

# Improving Neural Radiance Fields with Depth-aware Optimization for Novel View Synthesis

Shu Chen, Junyao Li, Yang Zhang, and Beiji Zou

**Abstract**—With dense inputs, Neural Radiance Fields (NeRF) is able to render photo-realistic novel views under static conditions. Although the synthesis quality is excellent, existing NeRF-based methods fail to obtain moderate three-dimensional (3D) structures. The novel view synthesis quality drops dramatically given sparse input due to the implicitly reconstructed inaccurate 3D-scene structure. We propose SfMNeRF, a method to better synthesize novel views as well as reconstruct the 3D-scene geometry. SfMNeRF leverages the knowledge from the self-supervised depth estimation methods to constrain the 3D-scene geometry during view synthesis training. Specifically, SfMNeRF employs the epipolar, photometric consistency, depth smoothness, and position-of-matches constraints to explicitly reconstruct the 3D-scene structure. Through these explicit constraints and the implicit constraint from NeRF, our method improves the view synthesis as well as the 3D-scene geometry performance of NeRF at the same time. In addition, SfMNeRF synthesizes novel sub-pixels in which the ground truth is obtained by image interpolation. This strategy enables SfMNeRF to include more samples to improve generalization performance. Experiments on two public datasets demonstrate that SfMNeRF surpasses state-of-the-art approaches. Code is available at <https://github.com/XTU-PR-LAB/SfMNeRF>.

**Index Terms**—neural radiance field, novel view synthesis, optical flow.

## I. INTRODUCTION

NEURAL Radiance Fields [1] have shown very impressive results for novel-view synthesis. NeRF employ a continuous five-dimensional (5D) function to implicitly encode the three-dimensional structure and appearance of a specific scene which is represented by a training multi-layer perceptron (MLP), and the novel view of the scene is synthesized by volumetric rendering. NeRF explicitly constrain the synthesized pixels to close the ground truth in which the employed volumetric rendering enables NeRF to implicitly constrain the 3D-scene geometry. However, the emitted colors and volume densities are entangled in the NeRF so that the NeRF require dense input views to eliminate the geometric ambiguity. When the input views are sparse, NeRF are prone to finding a degenerate solution to the image reconstruction objective [2]. On the other hand, NeRF fail to synthesize good novel views in the scenes with large texture-less regions such as couches and floors because it is unlikely to acquire sufficient correspondences cross view in NeRF.

S. Chen, J. Li and Y. Zhang are with the School of Computer Science, School of Cyberspace Security, Xiangtan University, Xiangtan, China and Key Laboratory of Intelligent Computing & Information Processing, Ministry of Education, Xiangtan, China (e-mail: csu\_cs@163.com; 351083542@qq.com; 460086012.com).

B. Zou is with School of Computer Science and engineering, Central South University, Changsha, China (e-mail: bjzou@csu.edu.cn).

The poor performance of NeRF in the sparse inputs or the scenes containing many texture-less regions comes from the implicitly estimated inaccurate 3D shape. To overcome this problem, some works leverage depth priors to improve the novel-view synthesis performance of neural radiance fields. These depth data are commonly estimated by running a structure-from-motion (SfM) approach [3], [4] from the input images. DS-NeRF [5] adopts the sparse 3D points from a SfM as the supervision in the NeRF optimization. However, the depth priors from SfM are commonly sparse and noisy. To obtain dense depth priors, NerfingMVS [6] employs a monocular depth network with the sparse depth from SfM reconstruction as the supervision to obtain the depth priors, and then using the adapted depth priors to guide volume sampling in the optimization of NeRF. Roessle et al. [7] propose a similar approach, they adopt depth completion to convert these sparse points into dense depth maps, which are used to guide NeRF optimization. However, this kind of solution still treats depth estimation and view synthesis as two separate processes that cannot benefit from each other.

Inspired by the self-supervised depth estimation approaches [9], [10], we integrate the novel-view synthesis and depth prediction into a single end-to-end procedure so they can benefit from each other. The employed explicit depth estimation can compensate for the weakly implicit 3D-geometry constraint in NeRF. With the improvement of depth estimation, the quality of novel-view synthesis is enhanced accordingly. Specifically, we employ the epipolar geometry to eliminate the pixels' depth uncertainty by constraining the corresponding points in another image of one point in one image must to be on a ray called the epipolar line. Furthermore, we leverage a patch photometric consistent loss across multiple views in which the image warping is implemented to ensure the identical region in different views has the same appearance. To further reduce the 3D-shape ambiguity, the surface smooth constraint in scenes and the position-of-matches constraint that the identical feature in different views has the same world coordinates are implemented. At the same time, we implement sub-pixel rendering in which the 2D image coordinates of pixels to be synthesized are not integer vectors but float vectors that are sampled based on the coordinates of image pixels with (0, 1) offset. The colors of the sampled sub-pixels are obtained by bilinear interpolation on the images. This new sampling strategy enables SfMNeRF to include more samples to improve generalization performance.

To sum up, our main contributions include:

- 1) By employing the cross-view consistent patch-based photometric loss, SfMNeRF explicitly constrains the 3D-scene

geometry that reduces the geometric ambiguity to some extent in NeRF.

2) The employed epipolar, smooth and position-of-matches constraints enable SfMNeRF to be aware of the scene’s structure, so the quality of novel-view synthesis improves.

3) The implemented sub-pixel rendering in SfMNeRF improves the generalization ability of the NeRF.

## II. RELATED WORK

We briefly review depth-aware NeRF and self-supervised depth prediction in this section.

### A. Depth-aware NeRF

NeRF have achieved the best results for novel-view synthesis under static conditions. NeRF implicitly represent a scene as a continuous five-dimensional (5D) function by training a MLP and then use volume rendering to synthesize a novel view. Many NeRF’s variants have been proposed, such as fast training and inference [11], [12], modelling non-rigid scenes [13], scalable [14], unbounded scenes [15], [16], editable [17], [18], handling reflections [19] and generalization [20], [21]. Although NeRF can achieve better performance for novel-view synthesis under rich texture and dense image inputs, they synthesize unplausible results due to the inherited geometric ambiguity when the scene is observed by sparse. Some works enforce regularization to improve the novel-view synthesis performance when the input is sparse, for example, by penalizing a semantic consistency loss [22] or introducing extra unobserved viewpoints [23]. However, existing approaches either regularize the depth with estimated sparse 3D point clouds from a SfM, or heavily rely on a extra multi-view dataset that might hard to collect or not be available. In contrast, our approach introduces more samples by sub-pixel rendering to enhance the generalization ability of the NeRF. Depth priors are also employed as the supervision in the NeRF optimization to guarantee a unique estimated 3D-scene geometry to handle the artifacts in sparse input scenarios. Here, we review the depth-aware Neural Radiance Fields in both scene-level and cross-scenes-level.

