_{j=1}^{i-1} \sigma_j \delta_j\right),$$

where \mathbf{c}_i and σ_i are the color and volume density of i -th sample along the ray and δ_i is the distance between adjacent samples. $\hat{\mathbf{C}}(\mathbf{r})$ denotes the final color of that pixel rendered by NeRF. In NeRF, a dense 3D scene is recovered implicitly in the form of neural network weights.

Multiplane Image. Multi-plane Image (MPI) represents the scene by dividing the 3D space into a collection of planes with RGB and density values in one camera frustum. In training, each batch consists of a pair of images $\mathbf{I}_s, \mathbf{I}_t \in \mathbb{R}^{H \times W \times 3}$ with corresponding camera intrinsic $\mathbf{K}_s, \mathbf{K}_t \in \mathbb{R}^{3 \times 3}$ and relative pose $\mathbf{P}_{s2t} = [\mathbf{R}_{s2t} \in \mathbb{R}^{3 \times 3}, \mathbf{t}_{s2t} \in \mathbb{R}^3]$ denoted as $\{(\mathbf{I}_s, \mathbf{K}_s), (\mathbf{I}_t, \mathbf{K}_t), \mathbf{P}\}$, subscripts s and t represent source and target viewpoint respectively. Depth for each plane is sampled $\{z = z_k | k = 1, 2, 3, \dots, N\}$ uniformly according to the scene bounds. An encoder-decoder based MPI-generator denoted as G_{MPI} is adopted to generate multiple planes of RGB and density at discrete depth as:

$$\{(\mathbf{c}_k, \sigma_k) | k = 1, 2, 3, \dots, N\} = G_{\text{MPI}}(\mathbf{I}_s). \quad (2)$$

Here the subscript k denotes the k -th plane. The rendering of MPI is performed in two steps: First, establish the correspondence between the pixel coordinates in the source and target plane through homography warping as:

$$[u_t, v_t, 1]^T = \mathbf{K}_t \left(\mathbf{R}_{s2t} - \frac{\mathbf{t}_{s2t} \mathbf{n}^T}{z_k} \right) \mathbf{K}_s^{-1} [u_s, v_s, 1]^T. \quad (3)$$

