

MobileNeRF: Exploiting the Polygon Rasterization Pipeline for Efficient Neural Field Rendering on Mobile Architectures

Zhiqin Chen^{1,2} Thomas Funkhouser¹ Peter Hedman¹ Andrea Tagliasacchi^{1,2,3}
 Google Research¹ Simon Fraser University² University of Toronto³

Abstract

Neural Radiance Fields (NeRFs) have demonstrated amazing ability to synthesize images of 3D scenes from novel views. However, they rely upon specialized volumetric rendering algorithms based on ray marching that are mismatched to the capabilities of widely deployed graphics hardware. This paper introduces a new NeRF representation based on textured polygons that can synthesize novel images efficiently with standard rendering pipelines. The NeRF is represented as a set of polygons with textures representing binary opacities and feature vectors. Traditional rendering of the polygons with a z-buffer yields an image with features at every pixel, which are interpreted by a small, view-dependent MLP running in a fragment shader to produce a final pixel color. This approach enables NeRFs to be rendered with the traditional polygon rasterization pipeline, which provides massive pixel-level parallelism, achieving interactive frame rates on a wide range of compute platforms, including mobile phones.

Project page: <https://mobile-nerf.github.io>

1. Introduction

Neural Radiance Fields (NeRF) [33] have become a popular representation for novel view synthesis of 3D scenes. They represent a scene using a multilayer perceptron (MLP) that evaluates a 5D implicit function estimating the density and radiance emanating from any position in any direction, which can be used in a volumetric rendering framework to produce novel images. NeRF representations optimized to minimize multi-view color consistency losses for a set of posed photographs have demonstrated remarkable ability to reproduce fine image details for novel views.

One of the main impediments to wide-spread adoption of NeRF is that it requires specialized rendering algorithms that are poor match for commonly available hardware. Traditional NeRF implementations use a volumetric rendering

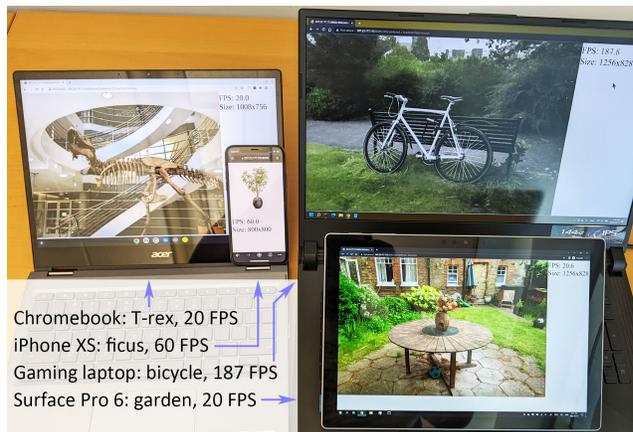


Figure 1. **Teaser** – We present a NeRF that can run on a variety of common devices at interactive frame rates.

algorithm that evaluates a large MLP at hundreds of sample positions along the ray for each pixel in order to estimate and integrate density and radiance. This rendering process is far too slow for interactive visualization.

Recent work has addressed this issue by “baking” NeRFs into a sparse 3D voxel grid [21, 51]. For example, Hedman et al. introduced Sparse Neural Radiance Grids (SNeRG) [21], where each active voxel contains an opacity, diffuse color, and learned feature vector. Rendering an image from SNeRG is split into two phases: the first uses ray marching to accumulate the precomputed diffuse colors and feature vectors along each ray, and the second uses a light-weight MLP operating on the accumulated feature vector to produce a view-dependent residual that is added to the accumulated diffuse color. This precomputation and deferred rendering approach increase the rendering speed of NeRF by three orders of magnitude. However, it still relies upon ray marching through a sparse voxel grid to produce the features for each pixel, and thus it cannot fully utilize the parallelism available in commodity graphics processing units (GPUs). In addition, SNeRG requires a significant amount of GPU memory to store the volumetric textures, which prohibits it from running on common mobile devices.

In this paper, we introduce MobileNeRF, a NeRF that

³Work done while at Google.

can run on a variety of common mobile devices at interactive frame rates. The NeRF is represented by a set of textured polygons, where the polygons roughly follow the surface of the scene, and the texture atlas stores opacity and feature vectors. To render an image, we utilize the classic polygon rasterization pipeline with Z-buffering to produce a feature vector for each pixel and pass it to a lightweight MLP running in a GLSL fragment shader to produce the output color. This rendering pipeline *does not* sample rays or sort polygons in depth order, and thus can model only binary opacities. However, it takes full advantage of the parallelism provided by z-buffers and fragment shaders in modern graphics hardware, and thus is $10\times$ faster than SNeRG with the same output quality on standard test scenes. Moreover, it requires only a standard polygon rendering pipeline, which is implemented and accelerated on virtually every computing platform, and thus it runs on mobile phones and other devices previously unable to support NeRF visualization at interactive rates.

Contributions. In summary, MobileNeRF:

- Is $10\times$ faster than the state-of-the-art (SNeRG), with the same output quality;
- Consumes less memory by storing *surface* textures instead of volumetric textures, enabling our method to run on integrated GPUs with limited memory and power;
- Runs on a web browser and is *compatible* with all devices we have tested, as our viewer is an HTML webpage;
- Allows real-time *manipulation* of the reconstructed objects/scenes, as they are simple triangle meshes.

2. Related work

Our work lies within the field of view-synthesis, which encompasses many areas of research: light fields, image-based rendering and neural rendering. To narrow the scope, we focus on methods that render output views in *real-time*.

Light fields [27] and Lumigraphs [19] store a dense grid of images, enabling real-time rendering of high quality scenes, albeit with limited camera freedom and significant storage overhead. Storage can be reduced by interpolating intermediate images with optical flow [5], representing the light field as a neural network [1], or by reconstructing a Multi-Plane Image (MPI) representation of the scene [15,32,37,47,54]. Multi-sphere images enable larger fields of view [2,6], but these representations still only support limited output camera motion

Other approaches leverage explicit 3D geometry to enable more camera freedom. While early methods applied view-dependent texturing to a 3D mesh [7, 12, 13], later methods incorporated convolutional neural networks as a post-processing step to improve quality [20, 31, 44]. Alternatively, the input geometry can be simplified into a col-

lection of textured planes with alpha [28]. Point-based representations further increase quality by jointly refining the scene geometry while training the post-processing network [24, 25, 40]. However, as this convolutional post-processing runs independently per output frame it often results in a lack of 3D consistency. Furthermore, unlike our work, they require powerful desktop GPUs and have not been demonstrated to run on a mobile device. Finally, unlike the vast majority of the methods above, our method does not need reconstructed 3D geometry as input.

