

FastNeRF: High-Fidelity Neural Rendering at 200FPS

Stephan J. Garbin* Marek Kowalski* Matthew Johnson
 Jamie Shotton Julien Valentin

Microsoft



Figure 1. **FastNeRF** renders high-resolution photorealistic novel views of objects at hundreds of frames per second. Comparable existing methods, such as NeRF, are orders of magnitude slower and can only render very low resolution images at interactive rates.

Abstract

Recent work on Neural Radiance Fields (NeRF) showed how neural networks can be used to encode complex 3D environments that can be rendered photorealistically from novel viewpoints. Rendering these images is very computationally demanding and recent improvements are still a long way from enabling interactive rates, even on high-end hardware. Motivated by scenarios on mobile and mixed reality devices, we propose FastNeRF, the first NeRF-based system capable of rendering high fidelity photorealistic images at 200Hz on a high-end consumer GPU. The core of our method is a graphics-inspired factorization that allows for (i) compactly caching a deep radiance map at each position in space, (ii) efficiently querying that map using ray directions to estimate the pixel values in the rendered image. Extensive experiments show that the proposed method is 3000 times faster than the original NeRF algorithm and at least an order of magnitude faster than existing work on accelerating NeRF, while maintaining visual quality and extensibility.

1. Introduction

Rendering scenes in real-time at photorealistic quality has long been a goal of computer graphics. Traditional approaches such as rasterization and ray-tracing often require

*Denotes equal contribution

Project website: <https://microsoft.github.io/FastNeRF>

significant manual effort in designing or pre-processing the scene in order to achieve both quality and speed. Recently, neural rendering [10, 27, 17, 25, 19] has offered a disruptive alternative: involve a neural network in the rendering pipeline to output either images directly [24, 18, 27] or to model implicit functions that represent a scene appropriately [5, 23, 29, 25, 40]. Beyond rendering, some of these approaches implicitly reconstruct a scene from static or moving cameras [25, 36, 1], thereby greatly simplifying the traditional reconstruction pipelines used in computer vision.

One of the most prominent recent advances in neural rendering is Neural Radiance Fields (NeRF) [25] which, given a handful of images of a static scene, learns an implicit volumetric representation of the scene that can be rendered from novel viewpoints. The rendered images are of high quality and correctly retain thin structures, view-dependent effects, and partially-transparent surfaces. NeRF has inspired significant follow-up work that has addressed some of its limitations, notably extensions to dynamic scenes [31, 7, 46], relighting [2, 3, 37], and incorporation of uncertainty [22].

One common challenge to all of the NeRF-based approaches is their high computational requirements for rendering images. The core of this challenge resides in NeRF’s volumetric scene representation. More than 100 neural network calls are required to render a single image pixel, which translates into several seconds being required to render low-resolution images on high-end GPUs. Recent explorations [35, 15, 28, 16] conducted with the aim of improving NeRF’s computational requirements reduced the

render time by up to $50\times$. While impressive, these advances are still a long way from enabling real-time rendering on consumer-grade hardware. Our work bridges this gap while maintaining quality, thereby opening up a wide range of new applications for neural rendering. Furthermore, our method could form the fundamental building block for neural rendering at high resolutions.

To achieve this goal, we use caching to trade memory for computational efficiency. As NeRF is fundamentally a function of positions $\mathbf{p} \in \mathbb{R}^3$ and ray directions $\mathbf{d} \in \mathbb{R}^2$ to color $\mathbf{c} \in \mathbb{R}^3$ (RGB) and a scalar density σ , a naïve approach would be to build a cached representation of this function in the space of its input parameters. Since σ only depends on \mathbf{p} , it can be cached using existing methodologies. The color \mathbf{c} , however, is a function of both ray direction \mathbf{d} and position \mathbf{p} . If this 5 dimensional space were to be discretized with 512 bins per dimension, a cache of around 176 terabytes of memory would be required – dramatically more than is available on current consumer hardware.

Ideally, we would treat directions and positions separately and thus avoid the polynomial explosion in required cache space. Fortunately this problem is not unique to NeRF; the rendering equation (modulo the wavelength of light and time) is also a function of \mathbb{R}^5 and solving it efficiently is the primary challenge of real-time computer graphics. As such there is a large body of research which investigates ways of approximating this function as well as efficient means of evaluating the integral. One of the fastest approximations of the rendering equation involves the use of spherical harmonics such that the integral results in a dot product of the harmonic coefficients of the material model and lighting model. Inspired by this efficient approximation, we factorize the problem into two independent functions whose outputs are combined using their inner product to produce RGB values. The first function takes the form of a Multi-Layer Perceptron (MLP) that is conditioned on position in space and returns a compact deep radiance map parameterized by D components. The second function is an MLP conditioned on ray direction that produces D weights for the D deep radiance map components.

This factorized architecture, which we call FastNeRF, allows for independently caching the position-dependent and ray direction-dependent outputs. Assuming that k and l denote the number of bins for positions and ray directions respectively, caching NeRF would have a memory complexity of $\mathcal{O}(k^3 l^2)$. In contrast, caching FastNeRF would have a complexity of $\mathcal{O}(k^3 * (1 + D * 3) + l^2 * D)$. As a result of this reduced memory complexity, FastNeRF can be cached in the memory of a high-end consumer GPU, thus enabling very fast function lookup times that in turn lead to a dramatic increase in test-time performance.

While caching does consume a significant amount of memory, it is worth noting that current implementations of

NeRF also have large memory requirements. A single forward pass of NeRF requires performing hundreds of forward passes through an eight layer 256 hidden unit MLP *per pixel*. If pixels are processed in parallel for efficiency this consumes large amounts of memory, even at moderate resolutions. Since many natural scenes (*e.g.* a living room, a garden) are sparse, we are able to store our cache sparsely. In some cases this can make our method actually more memory efficient than NeRF.

In summary, our main contributions are:

- The first NeRF-based system capable of rendering photorealistic novel views at 200FPS, thousands of times faster than NeRF.
- A graphics-inspired factorization that can be compactly cached and subsequently queried to compute the pixel values in the rendered image.
- A blueprint detailing how the proposed factorization can efficiently run on the GPU.

