

NerfAcc: Efficient Sampling Accelerates NeRFs

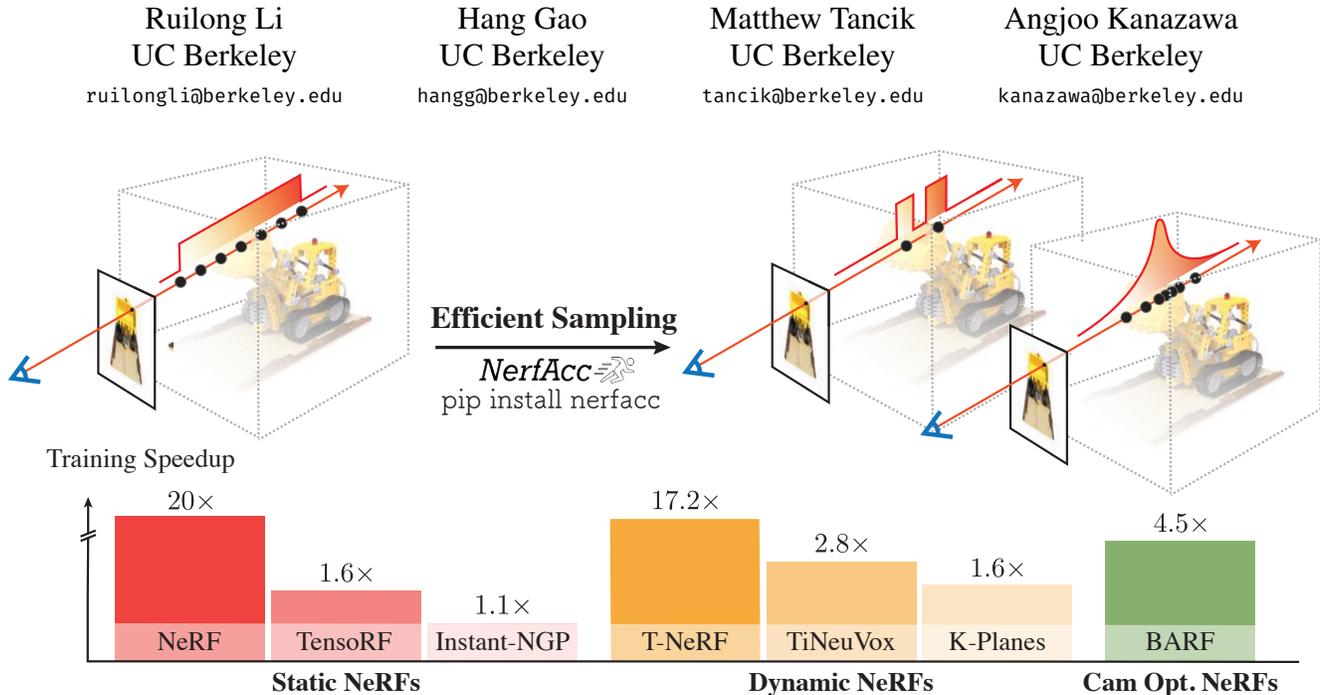


Figure 1: **NerfAcc Toolbox**. Our proposed toolbox, *NerfAcc*, integrates advanced efficient sampling techniques that lead to significant speedups in training various recent NeRF papers with minimal modifications to existing codebases.

Abstract

Optimizing and rendering Neural Radiance Fields is computationally expensive due to the vast number of samples required by volume rendering. Recent works have included alternative sampling approaches to help accelerate their methods, however, they are often not the focus of the work. In this paper, we investigate and compare multiple sampling approaches and demonstrate that improved sampling is generally applicable across NeRF variants under a unified concept of transmittance estimator. To facilitate future experiments, we develop NerfAcc, a Python toolbox that provides flexible APIs for incorporating advanced sampling methods into NeRF related methods. We demonstrate its flexibility by showing that it can reduce the training time of several recent NeRF methods by 1.5× to 20× with minimal modifications to the existing codebase. Additionally, highly customized NeRFs, such as Instant-NGP, can be implemented in native PyTorch using NerfAcc. Our code are open-sourced at <https://www.nerfacc.com>.

1. Introduction

Neural volume rendering has revolutionized the inverse rendering problem, with the Neural Radiance Field (NeRF) [16] being a key innovation. The continuous radiance field representation allows for rendering novel views of a scene from any camera position. However, the optimization of a NeRF can be computationally expensive due to the neural representation of radiance field and the large number of samples required by volume rendering. These challenges have limited practical applications of NeRF-based optimization and rendering.

Several recent works have successfully reduced the computational cost of neural volume rendering by proposing more efficient radiance field representations [17, 35, 36, 4, 24, 7]. While there are differences in the specific radiance field representations and their applications, most of these methods share a similar volume rendering pipeline which involves creating samples along the ray and accumulating them through alpha-composition.

However, compared to the considerable efforts focused on developing efficient radiance field representations, there has

been limited attention devoted to reducing the computational cost of neural volume rendering through *efficient sampling*. While a few recent works have included alternative sampling approaches to accelerate their methods [17, 35, 24, 2], these methods are often not the main focus of the paper. Moreover, implementing advanced sampling approaches typically requires non-trivial efforts. For example, Instant-NGP [17] and Plenoxels [35] both employ highly customized CUDA implementations to achieve spatial skipping during ray marching, which are tightly coupled with their respective radiance field implementations. Consequently, it can be challenging for researchers to benefit from these advanced sampling approaches in their own research.

In this paper, we investigate and compare several advanced sampling approaches from the literature and provide mathematical proofs demonstrating that they can all be viewed in a unified way of creating an *estimation of transmittance* for importance sampling. Our analysis shows that by understanding the spectrum of sampling through transmittance estimator, novel sampling strategies can be explored. Based on this, we decouple the sampling procedure from the neural volumetric rendering pipeline and demonstrate that improved sampling is generally applicable across different variants of NeRF. Furthermore, to facilitate future experiments, we propose NerfAcc, a plug-and-play toolbox that provides a flexible Python API for integrating advanced sampling approaches into NeRF-related methods, ready for researchers to incorporate into their own codebases. We demonstrate that with less than 100 lines of code change using NerfAcc, various NeRF methods [16, 4, 5, 7, 14, 21, 17] can enjoy $1.5\times$ to $20\times$ training speedup with better performance. Notably, using the NerfAcc library, one can train an Instant-NGP [17] model with pure Python code in the same amount of time, and achieve slightly better performance (+0.2dB) than the official pure CUDA implementation.