It is also possible to extract explicit triangle meshes via differentiable inverse-rendering [11, 16, 35]. DefTet [16] differentially renders a tetrahedral grid with occupancy and color at each vertex, and then compositing the interpolated values at all intersected faces along a ray. NVDiffRec [35] combines differentiable marching tetrahedra [42] with differentiable rasterization to perform full inverse rendering and extract triangle meshes, materials, and lighting from images. This representation enables elaborate editing and scene relighting. However, it incurs a significant loss in view-synthesis quality. Furthermore, while real-time rendering is possible with simple lighting, global illumination (GI) is computationally infeasible on mobile hardware. In contrast, our method simply caches the outgoing radiance, which does not need expensive compute to model GI effects, and also results in higher view-synthesis quality.

NeRF [33] represents the scene as a continuous field of opacity and view-dependent color, and produces images with volume rendering. This representation is 3D consistent and reaches high quality results [3,45]. However, rendering a NeRF involves evaluating a large neural network at multiple 3D locations per pixel, preventing real-time rendering.

Recent works have improved the training speed of NeRF. For example, by modeling the opacity and color of entire ray segments instead of just points [29] or by subdividing the scene and modeling each sub-region with a smaller neural network [38]. Recently, significant speed-ups have been achieved by decoding features fetched from a 3D embedding with a small neural network. This embedding can either be a dense voxel grid [23,43], a sparse voxel grid [41], a low-rank decomposition of a voxel grid [9], a point-based representation [50], or a multi-resolution hash map [34]. These 3D embeddings can also be used without a trained decoder, for example by directly storing diffuse colors [30] or by encoding view-dependent colors as spherical harmonics [41]. While these approaches drastically speed up training, they still require a large consumer GPU for rendering.

Rendering performance can further be increased by *post-processing* a trained NeRF. For example, by reducing the network queries per pixel with learned sampling [36], by evaluating the network for larger ray segments [48], or by subdividing the scene into smaller networks [38, 39, 49].

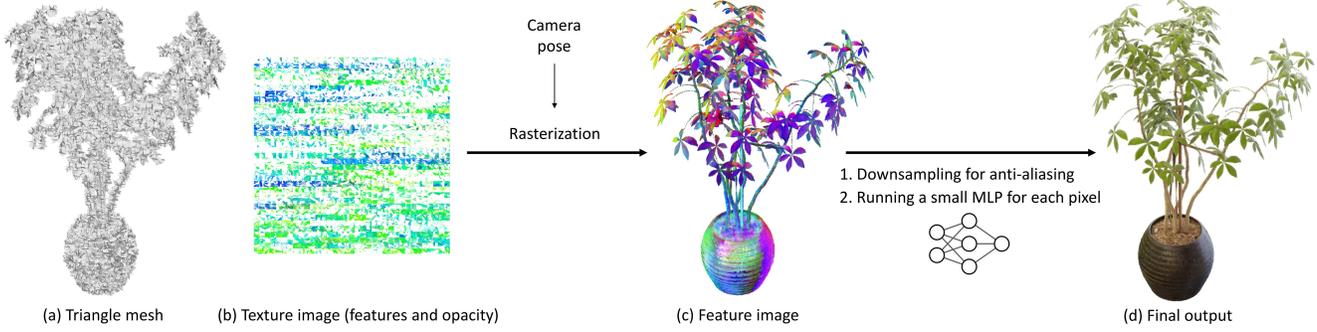


Figure 2. **Overview (rendering)** – We represent the scene as a triangle mesh textured by deep features. We first rasterize the mesh to a deferred rendering buffer. For each visible fragment, we execute a neural deferred shader that converts the feature and view direction to the corresponding output pixel color.

Alternatively, pre-computation can speed up rendering, by storing both scene opacity and a latent representation for view-dependent colors in a grid. FastNeRF [17] uses a dense voxel grid and represents view-dependence with a global spherical basis function. PlenOctrees [51] uses an octree representation, where each leaf node stores both opacity and spherical harmonics for colors. SNeRG [21] uses a sparse grid representation, and evaluates view-dependence as a post-process with a small neural network. Among these real-time methods, only SNeRG has been shown to work on lower-powered devices without access to CUDA. As our method directly targets rendering on low-powered hardware, we primarily compare with SNeRG in our experiments.

3. Method

Given a collection of (calibrated) images, we seek to optimize a representation for *efficient* novel-view synthesis. Our representation consists of a polygonal mesh (Figure 2a) whose texture maps (Figure 2b) store features and opacity. At rendering time, given a camera pose, we adopt a two-stage *deferred rendering* process:

- **Rendering Stage 1** – we rasterize the mesh to screen space and construct a *feature image* (Figure 2c), i.e. we create a deferred rendering buffer in GPU memory;
- **Rendering Stage 2** – we convert these features into a color image via a (neural) deferred renderer running in a fragment shader, i.e. a small MLP, which receives a feature vector and view direction and outputs a pixel color (Figure 2d).

Our representation is built in three *training* stages, gradually moving from a classical NeRF-like continuous representation towards a discrete one:

- **Training Stage 1 (Section 3.1)** – We train a NeRF-like model with *continuous* opacity, where volume rendering quadrature points are derived from the polygonal mesh;
- **Training Stage 2 (Section 3.2)** – We *binarize* the opac-

ities, as while classical rasterization can easily discard fragments, they cannot elegantly deal with semi-transparent fragments.

- **Training Stage 3 (Section 3.3)** – We *extract* a sparse polygonal mesh, bake opacities and features into texture maps, and store the weights of the neural deferred shader.

The mesh is stored as an OBJ file, the texture maps in PNGs, and the deferred shader weights in a (small) JSON file. As we employ the standard GPU rasterization pipeline, our real-time renderer is simply an HTML webpage.

As representing continuous signals with discrete representations can introduce aliasing, we also detail a simple, yet computationally efficient, anti-aliasing solution based on super-sampling (Section 3.4).

3.1. Continuous training (Training Stage 1)

As Figure 3 shows, our *training setup* consists of a polygonal mesh $\mathcal{M}=(\mathcal{T}, \mathcal{V})$ and three MLPs. The mesh topology \mathcal{T} is fixed, but the vertex locations \mathcal{V} and MLPs are optimized, similarly to NeRF, in an auto-decoding fashion by minimizing the mean squared error between predicted colors and ground truth colors of the pixels in the training images¹:

$$\mathcal{L}_{\mathbf{C}} = \mathbb{E}_{\mathbf{r}} \|\mathbf{C}(\mathbf{r}) - \mathbf{C}_{\text{gt}}(\mathbf{r})\|_2^2. \quad (1)$$

where the predicted color $\mathbf{C}(\cdot)$ is obtained by alpha-compositing the radiance \mathbf{c}_k along a ray $\mathbf{r}(t)=\mathbf{o} + t\mathbf{d}$, at the (depth sorted) quadrature points $\mathcal{K}=\{t_k\}_{k=1}^K$:

$$\mathbf{C}(\mathbf{r}) = \sum_{k=1}^K T_k \alpha_k \mathbf{c}_k, \quad T_k = \prod_{l=1}^{k-1} (1 - \alpha_l) \quad (2)$$

¹For real-world scenes, we further incorporate the distortion loss $\mathcal{L}_{\text{dist}}$ introduced by [3, Eq. 15] to suppress *floaters* and *background collapse*.