2. Related work

FastNeRF belongs to the family of Neural Radiance Fields methods [25] and is trained to learn an implicit, compressed deep radiance map parameterized by position and view direction that provides color and density estimates. Our method differs from [25] in the structure of the implicit model, changing it in such a way that positional and directional components can be discretized and stored in sparse 3D grids.

This also differentiates our method from models that use a discretized grid at training time, such as Neural Volumes [18] or Deep Reflectance Volumes [3]. Due to the memory requirements associated with generating and holding 3D volumes at training time, the output image resolution of [18, 3] is limited by the maximum volume size of 128^3 . In contrast, our method uses volumes as large as 1024^3 .

Subsequent to [18, 3], the problem of parameter estimation for an MLP with low dimensional input coordinates was addressed using Fourier feature encodings. After use in [41], this was popularized in [25, 49] and explored in greater detail by [40].

Neural Radiance Fields (NeRF) [25] first showed convincing compression of a light field utilizing Fourier features. Using an entirely implicit model, NeRF is not bound to any voxelized grid, but only to a specific domain. Despite impressive results, one disadvantage of this method is that a complex MLP has to be called for every sample along each ray, meaning hundreds of MLP invocations per image pixel.

One way to speed up NeRF’s inference is to scale to multiple processors. This works particularly well because every pixel in NeRF can be computed independently. For an 800×800 pixel image, JaxNeRF [6] achieves an inference

speed of 20.77 seconds on an Nvidia Tesla V100 GPU, 2.65 seconds on 8 V100 GPUS, and 0.35 seconds on 128 second generation Tensor Processing Units. Similarly, the volumetric integration domain can be split and separate models used for each part. This approach is taken in the Decomposed Radiance Fields method [35], where an integration domain is divided into Voronoi cells. This yields better test scores, and a speedup of up to $3\times$.

A different way to increase efficiency is to realize that natural scenes tend to be volumetrically sparse. Thus, efficiency can be gained by skipping empty regions of space. This amounts to importance sampling of the occupancy distribution of the integration domain. One approach is to use a voxel grid in combination with the implicit function learned by the MLP, as proposed in Neural Sparse Voxel Fields (NSVFs) [16] and Neural Geometric Level of Detail [38] (for SDFs only), where a dynamically constructed sparse octree is used to represent scene occupancy. As a network still has to be queried inside occupied voxels, however, NSVFs takes between 1 and 4 seconds to render an 800^2 image, with decreases to PSNR at the lower end of those timings. Our method is orders of magnitude faster in a similar scenario.

Another way to represent the importance distribution is via depth prediction for each pixel. This approach is taken in [28], which is concurrent to ours and achieves roughly 15FPS for 800^2 images at reduced quality or roughly half that for quality comparable to NeRF.

Orthogonal to this, AutoInt [15] showed that a neural network can be used to approximate the integrals along each ray with far fewer samples. While significantly faster than NeRF, this still does not provide interactive frame rates.

What differentiates our method from those described above is that FastNeRF’s proposed decomposition, and subsequent caching, lets us avoid calls to an MLP at inference time entirely. This makes our method faster in absolute terms even on a single machine.

It is worth noting that our method does not address training speed, as for example [39]. In that case, the authors propose improving the training speed of NeRF models by finding initialisation through meta learning.

Finally, there are orthogonal neural rendering strategies capable of fast inference, such as Neural Point-Based Graphics [11], Mixture of Volumetric Primitives [19] or Pulsar [12], which use forms of geometric primitives and rasterization. In this work, we only deal with NeRF-like implicit functional representations paired with ray tracing.

3. Method

In this section we describe FastNeRF, a method that is 3000 times faster than the original Neural Radiance Fields (NeRF) system [25] (Section 3.1). This breakthrough allows for rendering high-resolution photorealistic images at

over 200Hz on high-end consumer hardware. The core insight of our approach (Section 3.2) consists of factorizing NeRF into two neural networks: a position-dependent network that produces a deep radiance map and a direction-dependent network that produces weights. The inner product of the weights and the deep radiance map estimates the color in the scene at the specified position and as seen from the specified direction. This architecture, which we call FastNeRF, can be efficiently cached (Section 3.3), significantly improving test time efficiency whilst preserving the visual quality of NeRF. See Figure 2 for a comparison of the NeRF and FastNeRF network architectures.

3.1. Neural Radiance Fields

A Neural Radiance Field (NeRF) [25] captures a volumetric 3D representation of a scene within the weights of a neural network. NeRF’s neural network $\mathcal{F}_{NeRF} : (\mathbf{p}, \mathbf{d}) \mapsto (\mathbf{c}, \sigma)$ maps a 3D position $\mathbf{p} \in \mathbb{R}^3$ and a ray direction $\mathbf{d} \in \mathbb{R}^2$ to a color value \mathbf{c} and transparency σ . See the left side of Figure 2 for a diagram of \mathcal{F}_{NeRF} .

In order to render a single image pixel, a ray is cast from the camera center, passing through that pixel and into the scene. We denote the direction of this ray as \mathbf{d} . A number of 3D positions $(\mathbf{p}_1, \dots, \mathbf{p}_N)$ are then sampled along the ray between its near and far bounds defined by the camera parameters. The neural network \mathcal{F}_{NeRF} is evaluated at each position \mathbf{p}_i and ray direction \mathbf{d} to produce color \mathbf{c}_i and transparency σ_i . These intermediate outputs are then integrated as follows to produce the final pixel color $\hat{\mathbf{c}}$:

$$\hat{\mathbf{c}} = \sum_{i=1}^N T_i (1 - \exp(-\sigma_i \delta_i)) \mathbf{c}_i, \quad (1)$$

where $T_i = \exp(-\sum_{j=i}^{i-1} \sigma_j \delta_j)$ is the transmittance and $\delta_i = (\mathbf{p}_{i+1} - \mathbf{p}_i)$ is the distance between the samples. Since \mathcal{F}_{NeRF} depends on ray directions, NeRF has the ability to model viewpoint-dependent effects such as specular reflections, which is one key dimension in which NeRF improves upon traditional 3D reconstruction methods.