Depth-aware NeRF are expected to learn an accurate 3D shape to improve the novel-view synthesis performance [5], [6], [7], [24]. DS-NeRF [5] employs the sparse depth information from COLMAP [8] as the supervision to optimize the depth value of the pixel rendered. NerfingMVS [6] proposes a similar approach, but the depth priors are obtained by finetuning a monocular depth neural network on its sparse SfM reconstruction from the target scene. Dense Depth Priors [7] uses a depth completion network to convert the sparse depth data of each view obtained from the SfM into dense depth maps individually, which are used to guide NeRF optimization and supervise the depth of the pixel rendered. However, Dense Depth Priors process each view individually that it is not view-consistent. The depth priors in these approaches are commonly pre-obtained by an independent approach and the view synthesis and depth estimation are separated so that they cannot benefit from each other. In comparison, our approach

estimates the 3D-structure of the scene and synthesizes the novel view at the same time.

The per-scene NeRF optimize the representation of each scene individually, so it is time consuming and lack of generalization. To resolve these shortcomings, prior works also present generalizable radiance-field-based methods [20], [24], [25]. pixelNeRF [20] introduces a fully convolutional architecture to extract the feature maps of input images which are conditioned on the NeRF to learn a scene prior by training across multiple scenes. However, the features in pixelNeRF are aggregated from a multi-view 2D image. Point-NeRF [21] leverages deep multi-view stereo (MVS) techniques to reconstruct point clouds which are represented as anchors for feature extraction.

### B. Self-supervised Depth Prediction

Self-supervised depth prediction approaches aim to predict depth directly from monocular images, by enforcing a photometric loss on corresponding stereo images [26] or on temporally adjacent frames [9]. This kind of approach typically synthesizes a new view by image warping to serve as the supervisory signal. To eliminate the adverse effects of dynamic objects or occlusions, a mask explained for the motion is introduced to ignore certain regions that do not satisfy the static scene assumption. Zhou et al. [9] trained two networks, one for mask and another for depth estimation, to reconstruct the depth of scene from monocular images. In practice, the predicted mask is inaccurate that introduces more errors during training. Instead of being learned from a network, Godard et al. [26] introduce an auto-mask to eliminate training pixels that violate camera motion assumptions. Other methods alleviate this problem by either modeling the motion of individual objects [27], [28], estimating dense 3D translation field [29], or filtering by instance segmentation [30]. Garg et al. [31] introduce photometric consistency constraint to recover depth from stereo pairs, which is further improved by left-right consistency constraints [26]. Furthermore, some works [32], [33] use both temporal and spatial photometric warp errors to train the model. Inspired by self-supervised learning methods, in this work, we employed the patch-based multi-view consistent photometric constraint to guarantee an accurate 3D-scene’s shape.

## III. PRELIMINARY

A neural radiance field is represented by a continuous 5D function  $f$  that maps a 3D coordinate  $\mathbf{x} = (x, y, z)$  and viewing direction  $(\theta, \varphi)$  to a volume density  $\sigma$  and an emitted color  $\mathbf{c}$ . The continuous function is implicitly parameterized by a multi-layer perceptron, and the weights of the multi-layer perceptron are optimized to synthesize the input images of a specific scene.

$$f_w : (\beta(\mathbf{x}), \beta(\mathbf{d})) \rightarrow (\mathbf{c}, \sigma), \quad (1)$$

where  $w$  denotes the network weights, and  $\beta$  represents a predefined positional encoding applied to  $\mathbf{x}$  and  $\mathbf{d}$ .

Given  $m$  training images and the corresponding camera poses, a photometric loss is leveraged to optimize the NeRF as

$$L = \frac{1}{m} \sum_{i=1}^m \left\| I_i - \hat{I}_i \right\|_2^2, \quad (2)$$

where  $I_i$  and  $\hat{I}_i$  are the ground-truth color of image  $i$  and the corresponding synthesized image by volume rendering, respectively.

For each pixel of  $\hat{I}_i$ , casting a ray  $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$ ,  $\mathbf{o} \in \mathbb{R}^3$ ,  $\mathbf{d} \in S^2$ ,  $t \in [t_n, t_f]$  from the camera center  $\mathbf{o}$  through the pixel along direction  $\mathbf{d}$ , and its color  $\hat{\mathbf{c}}_\theta$  is computed using alpha compositing:

$$\hat{\mathbf{c}}_\theta(\mathbf{r}) = \int_{t_n}^{t_f} T(t) \sigma_\theta(\mathbf{r}(t)) \mathbf{c}_\theta(\mathbf{r}(t), \mathbf{d}) dt, \quad (3)$$

where  $T(t) = \exp\left(-\int_{t_n}^t \sigma_\theta(\mathbf{r}(s)) ds\right)$ , and  $\sigma_\theta(\cdot)$  and  $\mathbf{c}_\theta(\cdot, \cdot)$  indicate the volume density and color prediction of the radiance field, respectively.

#### IV. SFMNeRF

##### A. Overview

Fig. 1 shows the pipeline of our SfmNeRF framework for view synthesis. Firstly, a set of three images was selected from the dataset. In the set, one image is assumed to be the reference image, and the other two images selected are those that have an overlapping region with the reference image. Secondly, we extracted scale-invariant feature transform (SIFT) features in each image and obtained the matched SIFT correspondences between them. Finally, the matched SIFT correspondences were input into the neural radiance field to obtain their 3D coordinates. We used the positions of matched features loss to guarantee an accurate geometric shape based on the fact that the identical features in different views have the same world coordinates. Furthermore, we employed the epipolar constraint to enforce the correspondence of one point in reference image must lie in the epipolar line in another image by optimizing the minimum of the difference between the 3D coordinates of the point and the epipolar points. In addition, we leveraged the patch-based multi-view consistent photometric constraint to further constrain the estimated 3D-scene structure. We randomly selected a patch from the reference image and added random offset between (0, 1) to the image coordinates of each pixel in the patch to obtain sub-pixels; then the coordinates of sub-pixels  $(x, y)$  and the direction  $\mathbf{d}$  were input into the neural radiance field which output the color and depth of each sub-pixel. The ground-truth of each sub-pixel was obtained by a bilinear interpolator from the reference image. The pixels in the patch of the reference image were warped into the other view to synthesize a new patch, and a photometric reconstruction loss between the patch in the reference image and the new synthesized patch was implemented to constrain the estimated depth of each sub-pixel from the neural radiance field. In addition, the depth smooth loss was employed to further constrain the estimated 3D-scene geometry.

##### B. Sub-pixel Rendering

To achieve sub-pixel rendering, we first sample a patch from the reference image at random; the sampling region is denoted as  $\Omega_p$ . For any point  $p = (x, y) \in \Omega_p$  in the patch, we add an offset to it as

$$p = (x + x_{off}, y + y_{off}) | x_{off} \in (0, 1), y_{off} \in (0, 1). \quad (4)$$

The ground-truth color of the sub-pixel is obtained by linearly interpolating the color values of the four-pixel neighbors (top-left, top-right, bottom-left, and bottom-right) as

$$\hat{I}(p) = \sum_{i \in \{top, bottom\}, j \in \{left, right\}} w^{ij} I(p^{ij}), \quad (5)$$

where  $w^{ij}$  is determined according to linearly proportional to the Euclidean distance between  $p$  and  $p^{ij}$ , and  $\sum_{i,j} w^{ij} = 1$ .