Second, similar to NeRF, apply differentiable rendering to

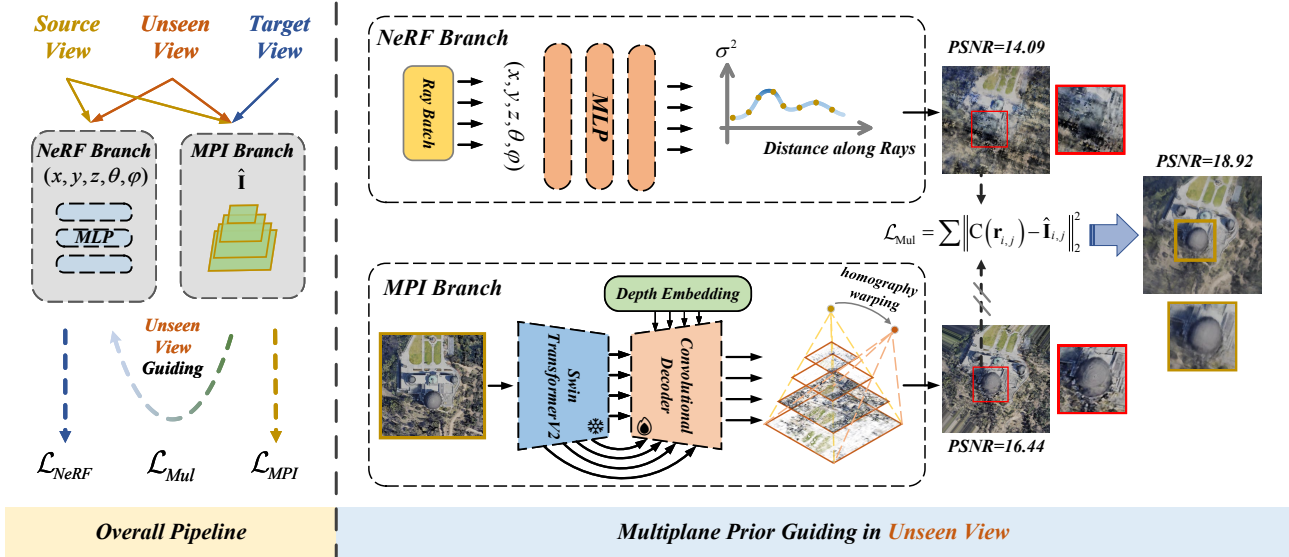


Figure 2. **Overall pipeline for training Multiplane Prior guided NeRF (MPNeRF).** Our novel MPNeRF architecture integrates a standard NeRF branch with an MPI branch, informed by a pre-trained SwinV2 Transformer. This design introduces a multiplane prior to guide the NeRF training, addressing the common challenges of rendering with sparse aerial data. The process begins by sampling three distinct views: a source and target view for training with known ground truth, and an unseen view from a novel viewpoint. The NeRF model is then refined using pseudo labels produced by the MPI branch, which are especially crucial for synthesizing views from previously unseen angles, as shown in the pipeline.

get target view 2D images \hat{I}_t :

$$\hat{I}_t = \sum_{k=1}^N T_k (1 - \exp(-\sigma_k \delta_k)) c_k, \quad (4)$$

$$T_i = \exp\left(-\sum_{l=1}^{k-1} \sigma_l \delta_l\right).$$

3.2. A Closer Look at The Behavior of NeRF & MPI

To better understand the behavior of NeRF and MPI, we conducted an investigation into their failure modes. In Figure 1, we visualize the rendering results of NeRF and MPI when encountering large camera movements. Our findings reveal that NeRF often produces blurry renderings, while MPI tends to exhibit overlapping ghosting effects and cropped corners.

Recall that NeRF represents the whole 3D scene continuously by encoding the volume density and color into an 8-layer MLP’s weights. In other words, NeRF utilizes a learning-based approach by forcing correct rendering from every angle of the scene with multi-view consistency. Such a model is highly compact when supervised with sufficient training views. When the supervised angle is limited, areas covered less (as in the non-overlapped camera frustum in Figure 1 (b)) are uncontrolled and may exhibit high-density values [25, 34, 55], leading to blurry or even collapsed results. Considering the structured nature of 3D aerial scenes, we recognize two key factors of aerial scenes: *aligned perspectives with predominant planarity*, and *consistent geomet-*

ric appearance. First, the typical flight paths of UAVs over these scenes predominantly capture landscapes aligned with the XY planes, providing a unique geometrical consistency. Second, objects in aerial scenes contain common visual characteristics, offering additional cues for scene interpretation and analysis.

In contrast to NeRF, MPI models the scene within each camera frustum and decomposes it into an explicit set of discrete 2D planes at fixed depths. This mirrors the overhead views and planar surfaces commonly found in aerial scenes. Also, the convolution-based or self-attention-based MPI-generator is inherently suited to carry prior knowledge of the scene. However, with insufficient supervision provided, the MPI for different camera frustums may not be properly calibrated. As a consequence, we observe the occurrence of overlapping ghosting effects in rendered unseen views. Additionally, when there is substantial camera movement, the corners of the target views may be excluded from the source views, resulting in invalid renderings. However, MPI is successful in preserving high-frequency details in the rendered image. We attribute this capability to the power of CNNs and the implicit encoding of prior knowledge in the MPI generator.

Due to the distinct failure mode of MPI and its favorable properties, we explore a strategy to enhance NeRF in a sparse aerial context.

3.3. Guiding NeRF with a Multiplane Prior

Based on the investigation of the different properties shown in NeRF and MPI when encountering sparse input. We turn to the task of few-shot aerial scene rendering and propose a simple yet effective strategy that treats the MPI as a bridge to convey information that is hard to learn by the traditional NeRF pipeline.

We formulate the proposed MPNeRF with a NeRF branch, and an MPI branch denoted as \mathbf{G}_{θ_1} and \mathbf{G}_{θ_2} . Given a batch contains images from source and target viewpoints alongside the corresponding camera parameters. To train the NeRF branch, we cast rays for the source viewpoints pixels using the camera parameters following [33]. The MSE loss is adopted to supervise the NeRF branch with the ground truth color $C(\mathbf{r})$:

$$\mathcal{L}_{\text{NeRF}} = \sum_{\mathbf{r} \in \mathbf{B}} \left\| \hat{C}(\mathbf{r}) - C(\mathbf{r}) \right\|_2^2, \quad (5)$$

where, B is the set of input rays during training. For the MPI branch, the encoder-decoder style MPI generator takes in images from the source view and outputs the corresponding MPI representation. In order to incorporate prior knowledge, we adopt a frozen Swin Transformer V2 model pre-trained with SimMIM [28] as a feature extractor to extract multi-scale features from aerial images. These features are fused to generate the final MPI representation. The loss function to optimize the MPI branch contains three components: L1 loss to match the synthesized target image $\hat{\mathbf{I}}_t$ to ground truth \mathbf{I}_t at a pixel level, SSIM loss to encourage structure consistency, and LPIPS loss [57] for perceptual consistency.