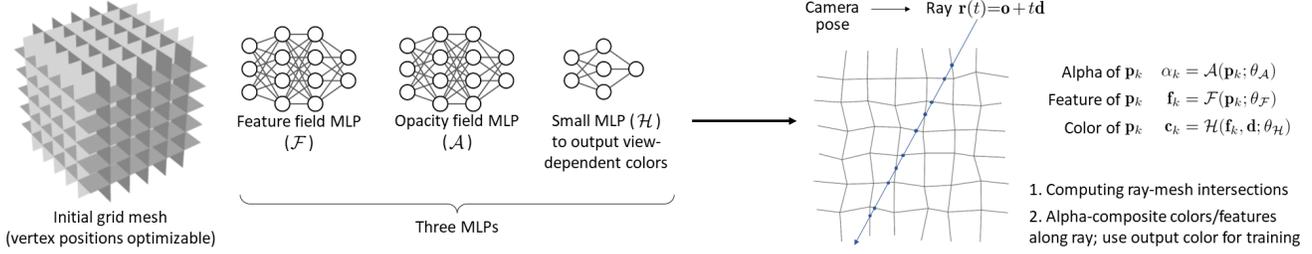


Figure 3. **Overview (train)** – We initialize the mesh as a regular grid, and use MLPs to represent features and opacity for any point on the mesh. For each ray, we compute its intersection points on the mesh, and alpha-composite the colors of those points to obtain the output color. In a later training stage, we enforce *binary* opacity, and perform super-sampling on features for anti-aliasing.

where *opacity* α_k and the view-dependent *radiance* \mathbf{c}_k are given by evaluating the MLPs at position $\mathbf{p}_k = \mathbf{r}(t_k)$:

$$\alpha_k = \mathcal{A}(\mathbf{p}_k; \theta_{\mathcal{A}}) \quad \mathcal{A} : \mathbb{R}^3 \rightarrow [0, 1] \quad (3)$$

$$\mathbf{f}_k = \mathcal{F}(\mathbf{p}_k; \theta_{\mathcal{F}}) \quad \mathcal{F} : \mathbb{R}^3 \rightarrow [0, 1]^8 \quad (4)$$

$$\mathbf{c}_k = \mathcal{H}(\mathbf{f}_k, \mathbf{d}; \theta_{\mathcal{H}}) \quad \mathcal{H} : [0, 1]^8 \times [-1, 1]^3 \rightarrow [0, 1]^3 \quad (5)$$

The small network \mathcal{H} is our *deferred neural shader*, which outputs the color of each fragment given the fragment feature and viewing direction. Finally, note that (2) does not perform compositing with volumetric density [33], but rather with opacity [1, Eq.8].

Polygonal mesh. Without loss of generality, we describe the polygonal mesh used in *Synthetic 360°* scenes, and provide the configurations for *Forward-Facing* and *Unbounded 360°* scenes in [supplementary](#) (Section G). 2D illustrations can be found in Figure 4. We first define a *regular grid* \mathcal{G} of size $P \times P \times P$ in the unit cube centered at the origin; see Figure 4a. We instantiate \mathcal{V} by creating one vertex per voxel, and \mathcal{T} by creating one quadrangle (two triangles) per grid edge connecting the vertices of the four adjacent voxels, akin to Dual Contouring [10, 22]. We locally parameterize vertex locations with respect to the voxel centers (and sizes), resulting in $\mathcal{V} \in [-.5, +.5]^{P \times P \times P \times 3}$ free variables. During optimization, we initialize the vertex locations to $\mathcal{V} = \mathbf{0}$, which corresponds to a regular Euclidean lattice, and we regularize them to prevent vertices from exiting their voxels, and to promote their return to their neutral position whenever the optimization problem is under-constrained:

$$\mathcal{L}_{\mathcal{V}} = \sum_{\mathbf{v} \in \mathcal{V}} (10^3 \mathcal{I}(\mathbf{v}) + 10^{-2}) \cdot \|\mathbf{v}\|_1, \quad (6)$$

where the indicator function $\mathcal{I}(\mathbf{v}) \equiv 1$ whenever \mathbf{v} is outside its corresponding voxel.

Quadrature. As evaluating the MLPs of our representation is computationally expensive, we rely on an acceleration grid to limit the cardinality $|\mathcal{K}|$ of quadrature points. First of all, quadrature points are only generated for the set of voxels that intersect the ray; see Figure 5a: Then, like InstantNGP [34], we employ an acceleration grid \mathcal{G} to prune

voxels that are unlikely to contain geometry; see Figure 5b. Finally, we compute intersections between the ray and the faces of \mathcal{M} that are incident to the voxel’s vertex to obtain the final set of quadrature points; see Figure 5c. We use the barycentric interpolation to back-propagate the gradients from the intersection point to the three vertices in the intersected triangle. For further technical details on the computation of intersections, we refer the reader to [supplementary](#) (Section F). In summary, for each input ray \mathbf{r} :

$$\tilde{\mathcal{B}} = \text{intersect}(\mathbf{r}, \mathcal{G}) \quad (7)$$

$$\mathcal{B} = \{b \in \tilde{\mathcal{B}} \mid \mathcal{G}[b] > \tau_{\mathcal{G}}\} \quad (8)$$

$$\mathcal{K} = \text{intersect}(\mathbf{r}, \{\mathbf{t} \in \mathcal{T} \mid \mathbf{t} \cap \mathcal{B}\}) \quad (9)$$

where $(a \cap b) = \text{true}$ if a intersects b , and the acceleration grid is supervised so to upper-bound² the alpha-compositing visibility $T_k \alpha_k$ across viewpoints during training.

$$\mathcal{L}_{\mathcal{G}}^{\text{bnd}} = \sum_k \max(\mathcal{N}[T_k \alpha_k] - \mathcal{G}[\mathbf{p}_k], 0) \quad (10)$$

where $\mathcal{N}[\cdot]$ is the stop-gradient operator that prevents the acceleration grid from (negatively) affecting the image reconstruction quality. This can be interpreted as a way to compute the so-called “surface field” *during* NeRF training, as opposed to *after* training as in nerf2nerf [18]. Similarly to Plenoxels [41], we additionally regularize the content of the grid by promoting its pointwise sparsity (i.e. lasso), and its spatial smoothness:

$$\mathcal{L}_{\mathcal{G}}^{\text{sparse}} = \|\mathcal{G}\|_1 \quad \mathcal{L}_{\mathcal{G}}^{\text{smooth}} = \|\nabla \mathcal{G}\|_2^2 \quad (11)$$

3.2. Binarized training (Training Stage 2)

Rendering pipelines implemented in typical hardware *do not* natively support semi-transparent meshes. Rendering semi-transparent meshes requires cumbersome (per-frame) sorting so to execute rendering in back-to-front order to guarantee correct alpha-compositing. We overcome this issue by converting the smooth opacity $\alpha_k \in [0, 1]$ from (3) to

²This loss performs a *stochastic upper-bound*, as we initialize $\mathcal{G}[\ast] = \mathbf{0}$, and $\mathcal{G}[\mathbf{p}_k]$ receives gradients whenever $T_k \alpha_k > \mathcal{G}[\mathbf{p}_k]$.