Our paper presents a unique contribution to the community. Unlike other papers that introduce novel algorithms, our work sheds light on the intricacies of various sampling approaches, which are often overlooked despite their significance. As far as we know, this work is the first paper that dives deep into this crucial aspect in the context of neural radiance field, and offer a novel, unified concept that allows researchers to view existing sampling algorithms in a fresh perspective. In addition, we provide a toolbox that facilitates research and development in this area. We believe translation of mathematical ideas into efficient, easy-to-modify implementation is fundamental to the research development. Overall, we hope that our proposed concept, along with the toolbox, can inspire new research ideas and accelerate progress in this field.

2. Related Works

NeRF Codebases. The recent explosion of NeRF-related research has led to numerous papers, many of which have released their own codebases [1, 2, 4, 7, 13, 16, 19, 37, 9]. These codebases address various tasks related to NeRF, including surface reconstruction [18, 32, 34], radiance field representation [1, 2, 17, 24, 35], dynamic modeling [5, 7, 13, 21], and camera optimization [14, 33]. However, each codebase is tailored to a specific task and supports only a single approach. While most of these methods share the same volume rendering pipeline of accumulating samples along the ray, transferring the implementation from one codebase to another requires non-trivial efforts. In this work, we address this common sampling problem by introducing NerfAcc, a plug-and-play toolbox that supports all the aforementioned tasks and methods, making it easy to integrate into any existing codebase.

NeRF Frameworks. Recently, several projects have been introduced to integrate different NeRF variants into a single framework, such as NeRF-Factory [11], Nerfstudio [29], and Kaolin-Wisp [28]. These frameworks have made significant progress in facilitating future NeRF-related research. NeRF-Factory offers a collection of NeRF variants [16, 39, 24, 35, 1, 2, 30] with original implementations and focuses on comprehensive benchmarking. Nerfstudio consolidates critical techniques introduced in existing literature [16, 1, 2, 17, 15] and provides modular components for the community to easily build on. Kaolin-Wisp builds upon the Kaolin [8] framework and implements a set of voxel-based NeRF papers [27, 17, 26]. However, these frameworks are designed to encourage researchers to develop within the framework, and do not benefit users working on their own codebases. Moreover, due to the high activity in the NeRF-related research, it is almost impossible to keep track of the latest developments and integrate advanced techniques into a single framework. Therefore, NerfAcc is designed as a standalone library that can be plugged into any codebase. It supports a wide range of NeRF related methods and can be easily maintained as new methods emerge.

3. Importance Sampling via Transmittance

Several advanced sampling approaches exist in the literature. For instance, Plenoxels [35] uses a sparse grid, Instant-NGP [17] uses an occupancy grid, NeRF [16] employs a coarse-to-fine strategy, Mip-NeRF 360 [2] proposes proposal networks. However, creating a toolbox that can support all of these approaches is not an easy task since they operate in completely different ways. In this section, we demonstrate that these approaches can all be viewed in a unified way of importance sampling. We also provide a mathematical proof that *transmittance is all you need for importance sampling*. Thus, each method essentially has its own way of creating an

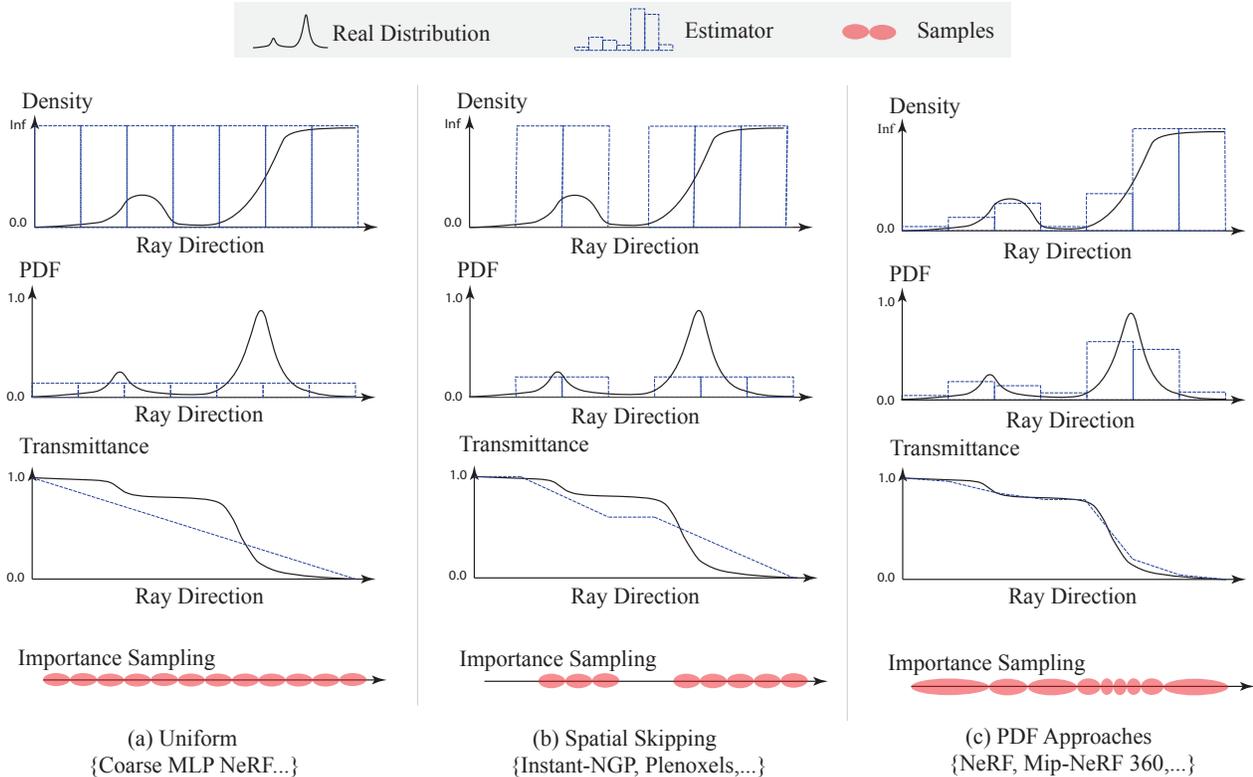


Figure 2: **Illustration of Sampling via Transmittance Estimator.** Although spatial skipping approaches (e.g., Occupancy Grid in Instant-NGP [17]) and PDF approaches (e.g., Proposal Network in Mip-NeRF 360 [2]) appear distinct from each other, they can both be viewed as constructing a transmittance estimator from which samples can be created via importance sampling.

estimation of the transmittance along the ray, which we refer to as the *transmittance estimator*. This observation enables us to unify different types of sampling approaches under the same framework and leads to our NerfAcc toolbox.

