

SurfelNeRF: Neural Surfel Radiance Fields for Online Photorealistic Reconstruction of Indoor Scenes

Yiming Gao Yan-Pei Cao Ying Shan
 ARC Lab, Tencent PCG

gaoyim9@mail3.sysu.edu.cn caoyanpei@gmail.com yingsshan@tencent.com

Abstract

Online reconstructing and rendering of large-scale indoor scenes is a long-standing challenge. SLAM-based methods can reconstruct 3D scene geometry progressively in real time but can not render photorealistic results. While NeRF-based methods produce promising novel view synthesis results, their long offline optimization time and lack of geometric constraints pose challenges to efficiently handling online input. Inspired by the complementary advantages of classical 3D reconstruction and NeRF, we thus investigate marrying explicit geometric representation with NeRF rendering to achieve efficient online reconstruction and high-quality rendering. We introduce SurfelNeRF, a variant of neural radiance field which employs a flexible and scalable neural surfel representation to store geometric attributes and extracted appearance features from input images. We further extend the conventional surfel-based fusion scheme to progressively integrate incoming input frames into the reconstructed global neural scene representation. In addition, we propose a highly-efficient differentiable rasterization scheme for rendering neural surfel radiance fields, which helps SurfelNeRF achieve $10\times$ speedups in training and inference time, respectively. Experimental results show that our method achieves the state-of-the-art 23.82 PSNR and 29.58 PSNR on ScanNet in feedforward inference and per-scene optimization settings, respectively.¹

1. Introduction

Large-scale scene reconstruction and rendering is a crucial but challenging task in computer vision and graphics with many applications. Classical visual simultaneous localization and mapping (SLAM) systems [6, 12, 16, 24, 39, 41] can perform real-time 3D scene reconstruction. However, they usually represent the scene geometry as solid surfaces and appearance as vertex color or texture maps; thus,

¹Project website: <https://gymat.github.io/SurfelNeRF-web>

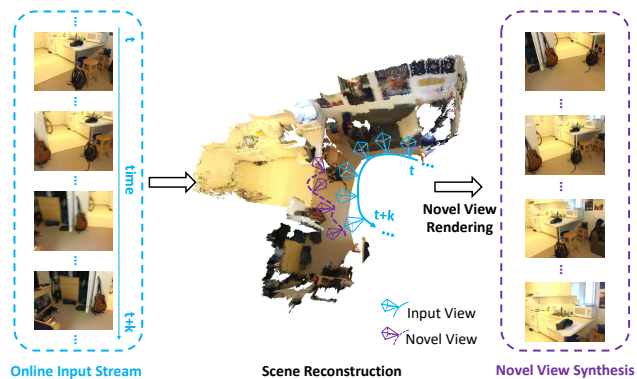


Figure 1. Examples to illustrate the task of online photorealistic reconstruction of an indoor scene. The *online photorealistic* reconstruction of large-scale indoor scenes: given an online input image stream of a previously unseen scene, the goal is to progressively build and update a scene representation that allows for high-quality rendering from novel views.

the reconstructed results fail to fully capture the scene content and cannot be used for photorealistic rendering. Recently, neural radiance fields (NeRF) and its variants [2, 3, 22, 27, 35, 42] have achieved unprecedented novel view synthesis quality on both object-centric and large-scale scenes. However, NeRFs suffer from long per-scene optimization time and slow rendering speed, especially for large-scale scenes. Although recent advances [7, 14, 19, 23, 33] achieve faster optimization and rendering via incorporating explicit representations, they still require gathering all input images in an *offline* fashion before optimizing each scene.

In this paper, we target the challenging task of *online photorealistic* reconstruction of large-scale indoor scenes: given an online input image stream of a previously unseen scene, the goal is to progressively build and update a scene representation that allows for high-quality rendering from novel views. The online setting can unlock a variety of real-time interactive applications, providing crucial immediate feedback to users during 3D capture. However, this task brings multiple extra requirements, including the scalabil-

Methods	Representation	Generalization	Real-time Rendering	Online Fusion	Scalability
DeepSurfel [21]	Surfels	✗	✓	✗	✗
Instant-NGP [23]	Hash Grids	✗	✓	✗	✗
PointNeRF [43]	Point Clouds	✓	✗	✗	✓
VBA [10]	B+ Trees	✗	✗	✓	✓
NeRFusion [50]	Voxel Grids	✓	✗	✓	✗
Ours	Surfels	✓	✓	✓	✓

Table 1. Comparison of representation and features with existing methods.

ity of the underlying scene representation, the ability to perform on-the-fly updates to the scene representation, and optimizing and rendering at interactive framerates. Recently, NeRFusion [50] followed NeuralRecon [34] to unproject input images into local sparse feature volumes, fusing them to a global volume via Gated Recurrent Units (GRUs), and then generating photorealistic results from the global feature volume via volume rendering. However, updating the sparse volumetric feature involves computationally heavy operations; the volume rendering is also very slow since it requires *hundreds* of MLP evaluations to render a pixel. Thus, although NeRFusion achieves efficient online scene reconstruction, it still needs dozens of seconds to render a frame. VBA [10] is another recent approach to online photorealistic scene capture, but it only applies to object-centric scenes. We compare with representation and key features used in online photorealistic rendering with existing methods, which is shown in Tab. 1.

We propose surfel-based neural radiance fields, SurfNeRF, for online photorealistic reconstruction and rendering of large-scale indoor scenes. Surfels (**surface elements**) [25] are point primitives containing geometric attributes, e.g., position, color, normal, and radius. We extend this representation to *neural surfels*, storing extra neural features that encode the neural radiance field of the target scene. Compared with volumetric representations, neural surfels are more compact and flexible and can easily scale to large scenes. Besides, we further employ a fast and differentiable rasterization process to render neural surfel radiance fields, which produces a pixel with only *a few* MLP evaluations based on the rasterized neural surfels.