Training a NeRF network requires a set of images of a scene as well as the extrinsic and intrinsic parameters of the cameras that captured the images. In each training iteration, a subset of pixels from the training images are chosen at random and for each pixel a 3D ray is generated. Then, a set of samples is selected along each ray and the pixel color $\hat{\mathbf{c}}$ is computed using Equation (1). The training loss is the mean squared difference between $\hat{\mathbf{c}}$ and the ground truth pixel value. For further details please refer to [25].

While NeRF renders photorealistic images, it requires calling \mathcal{F}_{NeRF} a large number of times to produce a single image. With the default number of samples per pixel $N = 192$ proposed in NeRF [25], nearly 400 million calls to \mathcal{F}_{NeRF} are required to compute a single high definition (1080p) image. Moreover the intermediate outputs of

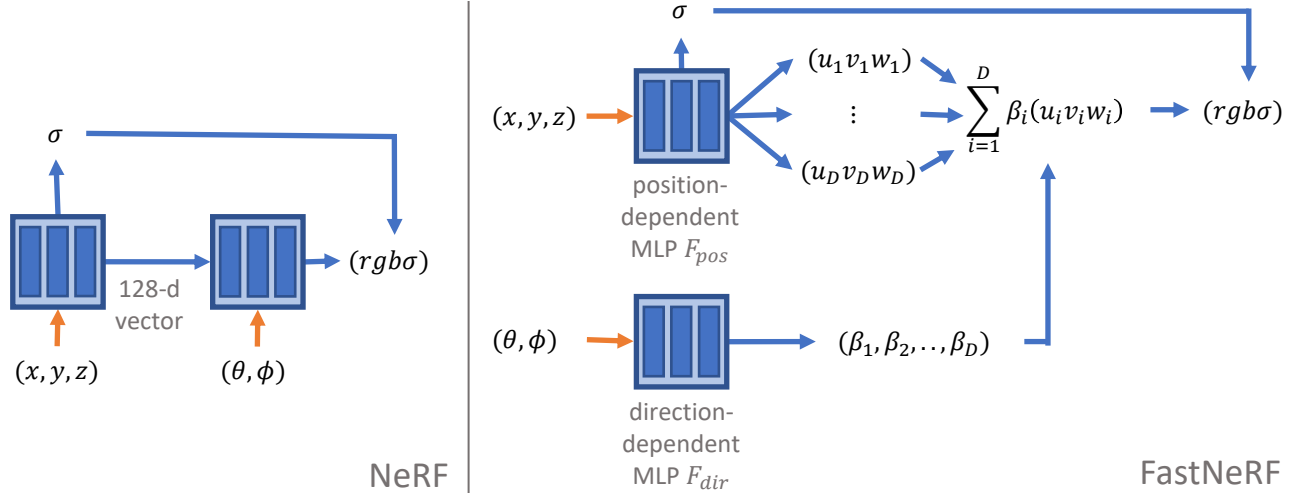


Figure 2. Left: NeRF neural network architecture. (x, y, z) denotes the input sample position, (θ, ϕ) denotes the ray direction and (r, g, b, σ) are the output color and transparency values. Right: our FastNeRF architecture splits the same task into two neural networks that are amenable to caching. The position-dependent network \mathcal{F}_{pos} outputs a deep radiance map $(\mathbf{u}, \mathbf{v}, \mathbf{w})$ consisting of D components, while the \mathcal{F}_{dir} outputs the weights for those components $(\beta_1, \dots, \beta_D)$ given a ray direction as input.

this process would take hundreds of gigabytes of memory. Even on high-end consumer GPUs, this constrains the original method to be executed over several batches even for medium resolution (800×800) images, leading to additional computational overheads.

3.2. Factorized Neural Radiance Fields

Taking a step away from neural rendering for a moment, we recall that in traditional computer graphics, the rendering equation [9] is an integral of the form

$$L_o(\mathbf{p}, \mathbf{d}) = \int_{\Omega} f_r(\mathbf{p}, \mathbf{d}, \boldsymbol{\omega}_i) L_i(\mathbf{p}, \boldsymbol{\omega}_i) (\boldsymbol{\omega}_i \cdot \mathbf{n}) d\boldsymbol{\omega}_i, \quad (2)$$

where $L_o(\mathbf{p}, \mathbf{d})$ is the radiance leaving the point \mathbf{p} in direction \mathbf{d} , $f_r(\mathbf{p}, \mathbf{d}, \boldsymbol{\omega}_i)$ is the reflectance function capturing the material properties at position \mathbf{p} , $L_i(\mathbf{p}, \boldsymbol{\omega}_i)$ describes the amount of light reaching \mathbf{p} from direction $\boldsymbol{\omega}_i$, and \mathbf{n} corresponds to the direction of the surface normal at \mathbf{p} . Given its practical importance, evaluating this integral efficiently has been a subject of active research for over three decades [9, 42, 8, 33]. One efficient way to evaluate the rendering equation is to approximate $f_r(\mathbf{p}, \mathbf{d}, \boldsymbol{\omega}_i)$ and $L_i(\mathbf{p}, \boldsymbol{\omega}_i)$ using spherical harmonics [4, 45]. In this case, evaluating the integral boils down to a dot product between the coefficients of both spherical harmonics approximations.

With this insight in mind, we can return to neural rendering. In the case of NeRF, $\mathcal{F}_{NeRF} : (\mathbf{p}, \mathbf{d})$ can also be interpreted as returning the radiance leaving point \mathbf{p} in direction \mathbf{d} . This key insight leads us to propose a new network architecture for NeRF, which we call FastNeRF. This novel architecture consists in splitting NeRF’s neural network \mathcal{F}_{NeRF} into two networks: one that depends only

on positions \mathbf{p} and one that depends only on the ray directions \mathbf{d} . Similarly to evaluating the rendering equation using spherical harmonics, our position-dependent \mathcal{F}_{pos} and direction-dependent \mathcal{F}_{dir} functions produce outputs that are combined using the dot product to obtain the color value at position \mathbf{p} observed from direction \mathbf{d} . Crucially, this factorization splits a single function that takes inputs in \mathbb{R}^5 into two functions that take inputs in \mathbb{R}^3 and \mathbb{R}^2 . As explained in the following section, this makes caching the outputs of the network possible and allows accelerating NeRF by 3 orders of magnitude on consumer hardware. See Figure 3 for a visual representation of the achieved speedup.

