

Analyzing the Internals of Neural Radiance Fields

Lukas Radl Andreas Kurz Michael Steiner Markus Steinberger

{lukas.radl, andreas.kurz, michael.steiner, steinberger}@icg.tugraz.at

Graz University of Technology, Austria

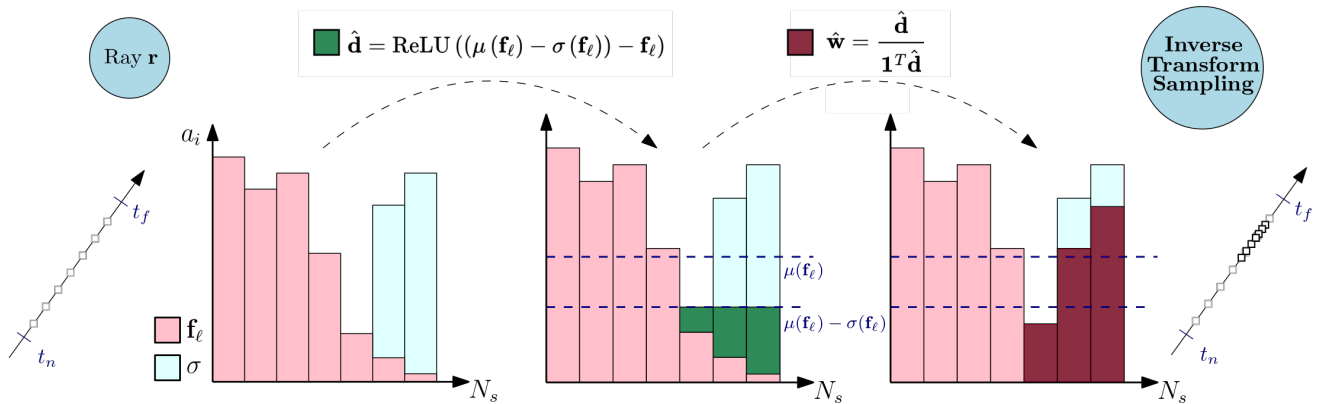


Figure 1. We show how density estimates derived from intermediate activations can accelerate inference for pre-trained Neural Radiance Fields by effectively reducing the capacity of large MLPs. Here, we show a small toy example for a ray with seven, uniformly placed samples between t_n and t_f . We obtain an activation feature vector \mathbf{f}_ℓ for layer ℓ , apply a function to obtain a density estimate $\hat{\mathbf{d}}$, leveraging the observation that minima in activation feature space indicate samples with high density σ . Finally, we perform inverse transform sampling with our weight estimate $\hat{\mathbf{w}}$.

Abstract

Modern Neural Radiance Fields (NeRFs) learn a mapping from position to volumetric density leveraging proposal network samplers. In contrast to the coarse-to-fine sampling approach with two NeRFs, this offers significant potential for acceleration using lower network capacity. Given that NeRFs utilize most of their network capacity to estimate radiance, they could store valuable density information in their parameters or their deep features. To investigate this proposition, we take one step back and analyze large, trained ReLU-MLPs used in coarse-to-fine sampling. Building on our novel activation visualization method, we find that trained NeRFs, Mip-NeRFs and proposal network samplers map samples with high density to local minima along a ray in activation feature space. We show how these large MLPs can be accelerated by transforming intermediate activations to a weight estimate, without any modifications to the training protocol or the network architecture. With our approach, we can reduce the computational requirements of trained NeRFs by up to 50% with only a slight hit in rendering quality. Extensive experimental evaluation on a variety

of datasets and architectures demonstrates the effectiveness of our approach. Consequently, our methodology provides valuable insight into the inner workings of NeRFs.

1. Introduction

Neural Fields [43] parameterize implicit functions over continuous domains using neural networks and have gained popularity due to their ability to represent various types of signals. To represent arbitrary high-frequency functions, positional encoding [20], Fourier features [34], Gaussian activations [6, 26], among other alternatives [7, 30, 31], are used to overcome the spectral bias [25]. Compared to their explicit counterparts, implicit representations are more memory-efficient and can achieve higher quality, but regrettably suffer from slower inference. In addition, as the signal is encoded in the weights and biases of a neural network, every change to a parameter has a non-local effect, which makes implicit representations hard to modify.

For novel view synthesis problems, the 5D plenoptic function mapping positions and directions to radiance and density can be learned from posed images leveraging a neural

field. NeRFs [20] have recently demonstrated impressive performance for real-world and synthetic datasets. However, the naïve approach of sampling the scene densely is slow and inefficient due to the prohibitive number of required samples. A coarse-to-fine sampling strategy is adopted to overcome these limitations, estimating the density distribution along a ray by uniformly sampling a coarse NeRF, which generates a new set of samples for a subsequent fine NeRF using inverse transform sampling. Although performance is improved substantially, the number of required samples still prohibits real-time rendering. Furthermore, as no ground-truth density information is available, the coarse NeRF is supervised with an MSE reconstruction loss.

Reducing the number of required samples for NeRFs is a fruitful endeavor to lower the computational requirements. Sampling networks [1, 3, 16, 22], which predict volumetric density along a ray to produce a set of samples for a subsequent shading network, prove highly effective for this task. The network capacity for a sampling network is significantly reduced, as solely predicting volumetric density is view-independent and thus easier to learn. For supervising these sampling networks ground-truth depth data [22], a pre-trained NeRF [24], multi-view stereo [17] or distillation from a concurrently trained shading network [1, 3, 16] may be used. Alternatively, occupancy grids [21] allow for efficient empty-space skipping by encoding the expected density within a discretized region.

A clear disadvantage of the coarse-to-fine NeRF pipeline is the use of the coarse model for density prediction only, even though it is optimized to reconstruct the ground-truth image. Sampling networks, on the other hand, map sample positions to volumetric density using a network with reduced capacity and can be supervised by the shading network. Given that sampling networks show that density can be modelled with lower capacity, a subset of a coarse NeRF should be able to model density accurately as well — particularly since recent work has found that NeRF-MLPs naturally partition themselves into structural and color layers [38].

To this end, we analyze the intermediate activations of coarse NeRFs: We derive a representation from these activations and show how it is tied to the density distribution along a view ray. Building on this crucial finding, we devise a novel method to extract density estimates given intermediate activations of coarse NeRFs. By elegantly combining activation analysis and our method to extract densities, our approach can reduce inference times while maintaining high quality at test-time. We demonstrate that our approach is not only applicable to NeRF and Mip-NeRF [2], but produces similar results when applied to proposal network samplers as well, with extensive evaluation supporting our claim. To the best of our knowledge, our approach is the first to use intermediate activations to accelerate inference for coordinate-based MLPs.