Inspired by classical real-time surfel-based geometric reconstruction methods [6, 15, 41], we propose an efficient neural radiance field fusion method to progressively build the scene representation by integrating neighboring neural surfels. Unlike point-based representations [1, 26, 30, 43] that are computationally heavy when finding neighboring points, it is easier to locate overlapping surfels and then merge neural features from multiview observations. By coupling the SurfNeRF representation, the efficient neural surfel fusion approach, and the fast neural surfel rasterization algorithm, we achieve high-quality, photorealistic 3D scene reconstruction in an online manner.

We conduct experiments on the large-scale indoor scene

dataset ScanNet [11], which contains complex scene structures and a large variety of scene appearances. We train the SurfNeRF end-to-end across the scenes on the ScanNet, obtaining a generalizable model that enables both feedforward inference on unseen data and per-scene fine-tuning. We demonstrate in experiments that the proposed SurfNeRF achieves favorably better rendering quality than the state-of-the-art approaches in both feedforward and fine-tuning settings while maintaining high training and rendering efficiency. We believe the proposed online photorealistic reconstruction framework has great potential in practical applications.

2. Related Work

Novel view synthesis. Novel view synthesis is a challenging task for computer vision and computer graphics. Previous methods utilize geometry-based view synthesis. Zhou *et al.* [51] proposes to use the multiplane image (MPI) representation for novel view synthesis. Recently, neural radiance fields [22] and neural-based geometry representation have attracted more interest due to high-quality rendering. Recent methods propose to associate neural features on geometry structure and then synthesize novel views via rendering rasterized features, such as point cloud [1, 30], texture [36] and surfel [21]. DeepSurfel performs online appearance fusion only on a fixed given geometry during online fusion, while our SurfNeRF performs online *appearance and geometry* integration. To render outputs, DeepSurfel rasterizes surfels and then passes them through 2D CNN filters, while our SurfNeRF performs neural volume rendering which could achieve higher-quality rendering. Some recent works [9, 13, 18, 45] focus on real-time neural rendering methods. However, these methods all require per-scene optimization, which makes them impractical to handle novel scenes. In contrast, our SurfNeRF is generalizable, which can reconstruct new scenes via direct feedforward inference.

Neural Radiance Fields. Recently, neural radiance fields (NeRF) methods [2, 22, 42] achieve promising results on novel view synthesis, but it requires expensive per-scene optimization. Some methods proposed to reduce the long per-scene optimization time via deferred neural rendering [13, 47], hash table [23] or geometry prior [20]. Recent,

generalization NeRF methods [5, 8, 37, 43, 46] are proposed to produce novel view without per-scene optimization by learning a radiance fields representation from offline given images of novel scenes. PixelNeRF [46] and IBRNet [37] render novel view images using volume rendering based on corresponding warped features from nearby given images. MVNeRF [8] reconstructs local neural volumes from local nearby views, and then conducts volume rendering based on volumes. FWD [5] renders novel views by warping all input images to the target view, and fusing and refining features of these warped inputs. However, these methods are generalizable on local small-scale scenes, and require taking all input images before rendering novel views in an offline fashion. PointNeRF [43] learns to reconstruct a point-based neural radiance field for the novel scene. However, point-based representation is redundant, especially on large-scale scenes, and the point clouds cannot be updated based on multi-view input. Unlike these works, our method can represent not only large-scale scenes but also update compact surfel representation in an online fashion.

Online scene reconstruction and rendering. Most online scene reconstruction methods [4, 34, 41, 52, 53] focus on 3D reconstruction only via TSDF-based fusion [4, 34], surfel-based fusion [41] and SLAM-based reconstruction [52]. However, they cannot render photorealistic novel view synthesis. Recently, some methods [10, 32, 50] employ radiance fields to tackle online scene representation for high-quality rendering. VBA [10] builds a neural dynamic B+Tree to represent scenes and apply volume rendering on it, but it still requires per-scene optimization. SRT [32] represents scenes with set-latent with a transformer encoder and decoder, but it cannot synthesize high-resolution images, which makes it not ready for real-world applications. NeRFusion [50] represents large-scale scenes with neural voxel grids, update them via GRU in an online fashion, and renders results via slow volume rendering. However, grid-based representation is wasteful for large-scale scenes, and is inefficient during features update due to features in empty areas are also updated. In contrast, our surfel representation is compact for large-scale scenes, and surfel-based fusion and corresponding rendering are more efficient.

3. Method

Unlike most existing methods [22, 23, 43] that build neural radiance fields in an offline fashion, we aim for the task of online photorealistic scene reconstruction. Given a sequence of images $\{I_n\}$, the goal is to progressively build the neural scene representation on the fly as the input images arrive.

To achieve this goal, we present SurfNeRF, an online neural reconstruction framework based on the neural surfel radiance fields representation. The insight of our method is that the classical explicit surfel representation is efficient

for both reconstruction and rendering, has compact memory footprints, has the scalability to large and complex scenes, and can be naturally extended to equip neural attributes. Meanwhile, one can integrate multi-frame information and perform incremental updates on neural surfels in a *geometrically meaningful* way, and the operations involved can be streamlined.

Figure 2 gives an overview of our proposed method. The framework first learns to generate per-view local neural surfel radiance fields \mathcal{S}_t for each input frame I_t . We then adopt the conventional surfel fusion process [6, 15, 41] to the neural setting, parameterizing the neural surfel fusion procedural as a neural network N_f . Through N_f , per-view representations $\{\mathcal{S}_t\}$ can be integrated into a global neural surfel radiance field \mathcal{S}_g on the fly, aggregating geometric and neural attributes. The neural surfel radiance fields can be efficiently rendered to photorealistic images using a differentiable rasterizer without extra optimization. The framework can be trained end-to-end on large-scale indoor scene datasets. Once trained, it can generalize to previously unseen data and perform online photorealistic reconstruction; if needed, the rendering quality can be further improved by a short period time of fine-tuning.