The position-dependent and direction-dependent functions of FastNeRF are defined as follows:

$$\mathcal{F}_{pos} : \mathbf{p} \mapsto \{\sigma, (\mathbf{u}, \mathbf{v}, \mathbf{w})\}, \quad (3)$$

$$\mathcal{F}_{dir} : \mathbf{d} \mapsto \boldsymbol{\beta}, \quad (4)$$

where $\mathbf{u}, \mathbf{v}, \mathbf{w}$ are D -dimensional vectors that form a deep radiance map describing the view-dependent radiance at position \mathbf{p} . The output of \mathcal{F}_{dir} , $\boldsymbol{\beta}$, is a D -dimensional vector of weights for the D components of the deep radiance map. The inner product of the weights and the deep radiance map

$$\mathbf{c} = (r, g, b) = \sum_{i=1}^D \beta_i (u_i, v_i, w_i) = \boldsymbol{\beta}^T \cdot (\mathbf{u}, \mathbf{v}, \mathbf{w}) \quad (5)$$

results in the estimated color $\mathbf{c} = (r, g, b)$ at position \mathbf{p} observed from direction \mathbf{d} .

When the two functions are combined using Equation (5), the resulting function $\mathcal{F}_{FastNeRF}(\mathbf{p}, \mathbf{d}) \mapsto (\mathbf{c}, \sigma)$ has the same signature as the function \mathcal{F}_{NeRF} used in NeRF. This effectively allows for using the FastNeRF architecture as a drop-in replacement for NeRF’s architecture at runtime.

3.3. Caching

Given the large number of samples that need to be evaluated to render a single pixel, the cost of computing \mathcal{F} dominates the total cost of NeRF rendering. Thus, to accelerate NeRF, one could attempt to reduce the test-time cost of \mathcal{F} by caching its outputs for a set of inputs that cover the space of the scene. The cache can then be evaluated at a fraction of the time it takes to compute \mathcal{F} .

For a trained NeRF model, we can define a bounding box \mathcal{V} that covers the entire scene captured by NeRF. We can then uniformly sample k values for each of the 3 world-space coordinates $(x, y, z) = \mathbf{p}$ within the bounds of \mathcal{V} . Similarly, we uniformly sample l values for each of the ray direction coordinates $(\theta, \phi) = \mathbf{d}$ with $\theta \in \langle 0, \pi \rangle$ and $\phi \in \langle 0, 2\pi \rangle$. The cache is then generated by computing \mathcal{F} for each combination of sampled \mathbf{p} and \mathbf{d} .

The size of such a cache for a *standard* NeRF model with $k = l = 1024$ and densely stored 16-bit floating point values is approximately 5600 Terabytes. Even for highly sparse volumes, where one would only need to keep 1% of the outputs, the size of the cache would still severely exceed the memory capacity of consumer-grade hardware. This huge memory requirement is caused by the usage of both \mathbf{d} and \mathbf{p} as input to $\mathcal{F}_{\text{NeRF}}$. A separate output needs to be saved for each combination of \mathbf{d} and \mathbf{p} , resulting in a memory complexity in the order of $\mathcal{O}(k^3 l^2)$.

Our FastNeRF architecture makes caching feasible. For $k = l = 1024$, $D = 8$ the size of two dense caches holding $\{\sigma, (\mathbf{u}, \mathbf{v}, \mathbf{w})\}$ and β would be approximately 54 GB. For moderately sparse volumes, where 30% of space is occupied, the memory requirement is low enough to fit into either the CPU or GPU memory of consumer-grade machines. In practice, the choice k and l depends on the scene size and the expected image resolution. For many scenarios a smaller cache of $k = 512, l = 256$ is sufficient, lowering the memory requirements further. Please see the supplementary materials for formulas used to calculate cache sizes for both network architectures.

4. Implementation

Aside from using the proposed FastNeRF architecture described in Section 3.2, training FastNeRF is identical to training NeRF [25]. We model \mathcal{F}_{pos} and $\mathcal{F}_{\text{view}}$ of FastNeRF using 8 and 4 layer MLPs respectively, with positional encoding applied to inputs [25]. The 8-layer MLP has 384 hidden units, and we use 128 for the view-predicting 4-layer network. The only exception to this is the *figus* scene, where we found a 256 unit MLP gave better results than a 384 unit MLP. While this makes the network larger than the baseline, cache lookup speed is not influenced by MLP complexity. Similarly to NeRF, we parameterize the view direction as a 3-dimensional unit vector. Other training-




Figure 3. Speed evaluation of our method and prior work [25, 35, 16, 28] on the Lego scene from the Realistic 360 Synthetic [25] dataset, rendered at 800×800 pixels. For previous methods, when numbers for the Lego scene were not available, we used an optimistic approximation.

time features and settings including the coarse/fine networks and sample count N follow the original NeRF work [25] (please see supplementary materials).

At test time, both our method and NeRF take a set of camera parameters as input which are used to generate a ray for each pixel in the output. A number of samples are then produced along each ray and integrated following Section 3.1. While FastNeRF can be executed using its neural network representation, its performance is greatly improved when cached. Once the cache is generated, we convert it to a sparse data structure using OpenVDB [26], which improves performance and reduces the memory footprint. To further improve performance, we use hardware accelerated ray tracing [32] to skip empty space, starting to integrate points along the ray only after the first hit with a collision mesh derived from the density volume, and visiting every voxel thereafter until a ray’s transmittance is saturated. The collision mesh can be computed from a sign-distance function derived from the density volume using Marching Cubes [20]. We use the same meshing and integration parameters for all scenes and datasets, please see the supplementary for details.

5. Experiments

We evaluate our method quantitatively and qualitatively on the Realistic 360 Synthetic [25] and Local Light Field Fusion (LLFF) [24] (with additions from [25]) datasets used in the original NeRF paper. While the NeRF synthetic dataset consists of 360 degree views of complex objects, the LLFF dataset consist of forward-facing scenes, with fewer images. In all comparisons with NeRF we use the same training parameters as described in the original paper [25].