According to NeRF [1], the rendering loss is formulated as

$$L_{ren} = \frac{1}{m} \sum_{p_i \in \Omega_p} \left\| \hat{I}(p_i) - \tilde{I}(p_i) \right\|_2^2, \quad (6)$$

where  $\hat{I}(p_i)$  is the ground-truth color of point  $p_i$  and  $\tilde{I}(p_i)$  is the corresponding synthesized pixel by volume rendering.

##### C. Positions of Matched Features Constraint

As introduced in SaNeRF [34], we also leverage the 3D coordinates of sparse keypoints to guide the geometry optimization in the NeRF, based on the fact that the identical keypoints in different views have the same world coordinates.

The 3D coordinates of the SIFT features are formulated as a weighted sum of all the samples volume densities  $\sigma_j$  output from the NeRF along the ray, defined as

$$\mathbf{x}_s = \sum_{j=1}^{N_c} w_j (\mathbf{o}' + t_j \mathbf{d}'), w_j = T_j (1 - \exp(-\sigma_j \delta_j)), \quad (7)$$

where  $\delta_j = t_{j-1} - t_j$  represents the distance between two consecutive samples and  $N_c$  denotes the number of samples.

We denote the 3D coordinates of matched correspondences as  $\{[\mathbf{x}_r^k, \mathbf{x}_i^k, \mathbf{x}_j^k] | k = 1, \dots, m\}$ , and the positions-of-matched features loss is formulated as

$$L_{3D} = \frac{1}{nml} \sum_{i=1}^n \sum_{j=1}^m \sum_{k=1}^l \left( \left\| \mathbf{x}_r^k - \mathbf{x}_i^k \right\|_2^2 + \left\| \mathbf{x}_r^k - \mathbf{x}_j^k \right\|_2^2 + \left\| \mathbf{x}_i^k - \mathbf{x}_j^k \right\|_2^2 \right), \quad (8)$$

where  $\mathbf{x}_r^k$ ,  $\mathbf{x}_i^k$ , and  $\mathbf{x}_j^k$  are the estimated 3D coordinates of the matched SIFT features  $k$  in the reference image and another two images according to (7), respectively.

##### D. Epipolar Constraint

Given the pose between two images, we can calculate the fundamental matrix. According to epipolar geometry, given one point in the reference image, we can derive a ray called epipolar line from the estimated fundamental matrix, the corresponding points in another image must to be on the epipolar line as shown in the Figure 2(b). In this work, we use the epipolar constraint to regularize the predicted depth of points from NeRF.

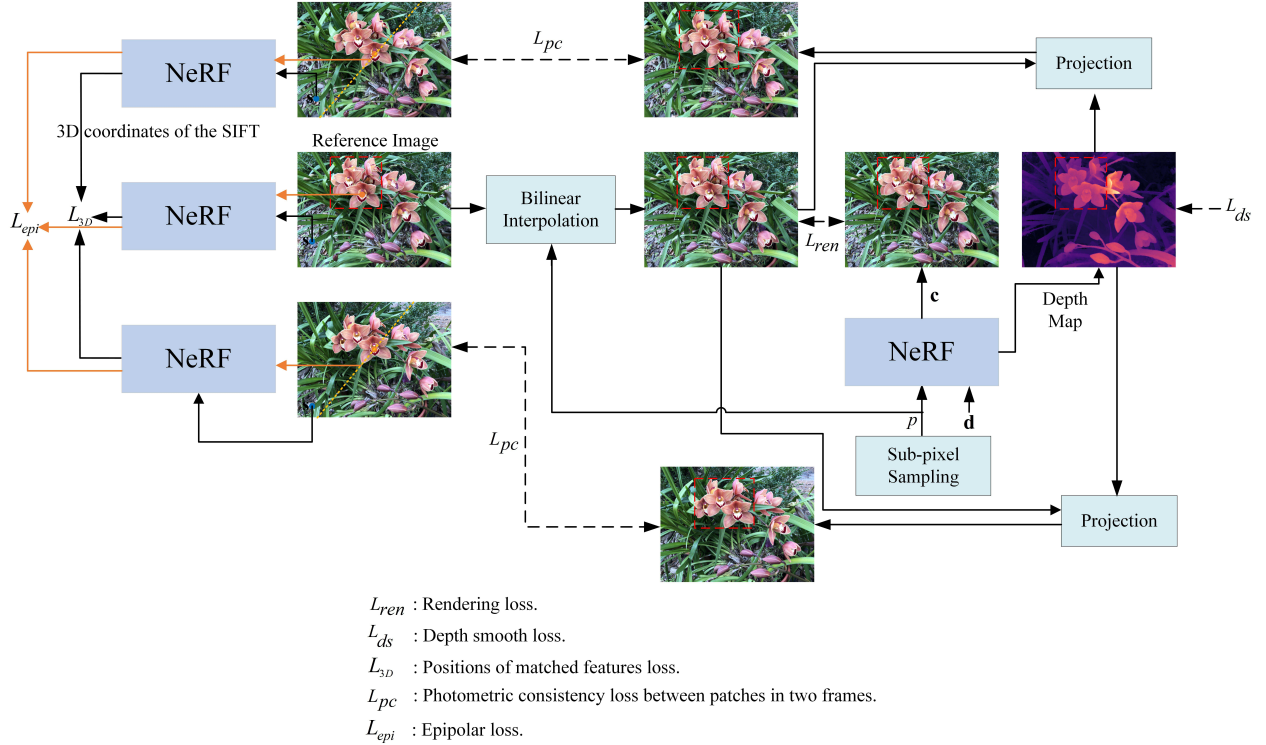


Fig. 1. Overview of SFMNeRF. The blue points in the other images are the matched features of the blue point in the reference image.

Given a point in the reference image, the 3D coordinates of this point is denoted as  $p_r$ , and the points in the epipolar line are denoted as  $\{p_i\}$ , we define the epipolar constraint as

$$L_{epi} = \min_i \left\{ \|p_r - p_i\|_2^2 \right\}, \quad (9)$$

where  $p_r$  and  $p_i$  are estimated according to (7).

Actually, the number of epipolar points is very large, it is necessary to eliminate the implausible points to improve the computational efficiency. On the other hand, some points may be occluded that it is not possible to find the corresponding points in the epipolar line. We use the color similarity to eliminate the implausible points, only the points in the epipolar line in which the color difference between them and the point in the reference image below a threshold are filtered as the epipolar points. We show the filtering visual results in the Figure 2.

### E. Photometric Consistency Constraint between Multiple Views

To eliminate the shape-radiance ambiguity in the NeRF [1], for each view, we explicitly leverage the patch-based photometric consistency constraint between multiple views to constrain the depth value of every pixel in the patch.

Firstly, the 3D coordinates of each pixel  $p$  in the patch are estimated by (7) and denoted as  $\mathbf{x}_p$ . Secondly, the pixel in the reference image is warped by a back-project process onto the other images (denoted as  $I_i$  and  $I_j$ , respectively) as the following equation, take image  $i$  for example.

$$\hat{p}_i = K T_{r \rightarrow i} \mathbf{x}_p, \quad (10)$$

where  $K$  is camera intrinsic matrix, and  $T_{r \rightarrow i} = [\mathbf{r}, \mathbf{t}]$  represents the relative pose between the reference image and the image  $i$  and defined as

$$\mathbf{r} = \mathbf{r}_i^{-1} \mathbf{r}_r, \mathbf{t} = \mathbf{r}_i^{-1} (\mathbf{t}_r - \mathbf{t}_i), \quad (11)$$

where  $[\mathbf{r}_i, \mathbf{t}_i]$ ,  $[\mathbf{r}_r, \mathbf{t}_r]$  are the global poses of image  $i$  and the reference image, respectively.

