

---

# Differentiable Rendering with Reparameterized Volume Sampling

---

Anonymous Author(s)

Affiliation

Address

email

## Abstract

1 We propose an alternative rendering algorithm for neural radiance fields based  
2 on importance sampling. In view synthesis, a neural radiance field approximates  
3 underlying density and radiance fields based on a sparse set of scene views. To  
4 generate a pixel of a novel view, it marches a ray through the pixel and computes a  
5 weighted sum of radiance emitted from a dense set of ray points. This rendering  
6 algorithm is fully differentiable and facilitates gradient-based optimization of the  
7 fields. However, in practice, only a tiny opaque portion of the ray contributes most  
8 of the radiance to the sum. Therefore, we can avoid computing radiance in the rest  
9 part. In this work, we use importance sampling to pick non-transparent points on  
10 the ray. Specifically, we generate samples according to the probability distribution  
11 induced by the density field. Our main contribution is the reparameterization of  
12 the sampling algorithm. It allows end-to-end learning with gradient descent as in  
13 the original rendering algorithm. With our approach, we can optimize a neural  
14 radiance field with just a few radiance field evaluations per ray. As a result, we  
15 alleviate the costs associated with the color component of the neural radiance field  
16 at the additional cost of the density sampling algorithm.

## 17 1 Introduction

18 We propose a volume rendering algorithm for learning 3D scenes and generating novel views.  
19 Recently, learning-based approaches led to significant progress in this area. As an early instance,  
20 [8] proposed to represent a scene via a density field and a radiance (color) field parameterized  
21 with an MLP. They run a differentiable volume rendering algorithm with the MLP-based fields and  
22 minimize the discrepancy between the produced images and a set of reference images to learn a  
23 scene representation. The algorithm we propose is a drop-in replacement for the volume rendering  
24 algorithm used in NeRF [8] and follow-ups.

25 The underlying model in NeRF generates an image point in the following way. It casts a ray from  
26 a camera through the point and defines the point color as a weighted sum along the ray. The sum  
27 aggregates the radiance of each ray point with weights induced by the density field. Each summand  
28 involves a costly neural network query, and model has a trade-off between rendering quality and  
29 computational load. NeRF obtained a better trade-off with a two-stage sampling algorithm used to get  
30 ray points with higher weights. The algorithm is reminiscent of importance sampling, yet it requires  
31 training an auxiliary model.

32 In this work we propose a rendering algorithm based on importance sampling. Our algorithm also  
33 acts in two stages. In the first stage, it marches through the ray to estimate density. In the second  
34 stage, it constructs a Monte-Carlo color approximation using the density to pick points along the ray.  
35 The resulting estimate is fully-differentiable and does not require any auxiliary models. Besides that,  
36 we only need a few samples to construct precise color approximation. An intuitive explanation is that

37 we only need to compute the radiance of the point where a ray hits a solid surface. In the experiments,  
 38 we query radiance for  $\times 16$  fewer ray points during training compared to baseline. Nevertheless, we  
 39 manage to obtain competitive model and rendering quality.

40 As a result, our algorithm is more suitable for recent solutions [10, 13, 12] that use distinct models to  
 41 parameterize radiance and density. Specifically, the first stage only queries the density field, whereas  
 42 the second stage only queries the radiance field. Compared to the standard rendering algorithm, the  
 43 second stage of our algorithm avoids redundant radiance queries and reduces the memory required  
 44 for rendering.

## 45 2 Neural Radiance Fields

46 Neural radiance fields represent 3D scenes with a scalar density field  $\sigma : \mathbb{R}^3 \rightarrow \mathbb{R}^+$  and a vector  
 47 radiance field  $c : \mathbb{R}^3 \times \mathbb{R}^3 \rightarrow \mathbb{R}^3$ . The scalar field  $\sigma$  represents volume density at each spatial  
 48 location  $\mathbf{x}$ , and  $c(\mathbf{x}, \mathbf{d})$  returns the light emitted from spatial location  $\mathbf{x}$  in direction  $\mathbf{d}$  represented as  
 49 a normalized three dimensional vector.