To assess the rendering quality of FastNeRF we compare its outputs to GT using Peak Signal to Noise Ratio (PSNR), Structural Similarity (SSIM) [43] and perceptual

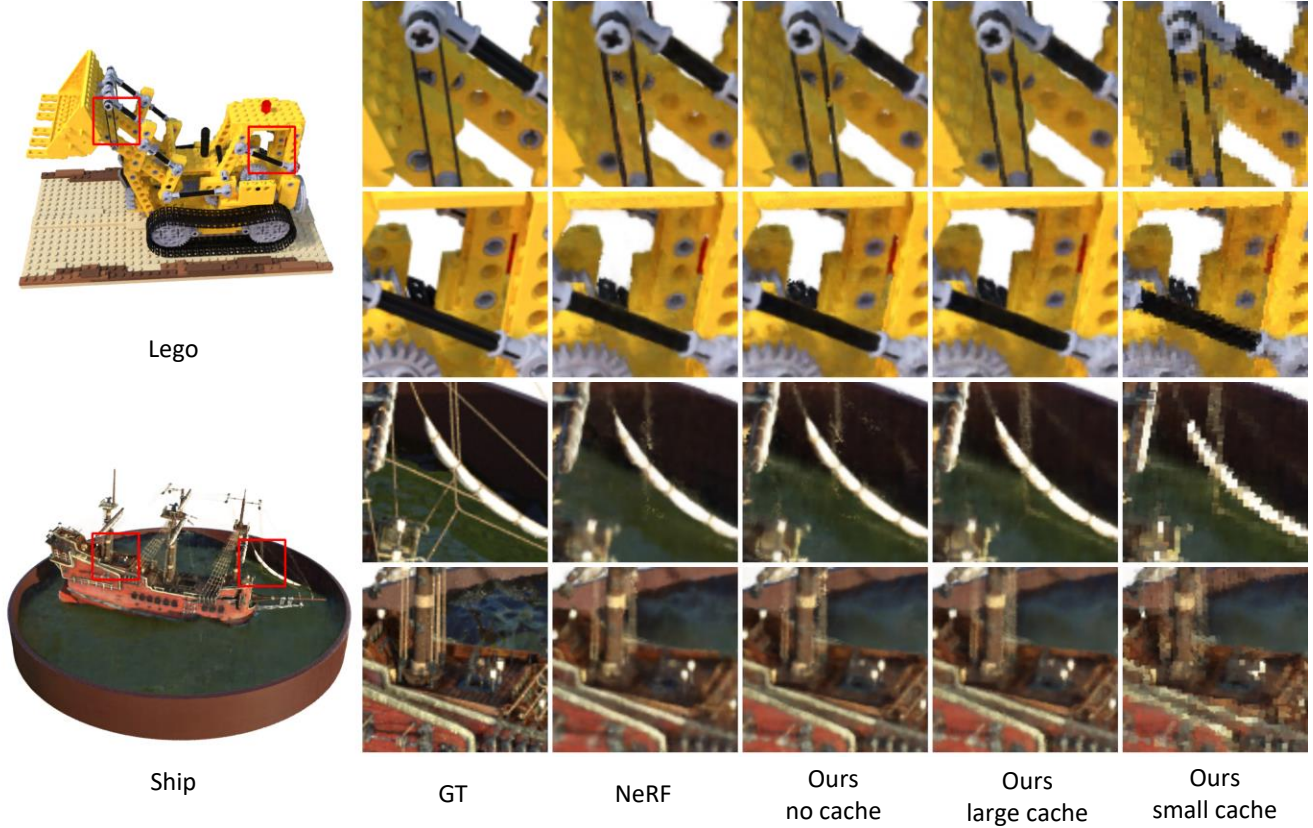


Figure 4. Qualitative comparison of our method vs NeRF on the dataset of [25] at 800^2 pixels using 8 components. *Small cache* refers to our method cached at 256^3 , and *large cache* at 768^3 . Varying the cache size allows for trading compute and memory for image quality resembling levels of detail (LOD) in traditional computer graphics.

Table 1. We compare NeRF to our method when not cached to a grid, and when cached at a high resolution in terms of PSNR, SSIM and LPIPS, and also provide the average speed in ms of our method when cached. LLFF denotes the dataset of [24] at 504×378 pixels, the other dataset is that of [25] at 800^2 pixels. We use 8 components and a 1024^3 cache for the NeRF Synthetic 360 scenes, and 6 components at 768^3 for LLFF. Please see the supplementary material for more detailed results.

Scene	NeRF			Ours - No Cache			Ours - Cache			Speed
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	
NeRF Synthetic	29.54dB	0.94	0.05	29.155dB	0.936	0.053	29.97dB	0.941	0.053	4.2ms
LLFF	27.72dB	0.88	0.07	27.958dB	0.888	0.063	26.035dB	0.856	0.085	1.4ms

LPIPS [48]. We use the same metrics in our ablation study, where we evaluate the effects of changing the number of components D and the size of the cache. All speed comparisons are run on one machine that is equipped with an Nvidia RTX 3090 GPU.

We keep the size of the view-dependent cache $\mathcal{F}_{\text{view}}^{\text{cache}}$ the same for all experiments at a base resolution of 384^3 . It is the base resolution of the RGBD cache $\mathcal{F}_{\text{pos}}^{\text{cache}}$ that represents the main trade-off in complexity of runtime and memory vs image quality. Across all results, grid resolution refers to the number of divisions of the longest side of the bounding volume surrounding the object of interest.

Rendering quality: Because we adopt the inputs, out-

puts, and training of NeRF, we retain the compatibility with ray-tracing [44, 9] and the ability to only specify loose volume bounds. At the same time, our factorization and caching do not affect the character of the rendered images to a large degree. Please see Figure 4 and Figure 5 for qualitative comparisons, and Table 1 for a quantitative evaluation. Note that some aliasing artefacts can appear since our model is cached to a grid *after* training. This explains the decrease across all metrics as a function of the grid resolution as seen in Table 1. However, both NeRF and FastNeRF mitigate multi-view inconsistencies, ghosting and other such artifacts from multi-view reconstruction. Table 1 demonstrates that at high enough cache resolutions, our method is capable