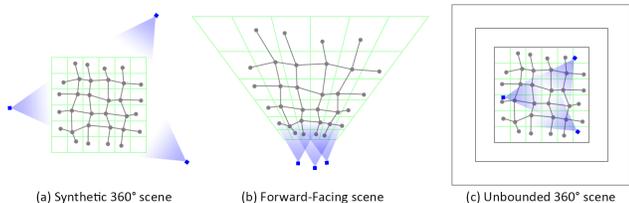


Figure 4. **Configurations of polygonal meshes** – The meshes employed for the different types of scenes. We sketch the distribution of camera poses in training views.

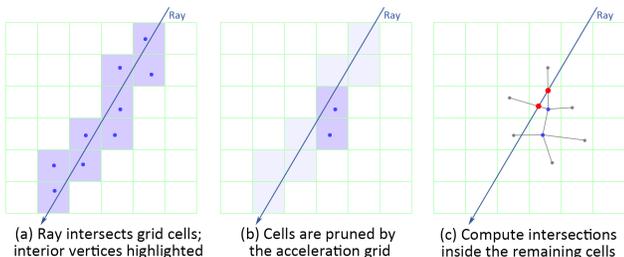


Figure 5. **Quadrature points** – are obtained by (a) identifying cells that intersect the ray; (b) pruning cells that do not contain geometry; and, (c) computing explicit intersections with the mesh.

a discrete/categorical opacity $\hat{\alpha}_k \in \{0, 1\}$. To optimize for discrete opacities via photometric supervision we employ a *straight-through estimator* [4]:

$$\hat{\alpha}_k = \alpha_k + \mathcal{V}[\mathbb{1}(\alpha_k > 0.5) - \alpha_k] \quad (12)$$

Please note that the gradients are transparently passed through the discretization operation (i.e. $\nabla \hat{\alpha} \equiv \nabla \alpha$), regardless of the values of α_k and the resulting $\hat{\alpha}_k \in \{0, 1\}$. To stabilize training, we then co-train the continuous and discrete models:

$$\mathcal{L}_C^{\text{bin}} = \mathbb{E}_{\mathbf{r}} \|\hat{\mathbf{C}}(\mathbf{r}) - \mathbf{C}_{\text{gt}}(\mathbf{r})\|_2^2 \quad (13)$$

$$\mathcal{L}_C^{\text{stage2}} = \frac{1}{2} \mathcal{L}_C^{\text{bin}} + \frac{1}{2} \mathcal{L}_C \quad (14)$$

where $\hat{\mathbf{C}}(\mathbf{r})$ is the output radiance corresponding to the discrete opacity model $\hat{\alpha}$:

$$\hat{\mathbf{C}}(\mathbf{r}) = \sum_{k=1}^K \hat{T}_k \hat{\alpha}_k \mathbf{c}_k, \quad \hat{T}_k = \prod_{l=1}^{k-1} (1 - \hat{\alpha}_l) \quad (15)$$

Once (14) has converged, we will apply a fine-tuning step to the weights in \mathcal{F} and \mathcal{H} by minimizing $\mathcal{L}_C^{\text{bin}}$, while fixing the weights of others.

3.3. Discretization (Training Stage 3)

After binarization and fine-tuning, we convert the representation into an explicit polygonal mesh (in OBJ format). We only store quads if they are at least partially visible in the training camera poses (i.e. non-visible quads are discarded). We then create a texture image whose size is proportional to the number of visible quads, and for each quad

we allocate a $K \times K$ patch in the texture, similarly to Disney’s Ptex [8]. We use $K=17$ in our experiments, so that the quad has a 16×16 texture with half-a-pixel boundary padding. We then iterate over the pixels of the texture, convert the pixel coordinate to 3D coordinates, and *bake* the values of the discrete opacity (i.e. (3) and (12)) and features (i.e. (4)) into the texture map. We quantize the $[0, 1]$ ranges to 8-bit integers, and store the texture into (losslessly compressed) PNG images. Our experiments show that quantizing the $[0, 1]$ range with 8-bit precision, which is not accounted for during back-propagation, does not significantly affect rendering quality.

3.4. Anti-aliasing

In classic rasterization pipelines, aliasing is an issue that ought to be considered to obtain high-quality rendering. While classical NeRF hallucinates smooth edges via semi-transparent volumes, as previously discussed, semi-transparency would require per-frame polygon sorting. We overcome this issue by employing anti-aliasing by super-sampling. While we could simply execute (5) four times/pixel and average the resulting color, the execution of the deferred neural shader \mathcal{H} is the computational bottleneck of our technique. We can overcome this issue by simply averaging the features, that is, *averaging the input* of the deferred neural shader, rather than averaging its output. We first rasterize features (at $2 \times$ resolution):

$$\mathbf{F}(\mathbf{r}) = \sum_k T_k \alpha_k \mathbf{f}_k, \quad (16)$$

and then average sub-pixel features to produce the anti-aliased representation we feed to our neural deferred shader:

$$\mathbf{C}(\mathbf{r}) = \mathcal{H}(\mathbb{E}_{\mathbf{r}_\delta \sim \mathbf{r}}[\mathbf{F}(\mathbf{r}_\delta)], \mathbb{E}_{\mathbf{r}_\delta \sim \mathbf{r}}[\mathbf{d}_\delta]) \quad (17)$$

where $\mathbb{E}_{\mathbf{r}_\delta \sim \mathbf{r}}$ computes the average between the sub-pixels (i.e. four in our implementation), and \mathbf{d}_δ is the direction of ray \mathbf{r}_δ . Note how with this change we only query \mathcal{H} once per output pixel. Finally, this process is analogously applied to (15) for discrete occupancies $\hat{\alpha}$. These changes for anti-aliasing are applied in training stage 2 (14).

3.5. Rendering

The result of the optimization process is a textured polygonal mesh (where texture maps store features rather than colors) and a small MLP (which converts view direction and features to colors). Rendering this representation is done in two passes using a deferred rendering pipeline:

1. we rasterize all faces of the textured mesh with a z-buffer to produce a $2M \times 2N$ feature image with 12 channels per pixel, comprising 8 channels of learned features, a binary opacity, and a 3D view direction;

- we synthesize an $M \times N$ output RGB image by rendering a textured rectangle that uses the feature image as its texture, with linear filtering to average the features for anti-aliasing. We apply the small MLP for pixels with non-zero alphas to convert features into RGB colors. The small MLP is implemented as a GLSL fragment shader.