50 For novel view synthesis, they adapt a volume rendering technique that computes a pixel color  
 51  $C(\mathbf{o}, \mathbf{d})$  (denoted with a capital letter). In particular, the expected color along a ray  $\mathbf{r} = \mathbf{o} + t\mathbf{d}$  going  
 52 from the camera through the pixel is

$$C(\mathbf{o}, \mathbf{d}) = \int_{t_n}^{+\infty} p_r(t) c(\mathbf{o} + t\mathbf{d}, \mathbf{d}) dt, \text{ for } p_r(t) = \sigma(\mathbf{o} + t\mathbf{d}) \exp\left(-\int_{t_n}^t \sigma(\mathbf{o} + s\mathbf{d}) ds\right). \quad (1)$$

53 Here,  $p_r(t)$  is a probability density function of a random variable  $T$  on a ray  $r$ . Intuitively,  $T$  is the  
 54 location on the ray where a portion of light coming into the point  $\mathbf{o}$  was emitted.

55 One way to approximate to the integral would be to cut off the integral at depth  $t_f$  and then use a grid  
 56  $t_n = t_0 < t_1 < \dots < t_m = t_f$  to compute the integral with a Riemann sum

$$\hat{C}_{\text{Riemann}}(\mathbf{o}, \mathbf{d}) = \sum_{i=1}^m (t_i - t_{i-1}) p_{r,i} c(\mathbf{o} + t_i \mathbf{d}, \mathbf{d}), \quad (2)$$

$$\text{where } p_{r,i} = \sigma(\mathbf{o} + t_i \mathbf{d}) \exp\left(-\sum_{j=1}^i (t_j - t_{j-1}) \sigma(\mathbf{o} + t_j \mathbf{d})\right). \quad (3)$$

57 Importantly, Eq 2 is fully differentiable and can be used as a part of gradient-based learning pipeline.

58 While such approximation works in practice, a faithful approximation requires a dense grid and  
 59 multiple evaluations of  $\sigma$  and  $c$ . Besides that, a common situation is when a ray intersects a solid  
 60 surface at some point  $s \in [t_n, t_f]$ . In this case, probability density  $p_r(t)$  will concentrate its mass  
 61 near  $s$  and will be close to zero in other parts of the ray. As a result, most of the summands in Eq. 2  
 62 will make negligible contribution to the sum.

63 Monte Carlo methods give another way to approximate the color. Given  $n$  i.i.d. samples  $t_1, \dots, t_n \sim$   
 64  $p_r(t)$ , the color estimate is gathered by

$$\hat{C}_{MC}(\mathbf{o}, \mathbf{d}) = \frac{1}{m} \sum_{i=1}^m c(\mathbf{o} + t_i \mathbf{d}, \mathbf{d}). \quad (4)$$

65 Due to the importance sampling with distribution  $p_r(t)$ , each term in Eq 4 contributes equally to the  
 66 sum as the samples come from regions with non-negligible density. Unlike the grid estimate in Eq. 2,  
 67 the Monte-Carlo estimate depends on the scene density  $\sigma$  implicitly and requires a custom gradient  
 68 estimate for the parameters of  $\sigma$ . For instance, NeRF addresses the issue via a hierarchical sampling  
 69 scheme. It trains a coarse model with a grid approximation to generate importance-weighted ray  
 70 locations for a separate fine-grained model.

71 In the next section, we propose a novel principled approach to training neural radiance  
 72 fields with importance-weighted color approximation as in Eq. 4.

73 **3 Learning with Stochastic Color Estimates**

74 In this section, we will discuss stochastic approximations to the expected color  $C(\mathbf{o}, \mathbf{d})$  in detail.  
 75 Recall that  $C(\mathbf{o}, \mathbf{d}) = \mathbb{E}_T c(\mathbf{o} + T\mathbf{d}, \mathbf{d})$ , where  $T$  is a random variable with density specified in Eq. 1.  
 76 Even though density  $p_r(t)$  involves an integral we cannot compute in closed form, below we first  
 77 assume that we have an algorithm to compute  $\int_{t_n}^t \sigma_r(s) ds$  used in  $p_r(t)$ .

78 Given a groundtruth expected color  $C_{gt}$ , optimization objective in NeRF captures the difference  
 79  $L(\hat{C}(\mathbf{o}, \mathbf{d}), C_{gt})$  between  $C_{gt}$  and the estimated color  $\hat{C}(\mathbf{o}, \mathbf{d})$ . To reconstruct a scene NeRF runs  
 80 a gradient based optimizer to minimize the objective averaged across multiple rays and multiple  
 81 viewpoints. Such approach works for grid estimate  $\hat{C}(\mathbf{o}, \mathbf{d}) = \hat{C}_{Riemann}(\mathbf{o}, \mathbf{d})$  that depends on  
 82 density  $\sigma_r$  explicitly, but Monte-Carlo estimate  $\hat{C}_{MC}(\mathbf{o}, \mathbf{d})$  of the expectation depends on  $\sigma$  implicitly  
 83 and a naive automatic differentiation algorithm will return zero gradients.