3.1. Formulation

Efficient sampling is a well-explored problem in Graphics [6], wherein the emphasis is on identifying regions that make the most significant contribution to the final rendering. This objective is generally accomplished through importance sampling, which aims to distribute samples based on the probability density function (PDF), denoted as $p(t)$, between the range of $[t_n, t_f]$. By computing the cumulative distribution function (CDF) through integration, i.e., $F(t) = \int_{t_n}^t p(v) dv$, samples are generated using the inverse transform sampling method:

$$t = F^{-1}(u) \quad \text{where} \quad u \sim \mathcal{U}[0, 1]. \quad (1)$$

In volumetric rendering, the contribution of each sample to the final rendering is expressed by the accumulation

weights $T(t)\sigma(t)$:

$$C(\mathbf{r}) = \int_{t_n}^{t_f} T(t) \sigma(t) c(t) dt \quad (2)$$

$$\text{where} \quad T(t) = \exp\left(-\int_{t_n}^t \sigma(s) ds\right).$$

Hence, the PDF for volumetric rendering is $p(t) = T(t)\sigma(t)$ and the CDF $F(t) = \int_{t_n}^t p(v) dv$ can be derived as a function of transmittance $T(t)$:

$$\begin{aligned} F(t) &= \int_{t_n}^t T(v)\sigma(v) dv \\ &= \int_{t_n}^t \frac{d}{dv} \left[-\exp\left(-\int_{t_n}^v \sigma(s) ds\right) \right] dv \\ &= 1 - \exp\left(-\int_{t_n}^t \sigma(s) ds\right) \\ &= 1 - T(t). \end{aligned} \quad (3)$$

Therefore, inverse sampling the CDF $F(t)$ is equivalent to inverse sampling the transmittance $T(t)$. Thus a transmittance estimator is sufficient to determine the optimal samples. Intuitively, this suggests to put more samples around regions

Sampling Methods	PDF Estimator $p(t_i)$	Transmittance Estimator $T(t_i)$	Update \mathcal{F}	Instantiation
Uniform	Constant	$1 - \frac{t-t_n}{t_f-t_n}$	-	Coarse MLP NeRF [16]
Spatial Skipping	$\frac{\mathbb{1}[\sigma(t_i) > \tau]}{\sum_{j=1}^n \mathbb{1}[\sigma(t_j) > \tau]}$	$1 - \frac{\sum_{j=1}^{i-1} \mathbb{1}[\sigma(t_j) > \tau]}{\sum_{j=1}^n \mathbb{1}[\sigma(t_j) > \tau]}$	EMA SGD	Instant-NGP [17] Plenoxels [35]
PDF Approaches	$\sigma(t_i) \exp(-\sigma(t_i) dt)$	$\exp(-\sum_{j=1}^{i-1} \sigma(t_j) dt)$	SGD SGD	Fine MLP NeRF [16] Mip-NeRF 360 [2]

Table 1: **Mathematical Formulations of Different Sampling Approaches.** We outline the significant mathematical distinctions among each sampling approach, under the perspective of transmittance estimator. See Section 3.1 for notations.

where the transmittance changes rapidly – and that is exactly what happens when a ray hit a surface. Implementation wise, this observation enables us to compute the CDF directly using $1 - T(t)$, instead of the computationally expensive integral $\int_{t_n}^t T(v)\sigma(v) dv$, which is the standard implementation adopted by many popular codebases [16, 1, 2, 29, 7].

While advanced transmittance estimator techniques, such as delta tracking [31] with Monte Carlo sampling, are utilized in production-level volumetric rendering in Graphics [6], NeRFs operate in a distinct setting in which the scene geometry is not predefined but optimized on the fly. During NeRF optimization, the radiance field changes between iterations, necessitating the dynamic update of the transmittance estimator at each step k :

$$\mathcal{F} : T(t)^{k-1} \mapsto T(t)^k. \quad (4)$$

This introduces additional challenges to efficient sampling because it becomes more difficult to accurately estimate the transmittance from a radiance field that is constantly changing. Current approaches employ either exponential moving average (EMA) or stochastic gradient descent (SGD) as the update function \mathcal{F} . However, we note that there may be other update functions that could be explored.

With these concepts in mind, let us now examine some of the existing approaches towards efficient sampling.

Uniform. If a transmittance estimator is not available, the only assumption we can make is that every point along the ray contributes equally to the final rendering. Mathematically, this assumption translates to a constant PDF and a linearly decaying transmittance $T(t) = 1 - (t-t_n)/(t_f-t_n)$. In this case, the sampling process is equivalent to uniformly sampling along the ray, i.e., $t_i = t_n + (t_f - t_n) \cdot u_i$. It is worth noting that every NeRF model that uses uniform sampling inherently assumes this linear transmittance decay, such as the coarse level in vanilla NeRF [16]. See Fig. 2(a) for the illustration.

Spatial Skipping. A more sophisticated approach to improve uniform sampling is to identify empty regions and skip them during sampling, as proposed in Instant-NGP’s Occupancy Grid [17] and Plenotrees’ Sparse Grid [36]. This technique binarizes the density along the ray with a conservative threshold τ : $\hat{\sigma}(t_i) = \mathbb{1}[\sigma(t_i) > \tau]$. Consequently,

the piece-wise constant PDF can be expressed as $p(t_i) = \hat{\sigma}(t_i) / \sum_{j=1}^n \hat{\sigma}(t_j)$, and the piece-wise linear transmittance estimator is $T(t_i) = 1 - \sum_{j=1}^{i-1} \hat{\sigma}(t_j) / \sum_{j=1}^n \hat{\sigma}(t_j)$. To update this estimator during optimization, Instant-NGP [17] directly updates the cached density with exponential moving average (EMA) over iteration k : $\sigma(t_i)^k = \gamma \cdot \sigma(t_i)^{k-1} + (1 - \gamma) \cdot \sigma(t_i)^k$. Meanwhile, Plenoxels [35] updates the density via the gradient descent through the rendering loss. See Fig. 2(b) for an illustration.