These rendering steps are implemented within the classic rasterization pipeline. Since z-buffering with binary transparency is order-independent, polygons *do not* need to be sorted into depth-order for each new view, and thus can be loaded into a buffer in the GPU once at the start of execution. Since the MLP for converting features to colors is very small, it can be implemented in a GLSL fragment shader [21], which is run in parallel for all pixels. These classical rendering steps are highly-optimized on GPUs, and thus our rendering system can run at interactive frame rates on a wide variety of devices; see Table 2. It is also easy to implement, since it requires only standard polygon rendering with a fragment shader. Our interactive viewer is an HTML webpage with Javascript, rendered by WebGL via the `three.js` library.

4. Experiments

We run a series of experiments to test how well MobileNeRF performs on a wide variety of scenes and devices. We test on three datasets: the 8 synthetic 360° scenes from NeRF [33], the 8 forward-facing scenes from LLFF [32], and 5 unbounded 360° outdoor scenes from Mip-NeRF 360 [3]. We compare with SNeRG [21], since, to our knowledge, it is the only NeRF model that can run in real-time on common devices. We also include extensive ablation studies to investigate the impact of different design choices.

4.1. Comparisons

To show the superior performance and compatibility of our method, we test our method and SNeRG on a variety of devices, as shown in Table 1. We report the rendering speed in Table 2. The rendering resolution is the same as the training images: 800×800 for synthetic, 1008×756 for forward-facing, and 1256×828 for unbounded. We test all methods on a chrome browser and rotate/pan the camera in a full circle to render 360 frames. Note that SNeRG is unable to represent unbounded 360° scenes due to its regular grid representation, and it does not run on phone or tablet due to compatibility or out-of-memory issues. We also report the GPU memory consumption and storage cost in Table 3. MobileNeRF requires 5x less GPU memory than SNeRG.

Rendering quality. We report the rendering quality in Table 4, while comparing with other methods using the common PSNR, SSIM [46], and LPIPS [53] metrics. Our method has roughly the same image quality as SNeRG, and

Device	Type	OS	GPU	Power
iPhone XS	Phone	iOS 15	Integrated GPU	6W
Pixel 3	Phone	Android 12	Integrated GPU	9W
Surface Pro 6	Tablet	Windows 10	Integrated GPU	15W
Chromebook	Laptop	Chrome OS	Integrated GPU	15W
Gaming laptop	Laptop	Windows 11	NVIDIA RTX 2070	115W
Desktop	PC	Ubuntu 16.04	NVIDIA RTX 2080 Ti	250W

Table 1. **Hardware specs** – of the devices used in our rendering experiments. The power is the max GPU power for discrete NVIDIA cards, and the combined max CPU and GPU power for integrated GPUs.

Dataset Method	Synthetic 360°		Forward-facing		Unbounded 360°
	Ours	SNeRG	Ours	SNeRG	Ours
iPhone XS	55.89	$0.0 \frac{2}{8}$	27.19 $\frac{2}{8}$	$0.0 \frac{2}{8}$	$22.20 \frac{3}{5}$
Pixel 3	37.14	$0.0 \frac{2}{8}$	12.40	$0.0 \frac{2}{8}$	9.24
Surface Pro 6	77.40	Unsupported	21.51	Unsupported	19.44
Chromebook	53.67	$22.62 \frac{1}{2012}$	19.44	$7.85 \frac{3}{2012}$	15.28
Gaming laptop	178.26	$8.30 \frac{1}{2012}$	57.72	3.63	55.32
Gaming laptop †	606.73	$43.87 \frac{1}{2012}$	250.17	26.01	192.59
Desktop †	744.91	207.26	349.34	50.71	279.70

Table 2. **Rendering speed** – on various devices in frames per second (FPS). The devices are on battery, except for the gaming laptop and the desktop which are plugged in, indicated with a †. The mobile devices (first four rows) have almost identical rendering speed when plugged in. With the notation $\frac{M}{N}$ we indicate that M out of N testing scenes failed to run due to out-of-memory errors.

Dataset Method	Synthetic 360°		Forward-facing		Unbounded 360°
	Ours	SNeRG	Ours	SNeRG	Ours
GPU memory	538.38	2707.25	759.25	4312.13	1162.20
Disk storage	125.75	86.75	201.50	337.25	344.60

Table 3. **Resources** – memory and disk storage (MB).

	Synthetic 360°			Forward-facing			Unbounded 360°		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
NeRF	31.00	0.947	0.081	26.50	0.811	0.250	-	-	-
JAXNeRF	31.65	0.952	0.051	26.92	0.831	0.173	21.46	0.458	0.515
NeRF++	-	-	-	-	-	-	22.76	0.548	0.427
SNeRG	30.38	0.950	0.050	25.63	0.818	0.183	-	-	-
Ours	30.90	0.947	0.062	25.91	0.825	0.183	21.95	0.470	0.470

Table 4. **Quantitative Analysis** – For NeRF [33] and NeRF++ [52], we dash entries where the original papers did not report quantitative performance. For SNeRG, while one could extend the method to include the unbounded design from [3], implementing this is far from trivial. Our method can be easily adapted to work across all modalities.

	Synthetic 360°		Forward-facing		Unbounded 360°	
	V	T	V	T	V	T
Number	494,289	224,341	830,076	338,535	1,436,033	608,785
Percentage	1.964%	1.783%	3.298%	2.690%	4.891%	4.147%

Table 5. **Polygon count** – Average number of vertices and triangles produced, and their percentage compared to all available vertices/triangles in the initial mesh.

better than NeRF. Visual results are shown in Figure 6 (a-c). Our method achieves image quality similar to SNeRG when the camera is at an appropriate distance. When the camera is zoomed in, SNeRG tends to render over-smoothed images.

