

Baking Neural Radiance Fields for Real-Time View Synthesis

Peter Hedman Pratul P. Srinivasan Ben Mildenhall Jonathan T. Barron Paul Debevec
Google Research
<https://nerf.live>

Abstract

Neural volumetric representations such as Neural Radiance Fields (NeRF) have emerged as a compelling technique for learning to represent 3D scenes from images with the goal of rendering photorealistic images of the scene from unobserved viewpoints. However, NeRF’s computational requirements are prohibitive for real-time applications: rendering views from a trained NeRF requires querying a multilayer perceptron (MLP) hundreds of times per ray. We present a method to train a NeRF, then precompute and store (i.e. “bake”) it as a novel representation called a Sparse Neural Radiance Grid (SNeRG) that enables real-time rendering on commodity hardware. To achieve this, we introduce 1) a reformulation of NeRF’s architecture, and 2) a sparse voxel grid representation with learned feature vectors. The resulting scene representation retains NeRF’s ability to render fine geometric details and view-dependent appearance, is compact (averaging less than 90 MB per scene), and can be rendered in real-time (higher than 30 frames per second on a laptop GPU). Actual screen captures are shown in our video.

1. Introduction

The task of view synthesis — using observed images to recover a 3D scene representation that can render the scene from novel unobserved viewpoints — has recently seen dramatic progress as a result of using neural volumetric representations. In particular, Neural Radiance Fields (NeRF) [31] are able to render photorealistic novel views with fine geometric details and realistic view-dependent appearance by representing a scene as a continuous volumetric function, parameterized by a multilayer perceptron (MLP) that maps from a continuous 3D position to the volume density and view-dependent emitted radiance at that location. Unfortunately, NeRF’s rendering procedure is quite slow: rendering a ray requires querying an MLP hundreds of times, such that rendering a frame at 800×800 resolution takes roughly a minute on a modern GPU. This prevents NeRF from being used for interactive view synthesis




Figure 1: Our method “bakes” NeRF’s continuous neural volumetric scene representation into a discrete Sparse Neural Radiance Grid (SNeRG) for real-time rendering on commodity hardware (~ 65 frames per second on a MacBook Pro in the example shown here). Our method is more than two orders of magnitude faster than prior work for accelerating NeRF’s rendering procedure and more than an order of magnitude faster than the next-fastest alternative (Neural Volumes) while achieving substantially higher-quality.

applications such as virtual and augmented reality, or even simply inspecting a recovered 3D model in a web browser.

We address the problem of rendering a trained NeRF in real-time, see Figure 1. Our approach accelerates rendering by three orders of magnitude, achieving 12 milliseconds per frame on a single GPU. We precompute and store (i.e. “bake”) a trained NeRF into a sparse 3D voxel grid data structure, which we call a Sparse Neural Radiance Grid (SNeRG). Each active voxel in a SNeRG contains opacity, diffuse color, and a learned feature vector that encodes

view-dependent effects. To render this representation, we first accumulate the diffuse colors and feature vectors along each ray. Next, we pass the accumulated feature vector through a lightweight MLP to produce a view-dependent residual that is added to the accumulated diffuse color.

We introduce two key modifications to NeRF to effectively bake it into this sparse voxel representation: 1) we design a “deferred” NeRF architecture that represents view-dependent effects with an MLP that only runs once per *pixel* (instead of once per 3D sample as done in NeRF), and 2) we regularize NeRF’s predicted opacity field during training to encourage sparsity, which improves both the storage cost and rendering time for the resulting SNeRG.

We demonstrate that our approach is able to increase the rendering speed of NeRF so that frames can be rendered in real-time, while retaining NeRF’s ability to represent fine geometric details and convincing view-dependent effects. Furthermore, our representation is compact, and requires less than 90 MB on average to represent a scene.

2. Related work

In this section, we review scene representations used for view synthesis with a specific focus on their ability to support real-time rendering, and discuss prior work in efficient representation and rendering of volumetric representations.

Scene Representations for View Synthesis The task of view synthesis, using observed images of an object or scene to render photorealistic images from novel unobserved viewpoints, has a rich history in graphics and computer vision. When the scene is captured by densely-sampled images, light field rendering techniques [9, 16, 23] can be used to efficiently render novel views by interpolating between sampled rays. Unfortunately, the sampling and storage requirements of light field interpolation techniques are typically intractable for significant viewpoint motion. Methods that aim to support free-viewpoint rendering from sparsely-sampled images typically reconstruct an explicit 3D representation of the scene [10]. One popular class of methods uses mesh-based representations, with either diffuse [29] or view-dependent [6, 10, 53] appearance. Recent methods trained deep networks to increase the quality of mesh renderings, improving robustness to errors in the reconstructed mesh geometry [17, 48]. Mesh-based approaches are naturally amenable to real-time rendering with modern rasterization pipelines. However, as gradient-based optimization of a rendering loss with a mesh representation is difficult, these methods struggle to reconstruct fine structures and detailed geometry.

Another popular class of methods uses discretized volumetric representations such as voxel grids [27, 41, 42, 45] or multiplane images [13, 37, 46, 56]. While volumetric approaches are better suited to gradient-based optimization, discretized voxel representations are fundamentally limited

by their cubic scaling. This restricts them to relatively low resolutions (for voxel grids), or rendering from a limited range of viewpoints (for multiplane images).

NeRF [31] replaces these discretized volumetric representations with an MLP that represents a scene as a *continuous neural volumetric function*, mapping from a 3D coordinate to the volume density and view-dependent emitted radiance at that position. NeRF has been remarkably successful for view synthesis, and follow-on works have extended it for generative modeling [7, 40], dynamic scenes [24, 35], non-rigidly deforming objects [14, 36], and relighting [3, 44]. NeRF is able to represent detailed geometry and realistic appearance extremely efficiently (NeRF uses approximately 5 MB of network weights to represent each scene), but this comes at the cost of slow rendering. NeRF needs to query its MLP hundreds of times per ray, and requires roughly a minute to render a single frame. We specifically address this issue, and present a method that enables a trained NeRF to be rendered in real-time.