3.1. Neural Surfel Radiance Field Construction

Surfel Representation. Adopted from classical 3D reconstruction research, each surfel s^n is equipped with the following geometric attributes: a position $\mathbf{p}^n \in \mathbb{R}^3$, normal $\mathbf{n}^n \in \mathbb{R}^3$, weight $w^n \in \mathbb{R}$ and radius $r^n \in \mathbb{R}$. The weight w^n measures the confidence of a surfel based on the distance of the current depth measurement from the camera center. The radius r^n of each surfel is designed to represent the local support area around the projected 3D position while minimizing visible holes. The weight and radius of each surfel are initialized following [41]. Besides geometric properties, we further associate each surfel with neural features $\mathbf{f}^n \in \mathbb{R}^N$ from input images for modeling semi-transparent structures and view-dependent effects. Although a similar representation has been explored in DeepSurfels [21], we construct neural surfels in a feedforward manner and incorporate them in progressive scene reconstruction.

Local Neural Surfel Radiance Field. Given an input frame I_t , camera parameters \mathbf{c}_t , and estimated depth D_t , the corresponding *local* neural surfel radiance field \mathcal{S}_t can be constructed by unprojecting pixels in the input image to 3D space using camera parameters, then estimating corresponding normal, weight, radius, and neural features. The depth D_t can be captured from sensors or estimated via off-the-shelf MVS-based methods [44], which might be incomplete and noisy. Thus, we employ a U-Net structure [29] to refine the depth following [5].

Neural Features on Surfels. We estimate neural features

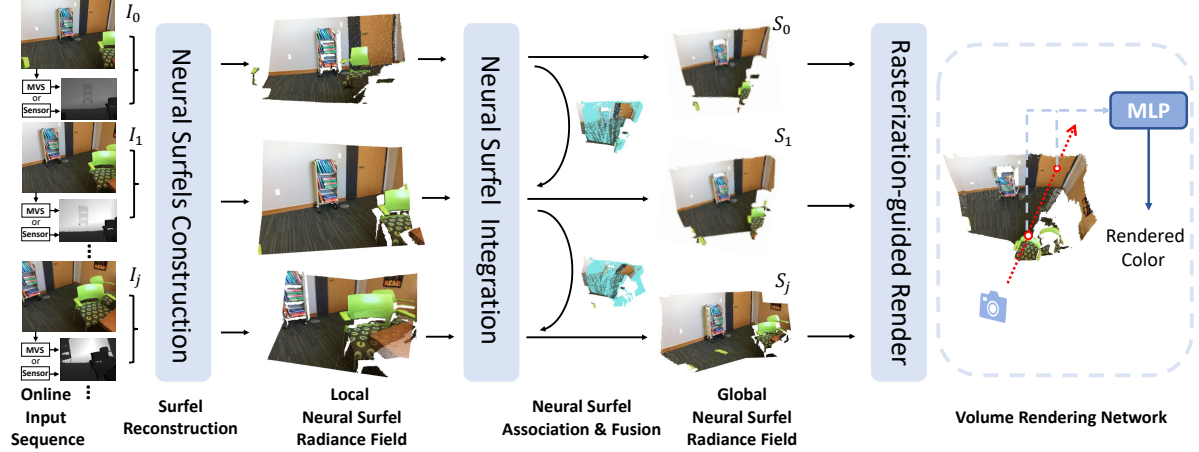


Figure 2. Overview of our SurfelNeRF. Given an online input stream of image sequences, we first reconstruct a surfel representation associated with neural features to build a local neural surfel radiance field for each input keyframe. Then the neural surfel radiance field integration is used to fuse the input local neural surfel radiance field into the global neural surfel radiance field by updating both surfel position and features. More specifically, input local neural surfels associated with global surfels are fused to global surfels with corresponding global ones, and remaining local surfels without corresponding ones are added to the global model. Furthermore, the novel views can be rendered from updated global surfels via our efficient rasterization-guided render. Our proposed rasterization-guided renderer renders color only on the intersection points of ray and surfels, which is faster than volume rendering.

f^n of surfels by extracting corresponding pixel features in the input image. We employ a 2D CNN to extract features of input images. Specifically, we use a MasNet following [34, 50] to extract multi-scale view-dependent features.

3.2. Online Integration of Neural Surfels

Based on the neural surfel representation, we adopt a geometry-guided fusion scheme to efficiently and progressively integrate local neural surfel radiance fields into a global model. Compared with previous volumetric fusion methods [10, 50], our surfel-based solution offers the advantages of 1) efficiency - the fusion operation only needs to be performed at surfel locations and 2) scalability - the memory footprint only grows linearly to the number of neural surfels.

In particular, our fusion scheme can be divided into two steps. Firstly, for each local neural surfel, we determine whether it should be added as a global neural surfel or should be merged with corresponding global surfels. Secondly, we insert it into the global model or perform updates to corresponding global surfels with local features.

Neural Surfel Association. Given an input image frame I_t and corresponding camera parameter, each surfel in global surfels can be projected as an ellipse onto the image plane of the current input camera view. For each pixel in the input image, we obtain top- K global surfels that are overlapped with it after rasterization. We discard irrelevant surfels with a normal filter. Then we determine the incremental surfel by comparing the depth distance between respective top- M global surfels from near to far step by step. If the min-

imal distance between a local surfel and its global surfel candidates is smaller than a threshold δ_{depth} , then we determine that this input surfel needs to be merged; otherwise, it should be directly added to global surfels. The overall workflow follows [41], and K is set to 8.

Neural Surfel Fusion. After the surfel association, there are a set of input surfels that need to be added directly and another set of surfels that will be merged with corresponding global surfels. For the former set of surfels, we directly construct neural surfels (Sec. 3.1) and then add them to global surfels. For the latter surfels, we use a GRU network to update neural features between the features $\mathbf{f}_t^{\text{merge}}$ of input surfels and features $\mathbf{f}_{t-1}^{\text{corrs}}$ of corresponding global surfels, which is given as

$$\mathbf{f}_t^{\text{corrs}} = \text{GRU}(\mathbf{f}_t^{\text{merge}}, \mathbf{f}_{t-1}^{\text{corrs}}), \quad (1)$$

where $\mathbf{f}_t^{\text{merge}}$ and $\mathbf{f}_{t-1}^{\text{corrs}}$ are input and hidden state of GRU respectively. We then update the geometric attributes of these global surfels via a weighted sum. The update rule of each respective pair of local surfel s_t^i and global surfel s_{t-1}^j can be expressed by

$$\begin{aligned} \mathbf{p}_t^j &= \frac{w_t^i \mathbf{p}_t^i + w_{t-1}^j \mathbf{p}_{t-1}^j}{w_t^i + w_{t-1}^j}, \\ \mathbf{n}_t^j &= \frac{w_t^i \mathbf{n}_t^i + w_{t-1}^j \mathbf{n}_{t-1}^j}{w_t^i + w_{t-1}^j}, \\ r_t^j &= \frac{w_t^i r_t^i + w_{t-1}^j r_{t-1}^j}{w_t^i + w_{t-1}^j}, \\ w_t^j &= w_t^i + w_{t-1}^j, \end{aligned} \quad (2)$$