Polygon count. Table 5 shows the average number of vertices and triangles produced by our method, and the per-

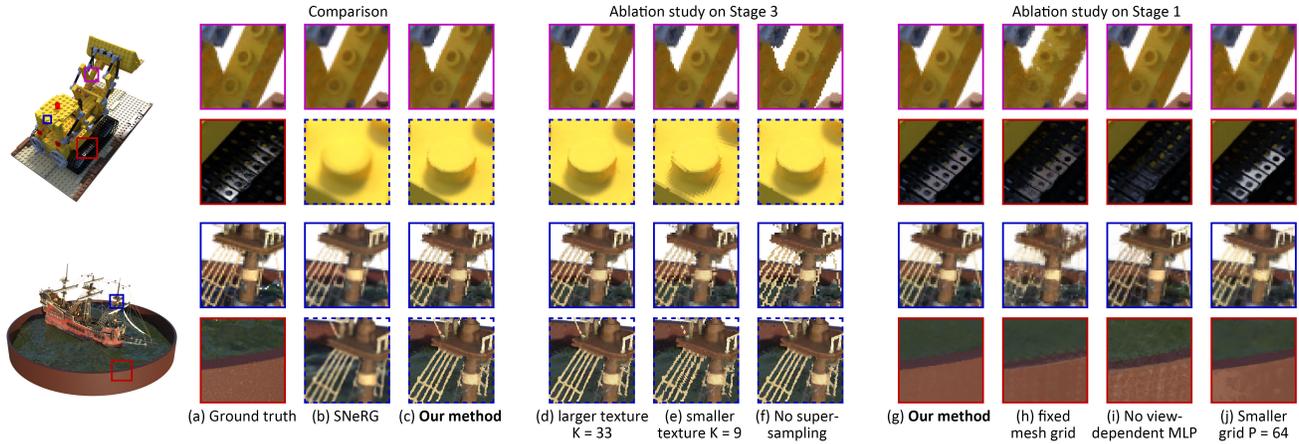


Figure 6. **Qualitative Results** – Comparisons to the state-of-the-art and ablation studies. With a *solid* line we denote zoom-ins of the rendered (800×800) image, while with a *dashed* line we move the camera to zoom-in onto the same detail.

	Synthetic 360°		Forward-facing	
	PSNR↑	SSIM↑	PSNR↑	SSIM↑
Stage 1, our method	32.13	0.955	26.57	0.839
Stage 1, fixed mesh grid	29.87	0.938	25.43	0.797
Stage 1, no view-dependent MLP	29.91	0.935	25.91	0.824
Stage 1, smaller grid $P=128 \rightarrow 64$	31.58	0.952	26.39	0.831
Stage 1, no acceleration grid	31.77	0.953	26.61	0.835
Stage 2, our method	31.01	0.948	26.32	0.833
Stage 2, no fine-tuning	30.80	0.946	26.25	0.832
Stage 2, only pseudo-gradients	29.70	0.935	26.01	0.820
Stage 2, binary loss	30.89	0.947	26.32	0.832
Stage 3, our method	30.90	0.947	25.91	0.825
Stage 3, larger texture $K=17 \rightarrow 33$	30.99	0.948	26.14	0.830
Stage 3, smaller texture $K=17 \rightarrow 9$	30.49	0.945	24.85	0.796
Stage 3, no supersampling	29.26	0.937	24.88	0.799

Table 6. **Ablation** – rendering quality.

	Speed in FPS			Space in MB	
	Pixel 3	Surface Pro 6	Gaming laptop ♣	GPU memory	Disk storage
Synthetic 360° scenes					
our method	37.14	77.40	606.73	538.38	125.75
Larger texture $K = 33$	32.48 $\frac{3}{4}$	59.15	589.20	1290.88	283.50
Smaller texture $K = 9$	37.74	94.62	617.74	336.63	67.00
No supersampling	51.81	113.41	649.86	440.25	125.75
No view-dependent MLP	52.16	96.76	638.30	538.38	125.75
Forward-facing scenes					
our method	12.40	21.51	250.17	759.25	201.50
Larger texture $K = 33$	12.88 $\frac{3}{4}$	18.79	241.52	2024.13	462.75
Smaller texture $K = 9$	12.70	23.61	257.64	394.13	105.75
No supersampling	16.97	42.11	413.02	645.00	201.50
No view-dependent MLP	23.72	28.06	385.65	759.25	201.50

Table 7. **Ablation** – rendering speed/memory.

centage compared to all available vertices/triangles in the initial mesh. As we only retain visible triangles, most vertices/triangles are removed in the final mesh.

Shading mesh. In Figure 2a and Figure 7, we show the extracted triangle meshes without the textures. Most triangle faces do not align with the actual object surface. This is perhaps due to the ambiguity that good rendering quality can be achieved despite how the triangles are aligned. For example, the results of our method after Stage 1 in Table 6 is similar to other methods in Table 4. Therefore, better regularization losses or training objectives need to be devised if one wishes to have better surface quality. However, optimizing vertices does improve the rendering quality, as shown in Figure 6h.

4.2. Ablation studies

In Table 6, we show the rendering quality of our method at each stage, and report our ablation studies. The rendering quality gradually drops after each stage, because each stage adds more constraints to the model. In Stage 1, the performance drops significantly if we use a fixed regular grid mesh instead of having optimizable mesh vertices, or

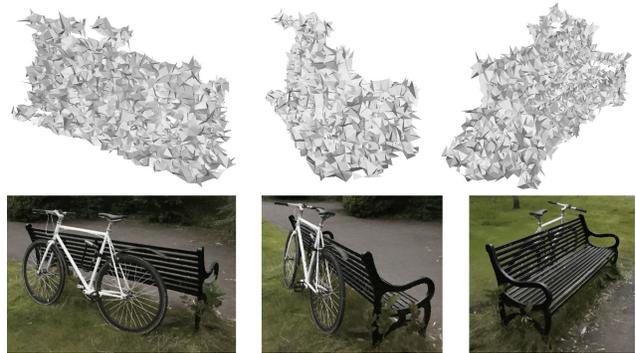


Figure 7. **Shading mesh** – not textured. The mesh corresponds to the bicycle (see Figure 1). We manually removed the background mesh to better show the geometry of the object. Zoom-in to see more details. In the bottom, we also show the rendered images of our method. Note how the coarse mesh is able to represent detailed structures such as the spokes of the wheels and the wires around the handles, thanks to high-resolution textures with transparencies.

if we forgo view-dependent effects by directly predicting the color and alpha of each point. The performance drops slightly if the grid is smaller ($P=64$ vs. 128). If we remove the acceleration grid, we are not able to quadruple the batch



Figure 8. **Limitations** – (a) the monitor/table are hollow, because the reflections are modelled as real objects behind the monitor and below the table. (b) our method generates scattered small fragments in the semi-transparent parts. (c) the camera is too close to the scene and details in the grass cannot be represented at the chosen texture resolution.