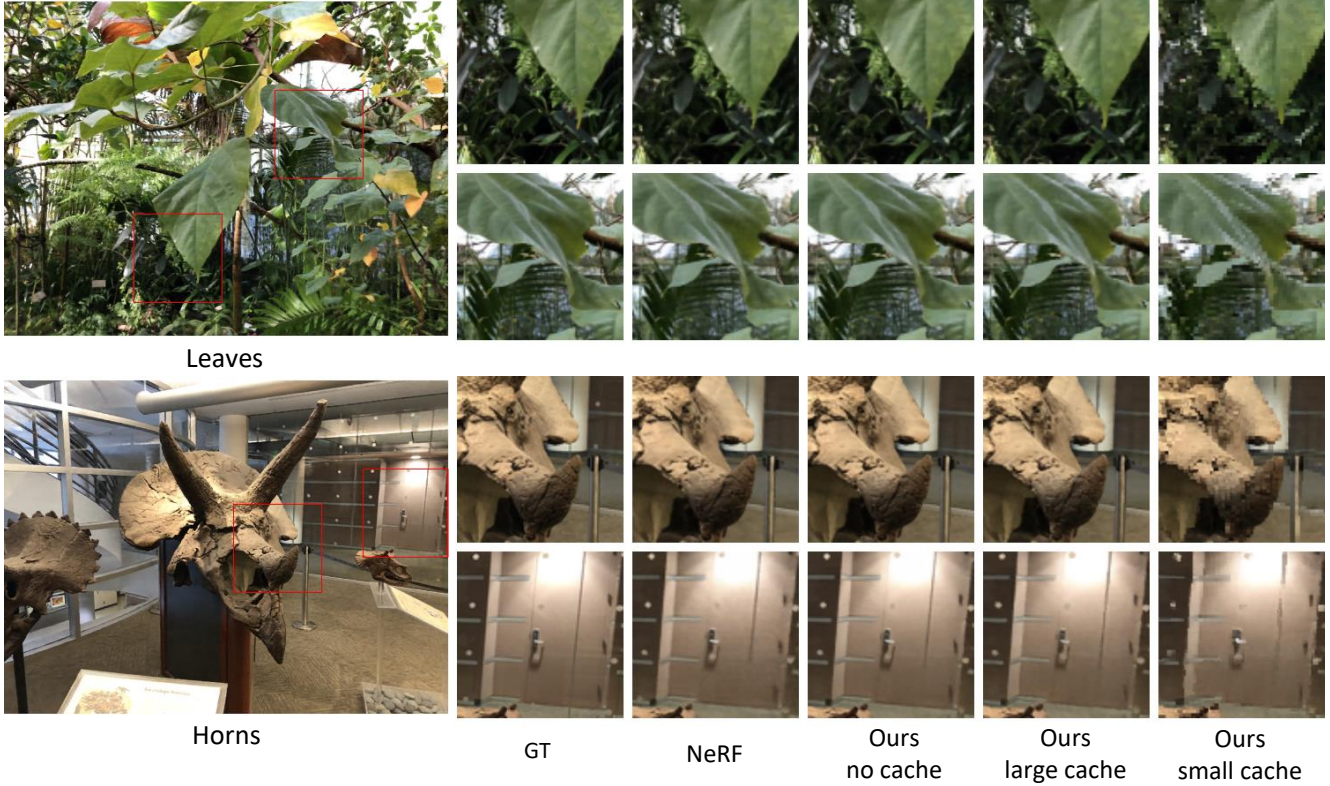


Figure 5. Qualitative comparison of our method vs NeRF on the dataset of [24] at 504×378 pixels using 6 factors. *Small cache* refers to our method cached at 256^3 , and *large cache* at 768^3 .

Table 2. Speed comparison of our method vs NeRF in ms. The *Chair* and *Lego* scenes are at rendered at 800^2 resolution, *Horns* and *Leaves* scenes at 504×378 . Our method never drops below 100FPS when cached, and is often significantly faster. Please note that our method is slower when not cached due to using larger MLPs; it performs identically to NeRF when using 256 hidden units. We do not compute the highest resolution cache for LLFF scenes because these are less sparse due to the use of normalized device coordinates.

Scene	NeRF	Ours - No Cache	256^3	384^3	512^3	768^3	1024^3	Speedup over NeRF
Chair	17.5K	28.2K	0.8	1.1	1.4	2.0	2.7	$6468\times - 21828\times$
Lego	17.5K	28.2K	1.5	2.1	2.8	4.2	5.6	$3118\times - 11639\times$
Horns*	3.8K	6.2K	0.5	0.7	0.9	1.2	-	$3183\times - 7640\times$
Leaves*	3.9K	6.3K	0.6	0.8	1.0	1.5	-	$2626\times - 6566\times$

of the same visual quality as NeRF. Figure 4 and Figure 5 further show that smaller caches appear ‘pixelated’ but retain the overall visual characteristics of the scene. This is similar to how different levels of detail are used in computer graphics [21], and an important innovation on the path towards neurally rendered worlds.

Challenging scenarios: We theorize that highly reflective surfaces could be a challenge for our method, as β might have to change with a high frequency depending on the view direction. In Section 1.1 of suppl. materials we discuss situations where the view-dependent network overfits to the training data and describe a simple solution.

Cache Resolution: As shown in Table 1, our method matches or outperforms NeRF on the dataset of synthetic objects at 1024^3 cache resolution. Because our method is fast enough to visit every voxel (as opposed to the fixed sample count used in NeRF), it can sometimes achieve better scores by not missing any detail. We observe that a cache of 512^3 is a good trade-off between perceptual quality, memory and rendering speed for the synthetic dataset. For the LLFF dataset, we found a 768^3 cache to work best. For the *Ship* scene, an ablation over the grid size, memory, and PSNR is shown in Table 3. For the view-filling LLFF scenes at 504×378 pixels, our method sees a slight

Table 3. Influence of the number of components and grid resolution on PSNR and memory required for caching the ship scene. Note how more factors increase grid sparsity. We find that 8 or 6 components are a reasonable compromise in practice.