84 In the rest of the section, we first introduce an algorithm to compute  $\hat{C}_{MC}(\mathbf{o}, \mathbf{d})$  and derive a gradient  
 85 estimate for the algorithm. Then, we conclude with a discussion our implementation of the estimate.  
 86 To ease the notation, we will also introduce  $\sigma_r(t) = \sigma(\mathbf{o} + t\mathbf{d})$  and  $c_r(t) = c(\mathbf{o} + t\mathbf{d}, \mathbf{d})$  to denote  
 87 fields restricted to a ray  $\mathbf{r} = \mathbf{o} + t\mathbf{d}$ .

88 **3.1 Estimate Reparameterization**

89 To make the dependence of  $\hat{C}(\mathbf{o}, \mathbf{d})$  on  $\sigma_r$  explicit, we change the variables in the expectation  
 90  $\mathbb{E}_T c_r(T)$ . For  $F(t) = 1 - \exp\left(-\int_{t_n}^t \sigma_r(s) ds\right)$  and  $y := F(t)$  we write

$$\mathbb{E}_T c_r(T) = \int_{t_n}^{+\infty} c_r(t) p_r(t) dt = \int_{y_n}^{y_f} c_r(F^{-1}(y)) dy. \quad (5)$$

91 Function  $F(t)$  acts as cumulative distribution function of the variable  $T$  with a single exception  
 92 that, in general,  $y_f = \lim_{t \rightarrow \infty} F(t) \neq 1$ . In volume rendering,  $F(t)$  is called the opacity function  
 93 with  $y_f$  being equal to pixel opaqueness. Bounds of integration are where  $y_n = F(t_n) = 0$  and  
 94  $y_f = \lim_{t \rightarrow +\infty} F(t)$ . For simplicity, below we replace  $y_f$  with  $F(t_f)$  where  $t_f$  is the maximum ray  
 95 depth.

96 In the right-hand side of Eq. 5, integration boundaries depend on the opacity  $F$  and, thus, on the  
 97 volume density  $\sigma_r$ . We further simplify the integral by changing the integration boundaries to  $[0, 1]$  :

$$\int_{y_n}^{y_f} c_r(F^{-1}(y)) dy = \int_0^1 (y_f - y_n) c_r(F^{-1}(y_n + (y_f - y_n)u)) du. \quad (6)$$

98 With this, we arrive to the following reparameterized Monte-Carlo estimate of the expected color  
 99 obtained with i.i.d  $U[0, 1]$  samples  $u_1, \dots, u_m$ :

$$\hat{C}_{MC}^R(\mathbf{o}, \mathbf{d}) := \frac{1}{m} \sum_{i=1}^m (y_f - y_n) c_r(F^{-1}(y_n + (y_f - y_n)u_i)). \quad (7)$$

100 In the above estimate sampling does not depend on volume density  $\sigma_r$  or color  $c_r$ . Essentially,  
 101 this is a reparameterized Monte-Carlo estimate that generates samples from  $p_r(t)$  using the inverse  
 102 cumulative distribution function  $F^{-1}(y_n + (y_f - y_n)u)$ .

103 We further improve the estimate using stratified sampling. To do this, we replace the uniform samples  
 104  $u_1, \dots, u_m$  with uniform independent samples within regular grid bins  $v_i \sim U[\frac{i-1}{m+1}, \frac{i}{m+1}]$ ,  $i =$   
 105  $1, \dots, m$  and derive a reparameterized (R) stratified (S) Monte Carlo estimate

$$\hat{C}_{SMC}^R(\mathbf{o}, \mathbf{d}) = \frac{1}{m} \sum_{i=1}^m (y_f - y_n) c_r(F^{-1}(y_n + (y_f - y_n)v_i)). \quad (8)$$

106 It is easy to show that both 7 and 8 are unbiased estimates of 1.

107 Next, we will discuss algorithms used to compute the inverse opacity function  $F^{-1}(y)$  and compute  
 108 the gradients of the function with automatic differentiation.