PDF Approaches. Another type of approach is to directly estimate the PDF along the ray with discrete samples. In vanilla NeRF [16], the coarse MLP is trained using volumetric rendering loss to output a set of densities $\sigma(t_i)$. This allows for the creation of a piece-wise constant PDF: $p(t_i) = \sigma(t_i) \exp(-\sigma(t_i) dt)$, and a piece-wise linear transmittance estimator $T(t_i) = \exp(-\sum_{j=1}^{i-1} \sigma(t_j) dt)$. This approach was further improved in Mip-NeRF 360 [16] with a PDF matching loss, which allows for the use of a much smaller MLP in the coarse level, namely Proposal Network, to speedup the PDF construction. In both cases, the transmittance estimator is updated through gradient descent. See Fig. 2(c) for an illustration.

Table 1 provides a mathematical summary and comparison of these approaches. Additionally, we present an illustration in Figure 2 to provide an intuitive comparison of these approaches from the perspective of PDF (second row) and transmittance (third row), as well as how the samples can be created from the transmittance estimator via importance sampling (last row). This visualization also reveals some pros and cons for each approach, which we will discuss in Section 3.3.

3.2. Design Spaces

Choice of Representations. The transmittance estimator can use either an explicit voxel [17, 35, 4], an MLP [1, 2, 16], or a hybrid representation [29]. Depending on whether the estimator is explicit or not, it can be updated with either rule-based EMA [17, 4] or gradient descent with some supervision [35, 2, 16, 29]. Generally, voxel-based estimators are faster than implicit (e.g., MLP-based) estimators but suffer more from aliasing issues. It is worth noting that the transmittance estimator’s representation can significantly benefit

from advances in radiance field representation. For example, the Nerfacto model [29] uses the most recent hybrid-representation HashEncoding [17] for both the radiance field and the sampling module, achieving the best quality-speed tradeoff in in-the-wild settings.

Handling Unbounded Scenes. So far, we have only discussed sampling within a bounded area $[t_n, t_f]$. For unbounded scenes, it is impossible to densely sample along the ray. Similar to the mipmaps used in graphics rendering, a general solution is to sample more coarsely as the ray goes further, as objects farther away appear in fewer pixels on the image plane. This can be achieved by creating a bijective mapping function $\Phi : s \in [s_n, s_f] \mapsto t \in [t_n, +\infty]$, and performing sampling in the s -space instead of the t -space. Several papers [2, 39, 23] that work on unbounded scenes have introduced different mapping functions Φ , to which we refer our readers for details.

3.3. Discussions

Pros and Cons. Sampling with uniform assumption is the easiest one to implement but with lowest efficiency in most cases. Spatial skipping is a more efficient technique since most of the 3D space is empty, but it still samples uniformly within occupied but occluded areas that contribute little to the final rendering (e.g., the last sample in Figure 2(b)). PDF-based estimators generally provide more accurate transmittance estimation, enabling samples to concentrate more on high-contribution areas (e.g., surfaces) and to be more spread out in both empty and occluded regions. However, this also means that samples are always spread out throughout the entire space without any skipping, as shown in Figure 2(c). Moreover, the current approaches all introduce aliasing effects to the volumetric rendering due to either (1) the piece-wise linearity assumption for estimating transmittance along the ray [1, 2, 16, 35, 17] as illustrated in Figure 2, or (2) the underlying voxel representation for the transmittance estimator [35, 17] as discussed in Section 3.2. A recent work, Zip-NeRF [3], addresses the aliasing problem tied to these two exact issues (called “z-aliasing” and “xy-aliasing” in their work), which are naturally revealed under our unified framework.

Implementation Difficulties. The current implementations for efficient sampling are all highly customized and tightly integrated with the specific radiance field proposed in each paper. For instance, spatial skipping is implemented with customized CUDA kernels in Instant-NGP [17] and Plenoxels [35]. Mip-NeRF 360 [2], K-planes [7], and Nerfacto [29] have a proposal network implemented but it is closely integrated with their repositories and can only support limited types of radiance fields that come with the repository. However, as shown before, the sampling process is independent of the radiance field representation, thus it should

Algorithm 1: NerfAcc Rendering Pipeline. In NerfAcc, the samples are created from a transmittance estimator, which can be updated during the NeRF training. See Section 4.2 for details.

```
# nerf: a radiance field model.
# r_o: ray origins. (n_rays, 3)
# r_d: ray normalized directions. (n_rays, 3)

# e.g., Prop Net, Occ Grid.
estimator = nerfacc.TransmittanceEstimator()

def density_fn(t0, t1, r_id):
    """Query density from nerf."""
    return nerf.density(r_o[r_id], r_d[r_id], t0, t1)

def rgb_density_fn(t0, t1, r_id):
    """Query rgb and density from nerf."""
    return nerf.forward(r_o[r_id], r_d[r_id], t0, t1)

# Efficient sampling.
# (t0, t1, r_id): packed samples. (all_samples,)
t0, t1, r_id = nerfacc.sampling(
    r_o, r_d, estimator, density_fn=density_fn
)

# Differentiable volumetric rendering.
color, opacity, depth, aux = nerfacc.rendering(
    t0, t1, r_id, rgb_density_fn=rgb_density_fn
)

# Update the transmittance estimator.
estimator.update_every_n_steps(t0, t1, r_id, aux)
# nerf, r_o and r_d all receive gradients.
F.mse_loss(color, color_gt).backward()
```

be easily transferable across different NeRF variants. Due to the various implementation details, it typically requires significant effort to correctly implement an efficient sampling approach from scratch. Therefore, having an implementation that is easily transferable from repository to repository would be valuable in supporting future research on NeRF.

Insights from Unified Formulation. Comprehending the sampling spectrum through the lens of Transmittance Estimator paves the way for investigating novel sampling strategies. For example, our framework reveals that the Occupancy Grid from Instant-NGP [17] and the Proposal Network from Mip-NeRF 360 [2] are not mutually exclusive but complementary, as both aim to estimate the transmittance along the ray. Therefore, combining them becomes straight-forward: one can first compute the transmittance with the occupancy grid, and then refine the estimated transmittance with a proposal network. This enables both skipping on the empty space and concentrating the samples onto the surface. We explore this approach in Section 4.4 and demonstrate that it overcomes the limitation of the proposal network approach, which always samples the entire space. Furthermore, this formulation could potentially shed light on questions such as how to enhance the sampling procedure with depth information or other priors, which we encourage readers to investigate further.