Factors	No Cache		256 ³		384 ³		512 ³		768 ³	
	PSNR↑	Memory	PSNR↑	Memory	PSNR↑	Memory	PSNR↑	Memory	PSNR↑	Memory
4	27.11dB	-	24.81dB	0.34GB	26.29dB	0.61GB	26.94dB	1.09GB	27.54dB	2.51GB
6	27.12dB	-	24.82dB	0.5GB	26.34dB	0.93GB	27.0dB	1.67GB	27.58dB	4.1GB
8	27.24dB	-	24.89dB	0.71GB	26.42dB	1.41GB	27.1dB	2.7GB	27.72dB	7.15GB
16	27.68dB	-	25.07dB	1.2GB	26.77dB	2.08GB	27.55dB	3.72GB	28.3dB	9.16GB



Figure 6. Face images rendered using FastNeRF combined with a deformation field network [30]. Thanks to the use of FastNeRF, expression-conditioned images can be rendered at 30FPS.

decrease in metrics when cached at 768³, but still produces qualitatively compelling results as shown in Figure 5, where intricate detail is clearly preserved.

Rendering Speed: When using a grid resolution of 768³, FastNeRF is on average more than 3000× faster than NeRF at largely the same perceptual quality. See Table 2 for a breakdown of run-times across several resolutions of the cache and Figure 3 for a comparison to other NeRF extensions in terms of speed. Note that we log the time it takes our method to fill an RGBA buffer of the image size with the final pixel values, whereas the baseline implementation needs to perform various steps to reshape samples back into images. Our CUDA kernels are not highly optimized, favoring flexibility over maximum performance. It is reasonable to assume that advanced code optimization and further compression of the grid values could lead to further reductions in compute time.

Number of Components: While we can see a theoretical improvement of roughly 0.5dB going from 8 to 16 components, we find that 6 or 8 components are sufficient for most scenes. As shown in Table 3, our cache can bring out fine details that are lost when only a fixed sample count is used, compensating for the difference. Table 3 also shows that more components tend to be increasingly sparse when cached, which compensates somewhat for the additional memory usage. Please see the supplementary material for more detailed results on other scenes.

6. Application

We demonstrate a proof-of-concept application of FastNeRF for a telepresence scenario by first gathering a dataset

of a person performing facial expressions for about 20s in a multi-camera rig consisting of 32 calibrated cameras. We then fit a 3D face model similar to FLAME [13] to obtain expression parameters for each frame. Since the captured scene is not static, we train a FastNeRF model of the scene jointly with a deformation model [30] conditioned on the expression data. The deformation model takes the samples p as input and outputs their updated positions that live in a canonical frame of reference modeled by FastNeRF.

We show example outputs produced using this approach in Figure 6. This method allows us to render 300×300 pixel images of the face at 30 fps on a single Nvidia Tesla V100 GPU – around 50 times faster than a setup with a NeRF model. At the resolution we achieve, the face expression is clearly visible and the high frame rate allows for real-time rendering, opening the door to telepresence scenarios. The main limitation in terms of speed and image resolution is the deformation model, which needs to be executed for a large number of samples. We employ a simple pruning method, detailed in the supplementary material, to reduce the number of samples. Finally, note that while we train this proof-of-concept method on a dataset captured using multiple cameras, the approach can be extended to use only a single camera, similarly to [7].

Since FastNeRF uses the same set of inputs and outputs as those used in NeRF, it is applicable to many of NeRF’s extensions. This allows for accelerating existing approaches for: reconstruction of dynamic scenes [30, 34, 14], single-image reconstruction [39], quality improvement [35, 47] and others [28]. With minor modifications, FastNeRF can be applied to even more methods allowing for control over illumination [2] and incorporation of uncertainty [22].

7. Conclusion

In this paper we presented FastNeRF, a novel extension to NeRF that enables the rendering of photorealistic images at 200Hz and more on consumer hardware. We achieve significant speedups over NeRF and competing methods by factorizing its function approximator, enabling a caching step that makes rendering memory-bound instead of compute-bound. This allows for the application of NeRF to real-time scenarios.