Recent works improve efficiency of NeRF’s neural volumetric rendering. AutoInt [25] designs a network architecture that automatically computes integrals along rays, which enables a piecewise ray-marching procedure with fewer MLP evaluations. Neural Sparse Voxel Fields [26] store a 3D voxel grid of latent codes, and sparsifies this grid during training to enable NeRF to skip free space during rendering. Decomposed Radiance Fields [38] represents a scene using a set of smaller MLPs instead of a single large MLP. However, these methods achieve moderate speedups of around 10× at best, and are therefore not suited for real-time rendering. In contrast to these methods, we specifically focus on accelerating the rendering of a NeRF after it has been trained, allowing us to leverage precomputation strategies that are difficult to incorporate during training.

Concurrent to our work, several approaches have explored real-time rendering of NeRF reconstructions. Rather than precomputing, Reiser et al. [39] represent the scene as a 3D grid of smaller MLPs. Neff et al. [32] predict sample locations with an MLP and can thus render images with much fewer samples per ray. It is also possible to bake view-dependent colors — the field of precomputed radiance transfer explores spherical basis functions [43, 33, 49] suitable for this purpose. Yu et al. [54] store the scene in a multi-resolution voxel octree and encode view-dependent color as spherical harmonics basis functions at every voxel. Garbin et al. [15] instead factorize the view-dependent computation as a dot product between a global spherical function and a per-sample feature vector. While fully baked representations [15, 54] perform the least computation per 3D sample point, our representation, KiloNeRF [39] and DONeRF [32] compress the scene more effectively, as they represent view-dependence with MLPs instead of precomputed spherical functions.

Efficient Volume Rendering Discretized volumetric rep-

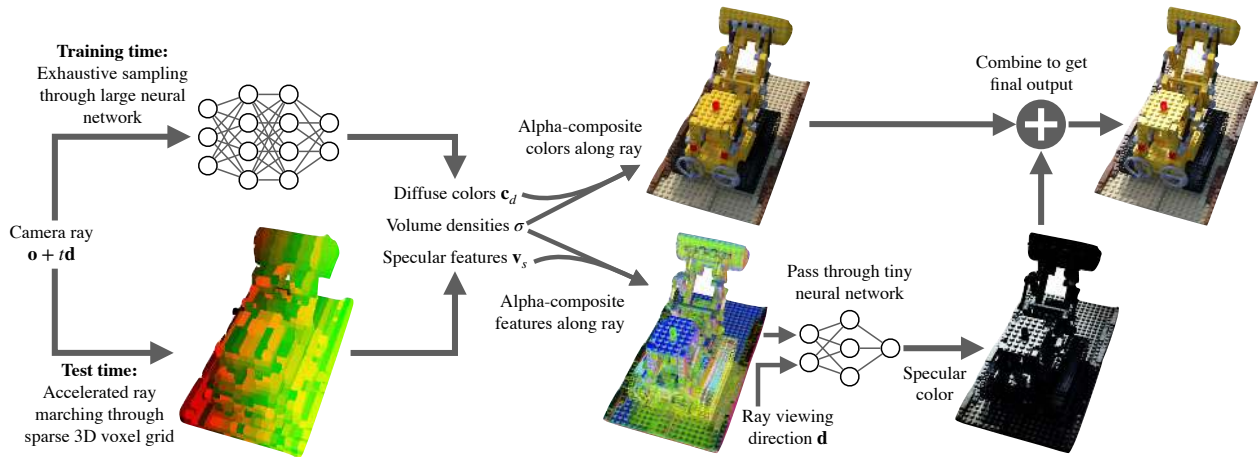


Figure 2: **Pipeline overview.** We jointly design a “deferred” NeRF as well as a procedure to precompute and store its the outputs in a sparse 3D voxel grid data structure for real-time rendering. At training time, we query the deferred NeRF’s MLP for the diffuse color, volume density, and feature-vector at any 3D location along a ray. At test time, we instead precompute and store these values in a sparse 3D voxel grid. Next, we alpha-composite the diffuse colors and feature vectors along the ray. Once the ray has terminated, we use another MLP to predict a view-dependent specular color from the accumulated feature vector, diffuse color, and the ray’s viewing direction. Note that this MLP is only run once per pixel.

representations have been used extensively in computer graphics to make rendering more efficient both in terms of storage and rendering speed. Early works [19, 22, 52] primarily focused on fast rendering of dense voxel grids. However, as shown by Laine and Karras [20], sparse voxel grids can be an effective and efficient representation for opaque surfaces. In situations where large regions of space share the same data value or a prefiltered representation is required to combat aliasing, hierarchical representations such as sparse voxel octrees [8] are a popular data structure. When a variable level-of-detail is not required, sparse voxel grids can be efficiently represented with hash tables [21, 34]. However, hashing each voxel independently can lead to incoherent memory fetches when traversing the representation during rendering. We make a deliberate trade-off to use a *block-sparse* representation, which improves memory coherence but slightly increases the size of our representation.

Our work aims to combine the reconstruction quality and view-dependence of NeRF with the speed of efficient volume rendering techniques. We achieve this by extending deferred neural rendering [48] to volumetric scene representations. This allows us to visualize trained NeRFs in real-time on commodity hardware, with minimal quality degradation.

3. Method Overview

Our overall goal is to design a practical representation that enables serving and real-time rendering of scenes reconstructed by NeRF. This implies three requirements: 1) Rendering a 800×800 resolution frame (the resolution used by NeRF) should require less than 30 milliseconds on com-

modity hardware. 2) The representation should be compressible to 100 MB or less. 3) The uncompressed representation should fit within GPU memory (approximately 4 GB) and should not require streaming.

