MixNeRF: Modeling a Ray with Mixture Density for Novel View Synthesis from Sparse Inputs

Seunghyeon Seo Donghoon Han* Yeonjin Chang* Nojun Kwak
Seoul National Univeristy
{zzzlssh, dhk1349, yjean8315, nojunk}@snu.ac.kr

Abstract

Neural Radiance Field (NeRF) has broken new ground in the novel view synthesis due to its simple concept and state-of-the-art quality. However, it suffers from severe performance degradation unless trained with a dense set of images with different camera poses, which hinders its practical applications. Although previous methods addressing this problem achieved promising results, they relied heavily on the additional training resources, which goes against the philosophy of sparse-input novel-view synthesis pursuing the training efficiency. In this work, we propose MixNeRF, an effective training strategy for novel view synthesis from sparse inputs by modeling a ray with a mixture density model. Our MixNeRF estimates the joint distribution of RGB colors along the ray samples by modeling it with mixture of distributions. We also propose a new task of ray depth estimation as a useful training objective, which is highly correlated with 3D scene geometry. Moreover, we remodel the colors with regenerated blending weights based on the estimated ray depth and further improves the robustness for colors and viewpoints. Our MixNeRF outperforms other state-of-the-art methods in various standard benchmarks with superior efficiency of training and inference.

1. Introduction

A photo-realistic view synthesis is one of the major research topics in computer vision. Recently, the coordinate-based neural representation [6, 25, 26, 30] has gained much popularity for the novel view synthesis task. Among them, Neural Radiance Field (NeRF) [28], which models a 3D scene by learning from a dense set of 2D images, enabled high-quality view synthesis with a simple concept and has become the prevailing mainstream. However, NeRF suffers from severe performance degradation in real-world applications, *e.g.* AR/VR, autonomous driving, and so on, where only a sparse set of views are available due to the burdensome task of collecting dense training images.

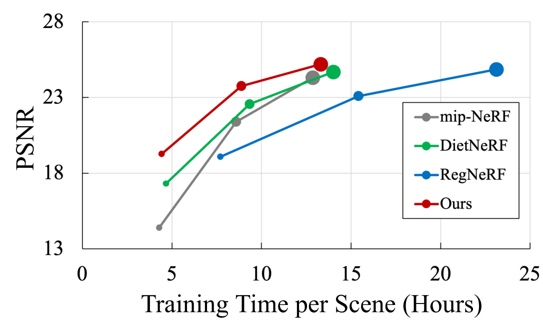

Figure 1. **Comparison with the vanilla mip-NeRF [1] and other regularization methods.** Given the same number of training batch and iterations, our MixNeRF outperforms mip-NeRF and DietNeRF [12] by a large margin with comparable or shorter training time. Compared to RegNeRF [29], ours achieves superior performance with about 42% shortened training time. The size of the circles are proportional to the number of input views, indicating 3/6/9-view, respectively. More details are provided in Sec. 4.2.

One of the key factors for a model's high-quality rendering with limited input views is its robustness in 3D geometry learning, *i.e.* accurate depth estimation for a scene. There are several works to address this problem and it can be classified into two major paradigms: *pre-training* and *regularization* approaches. For the pre-training approach [5,7,13,16,20,22,31,34,37,42], a general 3D geometry is trained by the multi-view images from a large-scale dataset and per-scene finetuning is optionally conducted in the test time. Although it has achieved promising results, it still requires the expensive cost for collecting a large-scale dataset across different scenes for pre-training and is not well-generalized for a novel domain in the test time.

Another line of research, the regularization approach [9, 12, 18, 29, 32, 39], performs per-scene optimization from scratch by applying regularization to prevent being overfitted from the limited inputs. Most existing methods of this kind depend heavily on the extra training resources for com-

*D. Han and Y. Chang equally contributed to this work.

This work was supported by NRF grant (2021R1A2C3006659) and IITP grants (2021-0-01343, 2022-0-00953) funded by Korean Government.

pensating a lack of supervisory signals, *e.g*. depth-map generation by running SfM [9, 32], unseen ray generation with arbitrary camera poses [18,29], leveraging external modules to exploit additional features [12,29,39], or so on. However, the additional training data might not always be available and the external modules should be pre-trained with a large-scale dataset. This is against the philosophy of novel view synthesis from sparse inputs, pursuing training efficiency.

In this work, we propose MixNeRF, an effective regularization approach for novel view synthesis from sparse inputs, modeling the colors along a ray with a mixture density model which represents a complex distribution with a mixture of component distributions. Since the concept of a mixture model is in line with that of alpha compositing in that both are represented by the weighted combination, we are able to regularize effectively both the colors and the densities of the samples along a ray by exploiting the blending weights as mixing coefficients for our mixture model. Furthermore, we propose a new auxiliary task of ray depth estimation for learning the 3D geometry which is crucial for the rendering quality. Since the estimated 3D geometry is highly correlated with the scene depth estimation, our proposed training objective acts as a useful supervisory signal. Finally, we regenerate the blending weights based on the estimated ray depth and remodel a ray. Since the estimated depth is not exactly the same, but nearly identical to the ground truth, it can play a role of pseudo geometry for adjacent points of the sample, like an unseen viewpoint. By remodeling the samples with the mixing coefficients based on the regenerated blending weights, we can further improve the robustness for shift of colors and viewpoints. Our main contributions are summarized as follows:

- Our method estimates the joint distribution of RGB color values along the ray samples by a mixture of distributions, learning the 3D geometry successfully with sparse views.
- We propose a ray depth estimation as an effective auxiliary task for few-shot novel view synthesis, playing a role of useful training objective.
- We use the regenerated blending weights based on the estimated ray depths for improving the robustness with negligible extra training cost.
- Our MixNeRF outperforms other state-of-the-art methods in the different standard benchmarks, showing much improved training and inference efficiency.