Similar to (5), the color of the projected point  $\hat{p}_i$  is obtained by a linear interpolator and denoted as  $\hat{I}(\hat{p}_i)$ .

Finally, the patch-based photometric consistency constraint is formulated as the combination of a photometric reconstruction loss  $L_{pr}$  and a structured similarity (SSIM) loss  $L_{SSIM}$ , and defined as  $L_{pc} = L_{pr} + L_{SSIM}$ .

The photometric reconstruction loss is defined as

$$L_{pr} = \sum_{p(k) \in \Omega_p} M(p(k)) \left| \hat{I}_r(p(k)) - \hat{I}_i(\hat{p}_i(k)) \right| + \sum_{p(k) \in \Omega_p} M(p(k)) \left| \hat{I}_r(p(k)) - \hat{I}_j(\hat{p}_j(k)) \right|, \quad (12)$$

where  $\hat{I}_r(p(k))$  is the pixel value of point  $k$  in the patch in the reference image, and  $\hat{I}_i(\hat{p}_i(k))$  is the synthesized pixel value of projected corresponding one in the image  $i$ .  $M(\cdot)$  represents the mask map for the patch to filter the pixels which do not have the correspondences in the other images. We pre-established a minimum rectangle to include all matched SIFT features in the reference image, if any point  $p(k)$  in the patch is in the minimum rectangle where its mask  $M(p(k))$  is set as one, otherwise set as zero.






Fig. 2. Implausible epipolar points elimination. (a) The reference image with a white point. (b) The epipolar line is depicted by the white line in another image and the white point is the corresponding point. (c) The obtained epipolar points after filtering which represented by the white points in the image.

The SSIM loss is defined as

$$L_{SSIM} = \frac{1}{2}M \left( \frac{1 - SSIM_{r \rightarrow i}}{2} + \frac{1 - SSIM_{r \rightarrow j}}{2} \right), \quad (13)$$

where  $SSIM_{r \rightarrow i}$  is the structured similarity between the patch in the reference image and the image  $i$ , and  $SSIM_{r \rightarrow j}$  is defined similarly.  $M$  is a binary mask that is set as one if the patch is in the minimum rectangle, otherwise it is set as zero.

#### F. Depth Smooth loss

Similar to [10], we enforce the depth smooth prior to the estimated depth maps in the NeRF [1]. Since large image gradients commonly mean depth discontinuities, we consider the edge constraint in our approach by enforcing the L1 norm of the depth gradients in the total loss, which is formulated as the weighted image gradients across adjacent pixels.

$$L_{ds} = \sum_{p(k) \in \Omega_p} \left| \nabla D(p(k)) \cdot \left( e^{-|\hat{I}(p(k))|} \right) \right|, \quad (14)$$

where  $\nabla$  represents the 2D differential operator and  $D(p(k))$  is the depth of pixel  $p(k)$ .

#### G. Training loss

The total loss is formulated as a combination of the aforementioned losses; each loss is controlled by a factor.

$$L_{total} = \lambda_{ren}L_{ren} + \lambda_{3D}L_{3D} + \lambda_{pr}L_{pr} + \lambda_{epi}L_{epi} + \lambda_{SSIM}L_{SSIM} + \lambda_{ds}L_{ds}. \quad (15)$$

## V. EXPERIMENT

### A. Experimental Details

We employed the commonly used deep learning library Pytorch to implement our approach. The performance of our system is evaluated on the LLFF-NeRF [1], [35] and ScanNet [36] datasets.

1) Datasets: The LLFF Dataset is established from eight scenes captured by a cellphone, with 20-62 images each. The resolution of each image in the dataset is  $4032 \times 3024$ . We downsized each image 1/8 scale to  $504 \times 378$  dimensions in

pixels due to the limited capacity of NVIDIA RTX 2080Ti, and held out 1/8 of these as the test set for novel view synthesis.

For the ScanNet dataset, we selected eight scenes in this dataset to evaluate our method as the experimental setup in [35]. In each scene, 40 images were selected to cover a local region, and all images were resized to  $648 \times 484$ . As introduced in NeRF [1], we held out 1/8 of these as the test set for novel view synthesis.

2) Training Details: We set  $\lambda_{ren} = 1.0$ ,  $\lambda_{3D} = 0.1$ ,  $\lambda_{pr} = 0.001$ ,  $\lambda_{epi} = 0.0001$ ,  $\lambda_{SSIM} = 0.01$  and  $\lambda_{ds} = 0.001$  in all the scenes of each dataset. For fair comparison, the same MLP architecture in NeRF [1] was employed in all experiments. The patch size was set as  $48 \times 48$ . The hierarchical sampling strategy in NeRF [1] was adopted and numbers of sampled points of both coarse sampling and importance sampling were set to 64. We use the Adam [37] optimizer to optimize our model with  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$  and the learning rate was set as  $1e-3$ . In all experiments, we used one NVIDIA RTX 2080Ti to training and testing.

3) Evaluation Metrics: We employed three kinds of metrics to evaluate the quality of novel-view rendering: Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index Measure (SSIM) [38] and Learned Perceptual Image Patch Similarity (LPIPS) [39].

### B. Comparison with State-of-the-Art

#### Results on LLFF-NeRF dataset.

We compared SfMNeRF to NeRF [1] and three recently proposed works: MVSNerF [41], TensorRF [40] and Plenoxels [12]. The quantitative comparisons for novel-view synthesis are shown in Table 1, and the visual results for novel-view synthesis and depth estimation are illustrated in Fig. 2. From Table 1, we notice that SfMNeRF achieved the best performance measured by PSNR in scenes (Fern, Leaves and Orchids) because the rich texture in these scenes is good for photometric consistency constraints between multiple views, and our approach achieved the best performance in six scenes on LPIPS.