Rendering a standard NeRF in real-time is completely intractable on current hardware. NeRF requires about 100 teraflops to render a single 800×800 frame, which results in a best-case rendering time of 10 seconds per frame on an NVIDIA RTX 2080 GPU with full GPU utilization. We must therefore exchange some of this computation for storage. However, we do not want to precompute and store the entire 5D view-dependent representation [16, 23], as that would require a prohibitive amount of GPU memory.

We propose a hybrid approach that precomputes and stores some content in a sparse 3D data structure but defers the computation of view-dependent effects to rendering time. We jointly design a reformulation of NeRF (Section 4) as well as a procedure to bake this modified NeRF into a discrete volumetric representation that is suited for real-time rendering (Section 5).

4. Modifying NeRF for Real-time Rendering

We reformulate NeRF in three ways: 1) we limit the computation of view-dependent effects to a single network evaluation *per ray*, 2) we introduce a small bottleneck in the network architecture that can be efficiently stored as 8 bit integers, and 3) we introduce a sparsity loss during training, which concentrates the opacity field around surfaces in the scene. Here, we first review NeRF’s architecture and rendering procedure before describing our modifications.

4.1. Review of NeRF

NeRF represents a scene as a continuous volumetric function parameterized by a MLP. Concretely, the 3D position $\mathbf{r}(t)$ and viewing direction \mathbf{d} along a camera ray $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$, are passed as inputs to an MLP with weights Θ to produce the volume density σ of particles at that location as well as the RGB color \mathbf{c} corresponding to the radiance emitted by particles at the input location along the input viewing direction:

$$\sigma(t), \mathbf{c}(t) = \text{MLP}_{\Theta}(\mathbf{r}(t), \mathbf{d}). \quad (1)$$

A key design decision made in NeRF is to architect the MLP such that volume density is only predicted as a function of 3D position, while emitted radiance is predicted as a function of both 3D position and 2D viewing direction.

To render the color $\hat{\mathbf{C}}(\mathbf{r})$ of a pixel, NeRF queries the MLP at sampled positions t_k along the corresponding ray and uses the estimated volume densities and colors to approximate a volume rendering integral using numerical quadrature, as discussed by Max [28]:

$$\hat{\mathbf{C}}(\mathbf{r}) = \sum_k T(t_k) \alpha(\sigma(t_k)\delta_k) \mathbf{c}(t_k), \quad (2)$$

$$T(t_k) = \exp\left(-\sum_{k'=1}^{k-1} \sigma(t_{k'})\delta_{k'}\right), \quad \alpha(x) = 1 - \exp(-x),$$

where $\delta_k = t_{k+1} - t_k$ is the distance between two adjacent points along the ray.

NeRF trains the MLP by minimizing the squared error between input pixels from a set of observed images (with known camera poses) and the pixel values predicted by rendering the scene as described above:

$$\mathcal{L}_r = \sum_i \|\mathbf{C}(\mathbf{r}_i) - \hat{\mathbf{C}}(\mathbf{r}_i)\|_2^2 \quad (3)$$

where $\mathbf{C}(\mathbf{r}_i)$ is the color of pixel i in the input images.

By replacing a traditional discrete volumetric representation with an MLP, NeRF makes a strong space-time trade-off: NeRF’s MLP requires multiple orders of magnitude less space than a dense voxel grid, but accessing the properties of the volumetric scene representation at any location requires an MLP evaluation instead of a simple memory lookup. Rendering a single ray that passes through the volume requires hundreds of these MLP queries, resulting in extremely slow rendering times. This tradeoff is beneficial during training; since we do not know where the scene geometry lies during optimization, it is crucial to use a compact representation that can represent highly-detailed geometry at arbitrary locations. However, after a NeRF has been trained, we argue that it is prudent to rethink this space-time tradeoff and bake the NeRF representation into a data structure that stores pre-computed values from the MLP to enable real-time rendering.

4.2. Deferred NeRF Architecture

NeRF’s MLP can be thought of as predicting a 256-dimensional feature vector for each input 3D location, which is then concatenated with the viewing direction and decoded into an RGB color. NeRF then accumulates these view-dependent colors into a single pixel color. However, evaluating an MLP at every sample along a ray to estimate the view-dependent color is prohibitively expensive for real-time rendering. Instead, we modify NeRF to use a strategy similar to deferred rendering [11, 48]. We restructure NeRF to output a diffuse RGB color \mathbf{c}_d and a 4-dimensional feature vector \mathbf{v}_s (which is constrained to $[0, 1]$ via a sigmoid so that it can be compressed, as discussed in Section 5.4) in addition to the volume density σ at each input 3D location:

$$\sigma(t), \mathbf{c}_d(t), \mathbf{v}_s(t) = \text{MLP}_{\Theta}(\mathbf{r}(t)). \quad (4)$$

To render a pixel, we accumulate the diffuse colors and feature vectors along each ray and pass the accumulated feature vector and color, concatenated to the ray’s direction, to a very small MLP with parameters Φ (2 layers with 16 channels each) to produce a view-dependent residual that we add to the accumulated diffuse color:

$$\hat{\mathbf{C}}_d(\mathbf{r}) = \sum_k T(t_k) \alpha(\sigma(t_k)\delta_k) \mathbf{c}_d(t_k), \quad (5)$$

$$\mathbf{V}_s(\mathbf{r}) = \sum_k T(t_k) \alpha(\sigma(t_k)\delta_k) \mathbf{v}_s(t_k), \quad (6)$$

$$\hat{\mathbf{C}}(\mathbf{r}) = \hat{\mathbf{C}}_d(\mathbf{r}) + \text{MLP}_{\Phi}(\mathbf{V}_s(\mathbf{r}), \mathbf{d}) \quad (7)$$

This modification enables us to precompute and store the diffuse colors and 4-dimensional feature vectors within our sparse voxel grid representation discussed below. Critically, we only need to evaluate the MLP_{Φ} to produce view-dependent effects once per *pixel*, instead of once per sample in 3D space as in the standard NeRF model.