2. Related Works

2.1. Neural Scene Representations

Recently, coordinate-based neural representations [6, 25, 26, 30] have gained a lot of popularity in the neural scene rendering [1, 2, 11, 15, 17, 21, 23, 28]. Among them, Neural Radiance Fields (NeRF) [28] have broken new ground in the novel view synthesis research due to its wide possibility with the simple concept and state-of-the-art quality. Since NeRF, several works have been followed to ameliorate its drawbacks and improve the performance. Mip-NeRF [1] tackled the problem of aliasing by introducing cone tracing method. Ref-NeRF [36] reparameterized NeRF from the view-dependent outgoing radiance to reflected radiance, leading to improvement for specular reflections.

However, these methods suffer from severe performance degradation unless trained with a set of dense images with different camera poses, which hinders their practical applications. In this work, we address the sparse input scenario which is closer to the real-world condition. We are able to perform high-quality view synthesis from sparse inputs by modeling a ray with a mixture density model and improve both the training and the inference efficiency.

2.2. Sparse Input Novel View Synthesis

One of the fundamental causes of performance degradation is the lack of 3D geometry information from training images, resulting in an inaccurate depth estimation. There are two major paradigms to tackle this problem in the novel view synthesis from sparse inputs: *pre-training* and *regularization* approaches. The former approach [5,7,13,16,20, 22, 31, 34, 37, 42] provides prior knowledges to conditional models through pre-training. The image features extracted by a CNN feature extractor [7, 42] or a 3D cost volume obtained by image warping [5, 16] are used for training a generalizable model. Although they achieved promising performances under the sparse input setting, a large-scale dataset of multi-view images with different scenes is required for pre-training, which is burdensome to collect. Furthermore, despite the lengthy pre-training phase, most of these methods require additional test-time fine-tuning and are apt to suffer from quality degradation on different data domains.

The regularization approach [9, 12, 18, 29, 32, 39] introduces extra supervision to regularize the color and the geometry without an expensive pre-training process. Additional training resources, *e.g*. external modules such as CLIP [12] or a pre-trained normalizing flow model [10], extra depth inputs obtained by running structure-from-motion (SfM), and additional rays of unseen viewpoints, are often used to provide abundant supervisory signals. However, the existing methods are overly dependent on the extra training resources which might not always be available, hampering data/time efficiency. Moreover, it goes against the philosophy of the sparse-input novel-view synthesis which pursues the training efficiency.

Our proposed method requires neither an external module nor an additional inference of extra supervisory signals, resulting in a more efficient training framework.

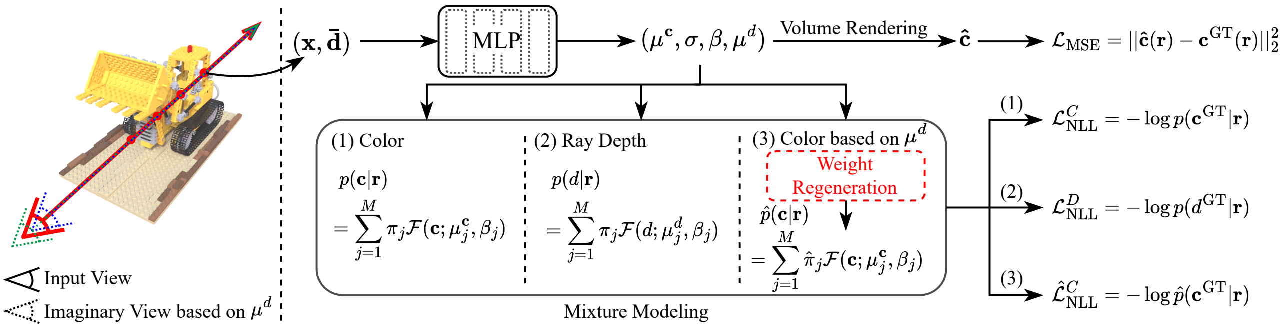

Figure 2. **Overview of MixNeRF.** Our method models the color and depth of a ray with a mixture of distributions, and remodels the color based on the estimated ray depth μ^d. The dotted rays indicate the imaginary views corresponding to μ^d. See Sec. 3 for more details.

2.3. Mixture Density Model

There exists a line of research utilizing a mixture density model in different tasks of computer vision [8, 19, 33, 35, 40, 41]. Among 3D vision tasks, Tosi *et al.* [33] proposed a novel stereo-matching framework, SMD-Nets, tackling the over-smoothing problem of output representations by leveraging a mixture density network [3]. Choi *et al.* [8] reformulated 3D bounding box regression as a density estimation problem using a Gaussian Mixture Model (GMM), achieving a more efficient 3D object detection framework with few heuristic design factors.

Although the mixture density model shows a great potential in 3D vision tasks, there has not been an attempt to utilize it in the NeRF framework. Our MixNeRF is able to learn the 3D geometry successfully, which is a critical factor for rendering quality under the sparse input setting, by modeling a ray with a mixture of distributions.

3. Method

In this work, we propose a novel training framework of neural radiance fields for novel view synthesis from sparse inputs. We build our MixNeRF upon mip-NeRF [1] which uses a multiscale scene representation (Sec. 3.1). Moreover, we leverage the mixture density model framework to learn 3D geometry efficiently. More specifically, we model the colors of samples along a ray by a mixture of Laplace distributions with the predicted weights as mixing coefficients, which contributes to learning a scene's geometry effectively with limited input views (Sec. 3.2). Furthermore, we estimate the depths of input rays as an auxiliary task and reuse it for producing blending weights once again as supplemental training resources, which enables robust rendering from unseen viewpoints with little additional burden for training (Sec. 3.3). In the training phase, our MixNeRF is not only trained to minimize the mean squared error (MSE) between predictions and GT colors, but also to maximize the likelihood of colors and depths for each ray (Sec. 3.4). Fig. 2 demonstrates an overview of our MixNeRF.