Figure 3. Our efficient rasterization-guided surfel rendering for neural surfel radiance fields is visualized using a toy model as an example. We begin by rasterizing neural surfels into the image plane, and then render the ray using shading points on the intersections between the ray and corresponding surfels. This method uses a combination of rasterization and ray tracing to render view-dependency results with surfel representation and 3D rendering.

where \mathbf{p}_t^j , \mathbf{n}_t^j , r_t^j and w_t^j are the updated position, normal, radius, and weight of the global surfels, respectively.

3.3. Rendering Neural Surfel Radiance Fields

We propose a fast rasterization-based rendering algorithm for neural surfel radiance fields (see Figure. 3). Given a neural surfel radiance field \mathcal{S} and the camera parameter \mathbf{c}_{tar} of a target view, we first conduct a differentiable rasterization process on the neural surfels to project them on the target image plane. For a pixel p_x in the target image, we gather the neural surfels $\{s_i\}_{i=1}^M$ that cover it and render its color $C(p_x)$ based on a few shading points $\{\mathbf{x}_i\}_{i=1}^M$ that are the intersections between $\{s_i\}_{i=1}^M$ and pixel ray, using the following volume rendering formula [22]:

$$C(p_x) = \sum \tau_i (1 - \exp(-\sigma_i \Delta_i)) c_i, \quad (3)$$

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

where τ indicates the accumulated transmittance; σ_i and c_i indicate the volume density and radiance on shading point \mathbf{x}_i ; Δ_j is the distance between adjacent shading points. For each shading point x_i , we regress its volume density σ_i and radiance c_i based on its view direction \mathbf{d} and shading point features $\mathbf{f}^i(\mathbf{x}_i)$ associated from surfel s_i , which can be formulated as:

$$\sigma_i, c_i = \text{Render}(x_i, \mathbf{f}^i(\mathbf{x}_i), \mathbf{d}), \quad (4)$$

where we use an MLP to parameterize the Render module. The ‘‘interpolated’’ surfel feature $\mathbf{f}^i(\mathbf{x}_i)$ is given as:

$$\mathbf{f}^i(\mathbf{x}_i) = \frac{\|\mathbf{x}^i - \mathbf{p}^i\|}{r^i} \mathbf{F}(\mathbf{f}^i, \mathbf{d}, \mathbf{n}^i, w^i), \quad (5)$$

where \mathbf{F} is also parameterized as an MLP-like network. Note that a recent concurrent work [48] also explores the

hybrid rendering process by combining rasterization and volume rendering, but using a different neural representation.

3.4. Training and Optimization

All modules in SurfelNeRF are differentiable, so they can be trained end-to-end on large-scale scene datasets with sequential images as input. We utilize L2 pixel-wise rendering loss to train our network supervised by the target ground truth image:

$$\mathcal{L}_{\text{render}} = \|C_{gt} - C_r\|_2^2, \quad (6)$$

where C_{gt} is the ground truth pixel color and C_r denotes the rendered pixel color output from SurfelNeRF. In addition, if we can acquire sensor depth D_i^{sen} sensor for image I_i , we further employ an L1 loss to train the depth refinement network:

$$\mathcal{L}_d = \|D_i^{\text{sen}} \odot M_i - D^i \odot M_i\|_1, \quad (7)$$

where D^i is the output from depth refinement network, and M_i is a binary mask indicating valid sensor depths. This supervision can make the depth refinement network to learn to match the sensor depths.

To sum up, the total loss to train SurfelNeRF end-to-end is given as

$$\mathcal{L} = \mathcal{L}_{\text{render}} + \lambda \cdot \mathbb{1}(\mathcal{L}_d), \quad (8)$$

where $\mathbb{1}$ denotes the indicator function whose value is set to 1 if sensor depth is available; λ is the balancing weight.

Training Details. Following [50], we uniformly sample keyframes from each image sequence as input for training SurfelNeRF. To train SurfelNeRF, we sample about 5% of the full sequence as keyframes, and other frames are used as supervision.

Per-scene Fine-tuning. After being trained on large-scale datasets, our SurfelNeRF is capable of constructing neural surfel radiance fields and rendering photorealistic novel views on unseen scenes in a feedforward manner. To further improve the rendering performance on test scenes, one can also perform a *short* period of per-scene optimization via fine-tuning the surfel representations and the rendering module. To make a fair comparison, we do not use the additional depths in the fine-tuning process and keep the depth refinement module fixed. Thanks to our efficient rendering scheme, SurfelNeRF takes less time when fine-tuning compared with existing methods [43, 50].

4. Experiments

4.1. Experimental Settings

Training Datasets. Our training data consists of large-scale indoor scenes from ScanNet [11] which provides the online sequence input of images with known camera parameters,

and the captured incomplete sensor depths. We focus on the online input of large-scale scenes in the main paper. Following [37, 50], we randomly sample 100 scenes from ScanNet for training.

Implementation Details. We train our SurfelNeRF on the large-scale training dataset mentioned above using the Adam [17] optimizer with an initial learning rate of 0.001 for total 200k iterations on 4 NVIDIA Tesla V100 32GB GPUs. For per-scene optimization settings, we fine-tune our method also using the Adam optimizer with an initial learning rate of 0.0002. All input images are resized to 640×480 . The λ in Eq. 8 is set to 0.1. The threshold of δ_{depth} is set to 0.1m in surfel fusion (Sec. 3.2). The maximum number of candidates M in Eq. 3 for rendering is set to 80. We implement our method using Pytorch and Pytorch3D.