#### Results on ScanNet dataset.

The quantitative comparisons for novel-view synthesis on the ScanNet dataset are shown in Table 2. From Table 2,

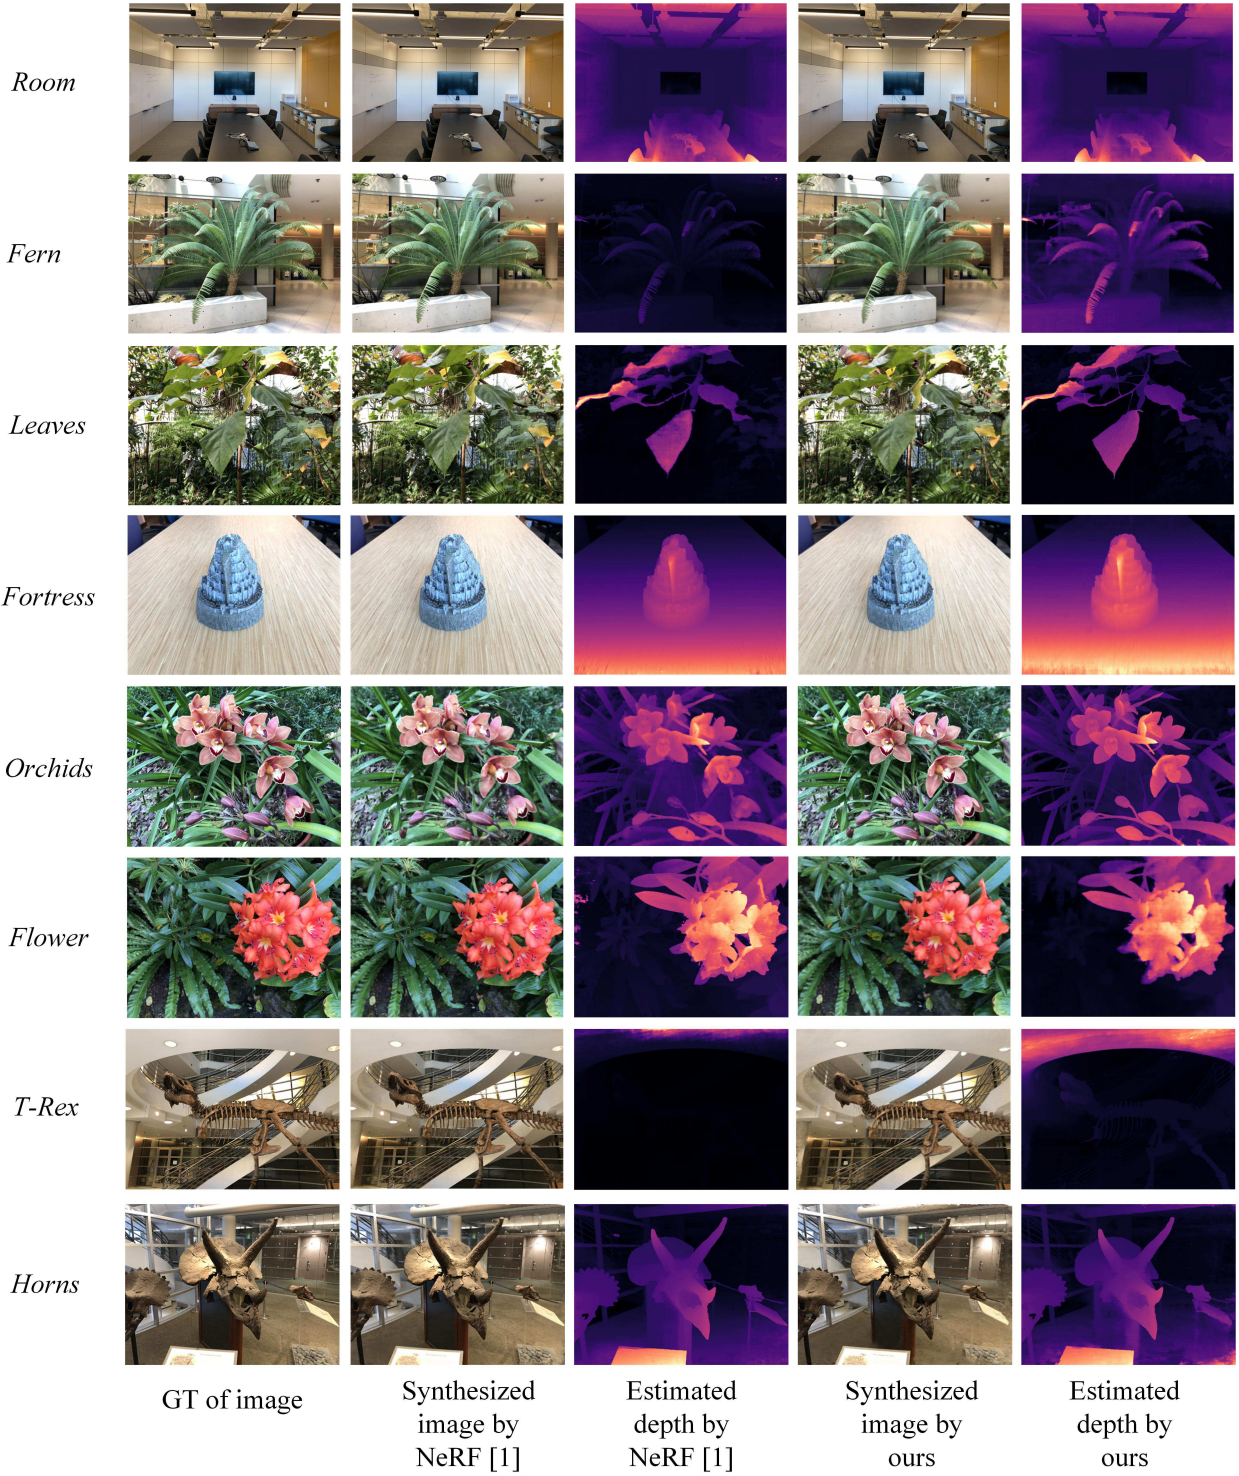


Fig. 3. Qualitative comparison between our SfMNeRF and other approaches on the LLFF-NeRF dataset.



Methods	Room	Fern	Leaves	Fortress	Orchids	Flower	T-Rex	Horns
<b>PSNR<math>\uparrow</math></b>								
NeRF [1] ECCV20	<b>32.70</b>	25.17	20.92	31.16	20.36	27.40	26.80	27.45
MVSNeRF [41] ICCV2021	26.95	21.15	17.51	26.03	17.85	24.74	23.20	23.57
TensoRF [40] ECCV22	32.35	25.27	21.30	<b>31.36</b>	19.87	<b>28.60</b>	<b>26.97</b>	<b>28.14</b>
Plenoxels [12] CVPR22	30.22	25.46	21.41	31.09	20.24	27.83	26.48	27.58
Ours	30.02	<b>25.79</b>	<b>21.46</b>	31.11	<b>20.81</b>	25.95	26.47	27.15
<b>SSIM<math>\uparrow</math></b>								
NeRF [1] ECCV20	0.948	0.792	0.690	0.881	0.641	0.827	0.880	0.828
MVSNeRF [41] ICCV2021	0.951	0.638	0.667	0.872	0.657	0.888	0.868	0.868
TensoRF [40] ECCV22	<b>0.952</b>	0.814	0.752	<b>0.897</b>	0.649	<b>0.871</b>	<b>0.900</b>	<b>0.877</b>
Plenoxels [12] CVPR22	0.937	0.832	0.760	0.885	0.687	0.862	0.890	0.857
Ours	0.953	<b>0.796</b>	0.730	0.881	<b>0.704</b>	0.791	0.895	0.840
<b>LPIS <math>v_{gg}\downarrow</math></b>								
NeRF [1] ECCV20	0.178	0.280	0.316	0.171	0.321	0.219	0.249	0.268
MVSNeRF [41] ICCV2021	0.172	0.238	0.313	0.208	0.274	0.196	0.184	0.237
TensoRF [40] ECCV22	0.167	0.237	0.217	0.148	0.278	<b>0.169</b>	0.221	0.196
Plenoxels [12] CVPR22	0.192	0.224	<b>0.198</b>	0.180	0.242	0.179	0.238	0.231
Ours	<b>0.049</b>	<b>0.219</b>	0.204	<b>0.072</b>	<b>0.160</b>	0.171	<b>0.097</b>	<b>0.113</b>