4. NerfAcc Toolbox

In this paper, we present *NerfAcc* toolbox, designed for Neural radiance field Acceleration. It provides efficient sampling for volumetric rendering, that is universally applicable and easily integrable for a diverse range of radiance fields [16, 4, 17, 14, 32]. In this section, we first introduce the design principles of this toolbox, along with critical implementation details. To demonstrate its flexibility, we further show that it can significantly speedup the training of various NeRF-related papers by $1.5\times$ to $20\times$ with only minor modifications to existing codebases.

4.1. Design Principles

This library is designed with these goals in mind:

- **Plug-and-play.** Our primary objective is to ease the challenges of implementing an efficient volumetric sampling technique for the research community. Therefore, NerfAcc is designed as a standalone library that can be easily installed from PyPI on both Windows and Linux platforms, and seamlessly integrated into any PyTorch codebase.
- **Efficiency & Flexibility.** To maximize the speed of the code, we fuse the operations into CUDA kernels as much as possible, while exposing flexible Python APIs to the users.
- **Radiance Field Complexity.** We target on supporting any radiance fields that are designed for per-scene optimization, including both density-based and SDF-based fields, for both static and dynamic scenes.

4.2. Implementation Details

NerfAcc incorporates two advanced sampling methods that can be decoupled from the radiance field representation, namely the Occupancy Grid from Instant-NGP [17] and the Proposal Network from Mip-NeRF 360 [2]. Algorithm 1 presents the pseudo code for volumetric rendering using the NerfAcc toolbox. In this section, we do not dive into the details of each algorithm, as we basically follow the original paper’s implementation. Instead, we discuss the implementation designs that are crucial to maintain high efficiency and flexibility of this toolbox.

Sample as Interval. In the NerfAcc toolbox, instead of representing each sample with a coordinate x , we represent the sample as an interval along the ray (t_0, t_1, r) , where t_0 and t_1 are the start and end of the interval along the r -th ray. This interval-based representation offers three key advantages. Firstly, representing a sample as an interval instead of a single point allows us to support research based on cone-based rays for anti-aliasing, such as Mip-NeRF [1] and Mip-NeRF 360 [2]. Secondly, since in almost all cases t_i does not require gradients, using (t_0, t_1, r) instead of $(\mathbf{x}_0, \mathbf{x}_1)$ to represent the interval allows for the detachment of the sam-

Methods	Dataset	Speedup	T (min) ↓	PSNR ↑	LPIPS ↓
TensoRF [4]	T&T	$1.5\times$	18.3	28.13	0.143
+ nerfacc (occ)			12.6	28.10	0.150
TensoRF [4]	Syn.	$1.6\times$	10.6	32.52	0.047
+ nerfacc (occ)			6.5	32.51	0.044
NeRF [16]	Syn.	$20\times$	>1000	31.00	0.081
+ nerfacc (occ) [†]			52.0	31.55	0.072
Instant-NGP [17]	Syn.	$1.0\times$	4.4	32.35	-
+ nerfacc (occ) [†]			4.4	32.55	0.056
+ nerfacc (prop) [†]			5.2	31.40	0.064
Instant-NGP [17]	360	$1.1\times$	5.3	25.93	-
+ nerfacc (occ) [†]			5.0	26.41	0.353
+ nerfacc (prop) [†]			4.9	27.58	0.292

(a) Static NeRFs.

Methods	Dataset	Speedup	T (min) ↓	PSNR ↑	LPIPS ↓
TiNeuVox [20]	Hyper.	$1.7\times$	56.3	24.19	0.425
+ nerfacc (occ)			33.0	24.19	0.434
+ nerfacc (prop)			34.3	24.26	0.398
TiNeuVox [5]	D-NeRF	$2.8\times$	11.8	31.14	0.050
+ nerfacc (occ)			4.2	31.75	0.038
K-Planes [7]	D-NeRF	$1.6\times$	63.9	30.28	0.043
+ nerfacc (occ)			38.8	30.35	0.042
T-NeRF [21]	D-NeRF	$20\times$	>1000	28.78	0.069
+ nerfacc (occ) [†]			58.0	32.22	0.040

(b) Dynamic NeRFs.

Methods	Dataset	Speedup	T (min) ↓	PSNR ↑	LPIPS ↓	E_R / E_T
BARF [20]	Syn.	$4.5\times$	586	28.83	0.054	0.19 / 0.74
+ nerfacc (occ)			130	30.11	0.044	0.07 / 0.35

(c) NeRFs for Camera Optimization. Rotation / translation errors are denoted as E_R/E_T . E_T is scaled by 100.

Table 2: Improvements on Various NeRFs. Experiments are conducted by replacing necessary code in the official repositories, except for those marked by [†] that are based on our re-implementations.

pling process from the differentiable computational graph, thereby maximizing its speed. Lastly, the ray id r attached to each sample enables support for various numbers of samples across rays with a packed tensor, which we will discuss in the next paragraph. A similar representation has been adopted in Nerfstudio [29] to support various radiance fields.

Packed Tensor. To support sampling with spatial skipping, it is necessary to consider that each ray may result in a different number of valid samples. Storing the data as a tensor with shape $(n_rays, n_samples, \dots)$ and an extra mask with shape $(n_rays, n_samples)$ to indicate which samples are valid leads to significant inefficient memory usage when large portions of space are empty. To address this, in NerfAcc, we represent samples as “packed tensors” with shape $(all_samples, \dots)$, in which only valid samples are stored (see Algo. 1). To keep track of the associated rays for each sample, we also host an integer tensor with shape