Metrics. Following previous works [50], we report the PSNR, SSIM [38] and LPIPS_{vgg} [49] to evaluate and compare the results. For fair comparisons with previous methods, we also evaluate our method on the same test set on the ScanNet dataset. The time cost of rendering and per-scene fine-tuning is measured on the same platform using one NVIDIA Tesla V100 32GB GPU.

4.2. Comparison with state-of-the-arts

In this section, we evaluate our method with other state-of-the-arts on the large-scale dataset, ScanNet. To achieve fair comparisons, we evaluate our methods strictly following the same training and evaluation scheme as NeRFusion [50] and NerfingMVS [40]. There are 8 selected testing scenes in ScanNet where 1/8 of images are held out for testing. Following [50], we report no per-scene optimization results via direct inference network and per-scene optimization results.

Baselines. In addition to comparing SurfelNeRF with NeRFusion [50], we also compare it with other state-of-the-art (SOTA) methods, such as fast per-scene optimization NeRF and neural point cloud representation, generalizable NeRF, scalable MLP-based representation NeRF, NeRF with additional depth supervision, and classical image-based rendering methods, including Instant-NGP [23], ADOP [31], PointNeRF [43], Mip-NeRF-360 [3], Dense Depth Prior NeRF(DDP-NeRF) [28], and IBRNet [37]. Among these methods, Mip-NeRF-360 is one of the state-of-the-art NeRF methods for large-scale scenes. DDP-NeRF employs an offline NeRF with additional depth supervision. PointNeRF achieves state-of-the-art performance thanks to point-based representation. ADOP achieves promising results on offline large-scale rendering. These two methods employ a similar representation to our method; thus we also compare with them. We run generalizable PointNeRF with the same No per-scene optimization and per-scene optimization settings using its open-source code.

The result is shown in Tab. 2 (a). For no per-scene

Methods	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	Time \downarrow
IBRNet [37]	21.19	0.786	0.358	-
NeRFusion [50]	22.99	0.838	0.335	38s \ddagger
PointNeRF [43]	20.47	0.642	0.544	30s
SurfelNeRF	23.82	0.845	0.327	0.2s

(a) Results of No per-scene optimization setting on ScanNet. And we report the average time to render an image. \ddagger indicates that the time of NeRFusion is reported by authors using a 2080Ti since they have not released executable code before submission.

Methods	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	Time \downarrow
Instant-NGP [23]	23.23	0.714	0.459	0.03
ADOP [31]	25.01	0.807	0.272	1s
NeRFingMVS [40]	26.37	0.903	0.245	-
IBRNet [37]	25.14	0.871	0.266	-
DDP-NeRF [28]	26.61	0.80	0.26	-
Mip-NeRF-360 [3]	27.85	0.86	0.25	-
NeRFusion [50]	26.49	0.915	0.209	38s
PointNeRF [43]	28.99	0.829	0.324	30s
SurfelNeRF	29.58	0.919	0.215	0.2s

(b) Per-scene optimization. Time \downarrow indicates the average time to render an image.

Table 2. Quantitative comparisons with SOTAs on the ScanNet dataset with No per-scene and per-scene optimization.

optimization settings, NeRFusion achieves PSNR of 22.99 thanks to the fusion between input features to global grid representations. Our SurfelNeRF outperforms NeRFusion since SurfelNeRF benefits from compact and precise feature fusion and the geometric guidance of surfels. For the per-scene optimization task, the result is shown in Tab. 2 (b). PointNeRF gains a large improvement after fine-tuning, which illustrates it achieves sub-optimal performance in no per-scene optimization settings due to its inability to fuse and update point features. Compared with DDP-NeRF, which employs additional depth supervision, our SurfelNeRF still outperforms it by taking advantage of surfel representation and its integration. Our method achieves state-of-the-art results again, which shows that our neural surfel representation can provide optimal initialization for fine-tuning.

The qualitative comparison is shown in Figure. 4. We compare our method with PointNeRF without per-scene optimization, and compare with per-scene optimization methods of PointNeRF and ADOP. Since NeRFusion has not release the executable code, we cannot generate qualitative results. As shown in Figure. 4, our results are visually much better than state-of-the-art methods, PointNeRF and ADOP.

4.3. Ablation Study

We investigate the effectiveness of designs in our SurfelNeRF by conducting the following ablation studies on