3.1. Preliminary: Neural Radiance Field

NeRF [28] represents a 3D scene with a continuous function, where a neural network $f(\cdot, \cdot)$ consisting of an MLP maps a 3D location $\mathbf{x} = (x, y, z)$ and viewing direction (θ, ϕ), which is expressed as a 3D Cartesian unit vector $\bar{\mathbf{d}}$ in practice, along rays to colors $\mathbf{c} = (r, g, b)$ and volume density σ:

$$f(\gamma(\mathbf{x}), \gamma(\bar{\mathbf{d}})) \to (\mathbf{c}, \sigma), \tag{1}$$

where $\gamma(\cdot)$ indicates the positional encoding applied to the inputs $(\mathbf{x}, \bar{\mathbf{d}})$. Following the volume rendering theory [24], a pixel on an image is rendered by alpha compositing the colors and densities along the ray $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$ cast from the camera origin $\mathbf{o}$, where $\mathbf{d}$ is the unnormalized direction vector, *i.e.* $\mathbf{d} = \|\mathbf{d}\|_2 \cdot \bar{\mathbf{d}}$. The volume rendering integrals are approximated by the quadrature rule in practice [28] as follows:

$$\begin{gathered}\hat{\mathbf{c}}(\mathbf{r}) = \sum_{i=1}^{N} T_i (1 - \exp(-\sigma_i \delta_i)) \mathbf{c}_i, \\ \text{where} \quad T_i = \exp(-\sum_{j=1}^{i-1} \sigma_j \delta_j).\end{gathered} \tag{2}$$

Note that N and $\delta_i = \|\mathbf{d}\|_2 \cdot (t_{i+1} - t_i)$ denote the number of samples and the interval between the i-th sample and its adjacent one, respectively. To improve rendering efficiency, the two-stage hierarchical sampling is performed: *coarse* and *fine* stage. The points are sampled uniformly along a ray in the coarse stage, and then more informed samples are generated in the fine stage based on the density estimated from the coarse stage. Finally, the radiance field is optimized by minimizing the MSE between the rendered color and ground truth color over the input images:

$$\mathcal{L}_{\text{MSE}} = \sum_{\mathbf{r} \in \mathcal{R}} ||\hat{\mathbf{c}}(\mathbf{r}) - \mathbf{c}^{\text{GT}}(\mathbf{r})||_2^2, \tag{3}$$

where $\mathcal{R}$ indicates a set of input rays.

Following RegNeRF [29], we adopt the mip-NeRF [1] representation for our MixNeRF. Mip-NeRF effectively alleviates the aliasing problem of NeRF by introducing a cone tracing method and an integrated positional encoding.

3.2. Modeling a Ray with Mixture Density Model

Given a set of input rays $\mathcal{R} = \{\mathbf{r}_1, \cdots, \mathbf{r}_K\}$ on training images with the ground truths $G = \{G_1, \cdots, G_K\}$ for each of K pixels, the i-th ground truth G_i consists of the RGB color values $\mathbf{c}_i^{\mathrm{GT}} = \{r_i, g_i, b_i\}$ and the unnormalized 3D ray vector $\mathbf{d}_i^{\mathrm{GT}}$, *i.e.* $G_i \triangleq \{\mathbf{c}_i^{\mathrm{GT}}, \mathbf{d}_i^{\mathrm{GT}}\}$. Note that $\mathbf{d}^{\mathrm{GT}}$ is the direction vector corresponding to $t = 1$ from the camera center. First, our MixNeRF estimates the distribution of the RGB color values $\mathbf{c}_i$ along the samples of the ray $\mathbf{r}_i$ on a pixel with a mixture model, which is derived from a weighted combination of component distributions. As shown in Fig. 2, in our model, $(\mathbf{c}, \sigma)$, the conventional outputs of NeRF for each sampled point $\mathbf{r}(t)$, are used as a location parameter $\mu^{\mathbf{c}}$ and to compute a mixing coefficient π, respectively. In addition to these, a scale parameter $\beta = \{\beta^r, \beta^g, \beta^b\}$ is also estimated in our model.

We assume that every element of $\mathbf{c}_i$ is independent of each other to simplify our mixture model formulation. Therefore, the j-th component's probability density function (pdf) corresponding to the j-th sampled point for the i-th ray $\mathbf{r}_i$ is as follows:

$$\begin{aligned} \mathcal{F}(\mathbf{c}; \mu_{ij}^{\mathbf{c}}, \beta_{ij}) &= \prod_{c \in \{r,g,b\}} \mathcal{F}(c; \mu_{ij}^{c}, \beta_{ij}^{c}) \\ &= \prod_{c \in \{r,g,b\}} \frac{1}{2\beta_{ij}^c} \exp\left(-\frac{|c - \mu_{ij}^c|}{\beta_{ij}^c}\right), \end{aligned} \tag{4}$$

where $\mathcal{F}$ denotes the Laplacian pdf. The pdf of our mixture model formed by the component distributions above is defined as:

$$p(\mathbf{c}|\mathbf{r}_i) = \sum_{j=1}^{M} \pi_{ij} \mathcal{F}(\mathbf{c}; \mu_{ij}^{\mathbf{c}}, \beta_{ij}), \tag{5}$$

where M denotes the number of mixture components which is the same as the number of samples along a ray. The mixture coefficient π_{ij} is derived from the density output σ_{ij} as follows:

$$\pi_j = \frac{w_j}{\sum_{m=1}^{M} w_m} = \frac{T_j(1 - \exp(-\sigma_j \delta_j))}{\sum_{m=1}^{M} T_m(1 - \exp(-\sigma_m \delta_m))}. \tag{6}$$

Note that we omitted ray index i for simplicity. Here, w_j and δ_j indicate the weight for the alpha compositing and the sample interval, respectively. Since the mixture components corresponding to the samples with higher weights, which contribute more to the alpha composition of the color than other samples, are likely to have higher π, we use the normalized weight as a mixing coefficient π so that $\sum_{j=1}^{M} \pi_j = 1$.