4.3. Opacity Regularization

Both the rendering time and required storage for a volumetric representation strongly depend on the sparsity of opacity within the scene. To encourage NeRF’s opacity field to be sparse, we add a regularizer that penalizes predicted density using a Cauchy loss during training:

$$\mathcal{L}_s = \lambda_s \sum_{i,k} \log\left(1 + \frac{\sigma(\mathbf{r}_i(t_k))^2}{c}\right), \quad (8)$$

where i indexes pixels in the input (training) images, k indexes samples along the corresponding rays, and hyperparameters λ_s and c control the magnitude and scale of the regularizer respectively ($\lambda_s = 10^{-4}$ and $c = 1/2$ in all experiments). To ensure that this loss is not unevenly applied due to NeRF’s hierarchical sampling procedure, we only compute it for the “coarse” samples that are distributed with

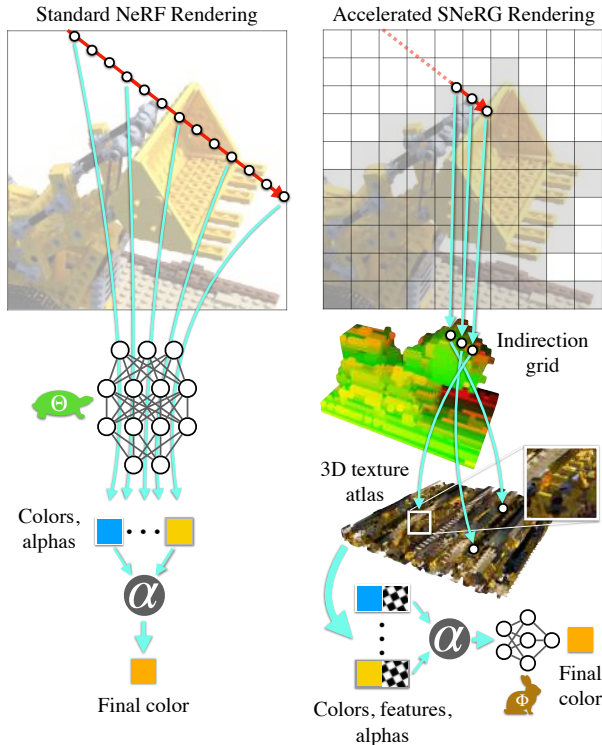


Figure 3: **Ray-marching for NeRF vs. SNeRG.** Left: To render a ray with NeRF, we densely sample points along the ray and pass the coordinates to a large MLP to compute colors and opacities, which are alpha-composited into a pixel color. Right: SNeRG significantly accelerates rendering by replacing compute-intensive MLP evaluations with lookups into a precomputed sparse 3D grid. We use an indirection grid to map occupied voxel blocks to locations within a compact 3D texture atlas. During ray-marching, we skip unoccupied blocks, and alpha-composite the diffuse colors and feature vectors fetched along the ray. We terminate ray-marching once the accumulated opacity saturates, and pass the accumulated color and features to a small MLP that evaluates the view-dependent component for the ray.

uniform density along each ray. We chose a Cauchy loss as it is highly robust to outliers (which is desirable as all “hard” surfaces with high density will be treated as outliers) while also being continuously differentiable near the origin — a desirable property in gradient-based optimization [2].

5. Sparse Neural Radiance Grids

We now convert a trained deferred NeRF model, described above, into a representation suitable for real-time rendering. The core idea is to trade computation for storage, significantly reducing the time required to render frames. In other words, we are looking to replace the MLP evaluations in NeRF with fast lookups in a precomputed data structure.

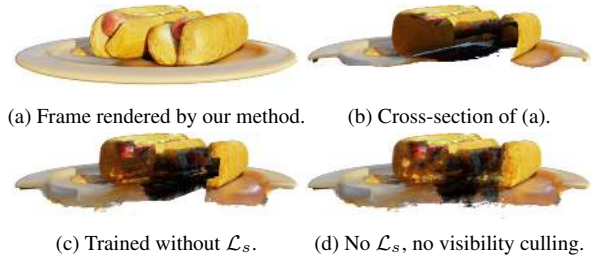


Figure 4: **Visualization of sparsity loss and visibility culling.** We render cross-sections of the *Hotdog* scene to inspect the effect of our sparsity loss \mathcal{L}_s and visibility culling. Our full method (b) represents the scene by only allocating content around visible scene surfaces. Removing either the sparsity loss alone (c) or both the sparsity loss and visibility culling (d) results in a much less compact representation.

We achieve this by precomputing and storing, i.e. *baking*, the diffuse colors c_d , volume densities σ , and 4-dimensional feature vectors v_s in a voxel grid data structure.

It is crucial for us to store this volumetric grid sparsely, as a dense voxel grid can easily fill up all available memory on a modern high-end GPU. By exploiting sparsity and only storing voxels that are both occupied and visible, we end up with a much more compact representation.

5.1. SNeRG Data Structure

Our data structure represents an N^3 voxel grid in a block-sparse format using two smaller dense arrays.

The first array is a 3D texture atlas containing densely-packed “macroblocks” of size B^3 each, corresponding to the content (diffuse color, feature vectors, and opacity) that actually exists in the sparse volume. Each voxel in the 3D atlas represents the scene at the full resolution of the dense N^3 grid, but the 3D texture atlas is much smaller than N^3 since it only contains the sparse “occupied” content. Compared to hashing-based data structures (where $B^3 = 1$), this approach helps keep spatially close content nearby in memory, which is beneficial for efficient rendering.

The second array is a low resolution $(N/B)^3$ indirection grid, which either stores a value indicating that the corresponding B^3 macroblock in the full voxel grid is empty, or stores an index pointing to the high-resolution content of that macroblock in the 3D texture atlas. This structure crucially lets us skip blocks of empty space during rendering, as we describe below.

5.2. Rendering

We render a SNeRG using a ray-marching procedure, as done in NeRF. The critical differences that enable real-time rendering are: 1) we precompute the diffuse colors and feature vectors at each 3D location, allowing us to look them up within our data structure instead of evaluating an MLP,

and 2) we only evaluate an MLP to produce view-dependent effects once per pixel, as opposed to once per 3D location.