TABLE I  
QUANTITATIVE COMPARISONS FOR NOVEL-VIEW SYNTHESIS ON LLFF-NeRF DATASET.

we observe that the performance of SfMNeRF was inferior to some supervised approaches, because the estimated geometric shape of the scene in those supervised approaches served as the supervision that constrained the NeRF optimization. Benefitting from the geometric constraints within the multi-views, our approach outperformed the original NeRF by a large margin. Fig. 3 illustrates the qualitative results about novel-view synthesis and depth estimation. SfMNeRF successfully recovered the accurate 3D-scene representation. In contrast, NeRF [1] failed to reconstruct the corrected 3D-scene structure, which in turn deteriorated the view synthesis.

### C. Ablation Study

We conducted extensive ablation studies on the ScanNet dataset to validate the effectiveness of the individual components in SfMNeRF for view synthesis. Table 3 shows the quantitative results. We explain items in Table 3 as follows:

*Basic*: Only the rendering loss is used.

*3D*: The positions-of-matched-features loss is included.

*ds*: The depth smooth loss is included.

*pr*: The photometric reconstruction loss between patches in different views is included.

*SSIM*: The SSIM loss between patches in different views is included.

*sub-pixel sampling*: The sub-pixel sampling strategy is implemented during training.

*epi*: The epipolar constraint is included.

From Table 3, we observe that each loss employed in SfMNeRF is able to improve novel-view synthesis accuracy, especially, the sub-pixel sampling strategy improved the results significantly.

## VI. CONCLUSION AND FUTURE WORK

In this paper, we analyzed the limitation of NeRF and proposed SfMNeRF, a neural radiance field is able to improve the quality of novel view synthesis by self-supervised

depth constraints. The employed depth priors are obtained by some constraints, such as patch-based photometric consistency constraint between multiple views, epipolar constraint and positions-of-matched-features constraint, without additional data. In this way, our approach learns a multi-view consistent geometry with depth constraints. The depth priors can eliminate the geometric ambiguity to some extent and improve the quality of novel-view synthesis accordingly. The employed sub-pixel sampling strategy introduces more samples in training that further implicitly constrain the 3D-scene geometry in NeRF. By employing explicit and implicit depth constraints, our approach improves the novel-view rendering quality of NeRF. In terms of comparison to other depth priors based NeRF, our approach does not acquire to estimate the sparse depth map by a SfM in advance. Our approach motivates the future research to further exploit the structural priors in multi-view inputs for view synthesis and other related tasks.

As with other photometry-based reconstruction methods, SfMNeRF suffers from the scenes with repeated structures which are likely to be caused by photometric ambiguity. In the future, we will investigate how to incorporate the rich priors of indoor datasets, and adapt the proposed approach to achieve a generalized NeRF trained on across large scale datasets.

### ACKNOWLEDGMENT

This research is supported in part by the National Key R&D Program of China (No. 2018AAA0102102).

### REFERENCES

- [1] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, "Nerf: Representing scenes as neural radiance fields for view synthesis," in Proc. Eur. Conf. Comput. Vis., Sep. 2020, pp. 405-421.
- [2] K. Zhang, G. Riegler, N. Snaveley, and V. Koltun, "NeRF++: Analyzing and improving neural radiance fields," arXiv, 2020.
- [3] Z. Zhou, F. Shi, J. Xiao, and W. Wu, "Non-rigid structure-from-motion on degenerate deformations with low-rank shape deformation model," IEEE Transactions on Multimedia, vol. 17, no. 2, pp. 171-185, 2015.