### 109 3.2 Implementation of Inverse Opacity for Volume Sampling

110 To compute the estimates in Eqs. eqs. (7) and (8), we need to compute the inverse opacity  $F^{-1}(y)$   
 111 along with its gradient. In practice, we start with a black-box density field  $\sigma_r(x)$  and compute  
 112 the induced density  $p_r(t)$  and opacity  $F(t)$  on a ray  $r$  via approximations. Assuming we have an  
 113 algorithm to compute  $\int_{t_n}^t \sigma_r(s)ds$ , below we show how to implement the inverse opacity  $F^{-1}$ .

114 We invert  $F(t) = 1 - \exp\left(-\int_{t_n}^t \sigma_r(s)ds\right)$  with binary search. Note that  $F(t)$  is a monotonic  
 115 function and for  $y \in (y_n, y_f) = (F(t_n), F(t_f))$  the inverse lies in  $(t_n, t_f)$ . To compute  $F^{-1}(y)$ , we  
 116 start with boundaries  $t_l = t_n$  and  $t_r = t_f$  and gradually decrease the gap between the boundaries  
 117 based on the comparison of  $F(\frac{t_l+t_r}{2})$  with  $y$ . Importantly, such procedure is easy to parallelize across  
 118 multiple inputs and multiple rays.

119 However, we cannot backpropagate through the binary search iterations and need a workaround to  
 120 compute the gradient  $\frac{\partial t}{\partial \theta}$  of  $t(\theta) = F^{-1}(y, \theta)$ . To do this, we compute differentials of the right and  
 121 the left hand side of equation  $y(\theta) = F(t, \theta)$

$$\frac{\partial y}{\partial \theta} d\theta = \frac{\partial F}{\partial t} \frac{\partial t}{\partial \theta} d\theta + \frac{\partial F}{\partial \theta} d\theta. \quad (9)$$

122 By the definition of  $F(t, \theta)$  we have

$$\frac{\partial F}{\partial t} = (1 - F(t, \theta))\sigma_r(t, \theta), \quad (10)$$

$$\frac{\partial F}{\partial \theta} = (1 - F(t, \theta)) \frac{\partial}{\partial \theta} \left( \int_{t_n}^t \sigma_r(s, \theta) ds \right). \quad (11)$$

123 We solve Eq. 9 for  $\frac{\partial t}{\partial \theta}$  and substitute the partial derivatives using Eqs. eqs. (10) and (11) to obtain the  
 124 final expression for the gradient

$$\frac{\partial t}{\partial \theta} = \frac{\frac{\partial y}{\partial \theta} - (1 - F(t, \theta)) \frac{\partial}{\partial \theta} \int_{t_n}^t \sigma_r(s, \theta) ds}{(1 - F(t, \theta))\sigma_r(t, \theta)}. \quad (12)$$

125 In our implementation, we use automatic differentiation to compute  $\partial y / \partial \theta$  and  $\frac{\partial}{\partial \theta} \int_{t_n}^t \sigma_r(s) ds$  so  
 126 combine the results as in Eq. 12.

### 127 3.3 Computing Opacity in Practice

128 To describe the sampling procedure, we assumed that we have an oracle for computing  $\int_{t_n}^t \sigma_r(s)ds$   
 129 along with its gradient. The integral is required to compute opacity  $F(t)$ . In this work, we consider an  
 130 arbitrary volumetric density  $\sigma(s)$  and approximate it with a linear spline on a ray  $r = o + td$  to sample  
 131 the points on the ray. Specifically, we take a grid  $t_0 < \dots < t_m$  and compute  $\sigma_r(t_0), \dots, \sigma_r(t_m)$   
 132 to construct the spline  $\hat{\sigma}_r(s)$  (Fig. 1). For the piecewise linear function  $\hat{\sigma}_r(x)$  we can compute the  
 133 integral  $\int_{t_n}^t \hat{\sigma}_r(s)ds$  in a closed form. Additionally, we can backpropagate the gradients through the  
 134 approximation to compute the gradients of knots  $\sigma_r(t_0), \dots, \sigma_r(t_m)$ . Thus, we obtain a differentiable  
 rendering algorithm for an arbitrary density field  $\sigma$ . Besides that, some recent works parameterize




Figure 1: Illustration of opacity inversion. We approximate an arbitrary density field  $\sigma$  with a linear spline(left). Then we use the spline to approximate opacity  $\hat{F}(t)$  and compute  $\hat{F}^{-1}(y)$  (right).