$$\begin{aligned} \mathcal{L}_{L1} &= \left\| \hat{\mathbf{I}}_t - \mathbf{I}_t \right\|_1, \\ \mathcal{L}_{\text{SSIM}} &= 1 - \text{SSIM}(\hat{\mathbf{I}}_t, \mathbf{I}_t), \\ \mathcal{L}_{\text{LPIPS}} &= \left\| \phi(\hat{\mathbf{I}}_t) - \phi(\mathbf{I}_t) \right\|_1. \end{aligned} \quad (6)$$

The overall loss function to optimize the MPI branch would then be a sum of these three losses:

$$\mathcal{L}_{\text{MPI}} = \mathcal{L}_{L1} + \mathcal{L}_{\text{SSIM}} + \mathcal{L}_{\text{LPIPS}}. \quad (7)$$

These conventional loss functions train both branches to give predictions based on training views. Based on the investigation in Sec. 3.2, we aim to guide the training process of the NeRF with a multiplane prior learned by the MPI branch. An intuitive choice is sampling a random number of pixels from an unseen view and matching the predicted color of two branches with an MSE loss.

$$\mathcal{L}_{\text{Mul}} = \sum \left\| \hat{C}(\mathbf{r}_{i,j}) - \hat{\mathbf{I}}_{i,j} \right\|_2^2. \quad (8)$$

i, j denotes the sampled pixel coordinates in the unseen viewpoint. In practice, this intuitive choice works surprisingly

well and the advantage of MPI is implicitly learned by NeRF. We give a further analysis of our design choice of \mathcal{L}_{Mul} in Sec. 4.5.

To summarize, the final objective functions of the NeRF branch \mathbf{G}_{θ_1} and MPI branch \mathbf{G}_{θ_2} are given as follows:

$$\begin{aligned} \mathcal{L}_{\mathbf{G}_{\theta_1}} &= \mathcal{L}_{\text{NeRF}} + \lambda \mathcal{L}_{\text{Mul}}, \\ \mathcal{L}_{\mathbf{G}_{\theta_2}} &= \mathcal{L}_{\text{MPI}}. \end{aligned} \quad (9)$$

In this way, the training experience of the MPI branch serves as a multiplane prior that guides the training process of NeRF. Even if the learned MPI is not entirely accurate, the NeRF branch benefits from this multiplane prior and thus avoids collapse during training.

4. Experiment

4.1. Implementation Details

Our method is implemented using PyTorch, and all experiments are conducted on a GeForce RTX 3090 GPU. For the NeRF branch, we use the original NeRF in [33]. During training, we randomly sample unseen views following the strategy proposed by [13]. The batch size is set to 1024 pixel rays in both source and unseen views. For each ray, we perform 64 coarse sampling and 32 fine sampling along the ray. For the MPI branch, we sample 16 layers of planes for each viewpoint. The optimization of the two branches of MPNeRF is performed using the Adam optimizer [18] with a learning rate of 5×10^{-4} . The hyperparameter λ in Eq. 9 is set to 1.

4.2. Datasets and Evaluation Metrics

The main experiments are conducted on 16 scenes collected by LEVIR-NVS [50]. These scenes contain various scenarios in common aerial imagery, including mountains, buildings, colleges, etc. 3 and 5 views are used for training and the rest for testing. Additional experiments and discussions can be found in the Appendix. In line with previous studies of few-shot neural rendering [13, 17, 34, 55], we report PSNR, SSIM and LPIPS [57].