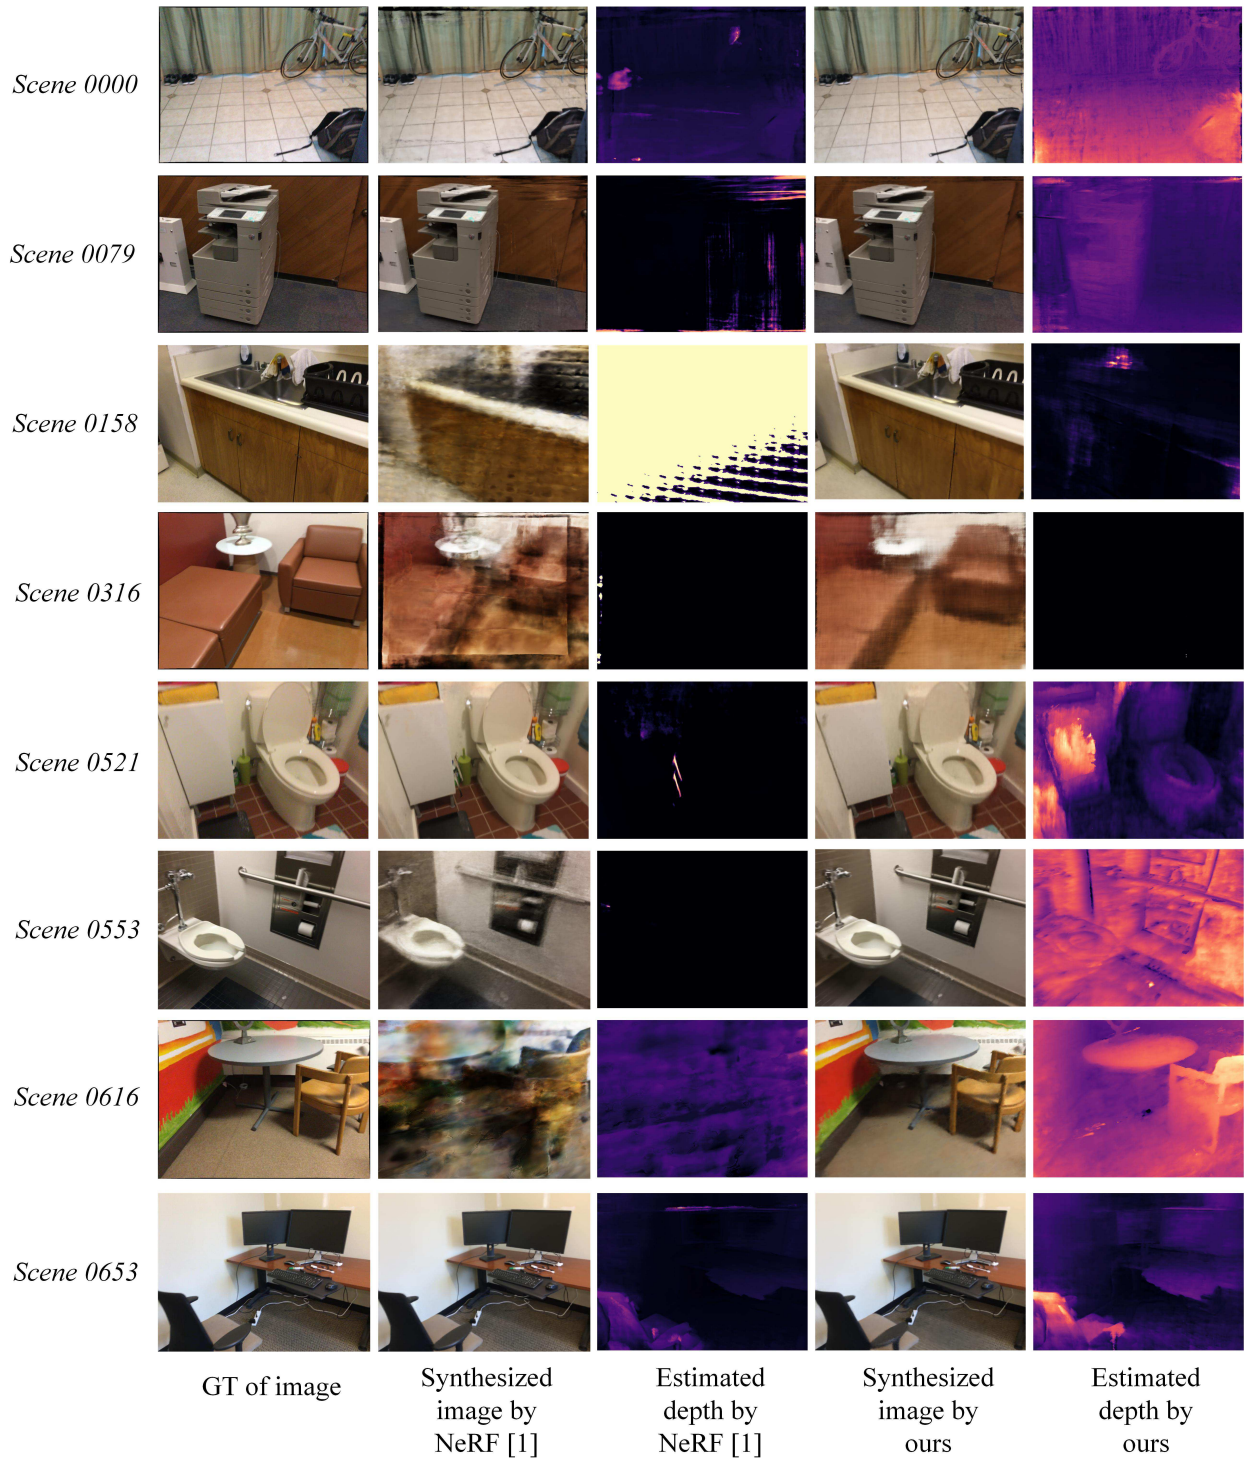


Fig. 4. Qualitative comparison between our SfMNeRF and other approaches on the ScanNet dataset.



Methods	Scene 0000	Scene 0079	Scene 0158	Scene 0316	Scene 0521	Scene 0553	Scene 0616	Scene 0653
<b>PSNR<math>\uparrow</math></b>								
NSVF [42] NeurIPS20	<b>23.36</b>	26.88	<b>31.98</b>	<b>22.29</b>	27.73	31.15	15.71	28.95
SVS [43] CVPR21	21.39	25.18	29.43	20.63	27.97	30.95	<b>21.38</b>	27.91
NeRF [1] ECCV20	18.75	25.48	29.19	17.09	24.41	30.76	15.76	30.89
NerfingMVS [6] ICCV21	22.10	<b>27.27</b>	30.55	20.88	<b>28.07</b>	<b>32.56</b>	18.07	31.43
Ours	20.97	25.53	31.45	16.11	27.25	30.92	20.76	<b>31.64</b>
<b>SSIM<math>\uparrow</math></b>								
NSVF [42] NeurIPS20	0.823	0.887	0.951	0.917	0.892	0.947	0.704	0.929
SVS [43] CVPR21	<b>0.914</b>	<b>0.923</b>	<b>0.953</b>	<b>0.941</b>	<b>0.924</b>	<b>0.968</b>	<b>0.899</b>	0.965
NeRF [1] ECCV20	0.751	0.896	0.928	0.828	0.871	0.950	0.699	0.953
NerfingMVS [6] ICCV21	0.880	0.916	0.948	0.899	0.901	0.965	0.748	<b>0.964</b>
Ours	0.750	0.832	0.945	0.715	0.798	0.892	0.749	0.926

TABLE II  
QUANTITATIVE COMPARISONS FOR NOVEL-VIEW SYNTHESIS ON SCANNET DATASET.

Methods	PSNR $\uparrow$	SSIM $\uparrow$
<i>Basic</i>	20.04	0.860
<i>Basic+3D</i>	22.32	0.848
<i>Basic+3D+ds</i>	22.35	0.815
<i>Basic+3D+ds+pr</i>	23.16	0.825
<i>Basic+3D+ds+pr+SSIM</i>	23.47	0.832
<i>Basic+3D+ds+pr+SSIM+sub-pixel sampling</i>	25.36	0.814
<i>Basic+3D+ds+pr+SSIM+sub-pixel sampling+epi</i>	25.58	0.826

TABLE III  
ABLATION STUDY ON SFMNeRF. THE EVALUATION IS PERFORMED ON SCANNET DATASET.