To estimate the color of each ray, we first march the ray through the indirection grid, skipping macroblocks that are marked as empty. For macroblocks that are occupied, we step at the voxel width $1/N$ through the corresponding block in the 3D texture atlas, and use trilinear interpolation to fetch values at each sample location. We further accelerate rendering and conserve memory bandwidth by only fetching features where the volume density is non-zero. We use standard alpha compositing to accumulate the diffuse color and features, terminating ray-marching once the opacity has saturated. Finally, we compute the view-dependent specular color for the ray by evaluating MLP_Φ with the accumulated color, feature vector and the ray’s viewing direction. We then add the resulting residual color to the accumulated diffuse color, as described in Equation 7.

5.3. Baking

To minimize storage cost and rendering time, our baking procedure aims to only allocate storage for voxels in the scene that are both non-empty and visible in at least one of the training views. We start by densely evaluating the NeRF network for the full N^3 voxel grid. We convert NeRF’s unbounded volume density values, σ , to traditional opacity values $\alpha = 1 - \exp(-\sigma v)$, where $v = 1/N$ is the width of a voxel. Next, we sparsify this voxel grid by culling empty space, i.e. macroblocks where the maximum opacity is low (below τ_α), and culling macroblocks for which the voxel visibilities are low (maximum transmittance T between the voxel and all training views is below τ_{vis}). In all experiments, we set $\tau_\alpha = 0.005$ and $\tau_{vis} = 0.01$. Finally, we compute an anti-aliased estimate for the content in the remaining macroblocks by densely evaluating the trained NeRF at 16 Gaussian distributed locations within each voxel ($\sigma = v/\sqrt{12}$) and averaging the resulting diffuse colors, feature vectors, and volume densities.

5.4. Compression

We quantize all values in the baked SNeRG representation to 8 bits and separately compress the indirection grid and the 3D texture atlas. We compress each slice of the indirection grid as a lossless PNG, and we compress the 3D texture atlas as either a set of lossless PNGs, a set of JPEGs, or as a single video encoded with H264. The quality versus storage tradeoff of this choice is evaluated in Table 3. For synthetic scenes, compressing the texture atlas results in approximately $80\times$, $100\times$, and $230\times$ compression rates for PNG, JPEG, and H264, respectively. We specifically choose a macroblock size of 32^3 voxels to align the 3D texture atlas macroblocks with the blocks used in image compression. This reduces the size of the compressed 3D texture atlas because additional coefficients are not needed to represent discontinuities between macroblocks.




Figure 5: **Impact of fine-tuning (FT) the view-dependent appearance network** (PSNR in parentheses). (a) Renderings from a SNeRG representation can be lower-quality than those from the deferred NeRF, primarily due the quantization from 32-bit floating point to 8-bit integer values (see Sections 5.3, 5.4). (b) We are able to regain most of the lost quality by fine-tuning the weights of the deferred shading network MLP_Φ (Section 5.5).

5.5. Fine-tuning

While the compression and quantization procedure described above is crucial for making SNeRG compact and easy to distribute, the quality of images rendered from the baked SNeRG is lower than the quality of images rendered from the corresponding deferred NeRF. Figure 5 visualizes how quantization affects view-dependent effects by biasing renderings towards a darker, diffuse-only color $\hat{C}_d(\mathbf{r})$.

Fortunately, we are able to recoup almost all of that lost accuracy by fine-tuning the weights of the deferred per-pixel shading MLP_Φ to improve the final rendering quality (Table 2). We optimize the parameters Φ to minimize the squared error between the observed input images used to train the deferred NeRF and the images rendered from our SNeRG. We use the Adam optimizer [18] with a learning rate of 3×10^{-4} and optimize for 100 epochs.

6. Implementation Details

Our deferred NeRF model is based on JAXNeRF [12], an implementation of NeRF in JAX [4]. We apply a positional encoding [47] to positions and view directions. We train all networks for 250k iterations with a learning rate which decays log-linearly from 2×10^{-3} to 2×10^{-5} . To improve stability we use JAXNeRF’s “warm up” functionality to reduce the learning rate to 2×10^{-4} for the first 2500 iterations, and clip gradients by value (at 0.01) and then by norm (also 0.01). We use a batch size of 8,192 for synthetic scenes and a batch size of 16,384 for real scenes.

As our rendering time is independent of MLP_Θ ’s model size, we can afford to use a larger network for our experiments. To this end, we base our method on the JAXNeRF+ model, which was trained with 576 samples per ray (192 coarse, 384 fine) and uses 512 channels per layer in MLP_Θ .

During baking, we set the voxel grid resolution to be

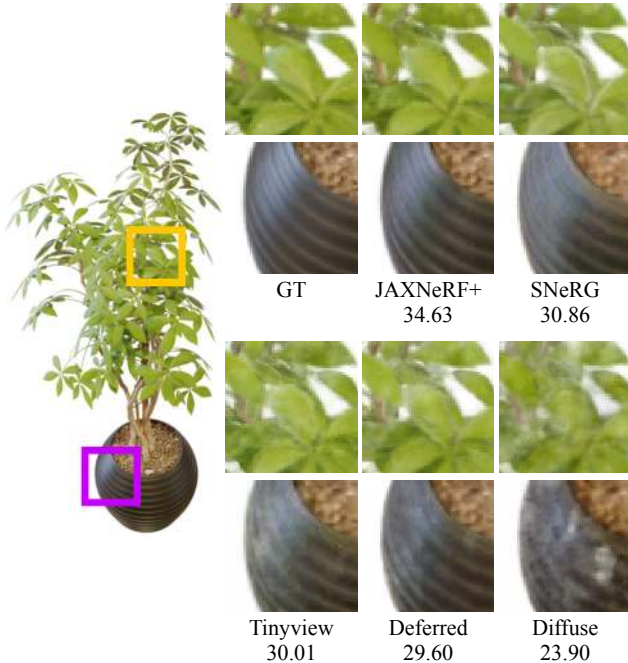


Figure 6: **Ablation study** showing visual examples from Table 2. Note the minimal difference in quality between the full JAXNeRF+ network and the successive approximations we make to speed up view-dependence (*Tinyview*, *Deferred NeRF* and *SNeRG*). For completeness, we show the floating alpha artifacts introduced by not modelling view-dependence at all (*Diffuse*).