135 density fields through voxel grids. For a voxel grid, when  $\sigma_r$  is a trilinear interpolation of the grid  
 136 values, we can compute the integral in a closed form.  
 137

138 **4 Related Work**

139 **Neural Radiance Fields & Efficient Sampling** Even in the original work on neural radiance  
 140 fields [8] the authors aimed to find an efficient sampling algorithm for volume rendering. Our  
 141 importance sampling approach is reminiscent of their hierarchical sampling solution. On the first  
 142 stage, they use an auxiliary model on a sparse grid. Then they use the predicted densities to generate  
 143 a dense grid with an importance sampling-like algorithm. As opposed to NeRF, we compute density  
 144 on a dense grid at the first stage and then use a sparse set of samples to evaluate radiance on the  
 145 second stage. Our algorithm also allows training without auxiliary models.

146 Several recent follow-up works also aimed to improve NeRF rendering time and overall efficiency.  
 147 Most of these works consider trainable encoding  $\theta$  and utilize some efficient data structure to make  
 148 each evaluation of multi-layered perceptron fast or avoid evaluating MLP at all. One of the earliest  
 149 work in this direction was NSVF [7]. The authors proposed to use octree to store point-based  
 150 embeddings and then estimate query point embedding with a trilinear interpolation and positional  
 151 encoding. During training, the octree gradually increased resolution in the regions of interest  
 152 and pruned the empty areas. However, this method still requires the time-consuming training of  
 153 MLPs. Voxel-based embedding structure was further studied in recent works and it was shown that  
 154 positional encoding doesn't affect model convergence - the network can be trained with fully trainable  
 155 embedding without any encoding. And also, what is more important, such a structure allows for  
 156 making neural network (MLP) shallower and consequently faster. Following this idea, DirectVoxGo  
 157 [12] proposes to avoid MLP at all in density computation, while Instant Neural Graphic Primitives  
 158 [10] uses it to solve hash collisions. When density field is a piecewise linear we can compute opacity  
 159 in a closed-form.

160 **Reparameterization Trick & Implicit Differentiation** Our solution is inspired by the literature on  
 161 deep latent variable models [6, 11] and approximate inference. In this area, models often contain  
 162 an internal sampling algorithm with parameters we need to optimize. The now-common approach  
 163 for continuous random variables is the reparameterization trick, which we apply in our setup. The  
 164 authors of [9] give a comprehensive overview of the area state.

165 A closely related work in the context of deep variable models is [4]. They were first to apply implicit  
 166 differentiation to estimate gradients for the reparameterization trick. While we use the implicit  
 167 differentiation to compute the gradient of binary search output, the same approach applies to other  
 168 iterative algorithms. The examples include ODE solves [3], fixed-point iterators [1] and optimization  
 169 algorithms.

170 **5 Experiments**

171 **5.1 Importance Sampling for a Single Ray**

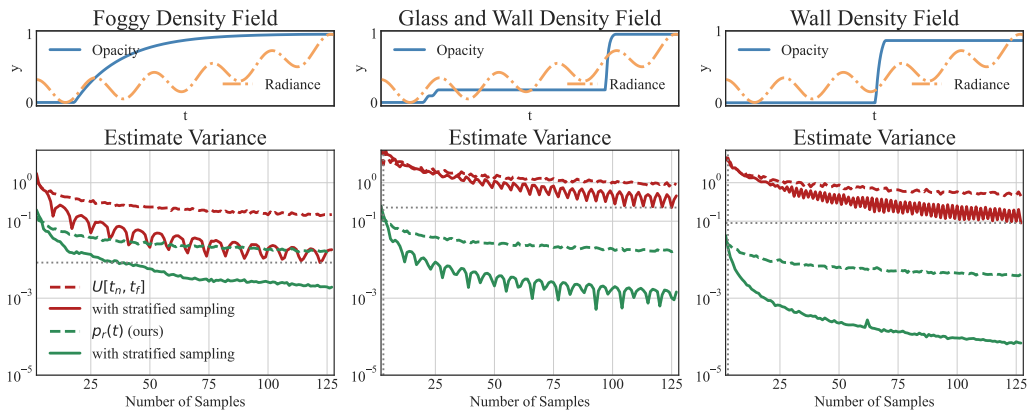


Figure 2: Color estimate variance compared for a varying number of samples. The upper plot illustrates underlying opacity function on a ray; the lower graph depicts variance in logarithmic scale. Our importance sampling approach (solid green) has significantly lower deviation than a stratified baseline (solid red) typically used in volume rendering.