4.3. Baseline Methods

We compare MPNeRF against various state-of-the-art methods including NeRF [33], Mip-NeRF [2], InfoNeRF [17], DietNeRF [13], PixelNeRF [56], RegNeRF [34] and FreeNeRF [55]. Among these methods, NeRF and Mip-NeRF are designed for dense view training, we mainly explore the performance gain achieved by MPNeRF. PixelNeRF aims to learn a generalized NeRF representation for all scenes and is pre-trained on the DTU dataset [14]. Since a large domain gap might exist when applied in aerial imagery, we report PixelNeRF’s results with and without additional fine-tuning per scene. Other methods are designed for few-shot

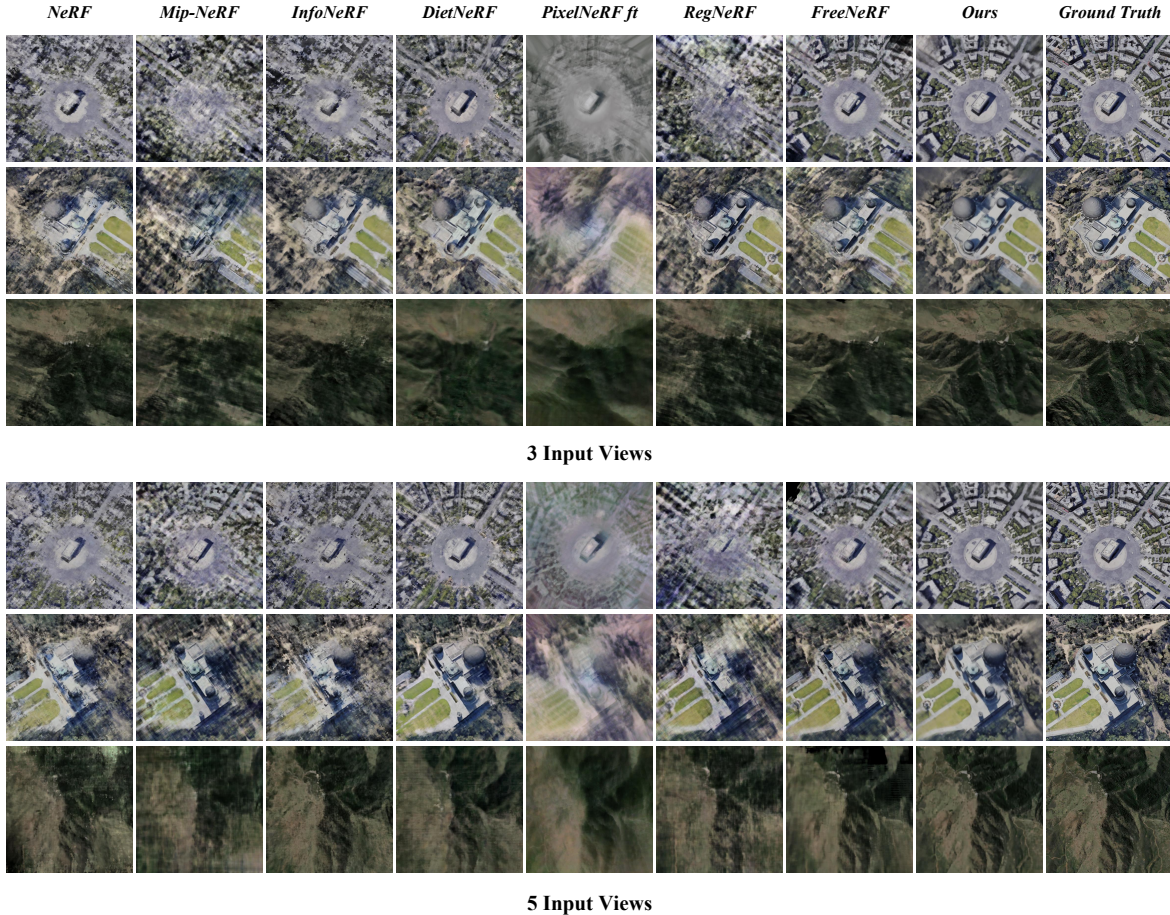


Figure 3. Visual comparisons on 3 selected scenes with 3 and 5 views. MPNeRF achieves photo-realistic quality in different scenes compared with ground-truth images on novel views.

neural rendering, and we conduct comparative experiments to investigate their performance when encountering aerial imagery.