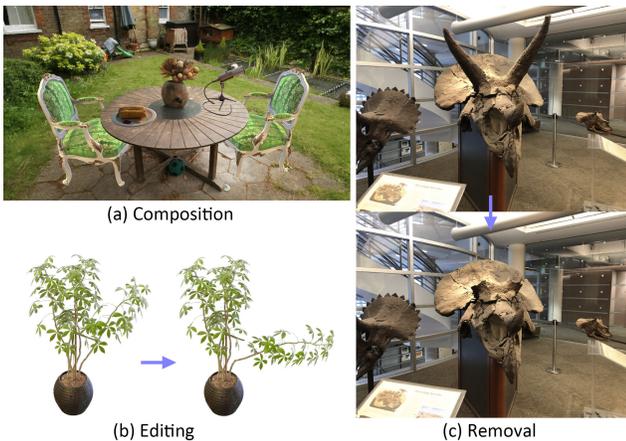


Figure 9. **Scene editing** – (a) four objects learned from the synthetic scenes are added into an unbounded scene. (b) a branch of the ficus is bent. (c) the horns are removed.

size during training; the performance drops if we train this model the same number of iterations as our method. Note that the PSNR of this model is higher on forward-facing scenes. This is because the acceleration grid will remove cells that are not visible in the training images, thus cannot “inpaint” the objects and may leave holes. In Stage 2, if we do not perform the fine-tuning step that only optimizes \mathcal{F} and \mathcal{H} and fix the weights of others, the performance drops. If we only use the binary opacity with pseudo-gradients by applying $\mathcal{L}_C^{\text{stage2}} = \mathcal{L}_C^{\text{bin}}$ instead of Eq. 14, the performance drops. If we use a binary loss on the predicted opacity, e.g., $\mathcal{L}_{\text{binary}} = -\sum |\alpha_k - 0.5|$, instead of using the pseudo-gradients with $\hat{C}(\mathbf{r})$, the performance drops slightly. In stage 3, when we use a larger texture size $K=33$ instead of 17, the performance improves, but the texture size will be quadrupled; the performance drops when we use a smaller texture size $K=9$. If we remove the super-

sampling step, the performance drops significantly. Visual results are shown in Figure 6. We omit some models because they do not have significant visual differences compared to our method. Notice the squared pixels of the texture images are clearly visible in the dashed-line boxes in (e) and almost invisible in (d). The aliasing artifacts are conspicuous in the solid-line boxes in (f). In Stage 1, if the grid vertices cannot be optimized, the results will be significantly worse, as shown in (h). Without the small MLP, the model cannot handle reflections, as shown in (i). Changing to a smaller grid size introduces some minor artifacts in (j). In Table 7, we show the rendering speed and space cost if we use a larger or smaller texture size, or if we remove the super-sampling step, or if we only perform the rasterization without using the small MLP to predict the view-dependent colors. One can find that the super-sampling step and the small MLP have the most significant impact.

5. Conclusions

We introduce MobileNeRF, an architecture that takes advantage of the classical rasterization pipeline (i.e. z-buffers and fragment shaders) to perform efficient rendering of surface-based neural fields on a wide range of compute platforms. It achieves frame rates an order of magnitude faster than the previous state-of-the-art (SNeRG) while producing images of equivalent quality.

Limitations. Our estimated *surface* may be incorrect, especially for scenes with specular surfaces and/or sparse views (Figure 8a); it uses *binary* opacities to avoid sorting polygons, and thus cannot handle scenes with semi-transparencies (Figure 8b); it uses fixed mesh and texture resolutions, which may be too coarse for close-up novel-view synthesis (Figure 8c); it models a radiance field without explicitly decomposing illumination and reflectance, and thus does not handle glossy surfaces as well as recent methods [45]. Extending the polygon rendering pipeline with efficient partial sorting, levels-of-detail, mipmaps, and surface shading should address some of these issues. Also, the current training speed of MobileNeRF is slow due to NeRF’s MLP backbone. The extension of MobileNeRF to fast-training architectures (e.g., Instant NGP [34]) constitutes an exciting avenue for future works.

The explicit mesh representation provided by MobileNeRF gives us direct editing control over the NeRF model without any complex architectural change (e.g. ControlNerf [26]), but in this paper we only superficially investigated these possibilities; see Figure 9 and the videos in the [supplementary](#) (Section B).

Acknowledgements. We thank the reviewers as well as Simon Kornblinth, Ting Chen, Daniel Rebain, Kevin Swersky, and David Fleet for their valuable feedback.