References

- [1] Mojtaba Berman, Karol Myszkowski, Hans-Peter Seidel, and Tobias Ritschel. X-fields: Implicit neural view-, light- and time-image interpolation. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia 2020)*, 39(6), 2020. [1](#)
- [2] Sai Bi, Zexiang Xu, Pratul Srinivasan, Ben Mildenhall, Kalyan Sunkavalli, Miloš Hašan, Yannick Hold-Geoffroy, David Kriegman, and Ravi Ramamoorthi. Neural reflectance fields for appearance acquisition, 2020. [1](#), [8](#)
- [3] Sai Bi, Zexiang Xu, Kalyan Sunkavalli, Miloš Hašan, Yannick Hold-Geoffroy, David Kriegman, and Ravi Ramamoorthi. Deep reflectance volumes: Relightable reconstructions from multi-view photometric images, 2020. [1](#), [2](#)
- [4] Brian Cabral, Nelson Max, and Rebecca Springmeyer. Bidirectional reflection functions from surface bump maps. *21(4):273–281*, Aug. 1987. [4](#)
- [5] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. *CoRR*, abs/1812.02822, 2018. [1](#)
- [6] Boyang Deng, Jonathan T. Barron, and Pratul P. Srinivasan. JaxNeRF: an efficient JAX implementation of NeRF, 2020. [2](#)
- [7] Guy Gafni, Justus Thies, Michael Zollhöfer, and Matthias Nießner. Dynamic neural radiance fields for monocular 4d facial avatar reconstruction. *arXiv preprint arXiv:2012.03065*, 2020. [1](#), [8](#)
- [8] Iliyan Georgiev, Jaroslav Křivánek, Tomáš Davidovič, and Philipp Slusallek. Light transport simulation with vertex connection and merging. *ACM Trans. Graph.*, 31(6):192:1–192:10, Nov. 2012. [4](#)
- [9] James T. Kajiya. The rendering equation. *SIGGRAPH Comput. Graph.*, 20(4):143–150, Aug. 1986. [4](#), [6](#)
- [10] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks, 2019. [1](#)
- [11] Maria Kolos, Artem Sevastopolsky, and Victor Lempitsky. Transpr: Transparency ray-accumulating neural 3d scene point renderer, 2020. [3](#)
- [12] Christoph Lassner and Michael Zollhöfer. Pulsar: Efficient sphere-based neural rendering, 2020. [3](#)
- [13] Tianye Li, Timo Bolkart, Michael J Black, Hao Li, and Javier Romero. Learning a model of facial shape and expression from 4d scans. *ACM Trans. Graph.*, 36(6):194–1, 2017. [8](#)
- [14] Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. *arXiv preprint arXiv:2011.13084*, 2020. [8](#)
- [15] David B Lindell, Julien NP Martel, and Gordon Wetzstein. Autoint: Automatic integration for fast neural volume rendering. *arXiv preprint arXiv:2012.01714*, 2020. [1](#), [3](#)
- [16] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields, 2021. [1](#), [3](#), [5](#)
- [17] Stephen Lombardi, Jason M. Saragih, Tomas Simon, and Yaser Sheikh. Deep appearance models for face rendering. *CoRR*, abs/1808.00362, 2018. [1](#)
- [18] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: Learning dynamic renderable volumes from images. *ACM Trans. Graph.*, 38(4):65:1–65:14, July 2019. [1](#), [2](#)
- [19] Stephen Lombardi, Tomas Simon, Gabriel Schwartz, Michael Zollhoefer, Yaser Sheikh, and Jason Saragih. Mixture of volumetric primitives for efficient neural rendering, 2021. [1](#), [3](#)
- [20] William E. Lorensen and Harvey E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. *SIGGRAPH Comput. Graph.*, 21(4):163–169, Aug. 1987. [5](#)
- [21] David Luebke, Martin Reddy, Jonathan D. Cohen, Amitabh Varshney, Benjamin Watson, and Robert Huebner. *Level of Detail for 3D Graphics*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2002. [7](#)
- [22] 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, 2021. [1](#), [8](#)
- [23] Lars M. Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. *CoRR*, abs/1812.03828, 2018. [1](#)
- [24] 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 Trans. Graph.*, 38(4), July 2019. [1](#), [5](#), [6](#), [7](#)
- [25] 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. *arXiv preprint arXiv:2003.08934*, 2020. [1](#), [2](#), [3](#), [5](#), [6](#)
- [26] Ken Museth. Vdb: High-resolution sparse volumes with dynamic topology. *ACM Trans. Graph.*, 32(3), July 2013. [5](#)
- [27] Oliver Nalbach, Elena Arabadzhiyska, Dushyant Mehta, H-P Seidel, and Tobias Ritschel. Deep shading: convolutional neural networks for screen space shading. In *Computer graphics forum*, volume 36, pages 65–78. Wiley Online Library, 2017. [1](#)
- [28] Thomas Neff, Pascal Stadlbauer, Mathias Parger, Andreas Kurz, Chakravarty R. Alla Chaitanya, Anton Kaplanyan, and Markus Steinberger. Donerf: Towards real-time rendering of neural radiance fields using depth oracle networks, 2021. [1](#), [3](#), [5](#), [8](#)
- [29] Jeong Joon Park, Peter Florence, Julian Straub, Richard A. Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. *CoRR*, abs/1901.05103, 2019. [1](#)
- [30] Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo Martin Brualla. Deformable neural radiance fields. *arXiv preprint arXiv:2011.12948*, 2020. [8](#)
- [31] Keunhong Park, Utkarsh Sinha, Jonathan T. Barron, Sofien Bouaziz, Dan B Goldman, Steven M. Seitz, and Ricardo Martin-Brualla. Deformable neural radiance fields. *arXiv preprint arXiv:2011.12948*, 2020. [1](#)
- [32] Steven G Parker, James Bigler, Andreas Dietrich, Heiko Friedrich, Jared Hoberock, David Luebke, David McAllister, Morgan McGuire, Keith Morley, Austin Robison, et al. Optix: a general purpose ray tracing engine. *Acm transactions on graphics (tog)*, 29(4):1–13, 2010. [5](#)

- [33] M. Pharr, W. Jakob, and G. Humphreys. *Physically Based Rendering: From Theory to Implementation*. Elsevier Science, 2016. 4
- [34] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. *arXiv preprint arXiv:2011.13961*, 2020. 8
- [35] Daniel Rebain, Wei Jiang, Soroosh Yazdani, Ke Li, Kwang Moo Yi, and Andrea Tagliasacchi. Derf: Decomposed radiance fields. *arXiv preprint arXiv:2011.12490*, 2020. 1, 3, 5, 8
- [36] Gernot Riegler and Vladlen Koltun. Stable view synthesis, 2020. 1
- [37] 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, 2020. 1
- [38] Towaki Takikawa, Joey Litalien, Kangxue Yin, Karsten Kreis, Charles Loop, Derek Nowrouzezahrai, Alec Jacobson, Morgan McGuire, and Sanja Fidler. Neural geometric level of detail: Real-time rendering with implicit 3d shapes, 2021. 3
- [39] Matthew Tancik, Ben Mildenhall, Terrance Wang, Divi Schmidt, Pratul P. Srinivasan, Jonathan T. Barron, and Ren Ng. Learned initializations for optimizing coordinate-based neural representations, 2020. 3, 8
- [40] 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, 2020. 1, 2
- [41] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. 2
- [42] Eric Veach and Leonidas J. Guibas. Optimally combining sampling techniques for monte carlo rendering. In *Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '95, page 419–428, New York, NY, USA, 1995. Association for Computing Machinery. 4
- [43] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 5
- [44] Turner Whitted. An improved illumination model for shaded display. In *ACM Siggraph 2005 Courses*, pages 4–es. 2005. 6
- [45] Lifan Wu, Guangyan Cai, Shuang Zhao, and Ravi Ramamoorthi. Analytic spherical harmonic gradients for real-time rendering with many polygonal area lights. *ACM Trans. Graph.*, 39(4), July 2020. 4
- [46] Wenqi Xian, Jia-Bin Huang, Johannes Kopf, and Changil Kim. Space-time neural irradiance fields for free-viewpoint video. *arXiv preprint arXiv:2011.12950*, 2020. 1
- [47] Kai Zhang, Gernot Riegler, Noah Snively, and Vladlen Koltun. Nerf++: Analyzing and improving neural radiance fields. *arXiv preprint arXiv:2010.07492*, 2020. 8
- [48] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. *CoRR*, abs/1801.03924, 2018. 5
- [49] Ellen D. Zhong, Tristan Bepler, Joseph H. Davis, and Bonnie Berger. Reconstructing continuous distributions of 3d protein structure from cryo-em images, 2020. 2