4.4. Comparative Results Analysis

Table 1 and Table 2 report the performance of MPNeRF and baseline methods in the 3-view and 5-view settings. Additionally, a qualitative comparison can be observed in Fig. 3. A very significant improvement can be found in all three metrics and rendering fidelity. The results demonstrate that PixelNeRF tends to produce blurry renderings, which we attribute to the poor localization of the CNN features. InfoNeRF and RegNeRF use local smoothness and sparsity to regularize NeRF explicitly. However, in scenarios with substantially limited information compared to the scene complexity, the performance of these methods could be compromised. DietNeRF implicitly distills the prior knowledge encoded in CLIP [36] and achieves better results. FreeNeRF investigates the frequency in NeRF training. By progressively learning each frequency component, FreeNeRF

has demonstrated remarkable effectiveness. Nonetheless, the progressive frequency regularization leads to relatively flat results, favoring PSNR but not metrics that consider local structures such as SSIM and LPIPS.

In fact, NeRF’s representation makes recovering 3D scenes from sparse inputs ill-posed. MPNeRF acquires superior results by the guiding of a multiplane prior to gaining a stronger understanding of local structures and semantics. In the more challenging scenes, such as Building in Figure 3, MPNeRF successfully avoids collapse during training.

4.5. Ablation Studies and Further Analyses

Ablation Analysis. We ablate the proposed multiplane prior to our method, and the results are shown in Table 3. Intuitively, it seems better to use MPI as a guide after fully training it, we first construct experiments where a two-stage training strategy is employed. We then assess the impact of SwinV2’s pre-trained weights on performance by removing them. Next, we evaluate the contribution of multi-scale features by disconnecting the skip connections in the MPI

generator. Finally, we integrate the multiplane directly within NeRF’s sampling space, omitting the separate MPI branch, to examine the inductive biases’ influence on performance. Employing MPI concurrent branch during training leads to slight improvements. We believe this is because the MPI’s training experience itself carries information. The exclusion of SwinV2’s pre-trained weights declines performance, affirming the value of the encoded prior knowledge. Similarly, omitting the multi-scale feature connection diminishes the fidelity of the rendered images. Most significantly, the absence of the MPI generator results in a marked decrease in all metrics. This suggests that uncalibrated MPI generated by self-attention and convolution is important to avoid degenerate solutions. Collectively, these findings demonstrate that each element of the proposed multiplane prior is crucial for the superior performance of MPNeRF.

Methods	PSNR	SSIM	LPIPS
Ours	21.72	0.80	0.19
training MPI beforehand	21.49	0.76	0.20
w/o pre-trained weights	20.46	0.75	0.24
w/o multi-scale feature	20.05	0.71	0.25
w/o MPI generator	16.38	0.42	0.52
Baseline NeRF	15.13	0.20	0.58

Table 3. Ablation analysis on the proposed Multiplane Prior.

Further Analyses on the design choice of the \mathcal{L}_{Mul} . One intuitive thought of designing \mathcal{L}_{Mul} is that geometry recovered by the MPI branch may provide more information than color alone. So we design two experiments, one is to match the expected depth of both branches as an auxiliary depth loss, and another is to model density on each ray as a distribution [17] and minimize the KL divergence. Another intuitive thought is that the choice of \mathcal{L}_{Mul} should reflect the local or nonlocal relationships within the pixels. Therefore, we adopt the recently proposed S3IM loss [53] to measure this relationship.

Design Choice	PSNR	SSIM	LPIPS
w/t depth matching	21.34	0.75	0.20
w/t ray matching	21.11	0.71	0.21
w/t relation matching	21.19	0.70	0.21
Ours	21.72	0.80	0.19

Table 4. Design choice of the \mathcal{L}_{Mul}

However, as shown in Table. 4, the result suggests that these intuitive designs worsen the results. The first two design involves direct supervision of the depth generated by the MPI generator. The last involves capturing the non-local relationships between the predictions of the NeRF and the MPI branch. Since the learned MPI is not entirely accurate, we believe the noise within pseudo-labels may compromise performance with these enhanced supervisions applied.