- [4] H. Zhou, X. Li, and A. H. Sadka, "Nonrigid structure-from-motion from 2-D images using markov chain monte carlo," *IEEE Transactions on Multimedia*, vol. 14, no. 1, pp. 168-177, 2011.
- [5] K. Deng, A. Liu, J. Zhu, and D. Ramanan, "Depth-supervised nerf: Fewer views and faster training for free," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2022, pp. 12872-12881.
- [6] Y. Wei, S. Liu, Y. Rao, W. Zhao, J. Lu, and J. Zhou, "NerfingMVS: Guided optimization of neural radiance fields for indoor multi-view stereo," in *Proc. Int. Conf. Comput. Vis.*, Oct. 2021, pp. 5590-5599.
- [7] B. Roessle, J. T. Barron, B. Mildenhall, P. P. Srinivasan, and M. Niebner, "Dense depth priors for neural radiance fields from sparse input views," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2022, pp. 12882-12891.
- [8] J. L. Schonberger and J. Frahm, "Structure from motion revisited," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jul. 2016, pp. 4104-4113.
- [9] T. Zhou, M. Brown, N. Snavely, and D. G. Lowe, "Unsupervised learning of depth and ego-motion from video," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2017, pp. 6612-6621.
- [10] S. Chen, Z. Pu, X. Fan, and B. Zou, "Fixing defect of photometric loss for self-supervised monocular depth estimation," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 32, no. 3, pp. 2043-2055, 2022.
- [11] T. Muller, A. Evans, C. Schied, and A. Keller, "Instant neural graphics primitives with a multiresolution hash encoding," *ACM Trans. Graphics*, vol. 41, no. 4, pp. 1-15, 2022.
- [12] A. Yu, S. Fridovich-Keil, M. Tancik, and Q. Chen, "Plenoxels: radiance fields without Neural networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2022, pp. 5491-5500.
- [13] A. Pumarola, E. Corona, G. Pons-Moll, and F. Moreno-Noguer, "D-NeRF: neural radiance fields for dynamic scenes," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2021, pp. 10313-10322.
- [14] M. Tancik, V. Casser, X. Yan, and S. Pradhan, "Block-NeRF scalable large scene neural view synthesis," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2022, pp. 8238-8248.
- [15] K. Zhang, G. Riegler, N. Snavely, and V. Koltun, "NeRF++: Analyzing and improving neural radiance fields," *arXiv*, 2020.
- [16] J. T. Barron, B. Mildenhall, D. Verbin, P. P. Srinivasan, and P. Hedman, "Mip-NeRF 360: unbounded anti-aliased neural radiance fields," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2022, pp. 5460-5469.
- [17] B. Yang, C. Bao, J. Zeng, H. Bao, Y. Zhang, Z. Cui, and G. Zhang, "NeuMesh: learning disentangled neural mesh-based implicit field for geometry and texture editing," in *Proc. Eur. Conf. Comput. Vis.*, Oct. 2022, pp. 597-614.
- [18] S. Liu, X. Zhang, Z. Zhang, R. Zhang, J. Zhu and B. Russell, "Editing conditional radiance fields," in *Proc. Int. Conf. Comput. Vis.*, Oct. 2021, pp. 5753-5763.
- [19] D. Verbin, P. Hedman, B. Mildenhall, T. Zickler, J. T. Barron, and P. P. Srinivasan, "Ref-nerf: Structured view-dependent appearance for neural radiance fields," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2022, pp. 5481-5490.
- [20] A. Yu, V. Ye, M. Tancik, and A. Kanazawa, "pixelNeRF Neural Radiance Fields from One or Few Images," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Nov. 2021, pp. 4576-4585.
- [21] Q. Xu, Z. Xu, J. Philip, S. Bi, Z. Shu, K. Sunkavalli, and U. Neumann, "Point-NeRF: point-based neural radiance fields," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2022, pp. 5428-5438.
- [22] A. Jain, M. Tancik, and P. Abbeel, "Putting nerf on a diet: Semantically consistent few-shot view synthesis," in *Proc. Int. Conf. Comput. Vis.*, 2021, pp. 5865-5874.
- [23] M. Niemeyer, J. T. Barron, B. Mildenhall, M. S. M. Sajjadi, A. Geiger and N. Radwan, "RegNeRF: regularizing neural radiance fields for view synthesis from sparse inputs," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2022, pp. 5470-5480.
- [24] Q. Wang, Z. Wang, K. Genova, P. Srinivasan, H. Zhou, J. T. Barron, R. Martin-Brualla, N. Snavely, and T. Funkhouser, "Ibrnet: Learning multi-view image-based rendering," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2021, pp. 4688-4697.
- [25] A. Trevithick and B. Ynag, "GRF: Learning a general radiance field for 3D representation and rendering," in *Proc. Int. Conf. Comput. Vis.*, 2021, pp. 15162-1572.
- [26] C. Godard, O. M. Aodha, and G. J. Brostow, "Unsupervised monocular depth estimation with left-right consistency," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Nov. 2017, pp. 6602-6611.
- [27] V. Casser, S. Pirk, R. Mahjourian, and A. Angelova, "Depth prediction without the sensors: Leveraging structure for unsupervised learning from monocular videos," in *Proc. the Thirty-Third AAAI Conference on Artificial Intelligence*, Jan. 2019, pp. 8001-8008.
- [28] Z. Yin and J. Shi, "GeoNet: unsupervised learning of dense depth,

- optical flow and camera pose,” in Proc. IEEE Conf. Comput. Vis. Pattern Recognit., Jun. 2018, pp. 1983-1992.
- [29] H. Li, A. Gordon, H. Zhao, V. Casser, and A. Angelova, “Unsupervised monocular depth learning in dynamic scenes,” arXiv preprint arXiv: 2010.16404, 2020.
- [30] S. Lee, S. Im, S. Lin, and I. S. Kweon, “Instance-wise Depth and motion learning from Monocular video,” arXiv preprint arXiv: 1912.09351, 2020.
- [31] R. Garg, V. K. BG, G. Carneiro, and I. Reid, “Unsupervised CNN for single view depth estimation: Geometry to the rescue,” in Proc. Eur. Conf. Comput. Vis., Oct. 2016, pp. 740-756.
- [32] R. Li, S. Wang, Z. Long, and D. Gu, “UnDeepVO: monocular visual odometry through unsupervised deep learning,” in Proc. IEEE Conf. on Robotics and Automation, May 2018, pp. 7286-7291.
- [33] H. Zhan, R. Garg, C. S. Weerasekera, K. Li, H. Agarwal, and I. Reid, “Unsupervised learning of monocular depth estimation and visual odometry with deep feature reconstruction,” in Proc. IEEE Conf. Comput. Vis. Pattern Recognit., Jun. 2018, pp. 340-349.
- [34] S. Chen, Y. Zhang, Y. Xu, and B. Zou, “Structure-aware NeRF without Posed Camera via Epipolar Constraint,” arXiv preprint arXiv: 2210.00183, 2022.
- [35] B. Mildenhall, P. P. Srinivasan, R. Ortiz-Cayon, N. K. Kalantari, R. Ramamoorthi, R. Ng, and A. Kar, “Local light Field fusion: Practical view synthesis with prescriptive sampling guidelines,” ACM Transactions on Graphics, vol. 38, no. 4, pp. 1-14, Aug. 2019.
- [36] A. Dai, A. X. Chang, M. Savva, M. Halber, T. Funkhouser, and M. Niesner, “ScanNet: Richly-annotated 3d reconstructions of indoor scenes,” in Proc. IEEE Conf. Comput. Vis. Pattern Recognit., Jun. 2017, pp. 5828-5839.
- [37] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” arXiv preprint arXiv: 1412.6980, 2014.
- [38] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, “Image quality assessment: from error visibility to structural similarity,” IEEE Transactions on Image Processing, vol. 13, no. 4, pp. 600-612, April. 2004.
- [39] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, “The unreasonable effectiveness of deep features as a perceptual metric,” in Proc. IEEE Conf. Comput. Vis. Pattern Recognit., Jun. 2018, pp. 586-595.
- [40] A. Chen, Z. Xu, A. Geiger, J. Yu, and H. Su, “TensorRF: Tensorial radiance fields,” in Proc. Eur. Conf. Comput. Vis., Sep. 2022, pp. 1-30.
- [41] A. Chen, Z. Xu, F. Zhao, X. Zhang, F. Xiang, J. Yu, and O. Wang, “MVSNeRF: fast generalizable radiance field reconstruction from multi-view stereo,” in Proc. Int. Conf. Comput. Vis, 2021, pp. 14104-14113.
- [42] L. Liu, J. Gu, K. Lin, T. Chua, and C. Theobalt, “Neural sparse voxel fields,” in Proc. Adv. Neural Inf. Process. Syst., Dec. 2020, pp. 15651-15663.
- [43] G. Riegler and V. Koltun, “Stable View Synthesis,” in Proc. IEEE Conf. Comput. Vis. Pattern Recognit., Jun. 2021, pp. 12211-12220.