The concept of a mixture model corresponds to that of alpha compositing in that a complex multimodal distribution is able to be represented by the weighted combination of component distributions with mixing coefficients π, like a pixel value derived from the weighted combination of estimated RGB values along ray samples with blending weights w. Motivated by this conceptual similarity, we are able to model a ray with a mixture of distributions successfully without any heuristic factors. The mixing coefficients derived from the blending weights provide effective supervisory signals toward the densities, which are the core factor for successfully learning 3D scene geometry with limited input views.

3.3. Depth Estimation by Mixture Density Model

We propose a ray's depth estimation as an effective auxiliary task for training our MixNeRF with sparse inputs. As demonstrated in Fig. 2, our MixNeRF estimates d, the ray's depth, which is defined as the length of the unnormalized ray direction vector $\mathbf{d}$, *i.e.* $d \triangleq \|\mathbf{d}\|$, along the ray samples. The ground truth G_i contains the ray direction values $\mathbf{d}_i$ as well, which are used in the form of 3D Cartesian unit vectors $\bar{\mathbf{d}}_i = \mathbf{d}_i / \|\mathbf{d}_i\|_2$ as an input viewing direction in practice. Like the RGB color values, the depths for each ray are modeled by our mixture model consisting of the Laplace distributions with the same scale parameters β and mixing coefficients π used above. The pdf of our mixture model for the depth of the i-th ray is as follows:

$$p(d|\mathbf{r}_i) = \sum_{j=1}^{M} \pi_{ij} \mathcal{F}(d; \mu_{ij}^{d}, \beta_{ij}). \tag{7}$$

Since the mixing coefficient π and parameter β are optimized through the supervision of the depth as well as the color values, it improves the robustness of our MixNeRF for slight changes of geometry. Also, considering that the successful depth estimation is crucial to the rendered images' quality in a NeRF model [9,18,29,32], our direct estimation of the ray's depth benefits a lot.

Blending weight regeneration. In addition, we exploit the estimated depth to regenerate the blending weights along the samples and model the RGB color values by a mixture of distributions once again. Since the estimated depth of each sample is trained to be nearly identical to the ground truth depth, but not exactly the same, it can play a role of pseudo geometry for adjacent points of the sample without any additional pre-generation process of extra training data, *e.g.* depth inputs made by SfM or rays from unobserved viewpoints. The new blending weight $\hat{w}_j$ of the j-th sample along a ray based on the estimated depth μ_j^d are

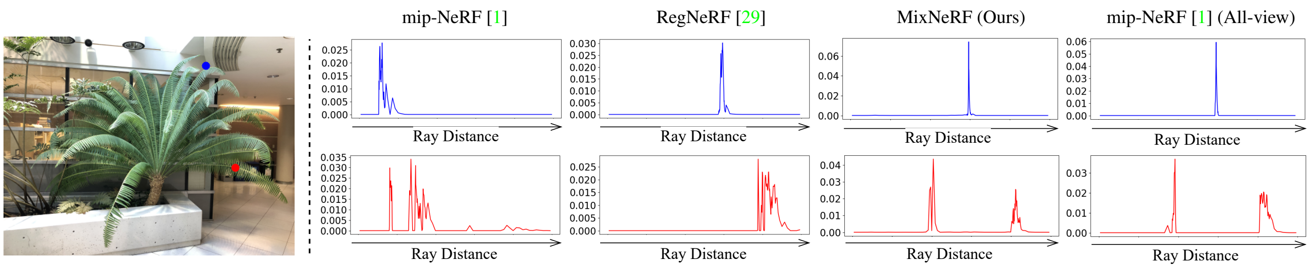

Figure 3. **Comparison of blending weight distributions.** Compared to the baselines, ours estimates the modes of weight distributions more accurately, leading to the precise 3D geometry. The ideal distributions on the blue and red points are unimodal and bimodal, respectively.

defined as follow:

$$\hat{w}_j = \hat{T}_j(1 - \exp(-\sigma_j \hat{\delta}_j)),\ \hat{\delta}_j = \mu^d_j(t_{j+1} - t_j), \quad (8)$$

in which we replace $\|\mathbf{d}\|_2$ in δ_j formulation with μ^d_j. Finally, we model the color values along a ray based on the new mixing coefficients $\hat{\pi}$ derived from $\hat{w}$ and the corresponding pdf is as follows:

$$\hat{p}(\mathbf{c}|\mathbf{r}_i) = \sum_{j=1}^{M} \hat{\pi}_{ij} \mathcal{F}(\mathbf{c}; \mu^{\mathbf{c}}_{ij}, \beta_{ij}). \quad (9)$$

Since the estimated ray depths are likely to be close enough to those of the ground truths, we use the same GT color values of input rays for modeling the mixture distribution based on the newly generated $\hat{\pi}$. It further improves the robustness for shift of colors and ray viewpoints by simply modeling a ray once again with regenerated blending weights, eliminating pre-generation and extra inference of unseen views without much computational overhead.