slightly larger than the size of the training images: 1000^3 for the synthetic scenes, and 1300^3 for the real datasets. We implement the SNeRG renderer in Javascript and WebGL using the THREE.js library. We load the indirection grid and 3D texture atlas into 8 bit 3D textures. The view-dependence MLP_{ϕ} is stored uncompressed and is implemented in a WebGL shader. In all experiments, we run this renderer on a 2019 MacBook Pro laptop equipped with a 85W AMD Radeon Pro 5500M GPU.

Training the JAXNeRF+ model takes approximately 24 hours on 16 TPU v2 accelerators. The baking process is faster and completes in less than two hours, where most of the time is spent on refinement, which is currently implemented on the CPU and could be significantly accelerated.

7. Experiments

We validate our design decisions with an extensive set of ablation studies and comparisons to recent techniques for accelerating NeRF. Our experiments primarily focus on free-viewpoint rendering of 360° scenes (scenes captured by inwards-facing cameras on the upper hemisphere). Though acceleration techniques already exist for the special case in which all cameras face the same direction (see




Figure 7: **Real 360° scene** (PSNR in parentheses). Note how our real-time method is able to model the mirror-like reflective surface of the garden spheres.

Broxton *et al.* [5]), 360° scenes represent a challenging and general use-case that has not yet been addressed. In Figure 7, we show an example of a real 360° scene and present more results in our supplement, including the forward-facing scenes from Local Light Field Fusion (LLFF) [30].

We evaluate according to three criteria: render-time performance (measured by frames per second as well as GPU memory consumption in gigabytes), storage cost (measured by megabytes required to store the compressed representation), and rendering quality (measured using the PSNR, SSIM [51], and LPIPS [55] quality metrics). It is important to account for power consumption when evaluating performance — algorithms that are fast on a high performance GPU are not necessarily fast on a laptop. We adopt the convention used by the high performance graphics community of measuring performance relative to power consumption, i.e. FPS per watt, or equivalently, frames per joule [1].

7.1. Ablation Studies

In Table 1, we ablate combinations of three components of our method that primarily affect speed and GPU memory usage. Ablation 1 shows that removing the view-dependence MLP has a minimal effect on runtime performance. Ablation 2 shows that removing the sparsity loss \mathcal{L}_s greatly increases (uncompressed) memory usage. Ablation 3 shows that switching from our “deferred” rendering back to NeRF’s approach of querying an MLP at each sample along the ray results in prohibitively large render times.

Table 2 and Figure 6 show the impact on rendering quality of each of our design decisions in building a representation suitable for real-time rendering. Although our simpli-

	MLP	\mathcal{L}_s	Defer	ms/frame ↓	GPU GB ↓
Ours	✓	✓	✓	11.9± 4.5	1.73 ± 1.48
1)		✓		9.2± 4.6	1.73 ± 1.48
2)	✓		✓	20.0*± 5.3	4.26 ± 1.56
3)	✓	✓		343.6± 247.5	1.73 ± 1.48

Table 1: Performance ablation study, including uncompressed GPU memory used during rendering, on the Synthetic 360° scenes. Ablation 2 ran out of memory on the *Ficus* scene, which biases the average runtime for that row.

	PSNR ↑	SSIM ↑	LPIPS ↓
JAXNeRF+	33.00	0.962	0.038
JAXNeRF+ Tinyview	31.65	0.954	0.047
JAXNeRF+ Deferred	30.55	0.952	0.049
SNeRG (PNG)	30.38	0.950	0.050
SNeRG (PNG, no \mathcal{L}_s)	30.22	0.949	0.050
SNeRG (PNG, no FT)	26.68	0.930	0.053
JAXNeRF+ Diffuse	27.39	0.927	0.068

Table 2: Quality ablation study on Synthetic 360° scenes.

fications of using a deferred rendering scheme (“Deferred”) and a smaller network architecture (“Tinyview”) for view-dependent appearance do slightly reduce rendering quality, they are crucial for enabling real-time rendering, as discussed above. Note that the initial impact on quality from quantizing and compressing our representation is significant. However, after fine tuning (“FT”), the final rendering quality of our SNeRG model remains competitive with the neural model from which it was derived (“Deferred”).

In Table 3 we explore the impact of various compression schemes on disk storage space requirements. Our sparse voxel grid benefits greatly from applying compression techniques such as JPEG or H264 to its 3D texture atlas, achieving a file size over 200× more compact than a naive 32 bit float array while sacrificing less than 1dB of PSNR. Because our sparsity loss \mathcal{L}_s concentrates opaque voxels around surfaces (see Figure 4), ablating it significantly increases model size. Our compressed SNeRG representations are small enough to be quickly loaded in a web page.

The positive impact of the sparsity loss \mathcal{L}_s is visible across these ablations. It more than doubles rendering speed, halves storage requirements of both the compressed representation on disk and the uncompressed representation in GPU memory, and minimally impacts rendering quality.

7.2. Baseline Comparisons

Table 4 shows that the quality of our method is comparable to the other methods listed here, while our run-time performance is an order of magnitude faster than the fastest competing approach (Neural Volumes [27]) and more than a thousand times faster than the slowest (NeRF). Note that we measure the run-time rendering performance of our method on a laptop with an 85W mobile GPU, while all other meth-

	PSNR ↑	SSIM ↑	LPIPS ↓	MB ↓
SNeRG (Float)	30.47	0.951	0.049	6883.6
SNeRG (PNG, no \mathcal{L}_s)	30.22	0.949	0.050	176.0
SNeRG (PNG)	30.38	0.950	0.050	86.7
SNeRG (JPEG)	29.71	0.939	0.062	70.9
SNeRG (H264)	29.86	0.938	0.065	30.2
JAXNeRF+	33.00	0.962	0.038	18.0
JAXNeRF	31.65	0.952	0.051	4.8

Table 3: Storage ablation study on Synthetic 360° scenes.