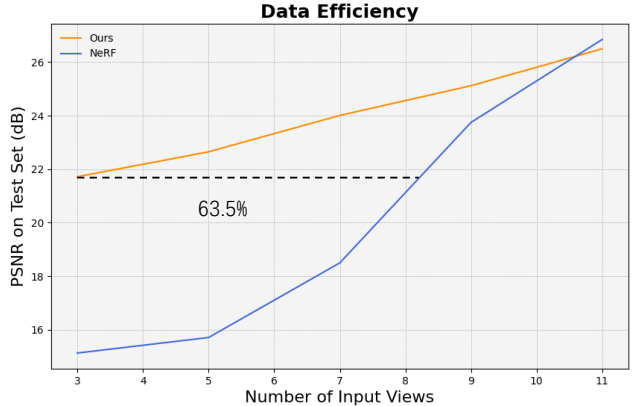


Figure 4. We investigate the data efficiency achieved by our method. Our method requires up to 63.5% training images to achieve a similar performance compared to a vanilla NeRF model.

Impact of Different Pre-trained Models. We perform a comparison study on three pre-trained vision transformers, i.e., CLIP [36], DINOv2 [35], and SimMIM [52]. We adopt the base model in our experiment. As shown in Table 5, all of these methods provide comparable results. The results show that the Swin Transformer pre-trained via SimMIM [52] outperforms others. We believe this can be attributed to the rich global and local details learned by SimMIM and the hierarchical structure of the Swin Transformer.

Pre-trained Image Encoder	PSNR	SSIM	LPIPS
CLIP [36]	20.15	0.70	0.25
DINOv2 [35]	20.01	0.68	0.25
SimMIM (Ours) [52]	21.72	0.80	0.19

Table 5. Impact of different pre-trained models.

Data Efficiency Since we aim to improve the capability of NeRF in aerial scenes when only sparse views are available, we investigate how much data MPNeRF can save to achieve similar rendering results compared to the original NeRF that requires dense view supervision. As shown in Fig. 4, the results show that our method requires up to 63.5% training images. This may help save energy and establish resource-efficient applications for UAVs based on NeRF.

5. Limitations and Conclusion

In this work, we introduce Multiplane Prior guided NeRF (MPNeRF), the first approach designed for few-shot aerial scene rendering. Through the guiding of the multiplane prior, MPNeRF effectively overcomes the typical pitfalls in sparse aerial scenes. We hope our work can provide insight into future NeRF-based applications in aerial scenes. However, further exploration of the guiding strategy design is needed. In particular, incorporating uncertainty prediction mechanisms or implementing grid-based representations holds promise for future research directions.