3.4. Total Loss

Our MixNeRF is trained to maximize the likelihood of $\mathbf{c}^{\text{GT}}$ and d^{GT} for a set of input rays $\mathcal{R}$ as well as to minimize the $\mathcal{L}_{\text{MSE}}$. Therefore, the loss functions can be simply defined to minimize the negative log-likelihood (NLL) of the ground truths as follows:

$$\begin{aligned} \mathcal{L}^C_{\text{NLL}} &= -\textstyle\sum_{\mathbf{r}\in\mathcal{R}} \log p(\mathbf{c}^{\text{GT}}|\mathbf{r}), \\ \mathcal{L}^D_{\text{NLL}} &= -\textstyle\sum_{\mathbf{r}\in\mathcal{R}} \log p(d^{\text{GT}}|\mathbf{r}), \\ \hat{\mathcal{L}}^C_{\text{NLL}} &= -\textstyle\sum_{\mathbf{r}\in\mathcal{R}} \log \hat{p}(\mathbf{c}^{\text{GT}}|\mathbf{r}), \end{aligned} \quad (10)$$

each of which corresponds to the NLL form of Eq. (5), Eq. (7) and Eq. (9), respectively. As a result, we define our total loss as:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{MSE}} + \lambda_C \mathcal{L}^C_{\text{NLL}} + \lambda_D \mathcal{L}^D_{\text{NLL}} + \hat{\lambda}_C \hat{\mathcal{L}}^C_{\text{NLL}}, \quad (11)$$

where λ_C, λ_D and $\hat{\lambda}_C$ are balancing terms for the losses. More details about training and implementation are provided in the supplementary material.

4. Experiments

4.1. Experimental Details

Datasets and metrics. We evaluate MixNeRF on the multiple standard benchmarks: DTU [14], LLFF [27] and Realistic Synthetic 360° [28]. DTU consists of images containing objects located on a white table with a black background. LLFF contains real forward-facing scenes and is usually used as an out-of-distribution test set for pre-training methods. We also compare our MixNeRF against other regularization methods on the Realistic Synthetic 360°, which provides 8 synthetic scenes each consisting of 400 images rendered from inward-facing cameras with various viewpoints. We follow the overall experimental protocols of [12, 18, 28, 42] for these datasets.

For the evaluation metrics, we adopt the mean of PSNR, structural similarity index (SSIM) [38], LPIPS perceptual metric [43], and the geometric average [1]. Kindly refer to the supp. mat. for more details about datasets and metrics.

Baselines. We compare our method against several representative pre-training and regularization approaches [5, 7, 12, 18, 29, 42] as well as the vanilla mip-NeRF [1]. We report the evaluation results from [29] for DTU and LLFF, which are superior to those from their original papers due to the improved training curriculum. The DTU dataset is used as a pre-training resources for PixelNeRF [42], MVSNeRF [5], and SRF [7], and the LLFF dataset serves as an out-of-domain test set. The regularization approaches and mip-NeRF are trained for each scene without pre-training. For Realistic Synthetic 360°, we train other regularization approaches [12, 18, 29] by their training schemes. Note that the pre-trained RealNVP [10] for training RegNeRF [29] is not publicly available and we report the results of RegNeRF trained without it on the Realistic Synthetic 360°. For the analysis of MixNeRF (Sec. 4.2), all models including ours are trained with the same batch size and iterations.

4.2. Analysis of MixNeRF

Benefit of mixture density model. We leverage a mixture density model, which represents a complex multimodal distribution with a weighted combination of component distri-

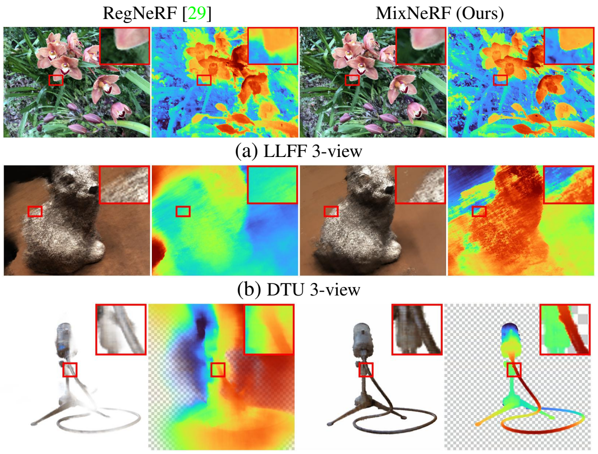

(a) LLFF 3-view

(b) DTU 3-view

(c) Realistic Synthetic 360° 4-view

Figure 4. **Comparison of estimated depth map.** We compare MixNeRF against RegNeRF, the state-of-the-art regularization approach. Our MixNeRF estimates more accurate depth maps and captures fine details better, leading to high-quality rendering with more distinct edges and less artifacts.

butions, to learn the distribution of density and colors along the ray samples effectively. Fig. 3 demonstrates the comparison of the blending weight distributions of cast rays on the LLFF fern scene of 3-view scenario. We compare ours against mip-NeRF and RegNeRF, with mip-NeRF trained from all training views as an ideal distribution. For the unimodal distribution in blue, mip-NeRF does not estimate the mode well and achieves degenerate geometry. However, RegNeRF and our MixNeRF show the unimodal weight distributions leading to higher-quality novel views, and especially our MixNeRF achieves the distribution with sharper mode than RegNeRF, which is more similar to that of mip-NeRF (All-view). In case of the bimodal-shaped distribution in red, our MixNeRF estimates the weight distribution successfully while both mip-NeRF and RegNeRF fail to estimate the accurate modes. Since the predicted 3D geometry is directly correlated with how well the density is estimated, our MixNeRF is able to learn the geometry more efficiently with limited input views through mixture density modeling.

Depth map estimation. We compare our MixNeRF against RegNeRF, which utilizes the prior of depth smoothness to learn 3D geometry. As illustrated in Fig. 4, our MixNeRF estimates more accurate depth maps with distinct edges while RegNeRF generates both RGB images and depth maps with blurry fine details. Especially for Realistic Synthetic 360°, we observe that RegNeRF fails to learn the geometry with its smoothing strategy and achieves degenerate results due to the overly strong prior of depth smoothness. However, since our MixNeRF learns the depth of a ray by leveraging a mixture density model without smoothing from additional unseen rays, the depth maps are predicted much more efficiently and precisely.