	PSNR ↑	SSIM ↑	LPIPS ↓	W ↓	FPS ↑	FPS/W ↑
JAXNeRF+ [12]	33.00	0.962	0.038	300	0.01	0.00002
NeRF [31]	31.00	0.947	0.081	300	0.03	0.00011
JAXNeRF [12]	31.65	0.952	0.051	300	0.05	0.00016
IBRNet [50]	28.14	0.942	0.072	300	0.18	0.00061
AutoInt [25]	25.55	0.911	0.170	300	0.38	0.00128
NSVF [26]	31.74	0.953	0.047	300	0.65	0.00217
NV [27]	26.05	0.893	0.160	300	3.33	0.01111
SNeRG (PNG)	30.38	0.950	0.050	85	84.06	0.98897

Table 4: Baseline comparisons on Synthetic 360° scenes.

ods are run on servers or workstations equipped with much more powerful GPUs (over 3× the power draw).

8. Conclusion

Our method renders Neural Radiance Fields in real-time by precomputing and storing a Sparse Neural Radiance Grid. This *SNeRG* uses a sparse voxel grid representation to store the precomputed scene geometry, but keeps storage requirements reasonable by maintaining a neural representation for view-dependent appearance. Rendering is accelerated by evaluating the view-dependent shading network only once per ray, achieving over 30 frames per second on a laptop GPU for typical NeRF scenes.

However, these design choices have some limitations. As we evaluate view-dependence once per ray, we will not be able to easily model scenes with both reflective and transparent surfaces, where the view-dependence for each surface should be modelled separately. Rendering will also become expensive in extremely large (e.g. city-scale) scenes, as our representation does not use multiple levels of detail. This could be addressed with an octree cone tracer [8].

In conclusion, we hope this ability to render neural volumetric representations such as NeRF in real time on commodity hardware will help increase the adoption of neural scene representations in vision and graphics applications.

Acknowledgements We thank Keunhong Park and Michael Broxton for their generous help with debugging, and Ryan Overbeck for last-minute JavaScript help. Many thanks to John Flynn, Dominik Kaeser, Keunhong Park, Ricardo Martin-Brualla, Hugues Hoppe, Janne Kontkanen, Utkarsh Sinha, Per Karlsson, and Mark Matthews for fruitful discussions, brainstorming, and testing out our viewer.

References

- [1] Tomas Akenine-Möller and Björn Johnson. Performance per what? *Journal of Computer Graphics Techniques*, 2012.
- [2] Jonathan T. Barron. A general and adaptive robust loss function. *CVPR*, 2019.
- [3] Sai Bi, Zexiang Xu, Pratul P. Srinivasan, Ben Mildenhall, Kalyan Sunkavalli, Miloš Hašan, Yannick Hold-Geoffroy, David Kriegman, and Ravi Ramamoorthi. Neural reflectance fields for appearance acquisition. *arXiv cs.CV arXiv:2008.03824*, 2020.
- [4] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018. <http://github.com/google/jax>.
- [5] 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.
- [6] Chris Buehler, Michael Bosse, Leonard McMillan, Steven Gortler, and Michael Cohen. Unstructured lumigraph rendering. *SIGGRAPH*, 2001.
- [7] Eric R. Chan, Marco Monteiro, Petr Kellnhofer, Jiajun Wu, and Gordon Wetzstein. pi-GAN: Periodic implicit generative adversarial networks for 3D-aware image synthesis. *CVPR*, 2021.
- [8] Cyril Crassin, Fabrice Neyret, Sylvain Lefebvre, and Elmar Eisemann. GigaVoxels: Ray-guided streaming for efficient and detailed voxel rendering. *Symposium on Interactive 3D Graphics and Games*, 2009.
- [9] Abe Davis, Marc Levoy, and Fredo Durand. Unstructured light fields. *Computer Graphics Forum*, 2012.
- [10] Paul Debevec, C. J. Taylor, and Jitendra Malik. Modeling and rendering architecture from photographs: a hybrid geometry- and image-based approach. *SIGGRAPH*, 1992.
- [11] Michael Deering, Stephanie Winner, Bic Schediwy, Chris Duffy, and Neil Hunt. The triangle processor and normal vector shader: A VLSI system for high performance graphics. *SIGGRAPH*, 1988.
- [12] Boyang Deng, Jonathan T. Barron, and Pratul P. Srinivasan. JaxNeRF: an efficient JAX implementation of NeRF, 2020. <http://github.com/google-research/google-research/tree/master/jaxnerf>.
- [13] John Flynn, Michael Broxton, Paul Debevec, Matthew DuVall, Graham Fyffe, Ryan Overbeck, Noah Snavely, and Richard Tucker. DeepView: View synthesis with learned gradient descent. *CVPR*, 2019.
- [14] Guy Gafni, Justus Thies, Michael Zollhöfer, and Matthias Nießner. Dynamic neural radiance fields for monocular 4D facial avatar reconstruction. *CVPR*, 2021.
- [15] Stephan J. Garbin, Marek Kowalski, Matthew Johnson, Jamie Shotton, and Julien P. C. Valentin. Fastnerf: High-fidelity neural rendering at 200fps. 2021.
- [16] Steven J. Gortler, Radek Grzeszczuk, Richard Szeliski, and Michael F. Cohen. The lumigraph. *SIGGRAPH*, 1996.
- [17] 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.
- [18] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *ICLR*, 2015.
- [19] Philippe Lacroute and Marc Levoy. Fast volume rendering using a shear-warp factorization of the viewing transformation. *SIGGRAPH*, 1994.
- [20] Samuli Laine and Tero Karras. Efficient sparse voxel octrees. *13D*, 2010.
- [21] Sylvain Lefebvre and Hugues Hoppe. Perfect spatial hashing. *ACM Transactions on Graphics*, 2006.
- [22] Marc Levoy. Efficient ray tracing of volume data. *ACM Transactions on Graphics*, 1980.
- [23] Marc Levoy and Pat Hanrahan. Light field rendering. *SIGGRAPH*, 1996.
- [24] Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. *CVPR*, 2021.
- [25] David B. Lindell, Julien N.P. Martel, and Gordon Wetzstein. AutoInt: Automatic integration for fast neural rendering. *CVPR*, 2021.
- [26] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. *NeurIPS*, 2020.
- [27] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: Learning dynamic renderable volumes from images. *SIGGRAPH*, 2019.
- [28] Nelson Max. Optical models for direct volume rendering. *IEEE TVCG*, 1995.
- [29] Michael Goesele, Michael Waechter, Nils Moehrl. Let there be color! Large-scale texturing of 3D reconstructions. *ECCV*, 2014.
- [30] Ben Mildenhall, Pratul P. Srinivasan, Rodrigo Ortiz-Cayon, Nima K. Kalantari, Ravi Ramamoorthi, Ren Ng, and Abhishek Kar. Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *ACM Transactions on Graphics*, 2019.
- [31] 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. *ECCV*, 2020.
- [32] Thomas Neff, Pascal Stadlbauer, Mathias Parger, Andreas Kurz, Joerg H. Mueller, Chakravarty R. Alla Chaitanya, Anton S. Kaplanyan, and Markus Steinberger. DONeRF: Towards Real-Time Rendering of Compact Neural Radiance Fields using Depth Oracle Networks. *Computer Graphics Forum*, 2021.
- [33] Ren Ng, Ravi Ramamoorthi, and Pat Hanrahan. All-frequency shadows using non-linear wavelet lighting approximation. *ACM Transactions on Graphics*, 2003.
- [34] Matthias Nießner, Michael Zollhofer, Shahram Izadi, and Marc Stamminger. Real-time 3D reconstruction at scale using voxel hashing. *ACM Transactions on Graphics*, 2013.