To summarize, our contributions are as follows:

- We present a novel method for visualizing and analyzing the activations of coordinate-based ReLU-MLPs.
- We present an approach for extracting a density estimate for fine re-sampling from activations in early layers of coarse NeRFs, significantly reducing the inference time with little decrease in quality. Consequently, our method provides valuable insight into the inner workings of NeRFs.
- We demonstrate the effectiveness of our method for several architectures with real-world and synthetic data.

2. Related Work

In the following, we review related work, focusing on real-time rendering and efficient sampling for NeRFs.

NeRFs for Real-time Rendering. Instead of the implicit representation of radiance fields used in NeRF [20], explicit representations with grid-like data structures [8, 33] or hybrid representation [4, 13, 21, 42, 44] are often used to accelerate inference and training. Most prominently of the aforementioned, Instant-NGP [21] replaces the fixed positional encoding used by NeRF with a learnable hash encoding, allowing for the use of much smaller MLPs and thus rapid convergence. Alternatively, 3D Gaussian Splatting [14] has recently demonstrated state-of-the-art rendering quality in addition to fast training and rendering using a fully explicit representation relying on anisotropic 3D Gaussians. Other works [27, 28] suggest a divide-and-conquer approach, utilizing many tiny MLPs responsible for a small region instead of a large MLP for the whole scene. Distillation of a NeRF into a render-friendly format has been explored extensively [9, 11], but these approaches are memory-intensive and require an additional distillation step after training. Recent work [15, 37] has also investigated approximate geometry equipped with neural features and small, global decoders. Finally, mesh-based view synthesis [5, 29, 45] has recently demonstrated real-time rendering on low-power devices, relying on approximate meshes extracted from radiance fields and the well-known graphics pipeline.

Efficient Sampling. Several works investigate a learnable proposal sampler, which can predict volumetric density [3] or sample locations [1]. DONeRF [22] learns a density volume as a multi-class classification task to guide sample placement with direct depth supervision. Building on DONeRF, AdaNeRF [16] mitigates the requirement of pre-training and depth supervision for the sampling network with a 4-phase training recipe. EfficientNeRF [12] proposes valid sampling, caching densities of sample positions and exclusively evaluating those with positive density — this approach shares similarities with our method in that it is motivated by the

density distribution of coarse NeRFs. TermiNeRF [24] distills a NeRF into a sampling and shading network. For their sampling network, they perform distribution matching, effectively reparameterizing NeRFs intervals to allow for fewer samples. Lin *et al.* [17] learn a proxy geometry with a cost volume for depth guidance, uniformly sampling a smaller depth range with a lower sample count. A better sampling strategy is also attained by unifying implicit surface models and radiance fields [23]. To enable mesh-based view synthesis, SDFs [45] or Binary Opacity Grids [29] can accurately locate surfaces. Finally, by restricting volume rendering to a small band around the surface [41], the number of samples can be reduced significantly.

3. Preliminaries

In the following section, we briefly summarize the most important architectures and concepts which build the foundation for our work.

3.1. NeRF

NeRFs [20] learn a function

$$\Theta_{\text{NeRF}} : \mathbb{R}^5 \rightarrow \mathbb{R}^4, (\mathbf{p}, \mathbf{d}) \mapsto (\mathbf{c}, \sigma), \quad (1)$$

where $\mathbf{p} = (x, y, z)$ and $\mathbf{d} = (\theta, \phi)$ are the sample position and viewing direction, $\mathbf{c} \in [0, 1]^3$ denotes the predicted output color and $\sigma \in \mathbb{R}$ is the volumetric density. NeRF uses a ReLU-MLP with 8 layers and 56 hidden units to predict a feature vector $\mathbf{f} \in \mathbb{R}^{256}$ and σ . This feature vector is then concatenated with \mathbf{d} and passed through a small ReLU-MLP to produce \mathbf{c} , which allows view-dependent phenomena to be modelled. For each pixel, we cast a ray $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$ from the origin \mathbf{o} in the direction of \mathbf{d} . We evaluate Θ_{NeRF} at different sample positions $t \in [t_n, t_f]$ between the near bound t_n and the far bound t_f . The color $\hat{\mathcal{C}}$ of a ray \mathbf{r} can be estimated using N_s samples with

$$\hat{\mathcal{C}}(\mathbf{r}) = \sum_{i=1}^{N_s} T_i (1 - \exp(-\sigma_i \delta_i)) \mathbf{c}_i, \quad (2)$$

where δ_i and T_i are given as

$$\delta_i = t_{i+1} - t_i, \quad T_i = \exp\left(-\sum_{j=1}^{i-1} \sigma_j \delta_j\right). \quad (3)$$

To increase rendering efficiency, two NeRFs are concurrently optimized. The coarse NeRF Θ_{coarse} places N_c samples uniformly for each ray \mathbf{r} , which produces a set of weights $\mathbf{w} \in \mathbb{R}^{N_c}$ with

$$\mathbf{w}_i = T_i (1 - \exp(-\sigma_i \delta_i)), \quad (4)$$

used to express the estimated coarse radiance as a weighted sum

$$\hat{\mathcal{C}}_c(\mathbf{r}) = \sum_{i=1}^{N_c} \mathbf{w}_i \mathbf{c}_i. \quad (5)$$

The weights \mathbf{w} are normalized such that $\mathbf{1}^T \mathbf{w} = 1$ to obtain a piecewise-constant PDF along \mathbf{r} . This distribution is used to produce a new set of N_f samples with inverse transform sampling. The final rendered color is produced by Θ_{fine} using all $N_c + N_f$ samples.

NeRF applies positional encoding $\gamma(\cdot)$ to \mathbf{p} and \mathbf{d} , allowing the network to model high-frequency details effectively [34]. The network is supervised with the MSE loss for both Θ_{coarse} and Θ_{fine} , i.e. $\sum_{\mathbf{r}} \|\hat{\mathcal{C}}_c(\mathbf{r}) - \mathcal{C}(\mathbf{r})\|_2^2 + \|\hat{\mathcal{C}}_f(\mathbf{r}) - \mathcal{C}(\mathbf{r})\|_2^2$, with $\mathcal{C}(\mathbf{r})$ denoting the ground truth color for \mathbf{r} and $\hat{\mathcal{C}}_f(\mathbf{r})$ denoting the estimated fine radiance of \mathbf{r} .