References

- [1] Brendan Alvey, Derek T Anderson, Andrew Buck, Matthew Deardorff, Grant Scott, and James M Keller. Simulated photo-realistic deep learning framework and workflows to accelerate computer vision and unmanned aerial vehicle research. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3889–3898, 2021. [1](#)
- [2] Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5855–5864, 2021. [2](#), [5](#), [6](#)
- [3] 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. [2](#), [3](#)
- [4] Wenzheng Chen, Huan Ling, Jun Gao, Edward Smith, Jaakko Lehtinen, Alec Jacobson, and Sanja Fidler. Learning to predict 3d objects with an interpolation-based differentiable renderer. *Advances in neural information processing systems*, 32, 2019. [2](#)
- [5] Xingyu Chen, Qi Zhang, Xiaoyu Li, Yue Chen, Ying Feng, Xuan Wang, and Jue Wang. Hallucinated neural radiance fields in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12943–12952, 2022. [3](#)
- [6] Abe Davis, Marc Levoy, and Fredo Durand. Unstructured light fields. In *Computer Graphics Forum*, pages 305–314. Wiley Online Library, 2012. [2](#)
- [7] Celso M de Melo, Antonio Torralba, Leonidas Guibas, James DiCarlo, Rama Chellappa, and Jessica Hodgins. Next-generation deep learning based on simulators and synthetic data. *Trends in cognitive sciences*, 2022. [1](#)
- [8] Kangle Deng, Andrew Liu, Jun-Yan Zhu, and Deva Ramanan. Depth-supervised nerf: Fewer views and faster training for free. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12882–12891, 2022. [2](#), [3](#)
- [9] Dawa Derksen and Dario Izzo. Shadow neural radiance fields for multi-view satellite photogrammetry. *Computer Vision and Pattern Recognition*, 2021. [3](#)
- [10] Thibaud Ehret, Roger Marí, and Gabriele Facciolo. Regularization of nerfs using differential geometry. *arXiv preprint arXiv:2206.14938*, 2022. [2](#), [3](#)
- [11] Steven J. Gortler, Radek Grzeszczuk, Richard Szeliski, and Michael F. Cohen. *The Lumigraph*. Association for Computing Machinery, New York, NY, USA, 1 edition, 2023. [2](#)
- [12] Philipp Henzler, Niloy J Mitra, and Tobias Ritschel. Learning a neural 3d texture space from 2d exemplars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8356–8364, 2020. [2](#)
- [13] 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](#)
- [14] Rasmus Jensen, Anders Dahl, George Vogiatzis, Engin Tola, and Henrik Aanaes. Large scale multi-view stereopsis evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 406–413, 2014. [5](#)
- [15] Daiju Kanaoka, Motoharu Sonogashira, Hakaru Tamukoh, and Yasutomo Kawanishi. Manifoldnerf: View-dependent image feature supervision for few-shot neural radiance fields. In *34th British Machine Vision Conference 2023, BMVC 2023, Aberdeen, UK, November 20-24, 2023*. BMVA, 2023. [2](#), [3](#)
- [16] Abhishek Kar, Christian Häne, and Jitendra Malik. Learning a multi-view stereo machine. *Advances in neural information processing systems*, 30, 2017. [2](#)
- [17] 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](#), [8](#)
- [18] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [5](#)
- [19] Kiriakos N Kutulakos and Steven M Seitz. A theory of shape by space carving. *International journal of computer vision*, 38:199–218, 2000. [2](#)
- [20] Minseop Kwak, Jiuhn Song, and Seungryong Kim. Geconerf: Few-shot neural radiance fields via geometric consistency. *arXiv preprint arXiv:2301.10941*, 2023. [3](#)
- [21] Obin Kwon, Jeongho Park, and Songhwa Oh. Renderable neural radiance map for visual navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9099–9108, 2023. [1](#)
- [22] Marc Levoy and Pat Hanrahan. Light field rendering. *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, 1996. [2](#)
- [23] Jiaxin Li, Zijian Feng, Qi She, Henghui Ding, Changhu Wang, and Gim Hee Lee. Mine: Towards continuous depth mpi with nerf for novel view synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12578–12588, 2021. [2](#)
- [24] Tzu-Mao Li, Miika Aittala, Frédo Durand, and Jaakko Lehtinen. Differentiable monte carlo ray tracing through edge sampling. *ACM Transactions on Graphics (TOG)*, 37(6): 1–11, 2018. [2](#)
- [25] Zhemin Li, Hongxia Wang, and Deyu Meng. Regularize implicit neural representation by itself. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10280–10288, 2023. [4](#)
- [26] Shichen Liu, Tianye Li, Weikai Chen, and Hao Li. Soft rasterizer: A differentiable renderer for image-based 3d reasoning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7708–7717, 2019. [2](#)
- [27] 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. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7824–7833, 2022. [3](#)
- [28] Ze Liu, Han Hu, Yutong Lin, Zhuliang Yao, Zhenda Xie, Yixuan Wei, Jia Ning, Yue Cao, Zheng Zhang, Li Dong, et al.