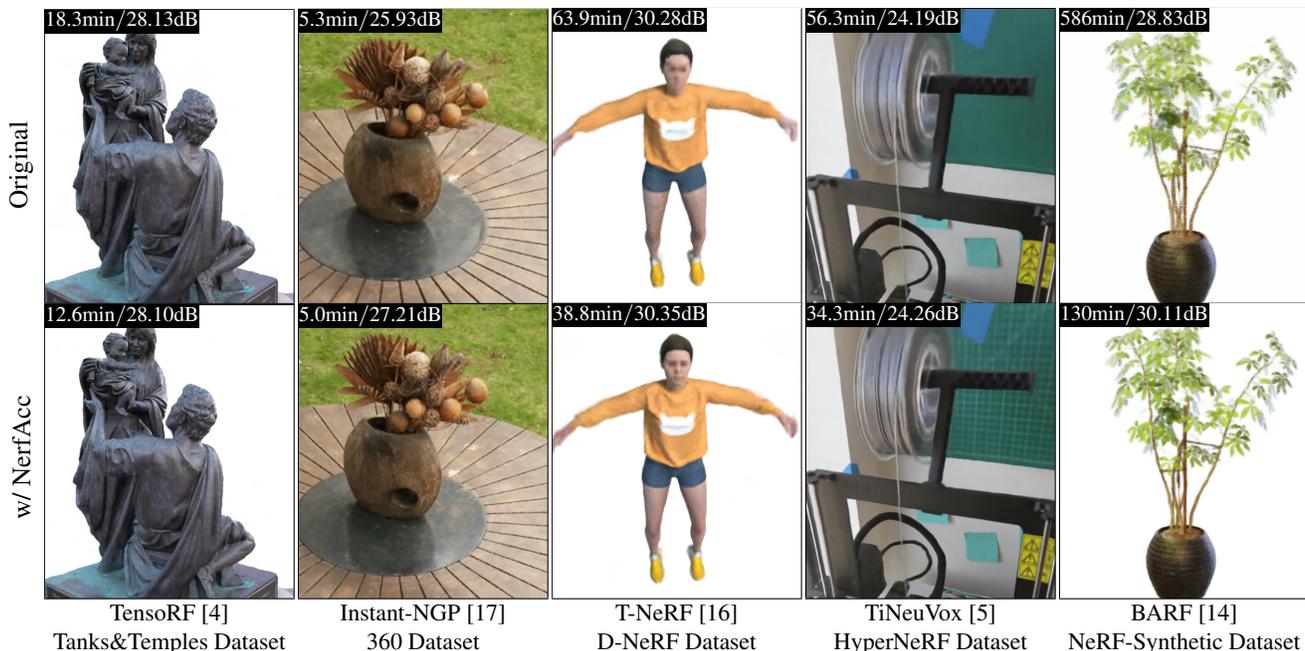


Figure 3: **Qualitative Results.** NerfAcc is able to significantly reduce the training time of various NeRF-related methods across multiple datasets, while also yielding slightly improved quality in the majority of cases. The training time and test PSNR are shown on the left corner of each image.

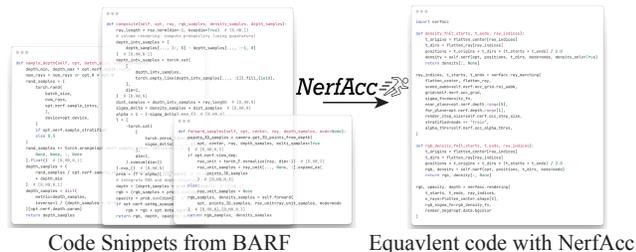


Figure 4: **Plug-and-play Example in BARF [14]’s Repository.** With around 50 lines of code change, both training speed and performance can be greatly improved with efficient sampling equipped in NerfAcc, as shown in Table 2b.

($n_rays, 2$), which stores the start index in the packed tensors and the number of samples on this ray. This approach is similar to that used in Instant-NGP [17] and PyTorch3D [22].

No Gradient Filtering. After importance sampling, inaccurate transmittance estimations can result in some samples lying in empty or occluded spaces, particularly in spatial skipping methods like Occupancy Grid. These samples can be filtered before being included in PyTorch’s differentiable computation graph by evaluating their transmittance using the radiance field with gradients disabled. As backward passes are not required during filtering, this is much faster ($\sim 10\times$) than keeping all samples in the computation graph. In practice, samples with transmittance below 10^{-4} are dis-

regarded during this process with almost no influence on the rendering quality. Note that this strategy is inspired from Instant-NGP [17]’s implementation.

4.3. Case Studies

We showcase the flexibility of NerfAcc on three types of NeRFs across seven papers: static NeRFs (NeRF [16], TensorRF [4], Instant-NGP [17]); dynamic NeRFs (D-NeRF [21], K-Planes [7] TiNeuVox [5]); and a NeRF variation for camera optimization (BARF [14]). Although many of these methods, *e.g.*, Instant-NGP, TensorRF, TiNeuVox and K-Planes, have already been highly optimized for efficiency, we are still able to accelerate their training by a large margin and achieve slightly better performance on nearly all cases. It is worth mentioning that experiments with TensorRF, TiNeuVox, K-Planes and BARF are conducted by integrating NerfAcc into the *official codebase* with around 100 lines of code change. The results of our experiments, including those of our baselines, are presented in Table 2a,2b, and2c, all of which were conducted under the same physical environment using a single NVIDIA RTX A5000 GPU to facilitate comparison, as per [25]. Aside from the experiments reported in this paper, NerfAcc has also been integrated into a few popular open-source projects such as nerfstudio [29] for density-based NeRFs, as well as sdfstudio [38] and instant-nsr-pl [10] for SDF-based NeRFs.

Static NeRFs. In this task, we experiment with three NeRF variants, including the original MLP-based NeRF [16],

TensoRF [4] and Instant-NGP [17]. We show that NerfAcc works with both MLP-based and Voxel-based radiance fields, on both bounded (NeRF-Synthetic dataset [16], Tank&Template dataset [12]) and unbounded scenes (360 Dataset [2]). It is worth to note that with NerfAcc, one can train an Instant-NGP model with pure Python code and achieve slightly better performance than the official pure CUDA implementation, as shown in Table 2a.

Dynamic NeRFs. In this task, we apply the NerfAcc toolbox to T-NeRF [21], K-Planes [7] and TiNeuVox [5], covering both the synthetic (D-NeRF [21]) and “in-the-wild” captures¹ (that accompany HyperNeRF [20]). When applying the occupancy grid approach to accelerate those dynamic methods, instead of representing a static scene with it, we share the occupancy grid across all frames. In other words, instead of using it to indicate the opacity of an area at a single timestamp, We use it to indicate the *maximum opacity at this area over all the timestamps*. This is not optimal but still makes the rendering very efficient, due to the fact there are limited movements in these datasets.