# of samples	128	64	32	16	8
mip-NeRF [1]	0.332	0.331	0.329	0.308	0.251
RegNeRF [29]	0.587	0.585	0.576	0.522	0.379
MixNeRF (Ours)	**0.629**	**0.629**	**0.620**	**0.580**	**0.468**

Table 1. **Comparison with baselines by the number of ray samples.** Our MixNeRF with 75% fewer samples (32-sample) outperforms RegNeRF with default 128-sample, and still achieves comparable results with only 16-sample ($\times 8$ reduction).

Efficiency in training and inference. Our MixNeRF improves the efficiency for both the training and the inference phases by learning the 3D geometry effectively without burdensome extra training resources. Fig. 1 illustrates that MixNeRF achieves superior performance with reduced training time among the vanilla mip-NeRF and two representative regularization methods on the LLFF. For a fair comparison, we compare the methods based on the identical JAX codebase [4] using the same batch size and iterations on 2 NVIDIA TITAN RTX. Although it takes a similar amount of time to train DietNeRF as MixNeRF, its performance is inferior significantly to ours in 3 and 6-view scenario. Compared to RegNeRF, ours outperforms it with about 42% shorter training time per scene under the same number of input view scenario, resulting from the elimination of extra inference for additional unseen rays. Furthermore, we also observe that our MixNeRF shows better data efficiency requiring up to about 60% fewer inputs than mip-NeRF to achieve comparable results, and outperforms mip-NeRF consistently in more than 9-view scenarios. It indicates that our proposed training strategy is effective in general scenarios as well as the sparse input setting. The related experimental results are provided in the suppl. material. For the inference efficiency, Tab. 1 demonstrates the SSIM results by the number of samples along a ray on the LLFF under the 3-view scenario. Our MixNeRF with 32-sample outperforms RegNeRF with default 128-sample, and still achieves comparable results with only 16-sample thanks to the capacity of our mixture model for representing the blending weight distributions successfully.

4.3. Ablation Study

We report the quantitative and qualitative results of our ablation study in Fig. 5. We observe that modeling a ray with mixture of distributions is helpful for improving performance under the sparse view setting ((1) $\rightarrow$ (2)). Also, our proposed ray depth estimation task contributes to further improving the rendering quality by generating more accurate depth maps ((2) $\rightarrow$ (3)). However, despite the well-estimated depth map, the RGB image suffers from the foggy artifacts upon the objects as shown in (3). By remodeling a ray through the weight regeneration process, our MixNeRF achieves high-quality of both RGB image and depth map ((3) $\rightarrow$ (5)). Since the regenerated weights are not helpful

	Approach	PSNR ↑ 3-view	PSNR ↑ 6-view	PSNR ↑ 9-view	SSIM ↑ 3-view	SSIM ↑ 6-view	SSIM ↑ 9-view	LPIPS ↓ 3-view	LPIPS ↓ 6-view	LPIPS ↓ 9-view	Average Error ↓ 3-view	Average Error ↓ 6-view	Average Error ↓ 9-view
LLFF.													
mip-NeRF [1]	N/A	14.62	20.87	24.26	0.351	0.692	0.805	0.495	0.255	0.172	0.246	0.114	0.073
PixelNeRF [42]	Pre-training	7.93	8.74	8.61	0.272	0.280	0.274	0.682	0.676	0.665	0.461	0.433	0.432
PixelNeRF *ft* [42]	Pre-training	16.17	17.03	18.92	0.438	0.473	0.535	0.512	0.477	0.430	0.217	0.196	0.163
SRF [7]	Pre-training	12.34	13.10	13.00	0.250	0.293	0.297	0.591	0.594	0.605	0.313	0.293	0.296
SRF *ft* [7]	Pre-training	17.07	16.75	17.39	0.436	0.438	0.465	0.529	0.521	0.503	0.203	0.207	0.193
MVSNeRF [5]	Pre-training	17.25	19.79	20.47	0.557	0.656	0.689	0.356	0.269	0.242	0.171	0.125	0.111
MVSNeRF *ft* [5]	Pre-training	17.88	19.99	20.47	0.584	0.660	0.695	0.327	0.264	0.244	0.157	0.122	0.111
DietNeRF [12]	Regularization	14.94	21.75	24.28	0.370	0.717	0.801	0.496	0.248	0.183	0.240	0.105	0.073
RegNeRF [29]	Regularization	19.08	23.10	24.86	0.587	0.760	0.820	0.336	0.206	0.161	0.146	0.086	0.067
MixNeRF (Ours)	Regularization	**19.27**	**23.76**	**25.20**	**0.629**	**0.791**	**0.833**	**0.236**	**0.115**	**0.087**	**0.124**	**0.066**	**0.052**
DTU.													
mip-NeRF [1]	N/A	8.68	16.54	23.58	0.571	0.741	0.879	0.353	0.198	0.092	0.323	0.148	0.056
PixelNeRF [42]	Pre-training	16.82	19.11	20.40	0.695	0.745	0.768	0.270	0.232	0.220	0.147	0.115	0.100
PixelNeRF *ft* [42]	Pre-training	**18.95**	20.56	21.83	0.710	0.753	0.781	0.269	0.223	0.203	0.125	0.104	0.090
SRF [7]	Pre-training	15.32	17.54	18.35	0.671	0.730	0.752	0.304	0.250	0.232	0.171	0.132	0.120
SRF *ft* [7]	Pre-training	15.68	18.87	20.75	0.698	0.757	0.785	0.281	0.225	0.205	0.162	0.114	0.093
MVSNeRF [5]	Pre-training	18.63	20.70	22.40	**0.769**	0.823	0.853	0.197	0.156	0.135	0.113	0.088	0.068
MVSNeRF *ft* [5]	Pre-training	18.54	20.49	22.22	**0.769**	0.822	0.853	0.197	0.155	0.135	0.113	0.089	0.069
DietNeRF [12]	Regularization	11.85	20.63	23.83	0.633	0.778	0.823	0.314	0.201	0.173	0.243	0.101	0.068
RegNeRF [29]	Regularization	18.89	22.20	24.93	0.745	**0.841**	**0.884**	**0.190**	0.117	0.089	**0.112**	0.071	0.047
MixNeRF (Ours)	Regularization	**18.95**	**22.30**	**25.03**	0.744	0.835	0.879	0.203	**0.102**	**0.065**	0.113	**0.066**	**0.042**