3.2. Mip-NeRF

Mip-NeRF [2] builds on NeRF and tackles the problem of aliasing when scene content is observed at various resolutions. The standard computer graphics solution of super-sampling is prohibitively expensive when considering NeRFs volume rendering. Instead, Mip-NeRF introduces a novel Integrated Positional Encoding (IPE), which is the expected positional encoding of all coordinates within a conical frustum. Each conical frustum is approximated with a multivariate Gaussian, and the IPE feature serves as the encoded input for Θ_{NeRF} . This modification allows the MLP to reason about the size and shape of the encoded region, efficiently combating aliasing which occurs for ambiguous point samples.

As the network can now learn a multi-scale scene representation, Θ_{NeRF} performs both coarse and fine sampling, reducing the model size by 50% — however, the network is still supervised with an MSE loss.

3.3. Mip-NeRF 360

NeRF and Mip-NeRF require that 3D coordinates exist in a bounded domain for sampling — an assumption which is violated in unbounded 3D scenes, where objects may exist at any distance from the camera. Therefore, Mip-NeRF 360 [3] proposes a spatial contraction $\text{contract}(\cdot)$, which maps coordinates from \mathbb{R}^3 to a ball of radius 2:

$$\text{contract}(\mathbf{x}) = \begin{cases} \mathbf{x} & \|\mathbf{x}\| \leq 1 \\ \left(2 - \frac{1}{\|\mathbf{x}\|}\right) \left(\frac{\mathbf{x}}{\|\mathbf{x}\|}\right) & \|\mathbf{x}\| > 1 \end{cases}. \quad (6)$$

To handle the increased capacity required for unbounded scenes and increase rendering efficiency, Mip-NeRF 360 concurrently optimizes two MLPs, Θ_{prop} and Θ_{NeRF} . Their proposal network sampler Θ_{prop} learns a function

$$\Theta_{\text{prop}} : \mathbb{R}^3 \rightarrow \mathbb{R}, \mathbf{p} \mapsto \sigma, \quad (7)$$

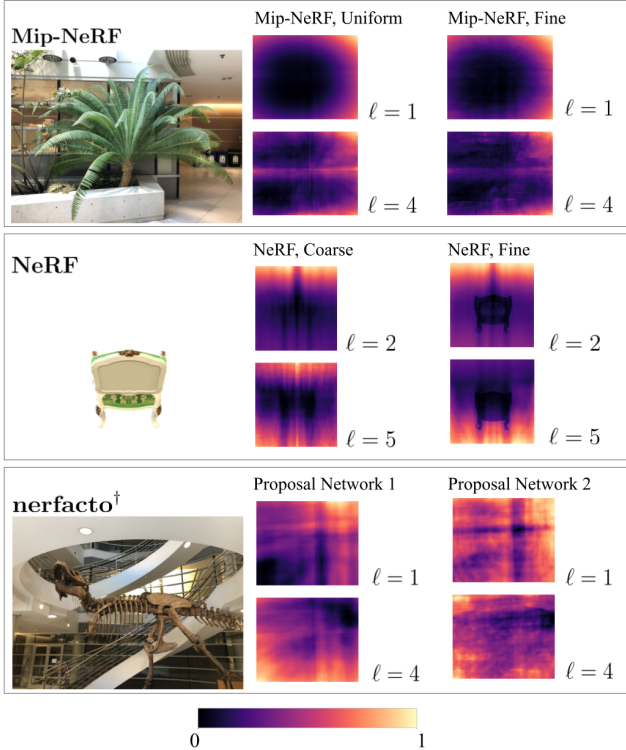


Figure 2. Ground-truth images and their corresponding normalized coarse and fine activations v_ℓ using the *magma* colormap reveal an interesting relationship between activations and outputs. With our visualization approach, we can infer some scene content using only v_ℓ . For each scene, we visualize activations for different layers ℓ .

and is optimized to bound the weights \mathbf{w} produced by Θ_{NeRF} . Two subsequent rounds of proposal sampling produce two weights $\{\hat{\mathbf{w}}_1, \hat{\mathbf{w}}_2\}$, which is efficient due to the small size of Θ_{prop} with 4 layers and 256 hidden units.

4. Activation Analysis

In this section, we present our method for visualizing and analyzing the intermediate activations for ReLU-MLPs. In addition, we show how we can leverage these intermediate activations to accelerate inference for pre-trained NeRFs.

4.1. Visualizing and Analyzing Activations

Benefiting from the inherent spatial structure baked into their building blocks, several works [32, 47] investigate the visualization of intermediate features for Convolutional Neural Networks (CNNs). These methods typically operate by back-projection from feature space to pixel space, allowing informative insight into the inner workings of these models. For MLPs, which are mostly used as universal function approximators, practitioners have not been as interested in their activations. To this end, we propose a novel method for visualizing activations of coordinate-based MLPs.




Figure 3. Intermediate activations allow for simple performance improvements. We can reduce the inference time of NeRFs in synthetic scenes if we perform the fine pass only if the condition $v_\ell < \mu(\mathbf{V}_\ell)$ is met.

We analyze the hidden layer activations of NeRF MLPs as follows: For each pixel (x, y) , a ray \mathbf{r} is generated, positional encoding $\gamma(\cdot)$ is applied and each sample along the ray is passed through Θ_{NeRF} . In the following, we describe the intermediate activations in terms of a single ray/pixel. After each linear layer, the activation $\mathbf{A}^{(\ell)}$ is of size $\mathbb{R}^{N_s \times N_h}$, where N_s and N_h denote the number of samples and the number of hidden units, respectively. To obtain a feature representation $\mathbf{f}_\ell \in \mathbb{R}^{N_s}$ for the activation in layer ℓ , we compute the mean over the dimension N_h :

$$\mathbf{f}_\ell = \frac{1}{N_h} \sum_{i=1}^{N_h} \mathbf{A}_i^{(\ell)}. \quad (8)$$

Further, we can produce a scalar $v_\ell \in \mathbb{R}$ for each ray \mathbf{r} for visualization purposes with

$$v_\ell = \frac{1}{N_s} \sum_{i=1}^{N_s} \sum_{j=1}^{N_h} \mathbf{A}_{i,j}^{(\ell)}. \quad (9)$$

We visualize the activations of different layers ℓ for some example views using Eqn. (9) in Fig. 2. As we can see, intermediate activations exhibit notable structural information. We note that when recording activations $\mathbf{A}^{(\ell)}$, we have the option to apply or not apply the activation function. Unless stated otherwise, we work with activations after the ReLU has been applied, thus $\mathbf{A}^{(\ell)} \in \mathbb{R}_+^{N_s \times N_h}$. Although unintuitive due to the information loss caused by the ReLU, early experiments showed that applying the activation function resulted in more consistent histograms for \mathbf{f}_ℓ .