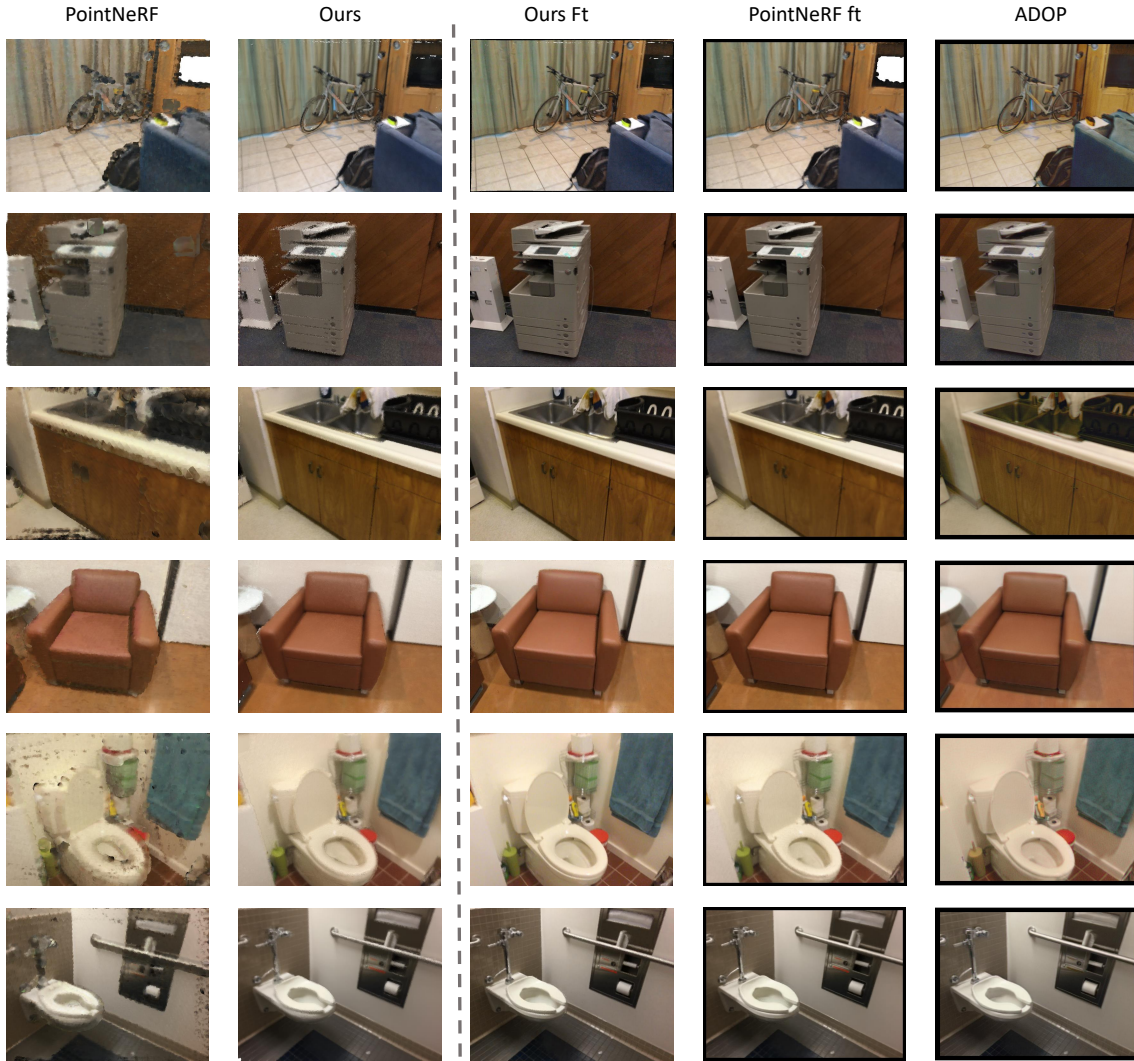


Figure 4. Qualitative comparisons of rendering quality. The 1-2 column are the results of no per-scene optimization. And the left three columns are per-scene optimization results. It is obvious that with or without per-scene optimization, our method can still achieve the most photorealistic rendering results. Since PointNeRF without a depth refinement network relies on captured depths only, it cannot handle the scenario that depth is incomplete. However, our method can tackle this problem with a lightweight depth refinement network. ADOP can produce geometrically correct results, but its color is less photorealistic.

ScanNet.

Fusion Schemes. Our surfel-based fusion is inspired by classical Surfel-based fusion in conventional 3D reconstruction [41]. The classical surfel-based fusion update features via the simple weighted sum of two overlapped surfels. We conduct an experiment that updates surfel features by weighted sum based on surfel weights, which is a similar operation of updating geometry attributes in Eq. 2. The results are shown in Tab. 3. Since the GRU module can learn to adaptively update features based on the global surfels and new incremental surfels, it can improve performance with negligible computation overhead.

Depth Refinement Module. To investigate the effective-

Fusion Scheme	Setting	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
Weighted Sum	No per-scene optimization	23.09	0.833	0.353
GRU		23.82	0.845	0.327
Weighted Sum	Per-scene optimization	28.54	0.884	0.293
GRU		29.58	0.919	0.215

Table 3. Ablation studies about fusion schemes in our SurfelNeRF.

ness of our depth refinement, we conduct experiments that reconstruct scenes with incomplete depth captured by sensors. The quantitative results are shown in Tab. 4, which show that PSNR performance drops about 4% without the

Depth Refinement	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
w/o	19.51	0.713	0.454
w	23.09	0.833	0.353

Table 4. Ablation studies about the effectiveness of depth refinement in *scene0000.01* in the ScanNet with the no per-scene optimization setting. The depth captured by the sensor in this scene is incomplete. If using captured depths only, some pixels cannot be projected to the 3D space, as the visualization results of PointNeRF shown in the first row in the qualitative results in Fig. 4.

Rendering Scheme	Setting	PSNR \uparrow	Time \downarrow
Volume Rendering	No per-scene optimization	23.80	30s
Ours	optimization	23.82	0.2s
Volume Rendering	Per-scene optimization	29.45	30s
Ours	optimization	29.58	0.2s

Table 5. Comparing with two rendering schemes, volume rendering [22] and our proposed rasterization-guided rendering, in the ScanNet dataset.

depth refinement. Depth quality affects surfel reconstruction geometry quality, but captured depth is always incomplete (visualized in the first row in Fig. 4). Thus, a depth refinement network is necessarily required to refine the input depth for better surfel geometry quality. The experimental result shows that the depth refinement network is effective in alleviating the noisy or incomplete captured or estimated depth. Moreover, we also conduct an experiment with worse depth quality that resizes input depth resolution to half size. Our performance just slightly drops 0.5% PSNR. These experiments show the effectiveness of our depth refinement module.

Rendering Schemes. Our proposed renderer renders a ray with a few shading points, which are the intersection points of ray and surfels, unlike volume rendering in NeRF [22] that requires hundreds of shading points each with one MLP evaluation. To validate the effectiveness of our differentiable rasterization scheme for rendering, we conduct an experiment that replaced our rendering scheme with NeRFs’ volume rendering. The result is shown in Tab. 5. Since online integration of surfels can limit the number of intersecting surfels, our rendering speed is persistently faster than volume rendering with fewer shading points. Specifically, The average and maximal number of rasterized surfels per pixel at novel views on ScanNet are 2.92 and 10, respectively.

Computational Cost and Memory Usage. To evaluate the computational cost and memory usage on a large-scale scene, we plot components of GPU processing time, surfel number, and memory usage of our method on a large three rooms scene measured on a V100 GPU implemented by Pytorch3D. During inference, our method renders 640×480




Figure 5. The computation cost, memory footprint, and the number of surfels grow according to input frames measured on a three rooms scene.

images at 5.2FPS on average. As shown in Fig. 5a and Fig. 5b, the number of surfels has a sub-linear growth thanks to the proposed surfel integration. The fluctuation of *Time (Rasterization)* is caused by the different numbers of corresponding surfels at novel views. Although the computational cost of rasterization depends on the number of rasterized surfels, but it is still faster than volume rendering [22]. These results show that our surfel representation is compact and efficient. The major computation cost comes from rasterization, and the cost of surfel construction as well as integration could be neglected.