Table 2. **Quantitative results on LLFF and DTU.** Our method achieves comparable or state-of-the-art performance across all scenarios. For LLFF, our MixNeRF outperforms both pre-training and regularization baselines without any burdensome extra training resources. Likewise, MixNeRF achieves competitive results against the state-of-the-art methods on DTU dataset. *ft* indicates fine-tuning.

	$\mathcal{L}^{C}_{\text{NLL}}$	$\mathcal{L}^{D}_{\text{NLL}}$	$\hat{\mathcal{L}}^{C}_{\text{NLL}}$	PSNR ↑	SSIM ↑	LPIPS ↓	Average ↓
(1)	Baseline (mip-NeRF† [1])			13.72	0.665	0.273	0.209
(2)	✓			17.44	0.729	**0.187**	0.130
(3)	✓	✓		18.35	0.739	0.209	0.121
(4)	✓		✓	17.45	0.728	0.191	0.134
(5)	✓	✓	✓	**18.95**	**0.744**	0.203	**0.113**

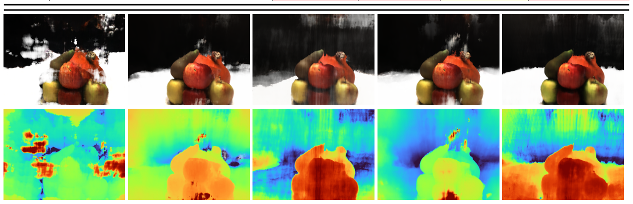

(1) (2) (3) (4) (5)

Figure 5. **Ablation study.** † uses a scene annealing strategy.

without a supervision toward the ray depth estimation task ((2) → (4)), our proposed auxiliary task of ray depth estimation is useful for learning 3D geometry and playing a role of additional training resources by weight regeneration process on-the-fly.

4.4. Comparison with other SOTA methods

LLFF. As shown in Tab. 2, the pre-training approaches except MVSNeRF are not able to achieve comparable results without fine-tuning for the 3-view scenario. The regularization approaches and vanilla mip-NeRF outperform the pre-training approaches in 6 and 9-view settings. Especially, RegNeRF improves the rendering quality by a large margin compared to mip-NeRF and DietNeRF, thanks to the regularization strategy using pre-generated rays from unseen viewpoints. Our MixNeRF achieves SOTA results across all scenarios and metrics without any pre-training process using a large-scale dataset or extra inference for pre-generated training resources. It improves the training efficiency, *e.g.* requiring about 42% shorter training time than RegNeRF in the same input view setting (see Fig. 1). Furthermore, as demonstrated in Fig. 6, we observe that our method achieves more realistic results with fine details than RegNeRF since MixNeRF learns the 3D geometry successfully without an explicit regularization for smoothing the depth.

	PSNR ↑ 4-view	PSNR ↑ 8-view	SSIM ↑ 4-view	SSIM ↑ 8-view	LPIPS ↓ 4-view	LPIPS ↓ 8-view	Average Err. ↓ 4-view	Average Err. ↓ 8-view
mip-NeRF [1]	14.12	18.74	0.720	0.830	0.380	0.240	0.220	0.120
DietNeRF [12]	15.42	21.31	0.730	0.847	0.314	0.153	0.201	0.086
InfoNeRF [18]	18.44	22.01	0.792	0.852	0.223	0.133	0.119	0.073
RegNeRF [29]	13.71	19.11	0.786	0.841	0.346	0.200	0.210	0.122
MixNeRF (Ours)	**18.99**	**23.84**	**0.807**	**0.878**	**0.199**	**0.103**	**0.113**	**0.060**

Table 3. **Quantitative results on Realistic Synthetic 360°.** MixNeRF outperforms other state-of-the-art regularization methods by a large margin.

DTU. Since the pre-training approaches are conditioned from the prior knowledge of DTU, they achieve competitive results with and without fine-tuning as demonstrated in Tab. 2. Our MixNeRF achieves comparable or better results than other baselines. Compared to the PSNR, our MixNeRF achieves worse quantitative results than RegNeRF for SSIM. We conjecture that since RegNeRF uses an explicit smoothing term for regularizing the depth, it can achieve slightly better quantitative results of SSIM which is a patch-wise evaluation metric. However, as shown in Fig. 6, our

(e) Realistic Synthetic 360° 4-view

Figure 6. **Qualitative results on LLFF (a,b), DTU (c,d), and Realistic Synthetic 360° (e).** More results are provided in suppl. material.

MixNeRF renders clearer images, capturing fine details better than other baselines despite a little worse quantitative results for some metrics. More qualitative results of LLFF and DTU are provided in the suppl. material.

Realistic Synthetic 360°. As shown in Tab. 3, our MixNeRF outperforms other regularization baselines across all settings and metrics by a large margin. Fig. 6e illustrates that other methods suffer from severe floating artifacts and degenerate colors in the 4-view scenario. Especially, the depth smoothing strategy of RegNeRF rather brings about the significant performance degradation. It implies that smoothing is not a fundamental solution universally effective for different datasets. Compared to the baselines, our MixNeRF achieves superior rendering quality with much less artifacts and more accurate geometry in both 4 and 8-view (see the suppl. material) scenarios.