NeRFs for Camera Optimization. In this task, we employed the NerfAcc toolbox to BARF [14] on the NeRF-Synthetic dataset with perturbed cameras. The goal is to jointly optimize the radiance field and camera extrinsic for multi-view images. We observed that the spatial skipping sampling provided by NerfAcc facilitated faster training and significantly improved both image quality and camera pose reconstruction. These improvements could be attributed to the sparsity enforced in our sampling procedure. This finding may provide interesting avenues for future research.

Analysis of Different Sampling Approaches. Results in Table 2a show that the choice between occupancy grid and proposal network sampling can noticeably affect run-time and performance on different datasets. As each approach relies on a distinct set of hyperparameters, a systematic comparison between the two methods is crucial by sweeping the hyperparameter space. We varied the resolution and marching step size for occupancy grid and the number of samples and size of the proposal network for the proposal network approach. We plot pareto curves for each approach for both the NeRF-Synthetic and Mip-NeRF 360 datasets in Fig.5. This analysis indicates that occupancy grid sampling is suitable for the NeRF-Synthetic dataset, whereas the proposal network approach performs better on the 360 dataset. This is likely because the NeRF-Synthetic dataset contains more empty space that can be skipped effectively using the occupancy grid approach. However, in the case of real, unbounded data, the use of the occupancy grid approach is limited by the bounding box and the lack of empty space to skip, making the proposal network approach more

¹These datasets teleport cameras and do not represent real captures [9].

effective. These experiments used the radiance field from Instant-NGP [17], with the same training recipes.

4.4. Combined Sampling

A benefit from the unified concept of transmittance estimator introduced in Section 3, is that it’s straight-forward to combine the two distinct sampling approaches, as both of them essentially provide an estimation of the transmittance that can be used for importance sampling. For example, we find that simply stacking an occupancy grid on top of the proposal network, can significantly reduce the number of rays and shrink the near-far plane for the remaining rays on the NeRF-Synthetic dataset. This leads to improvements in quality, from 31.40dB to 32.35dB, and a reduction in training time, from 5.2min to 4.3min, compared to using only the proposal network for importance sampling. Figure 6 shows an example with the FICUS scene, where the floaters are cleaned out with the combined sampling. This experiment is conducted using the HashEncoding from Instant-NGP [17] as the radiance field representation.

5. Conclusions

In conclusion, this paper highlights the significant impact of advanced sampling approaches on improving the efficiency of Neural Radiance Fields (NeRF) optimization and rendering. We demonstrate that advanced sampling can significantly speed up the training of various recent NeRF papers, while maintaining high-quality results. The development of NerfAcc, a flexible Python toolbox, enables researchers to incorporate advanced sampling methods into NeRF-related methods easily. The exploration and comparison of advanced sampling methods are important steps towards developing more efficient and accessible NeRF-based methods. The presented results also demonstrate the potential for further research to improve the performance of NeRF and other related techniques through advanced sampling strategies.

References

- [1] Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5855–5864, 2021.
- [2] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5470–5479, 2022.
- [3] Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Zip-nerf: Anti-aliased grid-based neural radiance fields. *arXiv*, 2023.

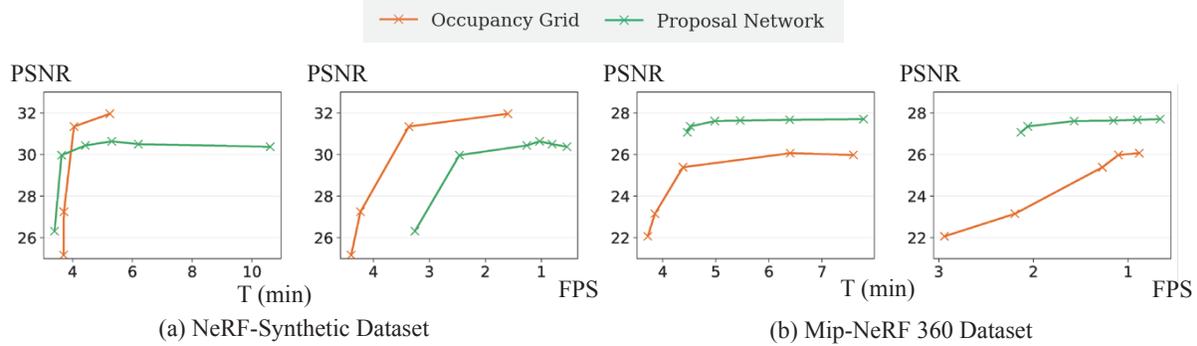


Figure 5: **Comparison between Different Sampling Approaches in terms of Training Time and Rendering FPS.** We sweep the hyper-parameter space for each sampling approach, and find out that occupancy grid approach gives the best performance-speed trade-off on the NeRF-Synthetic dataset, while the proposal network approach performs the best on the Mip-NeRF 360 dataset. Note T (min) denotes for *training* time and FPS is for *rendering* frames per second. All experiments use the HashEncoding from Instant-NGP [17] as the radiance field representation. Please see the supplementary materials for the hyperparameter space that we explored.

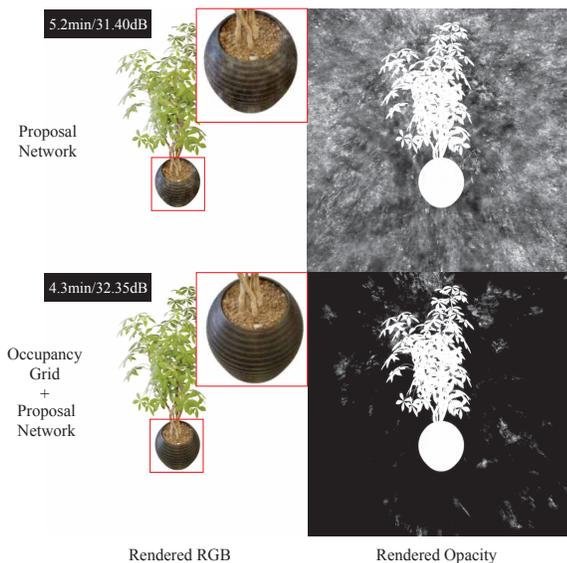


Figure 6: **Results of Combined Sampling on NeRF-Synthetic Dataset.** Benefit from our unified concept of transmittance estimator, we find it’s straight-forward to combine the two distinct sampling approaches.

[4] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. *arXiv preprint arXiv:2203.09517*, 2022.