4.2. Reducing Unnecessary Network Evaluations using Activations

As a first step towards speeding up the coarse-to-fine NeRF rendering pipeline, we analyze the densities σ predicted by Θ_{coarse} . Take, for example, any of the synthetic scenes from the Blender dataset [20]. Along rays \mathbf{r} in empty space, we do not expect any sample with significant volumetric density. We can sum the number of densities along each ray \mathbf{r} larger than a small threshold τ to obtain a mapping

ℓ	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	t [s]	Speedup	% fine Rays
1	29.66	0.93	0.04	34.69	40%	59%
2	30.45	0.94	0.03	34.02	43%	58%
3	31.21	0.95	0.02	33.11	47%	55%
4	30.71	0.95	0.03	34.76	43%	59%
NeRF	32.20	0.96	0.02	48.98	-	100%

Table 1. Quantitative results for our activation informed coarse-to-fine experiment for the *lego* scene from the Blender dataset [20]. Only performing the fine pass for rays \mathbf{r} where the summed activation is smaller than $\tau = \mu(\mathbf{V}_\ell)$ yields a significant performance gain with a slight drop in rendering quality.

$g : \mathbf{r} \mapsto \{0, \dots, N_c\}$. If for any of the rays the result is 0, the corresponding pixel likely belongs to the transparent background. We see in Fig. 3 that when the value of g is high, the value of v_ℓ is low, and vice-versa.

For each ray \mathbf{r} , we record the predicted coarse radiance \hat{C}_c and the activations v_ℓ from Θ_{coarse} . We choose a threshold τ , which indicates whether we observe scene content along \mathbf{r} . We choose $\tau = \mu(\mathbf{V}_\ell)$, the mean activation of the activation image \mathbf{V}_ℓ . We obtain a set of rays $P = \{\mathbf{r} : v_\ell < \tau\}$, which we treat as a mask. If $\mathbf{r} \in P$, we evaluate Θ_{fine} . Otherwise, we treat the coarse radiance \hat{C}_c as our final output. Using this approach, we can significantly reduce the number of evaluations for Θ_{fine} , which uses $N_c + N_f = 192$ samples, which results in a significant performance improvement. As we cover most of the object in the scene with our derived mask P , the image quality loss is negligible. We report the results for the *lego* scene from the Blender dataset [20] in Tab. 1 and visualize our approach in Fig. 3. This experiment demonstrates that early intermediate activations already contain notable density information and reveals an interesting connection between density and activations.

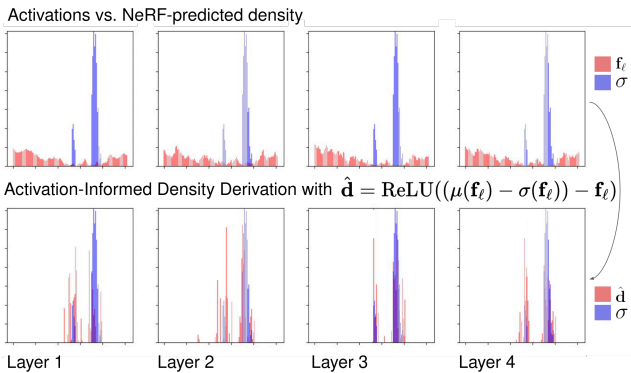


Figure 4. Visualization of our proposed approach for approximate density extraction for a real-world example: We visualize activation features \mathbf{f}_ℓ and densities σ for 128 uniform samples along an example ray, for a Mip-NeRF trained on the *chair* scene. Using Eqn. (10), a plausible density estimate $\hat{\mathbf{d}}$ is extracted from a activation feature \mathbf{f}_ℓ .

4.3. Approximating Density using Activations

Clearly, the approach from Sec. 4.2 heavily relies on the transparent background present in synthetic scenes. Real-world captures, on the other hand, often include highly detailed backgrounds. Therefore, we go beyond this basic approach by analyzing the activation features \mathbf{f}_ℓ along a ray \mathbf{r} . We accomplish this by examining histograms of the predicted density σ and the activation features \mathbf{f}_ℓ , as can be seen in Fig. 4. We notice a trend throughout multiple scenes: Activation feature space minima indicate samples with significant density. If we can find a function to transform \mathbf{f}_ℓ to a reasonable density estimate $\hat{\mathbf{d}}$, we can effectively replace the expensive evaluation for Θ_{coarse} with a pass through a smaller MLP. Ideally, this function should transform \mathbf{f}_ℓ to a vector of mostly zeros, except for locations where σ is significant. To this end, we propose 3 different functions f_i to map the intermediate activations directly to estimated densities $\hat{\mathbf{d}}$, which were handcrafted based on the observed relationship between activations and density:

$$f_1(\mathbf{f}_\ell) = \text{ReLU}((\mu(\mathbf{f}_\ell) - \sigma(\mathbf{f}_\ell)) - \mathbf{f}_\ell), \quad (10)$$

$$f_2(\mathbf{f}_\ell) = \text{ReLU}\left(\left(\mu(\mathbf{f}_\ell) - \frac{\sigma(\mathbf{f}_\ell)}{2}\right) - \mathbf{f}_\ell\right), \quad (11)$$

$$f_3(\mathbf{f}_\ell) = \text{ReLU}\left(\left(\mu(\mathbf{f}_\ell) - \frac{\sigma(\mathbf{f}_\ell)}{2}\right) - \mathbf{f}_\ell\right)^2, \quad (12)$$

where \mathbf{f}_ℓ denotes the intermediate activation feature for layer ℓ and $\mu(\cdot)$, $\sigma(\cdot)$ denote the mean and standard deviation along the last dimension (i.e. along the N_s samples of the ray \mathbf{r}). We visualize our approach for a toy example in Fig. 1.

Using Eqn. (10) results in fewer intervals with density along the ray, but is prone to failures if the distribution along a ray exhibits high standard deviation. On the other hand, Eqns. (11) and (12) extract density estimates more conservatively but might lead to insufficient detail, as the interval between fine samples grows larger.

Given a well-behaved distribution, we can estimate densities sufficiently close to σ such that we get good reconstruction quality when converting $\hat{\mathbf{d}}$ to a weight estimate $\hat{\mathbf{w}}$. Further, we can significantly reduce the required inference time by using only a small subset of the large Θ_{coarse} MLP. Clearly, the potential speedup is dependent on the number of samples N_s , which is large for NeRFs.