5. Conclusion

We have introduced MixNeRF, a novel regularization approach for training NeRF in the limited data scenario. However, previous approaches heavily depend on the extra training resources, which goes against the philosophy of sparse-input novel-view synthesis pursuing the efficiency of training. To overcome this bottleneck, we propose modeling a ray with mixture density, which enables effective learning of 3D geometry with sparse inputs. Furthermore, our novel training strategy, consisting of an auxiliary ray depth estimation and the following weight regeneration, further improves the rendering quality and better reconstructs 3D geometry by more accurate depth estimation without any extra training resources that should be prepared in advance. Our proposed MixNeRF outperforms both pre-training and regularization approaches across the multiple benchmarks with an enhanced efficiency of training and inference.

References

[1] 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. 1, 2, 3, 4, 5, 6, 7, 8

[2] Alexander Bergman, Petr Kellnhofer, and Gordon Wetzstein. Fast training of neural lumigraph representations using meta learning. *Advances in Neural Information Processing Systems*, 34:172–186, 2021. 2

[3] Christopher M Bishop. Mixture density networks. 1994. 3

[4] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, et al. Jax: composable transformations of python+ numpy programs. *Version 0.2*, 5:14–24, 2018. 6

[5] 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, 2, 5, 7

[6] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5939–5948, 2019. 1, 2

[7] Julian Chibane, Aayush Bansal, Verica Lazova, and Gerard Pons-Moll. Stereo radiance fields (srf): Learning view synthesis for sparse views of novel scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7911–7920, 2021. 1, 2, 5, 7

[8] Jaeseok Choi, Yeji Song, Yerim Kim, Jaeyoung Yoo, and Nojun Kwak. Md3d: Mixture-density-based 3d object detection in point clouds. *IEEE Access*, 2022. 3

[9] 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. 1, 2, 4

[10] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*, 2016. 2, 5

[11] Chen Gao, Yichang Shih, Wei-Sheng Lai, Chia-Kai Liang, and Jia-Bin Huang. Portrait neural radiance fields from a single image. *arXiv preprint arXiv:2012.05903*, 2020. 2

[12] 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, 5, 7, 8

[13] Wonbong Jang and Lourdes Agapito. Codenerf: Disentangled neural radiance fields for object categories. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12949–12958, 2021. 1, 2

[14] 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

[15] Yue Jiang, Dantong Ji, Zhizhong Han, and Matthias Zwicker. Sdfdiff: Differentiable rendering of signed distance fields for 3d shape optimization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1251–1261, 2020. 2

[16] Mohammad Mahdi Johari, Yann Lepoittevin, and François Fleuret. Geonerf: Generalizing nerf with geometry priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18365–18375, 2022. 1, 2

[17] Petr Kellnhofer, Lars C Jebe, Andrew Jones, Ryan Spicer, Kari Pulli, and Gordon Wetzstein. Neural lumigraph rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4287–4297, 2021. 2

[18] 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, 4, 5, 7, 8

[19] Chen Li and Gim Hee Lee. Generating multiple hypotheses for 3d human pose estimation with mixture density network. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9887–9895, 2019. 3

[20] 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. 1, 2

[21] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. *Advances in Neural Information Processing Systems*, 33:15651–15663, 2020. 2

[22] 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. 1, 2

[23] 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. 2

[24] Nelson Max. Optical models for direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 1(2):99–108, 1995. 3

[25] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4460–4470, 2019. 1, 2

[26] Mateusz Michalkiewicz, Jhony K Pontes, Dominic Jack, Mahsa Baktashmotlagh, and Anders Eriksson. Implicit surface representations as layers in neural networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4743–4752, 2019. 1, 2

[27] Ben Mildenhall, Pratul P Srinivasan, Rodrigo Ortiz-Cayon, Nima Khademi Kalantari, Ravi Ramamoorthi, Ren Ng, and Abhishek Kar. Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *ACM Transactions on Graphics (TOG)*, 38(4):1–14, 2019. 5

[28] 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, 3, 5

[29] 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, 4, 5, 6, 7, 8

[30] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 165–174, 2019. 1, 2

[31] Konstantinos Rematas, Ricardo Martin-Brualla, and Vittorio Ferrari. Sharf: Shape-conditioned radiance fields from a single view. In *International Conference on Machine Learning*, pages 8948–8958. PMLR, 2021. 1, 2

[32] Barbara Roessle, Jonathan T Barron, Ben Mildenhall, Pratul P Srinivasan, and Matthias Nießner. Dense depth priors for neural radiance fields from sparse input views. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12892–12901, 2022. 1, 2, 4

[33] Fabio Tosi, Yiyi Liao, Carolin Schmitt, and Andreas Geiger. Smd-nets: Stereo mixture density networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8942–8952, 2021. 3

[34] Alex Trevithick and Bo Yang. Grf: Learning a general radiance field for 3d representation and rendering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15182–15192, 2021. 1, 2

[35] Prune Truong, Martin Danelljan, Luc Van Gool, and Radu Timofte. Learning accurate dense correspondences and when to trust them. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5714–5724, 2021. 3

[36] 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. 2

[37] 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. 1, 2

[38] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 5

[39] 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. *arXiv preprint arXiv:2204.00928*, 2022. 1, 2

[40] Jaeyoung Yoo, Hojun Lee, Inseop Chung, Geonseok Seo, and Nojun Kwak. Training multi-object detector by estimating bounding box distribution for input image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3437–3446, 2021. 3

[41] Jaeyoung Yoo, Hojun Lee, Seunghyeon Seo, Inseop Chung, and Nojun Kwak. Sparse mdod: Training end-to-end multi-object detector without bipartite matching. *arXiv preprint arXiv:2205.08714*, 2022. 3

[42] 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, 2, 5, 7, 8

[43] 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