5. Limitations

Our method requires high-quality surfel geometry as neural surfel representation since SurfelNeRF relies on it to conduct surfel integration and rasterization-guided render. To obtain better geometry quality, we may require more accurate depth quality to construct surfel geometry, which necessitates extra depth supervision or the MVS depth estimator. Since captured or estimated depth quality is not exact, there may be geometry discontinuities or missing parts in the surfel representation. Consequently, there may be some missing parts or incorrect geometry in the rendered output (examples of missing parts are shown in the first row in Fig. 4 or video). As discussed in ablation studies of depth refinement above, inaccurate and noisy estimated depth as well as depth discontinuity could affect the construction of neural surfel representation. To alleviate this issue, we propose a depth refinement module (Tab. 4) to refine depth information.

6. Conclusion

In this paper, we propose SurfelNeRF for online photo-realistic reconstruction and rendering of large-scale indoor scenes. We employ neural surfels to store a neural radiance field, adopt an efficient neural radiance field fusion method to progressively build the scene representation, and propose a fast and differentiable rasterization-based rendering algorithm to tackle the challenging task. Experiments on the ScanNet dataset demonstrate that our method achieving favorably better results than the state-of-the-art methods.

References

- [1] Kara-Ali Aliev, Artem Sevastopolsky, Maria Kolos, Dmitry Ulyanov, and Victor Lempitsky. Neural point-based graphics. 2020. [2](#)
- [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. [1](#), [2](#)
- [3] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5470–5479, 2022. [1](#), [6](#)
- [4] Aljaz Bozic, Pablo Palafox, Justus Thies, Angela Dai, and Matthias Nießner. Transformerfusion: Monocular rgb scene reconstruction using transformers. *Advances in Neural Information Processing Systems*, 34:1403–1414, 2021. [3](#)
- [5] Ang Cao, Chris Rockwell, and Justin Johnson. Fwd: Real-time novel view synthesis with forward warping and depth. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15713–15724, 2022. [3](#)
- [6] Yan-Pei Cao, Leif Kobbelt, and Shi-Min Hu. Real-time high-accuracy three-dimensional reconstruction with consumer RGB-D cameras. *ACM Transactions on Graphics*, 37(5):171:1–171:16, 2018. [1](#), [2](#), [3](#)
- [7] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. *arXiv preprint arXiv:2203.09517*, 2022. [1](#)
- [8] 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. [3](#)
- [9] Zhiqin Chen, Thomas Funkhouser, Peter Hedman, and Andrea Tagliasacchi. Mobilenerf: Exploiting the polygon rasterization pipeline for efficient neural field rendering on mobile architectures. *arXiv preprint arXiv:2208.00277*, 2022. [2](#)
- [10] Ronald Clark. Volumetric bundle adjustment for online photorealistic scene capture. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6124–6132, 2022. [2](#), [3](#), [4](#)
- [11] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas A. Funkhouser, and Matthias Nießner. ScanNet: Richly-annotated 3d reconstructions of indoor scenes. In *IEEE CVPR*, pages 2432–2443, 2017. [2](#), [5](#)
- [12] Angela Dai, Matthias Nießner, Michael Zollhöfer, Shahram Izadi, and Christian Theobalt. Bundlefusion: Real-time globally consistent 3d reconstruction using on-the-fly surface reintegration. *ACM Transactions on Graphics*, 36(3):24:1–24:18, 2017. [1](#)
- [13] Peter Hedman, Pratul P Srinivasan, Ben Mildenhall, Jonathan T Barron, and Paul Debevec. Baking neural radiance fields for real-time view synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5875–5884, 2021. [2](#)
- [14] Animesh Karnewar, Tobias Ritschel, Oliver Wang, and Niloy Mitra. Relu fields: The little non-linearity that could. In *ACM SIGGRAPH 2022 Conference Proceedings*, SIGGRAPH '22, New York, NY, USA, 2022. Association for Computing Machinery. [1](#)
- [15] Maik Keller, Damien Lefloch, Martin Lambers, Shahram Izadi, Tim Weyrich, and Andreas Kolb. Real-time 3d reconstruction in dynamic scenes using point-based fusion. In *2013 International Conference on 3D Vision-3DV 2013*, pages 1–8. IEEE, 2013. [2](#), [3](#)
- [16] Jungeon Kim, Hyomin Kim, Hyeonsoo Nam, Jaesik Park, and Seungyong Lee. Textureme: High-quality textured scene reconstruction in real time. *ACM Trans. Graph.*, 41(3), mar 2022. [1](#)
- [17] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [6](#)
- [18] Zhong Li, Liangchen Song, Celong Liu, Junsong Yuan, and Yi Xu. Neulf: Efficient novel view synthesis with neural 4d light field. *arXiv preprint arXiv:2105.07112*, 2021. [2](#)
- [19] Jia-Wei Liu, Yan-Pei Cao, Weijia Mao, Wenqiao Zhang, David Junhao Zhang, Jussi Keppo, Ying Shan, Xiaohu Qie, and Mike Zheng Shou. Devrf: Fast deformable voxel radiance fields for dynamic scenes. *arXiv preprint arXiv:2205.15723*, 2022. [1](#)
- [20] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. *NeurIPS*, 2020. [2](#)
- [21] Marko Mihajlovic, Silvan Weder, Marc Pollefeys, and Martin R Oswald. Deepsurfels: Learning online appearance fusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14524–14535, 2021. [2](#), [3](#)
- [22] 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 *European conference on computer vision*, pages 405–421. Springer, 2020. [1](#), [2](#), [3](#), [5](#), [8](#)
- [23] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.*, 41(4):102:1–102:15, July 2022. [1](#), [2](#), [3](#), [6](#)
- [24] Richard A. Newcombe, Shahram Izadi, Otmar Hilliges, David Molyneaux, David Kim, Andrew J. Davison, Pushmeet Kohli, Jamie Shotton, Steve Hodges, and Andrew W. Fitzgibbon. Kinectfusion: Real-time dense surface mapping and tracking. In *IEEE ISMAR*, pages 127–136, 2011. [1](#)
- [25] Hanspeter Pfister, Matthias Zwicker, Jeroen Van Baar, and Markus Gross. Surfels: Surface elements as rendering primitives. In *Proceedings of the 27th annual conference on computer graphics and interactive techniques*, pages 335–342, 2000. [2](#)
- [26] Ruslan Rakhimov, Andrei-Timotei Ardelean, Victor Lempitsky, and Evgeny Burnaev. Npb++: Accelerating neural point-based graphics. In *Proceedings of the IEEE/CVF*

- Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15969–15979, June 2022. [2](#)
- [27] Konstantinos Rematas, Andrew Liu, Pratul P Srinivasan, Jonathan T Barron, Andrea Tagliasacchi, Thomas Funkhouser, and Vittorio Ferrari. Urban radiance fields. *arXiv preprint arXiv:2111.14643*, 2021. [1](#)
- [28] 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 (CVPR)*, June 2022. [6](#)
- [29] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. [3](#)
- [30] Darius Rückert, Linus Franke, and Marc Stamminger. Adop: Approximate differentiable one-pixel point rendering. *arXiv preprint arXiv:2110.06635*, 2021. [2](#)
- [31] Darius Rückert, Linus Franke, and Marc Stamminger. Adop: Approximate differentiable one-pixel point rendering. *ACM Transactions on Graphics (TOG)*, 41(4):1–14, 2022. [6](#)
- [32] Mehdi SM Sajjadi, Henning Meyer, Etienne Pot, Urs Bergmann, Klaus Greff, Noha Radwan, Suhani Vora, Mario Lučić, Daniel Duckworth, Alexey Dosovitskiy, et al. Scene representation transformer: Geometry-free novel view synthesis through set-latent scene representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6229–6238, 2022. [3](#)
- [33] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. *arXiv preprint arXiv:2111.11215*, 2021. [1](#)
- [34] Jiaming Sun, Yiming Xie, Linghao Chen, Xiaowei Zhou, and Hujun Bao. Neuralrecon: Real-time coherent 3d reconstruction from monocular video. In *IEEE CVPR*, pages 15598–15607, 2021. [2](#), [3](#), [4](#)
- [35] Matthew Tancik, Vincent Casser, Xinchun Yan, Sabeek Pradhan, Ben Mildenhall, Pratul P Srinivasan, Jonathan T Barron, and Henrik Kretzschmar. Block-nerf: Scalable large scene neural view synthesis. *arXiv preprint arXiv:2202.05263*, 2022. [1](#)
- [36] Justus Thies, Michael Zollhöfer, and Matthias Nießner. Deferred neural rendering: Image synthesis using neural textures. *ACM Transactions on Graphics (TOG)*, 38(4):1–12, 2019. [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. [3](#), [6](#)
- [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. [6](#)
- [39] Silvan Weder, Johannes Schonberger, Marc Pollefeys, and Martin R Oswald. Routedfusion: Learning real-time depth map fusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4887–4897, 2020. [1](#)
- [40] 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. [6](#)
- [41] Thomas Whelan, Stefan Leutenegger, Renato Salas-Moreno, Ben Glocker, and Andrew Davison. Elasticfusion: Dense slam without a pose graph. *Robotics: Science and Systems*, 2015. [1](#), [2](#), [3](#), [4](#), [7](#)
- [42] Yuanbo Xiangli, Linning Xu, Xingang Pan, Nanxuan Zhao, Anyi Rao, Christian Theobalt, Bo Dai, and Dahua Lin. Citynerf: Building nerf at city scale. *arXiv preprint arXiv:2112.05504*, 2021. [1](#), [2](#)
- [43] Qiangeng Xu, Zexiang Xu, Julien Philip, Sai Bi, Zhixin Shu, Kalyan Sunkavalli, and Ulrich Neumann. Point-nerf: Point-based neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5438–5448, 2022. [2](#), [3](#), [5](#), [6](#)
- [44] Yao Yao, Zixin Luo, Shiwei Li, Tian Fang, and Long Quan. Mvsnet: Depth inference for unstructured multi-view stereo. *European Conference on Computer Vision (ECCV)*, 2018. [3](#)
- [45] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. Plenotrees for real-time rendering of neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5752–5761, 2021. [2](#)
- [46] 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. [3](#)
- [47] Jian Zhang, Jinchi Huang, Bowen Cai, Huan Fu, Mingming Gong, Chaohui Wang, Jiaming Wang, Hongchen Luo, Rongfei Jia, Binqiang Zhao, et al. Digging into radiance grid for real-time view synthesis with detail preservation. In *European Conference on Computer Vision*, pages 724–740. Springer, 2022. [2](#)
- [48] Qiang Zhang, Seung-Hwan Baek, Szymon Rusinkiewicz, and Felix Heide. Differentiable point-based radiance fields for efficient view synthesis. *arXiv preprint arXiv:2205.14330*, 2022. [5](#)
- [49] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. [6](#)
- [50] Xiaoshuai Zhang, Sai Bi, Kalyan Sunkavalli, Hao Su, and Zexiang Xu. Nerfusion: Fusing radiance fields for large-scale scene reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5449–5458, 2022. [2](#), [3](#), [4](#), [5](#), [6](#)
- [51] Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. Stereo magnification: Learning view synthesis using multiplane images. *arXiv preprint arXiv:1805.09817*, 2018. [2](#)

- [52] 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 *IEEE CVPR*, pages 12786–12796, 2022. 3
- [53] Zi-Xin Zou, Shi-Sheng Huang, Yan-Pei Cao, Tai-Jiang Mu, Ying Shan, and Hongbo Fu. Mononeuralfusion: Online monocular neural 3d reconstruction with geometric priors. *arXiv preprint arXiv:2209.15153*, 2022. 3