Our Proposed Method for Weight Estimation. First, we extract the activation features \mathbf{f}_ℓ and apply a function f_i to transform this to a density estimate $\hat{\mathbf{d}}$ for each ray \mathbf{r} . To obtain our estimated weights $\hat{\mathbf{w}}$, which should conform to a piecewise-constant PDF along the ray \mathbf{r} , we use $\hat{w}_i = \hat{d}_i / \sum_{j=1}^{N_s} \hat{d}_j$ and subsequently perform inverse transform sampling, as done in NeRF [20] (c.f. Fig. 1).

5. Experiments

In the following, we apply our presented method to NeRF [20] and Mip-NeRF [2]. Additionally, we apply our method to proposal network samplers, utilizing a nerfacto model [35] with a Mip-NeRF 360 [3] proposal sampler, which we dub nerfacto[†]. We only apply our density extraction method to the first 3 layers: Recent work by Wang *et al.* [38] has found that NeRF MLPs naturally store structural information within the first MLP layers. In addition, when applying our method to earlier layers, we can accelerate inference more substantially.

Datasets. For our evaluation, we use four diverse datasets, which are well-established in the novel view synthesis literature: The **Blender** dataset [20] consists of 360° captures of synthetic objects with complex geometry, realistic non-Lambertian effects and transparent backgrounds in a bounded domain. **LLFF** [19] is a dataset of forward-facing, real-world scenes in high resolution. The **Mip-NeRF 360** dataset [3] contains 360° captures of real-world, unbounded scenes and is the most challenging benchmark in our evaluation. Finally, we also use 2 large-scale, bounded indoor scenes from **Deep Blending** [10] (*playroom & drjohnson*). For both LLFF and Mip-NeRF 360, we use the 4× down-sampled images. We report per-scene results in the supplementary material.

Implementation Details. We use the NeRF and Mip-NeRF implementations from Nerfstudio [35] and use the hyperparameter settings as reported in the respective works [2, 20]. To facilitate sampling in NDC, as NeRF does for LLFF [19], we implement a custom LLFF dataloader adapted from the nerf-pytorch codebase [46].

For our experiments with proposal network samplers, we adapt the nerfacto model from Nerfstudio [35]. This architecture unifies recent NeRF-related research [3, 18, 21, 36, 40] into a single, universally-applicable method. For our nerfacto[†] model, we replace the *HashMLPDensityField* with a 4-layer ReLU-MLP with 256 hidden units, following Mip-NeRF 360 [3], and use positional encoding [20]. Note that this configuration is slower than nerfacto due to the large MLP, however, it enables experiments with our proposed method for proposal network samplers.

5.1. Speeding up NeRF and Mip-NeRF

We evaluate our approach for NeRF and Mip-NeRF using the Blender dataset [20] and the LLFF dataset [19]. For the NeRF experiment, we could either record the activations from Θ_{coarse} or Θ_{fine} . However, only Θ_{coarse} is trained for uniform samples, hence we use the coarse MLP, which also performed better in our experiments.

		NeRF							
ℓ	f	Blender, 800 × 800				LLFF, 1008 × 756			
		PSNR	SSIM	LPIPS	Speedup	PSNR	SSIM	LPIPS	Speedup
1	f_1	27.48	0.92	0.09		25.11	0.75	0.23	
	f_2	28.69	0.92	0.08	24%	25.40	0.76	0.22	28%
	f_3	28.63	0.92	0.08		25.34	0.75	0.23	
2	f_1	27.99	0.92	0.08		24.72	0.73	0.26	
	f_2	29.05	0.93	0.07	20%	25.12	0.74	0.24	24%
	f_3	29.44	0.93	0.07		25.01	0.74	0.25	
3	f_1	24.17	0.88	0.12		24.21	0.71	0.28	
	f_2	26.23	0.90	0.11	18%	24.99	0.74	0.25	20%
	f_3	25.07	0.89	0.11		24.87	0.73	0.26	
NeRF		29.87	0.94	0.06	-	25.89	0.78	0.18	-
		Mip-NeRF							
1	f_1	27.54	0.92	0.09		25.07	0.74	0.24	
	f_2	28.75	0.92	0.08	58%	25.45	0.76	0.22	50%
	f_3	28.72	0.93	0.08		25.35	0.75	0.22	
2	f_1	26.67	0.91	0.09		25.02	0.74	0.25	
	f_2	29.04	0.93	0.07	49%	25.44	0.75	0.22	40%
	f_3	29.35	0.93	0.07		25.35	0.75	0.23	
3	f_1	23.47	0.90	0.10		24.21	0.71	0.28	
	f_2	27.47	0.92	0.09	39%	25.49	0.76	0.22	35%
	f_3	26.45	0.91	0.09		25.42	0.75	0.22	
Mip-NeRF		29.69	0.94	0.06	-	25.99	0.78	0.18	-
		nerfacto [†]							
1	f_1	25.67	0.91	0.10		24.46	0.81	0.15	
	f_2	26.20	0.91	0.09	83%	24.69	0.81	0.14	79%
	f_3	26.17	0.91	0.09		24.13	0.80	0.16	
2	f_1	26.24	0.91	0.09		24.21	0.80	0.16	
	f_2	26.56	0.91	0.09	29%	24.83	0.81	0.14	29%
	f_3	26.31	0.91	0.09		23.51	0.78	0.18	
3	f_1	26.16	0.91	0.09		23.21	0.78	0.19	
	f_2	25.58	0.91	0.10	-	23.78	0.79	0.17	-
	f_3	24.79	0.90	0.10		23.04	0.77	0.22	
nerfacto [†]		27.61	0.93	0.07	-	25.35	0.83	0.12	-

Table 2. Quantitative evaluation of our approach for NeRF, Mip-NeRF and nerfacto[†]. For each dataset, we report PSNR, SSIM [39] and LPIPS [48]. We highlight **best**, **second-best** and **third-best** and point out **failures** of our method.

We report quantitative results in Tab. 2 and show our visual results in Fig. 5. For NeRF, the performance gain is not as significant as for Mip-NeRF due to the different number of samples N_c, N_f : For Mip-NeRF, $N_c = N_f = 128$, contrary to $N_c = 64, N_f = 128 + 64$ for NeRF. Hence, our approach performs favorably for Mip-NeRF due to its more expensive coarse network evaluation. We find that Eqn. (11) for $\ell = 2$ performs best in our experiments. All run-time results for NeRF and Mip-NeRF were obtained on an NVIDIA Quadro RTX 8000.