172 We begin with an evaluation of color estimates in a one-dimensional setting. Our experiment models  
 173 light propagation on a single ray in three typical situations. The upper row of Fig. 2 defines a scalar  
 174 radiance field (orange)  $c(t)$  and opacity functions (blue)  $F(t)$  for

- 175 • "Foggy" density field. It models a semi-transparent volume. Similar fields occur after model  
 176 initialization during density field training;
- 177 • "Glass and wall" density field. Models light passing through nearly transparent volumes such  
 178 as glass. The light is emitted at three points: the inner and outer surface of the transparent  
 179 volume and an opaque volume near the end of the ray;
- 180 • "Wall" density field. Light is emitted from a single point on a ray. Such fields are most  
 181 common in applications.

182 For the three fields we estimated the expected radiance  $C = \int_{t_n}^{t_f} c(t)dF(t)$ . We considered two  
 183 baseline methods (both in red in Fig. 2): the first was a Monte Carlo estimate of  $C$  obtained with  
 184  $U[t_n, t_f]$  samples, the second was a stochastic modification of Eq. 2 using a grid  $t_n = t_0 < \dots <$   
 185  $t_m = t_f$ :

$$\hat{C} = \sum_{i=1}^m (t_i - t_{i-1})c(\tau_i) \left. \frac{dF}{dt} \right|_{t=\tau_i}, \text{ with independent } \tau_i \sim U[t_{i-1}, t_i]. \quad (13)$$

186 In other terms, the second baseline uses stratified sampling to reduce the baseline Monte Carlo  
 187 estimate variance. Eq 13 is an instance of a vanilla volume rendering algorithm one may encounter in  
 188 practice. We compared the baseline against estimate from Eq. eq. (7) and its stratified counterpart  
 189 from Eq. 8. All estimates are unbiased. Therefore, we only compared the estimates variances for a  
 190 varying number of samples  $m$ .

191 In all setups, our stratified estimate uniformly outperformed the baselines. For the most challenging  
 192 "foggy" field, approximately  $m = 32$  samples we required to match the baseline performance for  
 193  $m = 128$ . We matched the baseline with only a  $m = 4$  samples for other fields. Importance sampling  
 194 requires only a few points for degenerate distributions. In further experiments, we take  $m = 8, 32$  to  
 195 obtain a precise color estimate even when a model did not converge to a degenerate distribution.

## 196 5.2 Scene Reconstruction with Reparameterized Volume Sampling

197 Next, we apply our algorithm to 3D scene reconstruction based on a set of image projections. As a  
 198 benchmark, we use the common Lego dataset. The primary goal of the experiment is to demonstrate  
 199 computational advantages of our algorithm compared to a basic volume rendering baseline.

200 As a starting point, we took the original NeRF's MLP [8] with eight layers used to compute density  
 201 and radiance. Then we modified the architecture to use only three first layers to compute the density  
 202 field. When the density field is queried, we only compute the first three layers, while for the radiance  
 203 we compute the whole network. Even though such modification may have put additional limitations  
 204 on the density model, it illustrates the benefit of using fewer radiance queries. For density, we used  
 205 softplus activation to ensure its positivity, while for the radiance we used sigmoid activation to ensure  
 206 that the output will be a valid RGB image.

207 In our experiment, we did not reproduce the expensive hierarchical sampling used in NeRF and  
 208 trained a single model in all experiments. Our baseline calculated color using Eq. ???. We took a  
 209 dense grid of  $m = 128$  points along each ray and trained the model using Huber loss with the ground  
 210 truth colors and the predict colors. We additionally perturbed the grid to regularize the model. We  
 211 used Adam [5] optimizer for training and decayed the learning rate during 100 epochs of training  
 212 from  $3e-4$  to  $3e-7$  following MIP-NeRF's scheduler [2] with image batch size equal 8 and each  
 213 epoch consisting of 8000 batches. To form a training batch, for each image in an image batch we  
 214 selected 375 pixels and calculated loss over them.

215 We evaluated the importance sampling-based rendering algorithm with the same architecture and  
 216 hyperparameters as with the baseline model. We used the same algorithm to sample a dense grid of  
 217  $m = 128$  points to query the density field and construct an approximating spline. Then we calculated  
 218 color approximation with Eq. 8 with  $m' = \{8, 32\}$  samples from the inverse cumulative density  
 219 function approximated by the spline.