- Swin transformer v2: Scaling up capacity and resolution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12009–12019, 2022. 2, 5
- [29] Dominic Maggio, Marcus Abate, Jingnan Shi, Courtney Mario, and Luca Carlone. Loc-nerf: Monte carlo localization using neural radiance fields. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4018–4025. IEEE, 2023. 1
- [30] Roger Marí, Gabriele Facciolo, and Thibaud Ehret. Sat-NeRF: Learning multi-view satellite photogrammetry with transient objects and shadow modeling using rpc cameras. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1311–1321, 2022. 3
- [31] Ricardo Martin-Brualla, Noha Radwan, Mehdi SM Sajjadi, Jonathan T Barron, Alexey Dosovitskiy, and Daniel Duckworth. NeRF in the wild: Neural radiance fields for unconstrained photo collections. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7210–7219, 2021. 3
- [32] Christopher Maxey, Jaehoon Choi, Hyungtae Lee, Dinesh Manocha, and Heesung Kwon. Uav-sim: Nerf-based synthetic data generation for uav-based perception. *arXiv preprint arXiv:2310.16255*, 2023. 1
- [33] 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. 1, 2, 5, 6
- [34] 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, 4, 5, 6
- [35] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 8
- [36] 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. 7, 8
- [37] Steven M Seitz and Charles R Dyer. Photorealistic scene reconstruction by voxel coloring. *International Journal of Computer Vision*, 35:151–173, 1999. 2
- [38] Seunghyeon Seo, Yeonjin Chang, and Nojun Kwak. Flipnerf: Flipped reflection rays for few-shot novel view synthesis. *arXiv preprint arXiv:2306.17723*, 2023. 2, 3
- [39] Vincent Sitzmann, Justus Thies, Felix Heide, Matthias Nießner, Gordon Wetzstein, and Michael Zollhofer. Deepvoxels: Learning persistent 3d feature embeddings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2437–2446, 2019. 2
- [40] Edgar Sucar, Shikun Liu, Joseph Ortiz, and Andrew J Davison. imap: Implicit mapping and positioning in real-time. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6229–6238, 2021. 1
- [41] Matthew Tancik, Vincent Casser, Xinchun Yan, Sabeek Pradhan, Ben Mildenhall, Pratul P Srinivasan, Jonathan T Barron, and Henrik Kretschmar. Block-NeRF: Scalable large scene neural view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8248–8258, 2022. 3
- [42] Alex Trevithick and Bo Yang. Grf: Learning a general radiance field for 3d scene representation and rendering. 2020. 3
- [43] Richard Tucker and Noah Snavely. Single-view view synthesis with multiplane images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 551–560, 2020. 2
- [44] Michael Waechter, Nils Moehrle, and Michael Goesele. Let there be color! large-scale texturing of 3d reconstructions. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, pages 836–850. Springer, 2014. 2
- [45] 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. 2, 3
- [46] Peihao Wang, Xuxi Chen, Tianlong Chen, Subhashini Venugopalan, Zhangyang Wang, et al. Is attention all nerf needs? *arXiv preprint arXiv:2207.13298*, 2022. 2, 3
- [47] Qianqian Wang, Zhicheng Wang, Kyle Genova, Pratul P Srinivasan, Howard Zhou, Jonathan T Barron, Ricardo Martin-Brualla, Noah Snavely, and Thomas Funkhouser. Ibrnet: Learning multi-view image-based rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4690–4699, 2021. 2, 3
- [48] Yi Wei, Shaohui Liu, Yongming Rao, Wang Zhao, Jiwen Lu, and Jie Zhou. Nerfingmvs: Guided optimization of neural radiance fields for indoor multi-view stereo. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5610–5619, 2021. 3
- [49] Olivia Wiles, Georgia Gkioxari, Richard Szeliski, and Justin Johnson. Synsin: End-to-end view synthesis from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7467–7477, 2020. 2
- [50] Yongchang Wu, Zhengxia Zou, and Zhenwei Shi. Remote sensing novel view synthesis with implicit multiplane representations. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–13, 2022. 2, 5
- [51] Yuanbo Xiangli, Linning Xu, Xingang Pan, Nanxuan Zhao, Anyi Rao, Christian Theobalt, Bo Dai, and Dahua Lin. BungeeNeRF: Progressive neural radiance field for extreme multi-scale scene rendering. 2021. 3
- [52] 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. 3, 8
- [53] Zeke Xie, Xindi Yang, Yujie Yang, Qi Sun, Yixiang Jiang, Haoran Wang, Yunfeng Cai, and Mingming Sun. S3im:

- Stochastic structural similarity and its unreasonable effectiveness for neural fields. In *International Conference on Computer Vision*, 2023. [8](#)
- [54] Dejia Xu, Yifan Jiang, Peihao Wang, Zhiwen Fan, Humphrey Shi, and Zhangyang Wang. Sinnerf: Training neural radiance fields on complex scenes from a single image. In *European Conference on Computer Vision*, pages 736–753. Springer, 2022. [2](#)
- [55] Jiawei Yang, Marco Pavone, and Yue Wang. Freenerf: Improving few-shot neural rendering with free frequency regularization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8254–8263, 2023. [2](#), [4](#), [5](#), [6](#)
- [56] 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. [2](#), [3](#), [5](#), [6](#)
- [57] 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](#)
- [58] Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. Stereo magnification: learning view synthesis using multiplane images. *ACM Transactions on Graphics (TOG)*, 37(4):1–12, 2018. [2](#)
- [59] Zihan Zhu, Songyou Peng, Viktor Larsson, Weiwei Xu, Hujun Bao, Zhaopeng Cui, Martin R Oswald, and Marc Pollefeys. Nice-slam: Neural implicit scalable encoding for slam. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12786–12796, 2022. [1](#)