5.2. Speeding up Proposal Networks

First, we evaluate our approach for nerfacto[†] using the Blender dataset [20] and the LLFF dataset [19]. For our experiments, we approximate the first iteration of proposal sampling: As the proposal sampler iterations use $\{256, 96\}$

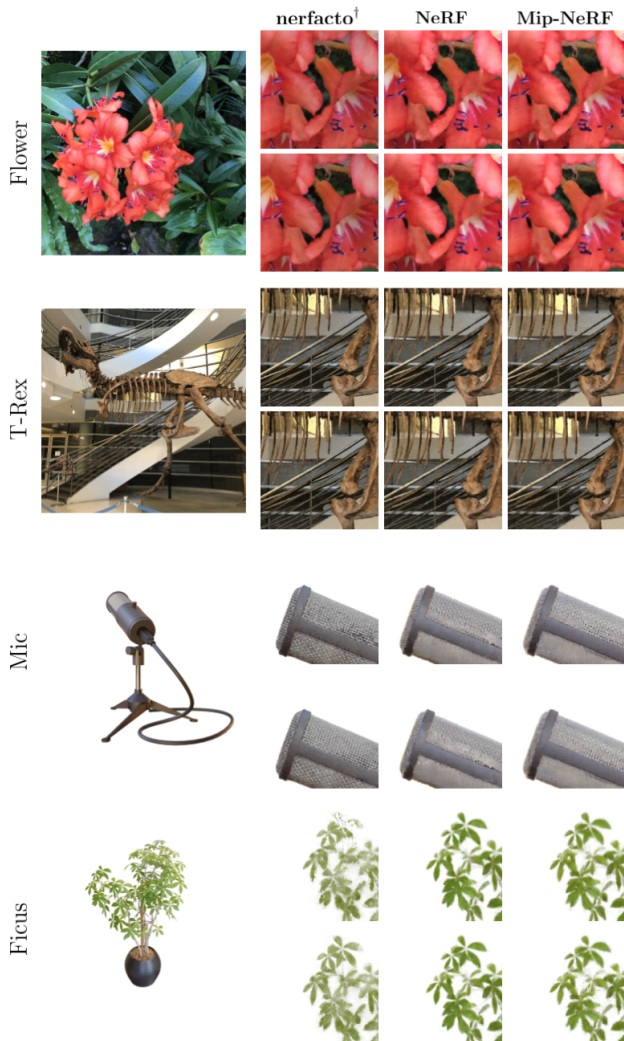


Figure 5. Qualitative results for our approach (top row) compared to baseline methods (bottom row) for synthetic and real-world scenes. As can be seen in the zoomed-in views, our best renderings are virtually indistinguishable from the baseline in most configurations.

samples, respectively, approximating the first round yields more performance improvements¹. In addition, our functions f_i are biased towards uniform or piecewise-linear sample placement, and lead to unfavorable results when interval sizes vary significantly.

We report quantitative results in Tab. 2 and show our visual results in Fig. 5. For the nerfacto[†] model, we observe a huge performance gain for the first layer, which declines sharply — this is due to the large number of samples for the first proposal sampler iteration and the small shading network. We find that Eqn. (11) applied in layer $\ell = 2$ performed best when considering standard image metrics.

¹Please note that nerfacto[†] is significantly slower than nerfacto due to the use of an MLP.

	Mip-NeRF 360				Deep Blending				
	PSNR	SSIM	LPIPS	Speedup	PSNR	SSIM	LPIPS	Speedup	
1	f_1	22.80	0.687	0.282	76 %	28.01	0.858	0.237	72 %
	f_2	22.92	0.691	0.279		28.01	0.859	0.237	
	f_3	22.80	0.687	0.283		27.98	0.858	0.239	
2	f_1	22.71	0.683	0.289	28 %	27.97	0.858	0.239	27 %
	f_2	22.66	0.681	0.292		27.95	0.856	0.243	
	f_3	22.67	0.680	0.297		27.94	0.856	0.242	
3	f_1	22.44	0.675	0.305	-	27.77	0.849	0.256	-
	f_2	22.21	0.669	0.309		27.48	0.844	0.264	
	f_3	22.39	0.669	0.316		27.78	0.847	0.262	
nerfacto [†]	24.20	0.732	0.237	-	29.26	0.880	0.207	-	

Table 3. Quantitative evaluation of our approach for nerfacto[†] for Mip-NeRF 360 [3] and Deep Blending [10]. For each dataset, we report PSNR, SSIM [39] and LPIPS [48]. We highlight **best**, **second-best** and **third-best** for each image metric.

The reported run-time results for nerfacto[†] were obtained on an NVIDIA RTX 2070 Super.

Further, we evaluate our method on unbounded, real-world scenes, leveraging the Mip-NeRF 360 [3] and Deep Blending [10] datasets. We adapt nerfacto[†] in the following ways: We use 7×10^5 training iterations, disable the appearance embedding, camera optimizer, distortion loss and the collider. We present our results in Tab. 3 and show example outputs in Fig. 6. Compared to the results presented for the LLFF dataset [19], our method performs less favorably in this scenario. We attribute this to the large sampling domain in these datasets, which requires more precise density estimates. Performance-wise, we observe the same trends as for the previous experiments. We used an NVIDIA RTX 4090 for these two datasets.

5.3. Ablation studies

To support our quantitative evaluation, we conduct additional ablation studies to analyze the effect of varying network capacity and the use of the ReLU function in our proposed functions f_i . More details and ablation studies are provided in the supplementary material.

Varying Network Capacity. To investigate the effect of network capacity, we vary the number of layers for the proposal network sampler of nerfacto[†] using the Mip-NeRF 360 dataset [3]. We report results in Tab. 4 — remarkably, lower network capacity allows for better visual quality retention for our best configuration. As the decomposition property of NeRFs [38] does not apply to proposal network samplers, lower network capacity forces the proposal network sampler to learn a more informative latent representation of density. We report results for all configurations of ℓ and f_i in the supplementary material — the configuration $\ell = 1$ and f_2 performs best for all tested network capacities.



Figure 6. Example renderings from our our method compared to nerfacto[†] (Baseline) for the Mip-NeRF 360 dataset [3]. In most configurations, our approach is able to retain high visual fidelity on this challenging benchmark.

#Layers	nerfacto [†]			Ours, $\ell = 1, f_2$		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
2	24.201	0.726	0.247	23.686	0.704	0.267
3	24.272	0.732	0.237	23.310	0.700	0.268
4	24.201	0.732	0.237	22.925	0.691	0.279
5	24.176	0.729	0.240	22.941	0.692	0.276
6	24.247	0.732	0.234	22.891	0.690	0.278

Table 4. Effect of varying network capacity for nerfacto[†] using the Mip-NeRF 360 dataset [3]: When we decrease capacity and use the previous best configuration of $\ell = 1$ and f_2 , our method retains more quality and exhibits better image metrics.