- [35] Julian Ost, Fahim Mannan, Nils Thuerey, Julian Knodt, and Felix Heide. Neural scene graphs for dynamic scenes. *CVPR*, 2021.
- [36] Keunhong Park, Utkarsh Sinha, Jonathan T. Barron, Sofien Bouaziz, Dan B Goldman, Steven M. Seitz, and Ricardo Martin-Brualla. Deformable neural radiance fields. *arXiv cs.CV 2011.12948*, 2020.
- [37] Eric Penner and Li Zhang. Soft 3D reconstruction for view synthesis. *ACM Transactions on Graphics*, 2017.
- [38] Daniel Rebain, Wei Jiang, Soroosh Yazdani, Ke Li, Kwang Moo Yi, and Andrea Tagliasacchi. DeRF: Decomposed radiance fields. *CVPR*, 2021.
- [39] Christian Reiser, Songyu Peng, Yiyi Liao, and Andreas Geiger. Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps. 2021.
- [40] Katja Schwarz, Yiyi Liao, Michael Niemeyer, and Andreas Geiger. GRAF: Generative radiance fields for 3D-aware image synthesis. *NeurIPS*, 2020.
- [41] Steven M. Seitz and Charles R. Dyer. Photorealistic scene reconstruction by voxel coloring. *IJCV*, 1999.
- [42] Vincent Sitzmann, Michael Zollhoefer, and Gordon Wetzstein. Scene representation networks: Continuous 3D-structure-aware neural scene representations. *NeurIPS*, 2019.
- [43] Peter-Pike Sloan, Jan Kautz, and John Snyder. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. *ACM Transactions on Graphics*, 2002.
- [44] Pratul P. Srinivasan, Boyang Deng, Xiuming Zhang, Matthew Tancik, Ben Mildenhall, and Jonathan T. Barron. NeRV: Neural reflectance and visibility fields for relighting and view synthesis. *CVPR*, 2021.
- [45] Pratul P. Srinivasan, Ben Mildenhall, Matthew Tancik, Jonathan T. Barron, Richard Tucker, and Noah Snavely. Lighthouse: Predicting lighting volumes for spatially-coherent illumination. *CVPR*, 2020.
- [46] Pratul P. Srinivasan, Richard Tucker, Jonathan T. Barron, Ravi Ramamoorthi, Ren Ng, and Noah Snavely. Pushing the boundaries of view extrapolation with multiplane images. *CVPR*, 2019.
- [47] Matthew Tancik, Pratul P. Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan T. Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *NeurIPS*, 2020.
- [48] Justus Thies, Michael Zollhoefer, and Matthias Nießner. Deferred neural rendering: Image synthesis using neural textures. *ACM Transactions on Graphics*, 2019.
- [49] Yu-Ting Tsai and Zen-Chung Shih. All-frequency precomputed radiance transfer using spherical radial basis functions and clustered tensor approximation. *ACM Transactions on Graphics*, 2006.
- [50] 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. *CVPR*, 2021.
- [51] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE TIP*, 2004.
- [52] Lee A Westover. Splatting: A parallel, feed-forward volume rendering algorithm. Technical report, University of North Carolina at Chapel Hill, USA, 1991.
- [53] Daniel Wood, Daniel Azuma, Wyvern Aldinger, Brian Curless, Tom Duchamp, David Salesin, and Werner Stuetzle. Surface light fields for 3D photography. *SIGGRAPH*, 2000.
- [54] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. Plenotrees for real-time rendering of neural radiance fields. 2021.
- [55] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. *CVPR*, 2018.
- [56] Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. Stereo magnification: Learning view synthesis using multiplane images. *ACM Transactions on Graphics*, 2018.