[5] Jiemin Fang, Taoran Yi, Xinggang Wang, Lingxi Xie, Xiopeng Zhang, Wenyu Liu, Matthias Nießner, and Qi Tian. Fast dynamic radiance fields with time-aware neural voxels. In *SIGGRAPH Asia 2022 Conference Papers*, 2022.

[6] Julian Fong, Magnus Wrenninge, Christopher Kulla, and Ralf Habel. Production volume rendering: Siggraph 2017 course. In *ACM SIGGRAPH 2017 Courses*, pages 1–79. 2017.

[7] Sara Fridovich-Keil, Giacomo Meanti, Frederik Warburg,

Benjamin Recht, and Angjoo Kanazawa. K-planes: Explicit radiance fields in space, time, and appearance. *arXiv preprint arXiv:2301.10241*, 2023.

[8] Clement Fuji Tsang, Maria Shugrina, Jean Francois Lafleche, Towaki Takikawa, Jiehan Wang, Charles Loop, Wenzheng Chen, Krishna Murthy Jatavallabhula, Edward Smith, Artem Rozantsev, Or Perel, Tianchang Shen, Jun Gao, Sanja Fidler, Gavriel State, Jason Gorski, Tommy Xiang, Jianing Li, Michael Li, and Rev Lebaredian. Kaolin: A pytorch library for accelerating 3d deep learning research. <https://github.com/NVIDIAGameWorks/kaolin>, 2022.

[9] Hang Gao, Ruilong Li, Shubham Tulsiani, Bryan Russell, and Angjoo Kanazawa. Monocular dynamic view synthesis: A reality check. In *Advances in Neural Information Processing Systems*, 2022.

[10] Yuan-Chen Guo. Instant neural surface reconstruction, 2022.

[11] Yoonwoo Jeong, Seungjoo Shin, and Kibaek Park. Nerf-factory: An awesome pytorch nerf collection, 2022.

[12] Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics*, 36(4), 2017.

[13] Ruilong Li, Julian Tanke, Minh Vo, Michael Zollhofer, Jurgen Gall, Angjoo Kanazawa, and Christoph Lassner. Tava: Template-free animatable volumetric actors. *arXiv preprint arXiv:2206.08929*, 2022.

[14] Chen-Hsuan Lin, Wei-Chiu Ma, Antonio Torralba, and Simon Lucey. Barf: Bundle-adjusting neural radiance fields. In *IEEE International Conference on Computer Vision (ICCV)*, 2021.

[15] Ricardo Martin-Brualla, Noha Radwan, Mehdi SM Sajjadi, Jonathan T Barron, Alexey Dosovitskiy, and Daniel Duckworth. Nerf in the wild: Neural radiance fields for unconstrained photo collections. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7210–7219, 2021.

[16] 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. *Communications of the ACM*, 65(1):99–106, 2021.
- [17] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *arXiv preprint arXiv:2201.05989*, 2022.
- [18] Michael Oechsle, Songyou Peng, and Andreas Geiger. Unisurf: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5589–5599, 2021.
- [19] Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5865–5874, 2021.
- [20] Keunhong Park, Utkarsh Sinha, Peter Hedman, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Ricardo Martin-Brualla, and Steven M Seitz. Hypernerf: A higher-dimensional representation for topologically varying neural radiance fields. *arXiv preprint arXiv:2106.13228*, 2021.
- [21] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10318–10327, 2021.
- [22] Nikhila Ravi, Jeremy Reizenstein, David Novotny, Taylor Gordon, Wan-Yen Lo, Justin Johnson, and Georgia Gkioxari. Accelerating 3d deep learning with pytorch3d. *arXiv preprint arXiv:2007.08501*, 2020.
- [23] Christian Reiser, Richard Szeliski, Dor Verbin, Pratul P Srinivasan, Ben Mildenhall, Andreas Geiger, Jonathan T Barron, and Peter Hedman. Merf: Memory-efficient radiance fields for real-time view synthesis in unbounded scenes. *arXiv preprint arXiv:2302.12249*, 2023.
- [24] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5459–5469, 2022.
- [25] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Improved direct voxel grid optimization for radiance fields reconstruction. *arXiv preprint arXiv:2206.05085*, 2022.
- [26] Towaki Takikawa, Alex Evans, Jonathan Tremblay, Thomas Müller, Morgan McGuire, Alec Jacobson, and Sanja Fidler. Variable bitrate neural fields. In *ACM SIGGRAPH 2022 Conference Proceedings*, pages 1–9, 2022.
- [27] 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. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11358–11367, 2021.
- [28] Towaki Takikawa, Or Perel, Clement Fuji Tsang, Charles Loop, Joey Litalien, Jonathan Tremblay, Sanja Fidler, and Maria Shugrina. Kaolin wisp: A pytorch library and engine for neural fields research. <https://github.com/NVIDIAGameWorks/kaolin-wisp>, 2022.
- [29] Matthew Tancik, Ethan Weber, Evonne Ng, Ruilong Li, Brent Yi, Justin Kerr, Terrance Wang, Alexander Kristoffersen, Jake Austin, Kamyar Salahi, et al. Nerfstudio: A modular framework for neural radiance field development. *arXiv preprint arXiv:2302.04264*, 2023.
- [30] 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 *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5481–5490. IEEE, 2022.
- [31] John Von Neumann. 13. various techniques used in connection with random digits. *Appl. Math Ser.*, 12(36-38):3, 1951.
- [32] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *arXiv preprint arXiv:2106.10689*, 2021.
- [33] Zirui Wang, Shangzhe Wu, Weidi Xie, Min Chen, and Victor Adrian Prisacariu. Nerf-: Neural radiance fields without known camera parameters. *arXiv preprint arXiv:2102.07064*, 2021.
- [34] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. *Advances in Neural Information Processing Systems*, 34:4805–4815, 2021.
- [35] Alex Yu, Sara Fridovich-Keil, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. *arXiv preprint arXiv:2112.05131*, 2021.
- [36] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. Plenotrees for real-time rendering of neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5752–5761, 2021.
- [37] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from one or few images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4578–4587, 2021.
- [38] Zehao Yu, Anpei Chen, Bozidar Antic, Songyou Peng, Apratim Bhattacharyya, Michael Niemeyer, Siyu Tang, Torsten Sattler, and Andreas Geiger. Sdfstudio: A unified framework for surface reconstruction, 2022.
- [39] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. Nerf++: Analyzing and improving neural radiance fields. *arXiv preprint arXiv:2010.07492*, 2020.