ReLU or no ReLU. We also conduct an additional ablation study to test the effectiveness of applying the ReLU in Eqns. (10), (11) and (12): We remove the ReLU for all f_i , use the nerfacto[†] model for the Mip-NeRF 360 dataset [3] and report the results in Tab. 5.

As we can see, not applying the ReLU leads to worse outputs for every configuration we tested. Empirically, this is caused by the more inconsistent histograms and larger standard deviation $\sigma(f_\ell)$. We note that the configuration $\ell = 1, f_2$ performs best for both variants of our method.

6. Limitations

Our method is not without its limitations. For certain test set views, using Eqn. (10) results in rays \mathbf{r} with $\hat{\mathbf{w}} = \mathbf{0}$, which does not produce a good set of samples for the subsequent shading pass. We indicate these cases in Tab. 2 and show examples in Fig. 7. These failures are particularly evident for Mip-NeRF, as the fine pass does not include the uniformly

		with ReLU			without ReLU		
		PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
1	f_1	22.796	0.687	0.282	22.347	0.666	0.291
	f_2	22.925	0.691	0.279	22.565	0.671	0.284
	f_3	22.804	0.687	0.283	22.535	0.670	0.285
2	f_1	22.708	0.683	0.289	21.527	0.637	0.327
	f_2	22.656	0.681	0.292	21.536	0.637	0.327
	f_3	22.666	0.680	0.297	21.501	0.635	0.329
3	f_1	22.443	0.675	0.305	21.598	0.635	0.334
	f_2	22.214	0.669	0.309	21.618	0.637	0.332
	f_3	22.390	0.669	0.316	21.570	0.636	0.333

Table 5. Ablation study on the effectiveness of using the ReLU function. We evaluate our method with f_i , either using the ReLU or not, for the Mip-NeRF 360 dataset [3] and report averaged scores. With the ReLU, we obtain better performance for all configurations.

distributed samples. Eqns. (11) and (12) mitigate this issue effectively, but might lead to outputs with less detail due to their more conservative formulation.

Although our approach performs well for a variety of architectures and datasets, extending it to hybrid or explicit representations such as [8, 21, 33, 44] is not trivial, as these approaches replace the MLP with more efficient alternatives. Further, our handcrafted functions for density estimation were created intuitively and are biased towards uniform or piecewise-linear sample placement.



Figure 7. Examples of failures of our method. Using Eqn. (10) may lead to rays \mathbf{r} with $\hat{\mathbf{w}} = \mathbf{0}$, producing wrong samples for the subsequent fine pass.

7. Conclusion and Outlook

We have presented a generally applicable framework to analyze and visualize the intermediate activations of coordinate-based ReLU-MLPs. Building on our findings, we proposed a novel method to obtain density estimates using our derived activation features. We demonstrated that our approach produces faithful renderings with improved performance and provided valuable insight into the internals of neural fields. We believe our concepts could be applied to other tasks involving coordinate-based MLPs.

References

- [1] Relja Arandjelović and Andrew Zisserman. NeRF in detail: Learning to sample for view synthesis. *arXiv CoRR*, abs/2106.05264, 2021. [2](#)
- [2] Jonathan T. Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P. Srinivasan. Mip-NeRF: A Multiscale Representation for Anti-Aliasing Neural Radiance Fields. In *ICCV*, 2021. [2](#), [3](#), [6](#)
- [3] Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Mip-NeRF 360: Unbounded Anti-Aliased Neural Radiance Fields. In *CVPR*, 2022. [2](#), [3](#), [6](#), [7](#), [8](#)
- [4] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. TensorRF: Tensorial Radiance Fields. In *ECCV*, 2022. [2](#)
- [5] Zhiqin Chen, Thomas Funkhouser, Peter Hedman, and Andrea Tagliasacchi. MobileNeRF: Exploiting the Polygon Rasterization Pipeline for Efficient Neural Field Rendering on Mobile Architectures. In *CVPR*, 2023. [2](#)
- [6] Shin-Fang Chng, Sameera Ramasinghe, Jamie Sherrah, and Simon Lucey. Gaussian Activated Neural Radiance Fields for High Fidelity Reconstruction & Pose Estimation. In *ECCV*, 2022. [1](#)
- [7] Rizal Fathony, Anit Kumar Sahu, Devin Willmott, and J Zico Kolter. Multiplicative Filter Networks. In *ICLR*, 2021. [1](#)
- [8] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance Fields without Neural Networks. In *CVPR*, 2022. [2](#), [8](#)
- [9] Stephan J. Garbin, Marek Kowalski, Matthew Johnson, Jamie Shotton, and Julien Valentin. FastNeRF: High-Fidelity Neural Rendering at 200FPS. In *ICCV*, 2021. [2](#)
- [10] Peter Hedman, Julien Philip, True Price, Jan-Michael Frahm, George Drettakis, and Gabriel Brostow. Deep Blending for Free-viewpoint Image-based Rendering. *ACM TOG*, 37(6), 2018. [6](#), [7](#)
- [11] Peter Hedman, Pratul P. Srinivasan, Ben Mildenhall, Jonathan T. Barron, and Paul E. Debevec. Baking Neural Radiance Fields for Real-Time View Synthesis. In *ICCV*, 2021. [2](#)
- [12] Tao Hu, Shu Liu, Yilun Chen, Tiancheng Shen, and Jiaya Jia. EfficientNeRF: Efficient Neural Radiance Fields. In *CVPR*, 2022. [2](#)
- [13] Animesh Karnewar, Tobias Ritschel, Oliver Wang, and Niloy J. Mitra. ReLU Fields: The Little Non-linearity That Could. In *SIGGRAPH*, 2022. [2](#)
- [14] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM TOG*, 42(4), 2023. [2](#)
- [15] Jonas Kulhanek and Torsten Sattler. Tetra-NeRF: Representing Neural Radiance Fields Using Tetrahedra. In *ICCV*, 2023. [2](#)
- [16] Andreas Kurz, Thomas Neff, Zhaoyang Lv, Michael Zollhöfer, and Markus Steinberger. AdaNeRF: Adaptive Sampling for Real-time Rendering of Neural Radiance Fields. In *ECCV*, 2022. [2](#)
- [17] Haotong Lin, Sida Peng, Zhen Xu, Yunzhi Yan, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Efficient Neural Radiance Fields with Learned Depth-Guided Sampling. In *SIGGRAPH Asia*, 2022. [2](#), [3](#)
- [18] Ricardo Martin-Brualla, Noha Radwan, Mehdi S. M. Sajjadi, Jonathan T. Barron, Alexey Dosovitskiy, and Daniel Duckworth. NeRF in the Wild: Neural Radiance Fields for Unconstrained Photo Collections. In *CVPR*, 2021. [6](#)
- [19] 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 TOG*, 38(4), 2019. [6](#), [7](#)
- [20] 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](#), [3](#), [4](#), [5](#), [6](#)
- [21] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant Neural Graphics Primitives with a Multiresolution Hash Encoding. *ACM TOG*, 41(4), 2022. [2](#), [6](#), [8](#)
- [22] 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. *Comput. Graph. Forum*, 40(4):45–59, 2021. [2](#)
- [23] Michael Oechsle, Songyou Peng, and Andreas Geiger. UNISURF: Unifying Neural Implicit Surfaces and Radiance Fields for Multi-View Reconstruction. In *ICCV*, 2021. [3](#)
- [24] Martin Píala and Ronald Clark. TerminiNeRF: Ray Termination Prediction for Efficient Neural Rendering. In *3DV*, 2021. [2](#), [3](#)
- [25] Nasim Rahaman, Aristide Baratin, Devansh Arpit, Felix Draxler, Min Lin, Fred Hamprecht, Yoshua Bengio, and Aaron Courville. On the Spectral Bias of Neural Networks. In *ICLR*, 2019. [1](#)
- [26] Sameera Ramasinghe and Simon Lucey. Beyond Periodicity: Towards a Unifying Framework for Activations in Coordinate-MLPs. In *ECCV*, 2022. [1](#)
- [27] Daniel Rebain, Wei Jiang, Soroosh Yazdani, Ke Li, Kwang M. Yi, and Andrea Tagliasacchi. DeRF: Decomposed Radiance Fields. In *CVPR*, 2021. [2](#)
- [28] Christian Reiser, Songyou Peng, Yiyi Liao, and Andreas Geiger. KiloNeRF: Speeding up Neural Radiance Fields with Thousands of Tiny MLPs. In *ICCV*, 2021. [2](#)
- [29] Christian Reiser, Stephan Garbin, Pratul P. Srinivasan, Dor Verbin, Richard Szeliski, Ben Mildenhall, Jonathan T. Barron, Peter Hedman, and Andreas Geiger. Binary Opacity Grids: Capturing Fine Geometric Detail for Mesh-Based View Synthesis. *arXiv CoRR*, abs/2402.12377, 2024. [2](#), [3](#)
- [30] Vishwanath Saragadam, Daniel LeJeune, Jasper Tan, Guha Balakrishnan, Ashok Veeraraghavan, and Richard G Baraniuk. WIRE: Wavelet Implicit Neural Representations. In *CVPR*, 2023. [1](#)
- [31] Vincent Sitzmann, Julien N.P. Martel, Alexander W. Bergman, David B. Lindell, and Gordon Wetzstein. Implicit Neural Representations with Periodic Activation Functions. In *NeurIPS*, 2020. [1](#)