Model		PSNR ( $\uparrow$ )	SSIM ( $\uparrow$ )	LPIPS (alex) [14]( $\downarrow$ )
Baseline		27.247	0.904	0.1138
Splines, #pts in estimation 8				
Training	Validation			
8 pts	1 pts	23.377	0.822	0.1819
8 pts	2 pts	25.193	0.858	0.1449
8 pts	4 pts	26.210	0.883	0.1215
8 pts	8 pts	26.502	0.892	0.1243
8 pts	16 pts	26.570	0.894	0.1333
8 pts	32 pts	26.585	0.895	0.1369
32 pts	1 pts	22.519	0.805	0.2050
32 pts	2 pts	24.902	0.846	0.1523
32 pts	4 pts	26.523	0.881	0.1181
32 pts	8 pts	27.100	0.897	0.1083
32 pts	16 pts	27.252	0.902	0.1167
32 pts	32 pts	27.286	0.904	0.1223

Table 1: Ablation study and comparison with the baseline. Metrics are calculated over test views for Lego scene [8]

220 First, we compared the rendering quality of our algorithm against the baseline. Tab. 1 contains the  
 221 quantitative results and figs. 3 and 4 contain qualitative results. From the rendering quality viewpoint  
 222 (1), with  $m' = 32$  samples, our model works on par with the baseline, while with  $m' = 8$  samples it  
 223 has slightly worse performance. Though we did not aim to reproduce the state-of-the-art results, we  
 224 speculate that a better density model could improve the results even further. In Fig. ??, we compared  
 225 the rendering performance of importance sampling for varying  $m'$ . Our algorithm produced sensible  
 226 renders even for  $m' = 1$ , however noise artifacts only disappeared for  $m' = 32$ . Fig. 4 shows a  
 227 stratified estimate renders (Eq. 8) along with a Monte Carlo renders (Eq. 7) for  $m' = 32$ . With  
 228 the same rendering complexity, the variance reduction obtained via stratified sampling purges the  
 229 rendering artifacts that a naive Monte Carlo estimate has.

Model	Iter/sec ( $\uparrow$ )	Mem Usage ( $\downarrow$ )
Baseline	3.90	8.5 Gb
Splines 1 pts	4.89	1.8 Gb
Splines 2 pts	5.06	1.8 Gb
Splines 4 pts	4.88	2.1 Gb
Splines 8 pts	4.53	2.2 Gb
Splines 16 pts	3.81	2.5 Gb
Splines 32 pts	2.98	2.8 Gb

Table 2: Speed & memory estimates. Iteration time is measured during training on GTX 1080 ti, memory usage is measured during inference with batch size equal 1024

230 Besides the rendering quality, we estimated the training speed and memory footprint of our algorithm  
 231 in Tab. 2. For  $m' = 8$  training iterations were on average  $\times 1.2$  faster, while for  $m' = 32$  training  
 232 iterations took  $\times 1.3$  more time. The difference occurred due to a varying number of radiance queries.  
 233 For a memory footprint viewpoint, our algorithm used  $\times 3.0$  and  $\times 3.9$  less memory for  $m' = 32$  and  
 234  $m = 8$  correspondingly. With this, important sampling leaves room for further optimization as it  
 235 allows to work with bigger batches with a moderate variability in rendering speed and quality.




Figure 3: Rendering results with a different number of samples in the stratified estimate. From left to right and from top to down: 1, 2, 8, 32 points estimates, Baseline and Target for reference.




Figure 4: Comparison of rendering results from different viewing angles with Monte-Carlo estimate (top row) and stratified Monte-Carlo estimate (bottom row), both with 32 points along each ray



## 236 6 Conclusion

237 We proposed an alternative to classic volume rendering algorithms used in 3D scene reconstruction.  
238 For a synthetic experiment and in full-scale reconstruction task we achieve better estimation results  
239 in terms of variance with a significantly smaller computation footprint. In particular, our algorithm  
240 allows for significant memory reductions and even increased inference time. At the same time, we  
241 demonstrate competitive rendering quality. We believe that our approach is a promising alternative to  
242 standard volume rendering techniques.

### 243 6.1 Broader Impact

244 We hypothesize that models like NeRFs may be used in online stores for a better user experience.  
245 Then people will choose more suitable products. We are not aware of any possibilities to use this in a  
246 negative way. Furthermore, we are sure that the efficient sampling we proposed for 3D rendering  
247 may reduce computation costs and therefore environmental damage.