References

- [1] Benjamin Attal, Jia-Bin Huang, Michael Zollhöfer, Johannes Kopf, and Changil Kim. Learning neural light fields with ray-space embedding networks. *CVPR*, 2022. 2, 4
- [2] Benjamin Attal, Selena Ling, Aaron Gokaslan, Christian Richardt, and James Tompkin. MatryODShka: Real-time 6DoF video view synthesis using multi-sphere images. *ECCV*, 2020. 2
- [3] Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. *CVPR*, 2022. 2, 3, 6
- [4] Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013. 5
- [5] Tobias Bertel, Mingze Yuan, Reuben Lindroos, and Christian Richardt. OmniPhotos: Casual 360° VR photography. *ACM Transactions on Graphics*, 2020. 2
- [6] Michael Broxton, John Flynn, Ryan Overbeck, Daniel Erickson, Peter Hedman, Matthew DuVall, Jason Dourgarian, Jay Busch, Matt Whalen, and Paul Debevec. Immersive light field video with a layered mesh representation. *ACM Transactions on Graphics*, 2020. 2
- [7] Chris Buehler, Michael Bosse, Leonard McMillan, Steven Gortler, and Michael Cohen. Unstructured lumigraph rendering. In *Proceedings of Computer graphics and interactive techniques*, 2001. 2
- [8] Brent Burley and Dylan Lacewell. Ptex: Per-face texture mapping for production rendering. In *Computer Graphics Forum*, 2008. 5
- [9] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. *ECCV*, 2022. 2
- [10] Zhiqin Chen, Andrea Tagliasacchi, Thomas Funkhouser, and Hao Zhang. Neural dual contouring. *ACM Transactions on Graphics*, 2022. 4
- [11] Forrester Cole, Kyle Genova, Avneesh Sud, Daniel Vlasic, and Zhoutong Zhang. Differentiable surface rendering via non-differentiable sampling. In *ICCV*, 2021. 2
- [12] Abe Davis, Marc Levoy, and Fredo Durand. Unstructured light fields. *Computer Graphics Forum*, 2012. 2
- [13] Paul Debevec, Yizhou Yu, and George Borshukov. Efficient view-dependent image-based rendering with projective texture-mapping. In *Eurographics Workshop on Rendering Techniques*, 1998. 2
- [14] Boyang Deng, Jonathan T. Barron, and Pratul P. Srinivasan. JaxNeRF: an efficient JAX implementation of NeRF, 2020. 3, 4, 5
- [15] John Flynn, Michael Broxton, Paul Debevec, Matthew DuVall, Graham Fyffe, Ryan Overbeck, Noah Snively, and Richard Tucker. DeepView: View synthesis with learned gradient descent. *CVPR*, 2019. 2
- [16] Jun Gao, Wenzheng Chen, Tommy Xiang, Clement Fuji Tsang, Alec Jacobson, Morgan McGuire, and Sanja Fidler. Learning deformable tetrahedral meshes for 3d reconstruction. In *NeurIPS*, 2020. 2
- [17] Stephan J. Garbin, Marek Kowalski, Matthew Johnson, Jamie Shotton, and Julien P. C. Valentin. Fastnerf: High-fidelity neural rendering at 200fps. In *ICCV*, 2021. 3
- [18] Lily Goli, Daniel Rebain, Sara Sabour, Animesh Garg, and Andrea Tagliasacchi. nerf2nerf: Pairwise registration of neural radiance fields. *ICRA*, 2023. 4
- [19] Steven J. Gortler, Radek Grzeszczuk, Richard Szeliski, and Michael F. Cohen. The lumigraph. *SIGGRAPH*, 1996. 2
- [20] Peter Hedman, Julien Philip, True Price, Jan-Michael Frahm, George Drettakis, and Gabriel Brostow. Deep blending for free-viewpoint image-based rendering. *ACM Transactions on Graphics*, 2018. 2
- [21] Peter Hedman, Pratul P. Srinivasan, Ben Mildenhall, Jonathan T. Barron, and Paul Debevec. Baking neural radiance fields for real-time view synthesis. *ICCV*, 2021. 1, 3, 6, 4
- [22] Tao Ju, Frank Losasso, Scott Schaefer, and Joe Warren. Dual contouring of Hermite data. *ACM Transactions on graphics*, 2002. 4
- [23] Animesh Karnewar, Tobias Ritschel, Oliver Wang, and Niloy J. Mitra. Relu fields: The little non-linearity that could. *ACM Transactions on Graphics*, 2022. 2
- [24] Georgios Kopanas, Julien Philip, Thomas Leimkühler, and George Drettakis. Point-based neural rendering with per-view optimization. In *Computer Graphics Forum*, 2021. 2
- [25] Christoph Lassner and Michael Zollhofer. Pulsar: Efficient sphere-based neural rendering. *CVPR*, 2021. 2
- [26] Verica Lazova, Vladimir Guzov, Kyle Olszewski, Sergey Tulyakov, and Gerard Pons-Moll. Control-nerf: Editable feature volumes for scene rendering and manipulation. *arXiv preprint arXiv:2204.10850*, 2022. 8
- [27] Marc Levoy and Pat Hanrahan. Light field rendering. *SIGGRAPH*, 1996. 2
- [28] Zhi-Hao Lin, Wei-Chiu Ma, Hao-Yu Hsu, Yu-Chiang Frank Wang, and Shenlong Wang. Neurmips: Neural mixture of planar experts for view synthesis. *CVPR*, 2022. 2
- [29] David B. Lindell, Julien N.P. Martel, and Gordon Wetzstein. Autoint: Automatic integration for fast neural rendering. *CVPR*, 2021. 2
- [30] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: Learning dynamic renderable volumes from images. *SIGGRAPH*, 2019. 2
- [31] Ricardo Martin-Brualla, Rohit Pandey, Shuoran Yang, Pavel Pidlypenskyi, Jonathan Taylor, Julien Valentin, Sameh Khamis, Philip Davidson, Anastasia Tkach, Peter Lincoln, Adarsh Kowdle, Christoph Rhemann, Dan B Goldman, Cem Keskin, Steve Seitz, Shahram Izadi, and Sean Fanello. Lookingood: Enhancing performance capture with real-time neural re-rendering. *ACM Transactions on Graphics*, 2018. 2
- [32] 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*, 2019. 2, 6

- [33] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. 1, 2, 4, 6, 3
- [34] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics*, 2022. 2, 4, 8
- [35] Jacob Munkberg, Jon Hasselgren, Tianchang Shen, Jun Gao, Wenzheng Chen, Alex Evans, Thomas Müller, and Sanja Fidler. Extracting triangular 3d models, materials, and lighting from images. In *CVPR*, 2022. 2
- [36] Thomas Neff, Pascal Stadlbauer, Mathias Parger, Andreas Kurz, Joerg H Mueller, Chakravarty R Alla Chaitanya, Anton Kaplanyan, and Markus Steinberger. Donerf: Towards real-time rendering of compact neural radiance fields using depth oracle networks. In *Computer Graphics Forum*, 2021. 2
- [37] Eric Penner and Li Zhang. Soft 3D reconstruction for view synthesis. *ACM Transactions on Graphics*, 2017. 2
- [38] Daniel Rebain, Wei Jiang, Soroosh Yazdani, Ke Li, Kwang Moo Yi, and Andrea Tagliasacchi. DeRF: Decomposed radiance fields. *CVPR*, 2021. 2
- [39] Christian Reiser, Songyou Peng, Yiyi Liao, and Andreas Geiger. Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps. In *ICCV*, 2021. 2
- [40] Darius Rückert, Linus Franke, and Marc Stamminger. Adop: Approximate differentiable one-pixel point rendering. *arXiv preprint arXiv:2110.06635*, 2021. 2
- [41] Sara Fridovich-Keil and Alex Yu, Matthew Tancik, Qinlong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *CVPR*, 2022. 2, 4
- [42] Tianchang Shen, Jun Gao, Kangxue Yin, Ming-Yu Liu, and Sanja Fidler. Deep marching tetrahedra: a hybrid representation for high-resolution 3d shape synthesis. In *NeurIPS*, 2021. 2
- [43] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. *CVPR*, 2022. 2
- [44] Justus Thies, Michael Zollhöfer, and Matthias Nießner. Deferred neural rendering: Image synthesis using neural textures. *ACM Transactions on Graphics*, 2019. 2
- [45] 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. *CVPR*, 2022. 2, 8
- [46] Zhou 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*, 2004. 6
- [47] Suttisak Wizadwongsa, Pakkapon Phongthawee, Jiraphon Yenphraphai, and Supasorn Suwajanakorn. Nex: Real-time view synthesis with neural basis expansion. *CVPR*, 2021. 2
- [48] Liwen Wu, Jae Yong Lee, Anand Bhattad, Yuxiong Wang, and David Forsyth. Diver: Real-time and accurate neural radiance fields with deterministic integration for volume rendering. *CVPR*, 2022. 2
- [49] Xiuchao Wu, Jiamin Xu, Zihan Zhu, Hujun Bao, Qixing Huang, James Tompkin, and Weiwei Xu. Scalable neural indoor scene rendering. *ACM Transactions on Graphics*, 2022. 2
- [50] Qiangeng Xu, Zexiang Xu, Julien Philip, Sai Bi, Zhixin Shu, Kalyan Sunkavalli, and Ulrich Neumann. Point-nerf: Point-based neural radiance fields. *CVPR*, 2022. 2
- [51] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. PlenOctrees for real-time rendering of neural radiance fields. In *ICCV*, 2021. 1, 3
- [52] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. Nerf++: Analyzing and improving neural radiance fields. *arXiv preprint arXiv:2010.07492*, 2020. 6, 5
- [53] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 6
- [54] Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. Stereo magnification: Learning view synthesis using multiplane images. *ACM Transactions on Graphics*, 2018. 2