- [32] Jost T. Springenberg, Alexey Dosovitskiy, Thomas Brox, and Martin Riedmiller. Striving for Simplicity: The All Convolutional Net. In *ICLR*, 2015. 4
- [33] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct Voxel Grid Optimization: Super-fast Convergence for Radiance Fields Reconstruction. In *CVPR*, 2022. 2, 8
- [34] 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. In *NeurIPS*, 2020. 1, 3
- [35] Matthew Tancik, Ethan Weber, Evonne Ng, Ruilong Li, Brent Yi, Justin Kerr, Terrance Wang, Alexander Kristoffersen, Jake Austin, Kamyar Salahi, Abhik Ahuja, David McAllister, and Angjoo Kanazawa. Nerfstudio: A Modular Framework for Neural Radiance Field Development. In *SIGGRAPH*, 2023. 6
- [36] Dor Verbin, Peter Hedman, Ben Mildenhall, Todd Zickler, Jonathan T. Barron, and Pratul P. Srinivasan. Ref-NeRF: Structured View-Dependent Appearance for Neural Radiance Fields. In *CVPR*, 2022. 6
- [37] Ziyu Wan, Christian Richardt, Aljaž Božič, Chao Li, Vijay Rengarajan, Seonghyeon Nam, Xiaoyu Xiang, Tuotuo Li, Bo Zhu, Rakesh Ranjan, and Jing Liao. Learning Neural Duplex Radiance Fields for Real-Time View Synthesis. In *CVPR*, 2023. 2
- [38] Shengze Wang, Alexey Supikov, Joshua Ratcliff, Henry Fuchs, and Ronald Azuma. INV: Towards Streaming Incremental Neural Videos. *arXiv CoRR*, abs/2302.01532, 2023. 2, 6, 7
- [39] Zhou Wang, Alam C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image Quality Assessment: From Error Visibility to Structural Similarity. *IEEE TIP*, 13(4):600–612, 2004. 6, 7
- [40] Zirui Wang, Shangzhe Wu, Weidi Xie, Min Chen, and Victor Adrian Prisacariu. NeRF—: Neural Radiance Fields Without Known Camera Parameters. *arXiv CoRR*, abs/2102.07064, 2021. 6
- [41] Zian Wang, Tianchang Shen, Merlin Nimier-David, Nicholas Sharp, Jun Gao, Alexander Keller, Sanja Fidler, Thomas Müller, and Zan Gojcic. Adaptive Shells for Efficient Neural Radiance Field Rendering. *ACM TOG*, 42(6), 2023. 3
- [42] 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. In *CVPR*, 2022. 2
- [43] Yiheng Xie, Towaki Takikawa, Shunsuke Saito, Or Litany, Shiqin Yan, Numair Khan, Federico Tombari, James Tompkin, Vincent Sitzmann, and Srinath Sridhar. Neural Fields in Visual Computing and Beyond. *Comput. Graph. Forum*, 41(2):641–676, 2022. 1
- [44] Qiangeng Xu, Zexiang Xu, Julien Philip, Sai Bi, Zhixin Shu, Kalyan Sunkavalli, and Ulrich Neumann. Point-NeRF: Point-based Neural Radiance Fields. In *CVPR*, 2022. 2, 8
- [45] Lior Yariv, Peter Hedman, Christian Reiser, Dor Verbin, Pratul P Srinivasan, Richard Szeliski, Jonathan T Barron, and Ben Mildenhall. BakedSDF: Meshing Neural SDFs for Real-Time View Synthesis. In *SIGGRAPH*, 2023. 2, 3
- [46] Lin Yen-Chen. NeRF-pytorch. <https://github.com/yenchenlin/nerf-pytorch/>, 2020. 6
- [47] Matthew D. Zeiler and Rob Fergus. Visualizing and Understanding Convolutional Networks. In *ECCV*, 2014. 4
- [48] 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, 7