## 248 References

- 249 [1] Shaojie Bai, J Zico Kolter, and Vladlen Koltun. Deep equilibrium models. *Advances in Neural*  
250 *Information Processing Systems*, 32, 2019.
- 251 [2] Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla,  
252 and Pratul P Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance  
253 fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages  
254 5855–5864, 2021.
- 255 [3] Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary  
256 differential equations. *Advances in neural information processing systems*, 31, 2018.
- 257 [4] Mikhail Figurnov, Shakir Mohamed, and Andriy Mnih. Implicit reparameterization gradients.  
258 *Advances in Neural Information Processing Systems*, 31, 2018.
- 259 [5] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint*  
260 *arXiv:1412.6980*, 2014.
- 261 [6] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint*  
262 *arXiv:1312.6114*, 2013.
- 263 [7] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse  
264 voxel fields. *Advances in Neural Information Processing Systems*, 33:15651–15663, 2020.
- 265 [8] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi,  
266 and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *European*  
267 *conference on computer vision*, pages 405–421. Springer, 2020.
- 268 [9] Shakir Mohamed, Mihaela Rosca, Michael Figurnov, and Andriy Mnih. Monte carlo gradient  
269 estimation in machine learning. *J. Mach. Learn. Res.*, 21(132):1–62, 2020.
- 270 [10] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics  
271 primitives with a multiresolution hash encoding. *arXiv preprint arXiv:2201.05989*, 2022.
- 272 [11] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation  
273 and approximate inference in deep generative models. In *International conference on machine*  
274 *learning*, pages 1278–1286. PMLR, 2014.
- 275 [12] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct voxel grid optimization: Super-fast  
276 convergence for radiance fields reconstruction. *arXiv preprint arXiv:2111.11215*, 2021.
- 277 [13] Alex Yu, Sara Fridovich-Keil, Matthew Tancik, Qinhong Chen, Benjamin Recht, and  
278 Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. *arXiv preprint*  
279 *arXiv:2112.05131*, 2021.
- 280 [14] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreason-  
281 able effectiveness of deep features as a perceptual metric. In *CVPR*, 2018.

282 **Checklist**

- 283 1. For all authors...
- 284 (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s  
285 contributions and scope? [Yes] See Section 1.
- 286 (b) Did you describe the limitations of your work? [Yes] See Section 1.
- 287 (c) Did you discuss any potential negative societal impacts of your work? [Yes] See  
288 Section 6.1
- 289 (d) Have you read the ethics review guidelines and ensured that your paper conforms to  
290 them? [Yes]
- 291 2. If you are including theoretical results...
- 292 (a) Did you state the full set of assumptions of all theoretical results? [Yes] We listed our  
293 assumptions in section
- 294 (b) Did you include complete proofs of all theoretical results? [Yes] Proofs in section are  
295 complete and rely on basic math.
- 296 3. If you ran experiments...
- 297 (a) Did you include the code, data, and instructions needed to reproduce the main experi-  
298 mental results (either in the supplemental material or as a URL)? [No] But we plan to  
299 release code soon, within supplementary materials
- 300 (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they  
301 were chosen)? [Yes] See Section 5.2
- 302 (c) Did you report error bars (e.g., with respect to the random seed after running experi-  
303 ments multiple times)? [No] We have rerun our experiments several times and so  
304 that results are pretty similar. We plan to continue experiments and made more run  
305 with other MLP architectures as well as voxel-based models.
- 306 (d) Did you include the total amount of compute and the type of resources used (e.g., type  
307 of GPUs, internal cluster, or cloud provider)? [Yes] See Section 5.2
- 308 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- 309 (a) If your work uses existing assets, did you cite the creators? [Yes] Yes, we cite Pytorch,  
310 pytorch3d and NeRF’s creators as well as many other relevant. See Section 5.2
- 311 (b) Did you mention the license of the assets? [No] We use opensource assets and cite/share  
312 links to them.
- 313 (c) Did you include any new assets either in the supplemental material or as a URL? [No]  
314 But we plan to release code soon, within supplementary materials
- 315 (d) Did you discuss whether and how consent was obtained from people whose data you’re  
316 using/curating? [N/A] Not applicable, data is artificially generated
- 317 (e) Did you discuss whether the data you are using/curating contains personally identifiable  
318 information or offensive content? [N/A] Not applicable, data is artificially generated
- 319 5. If you used crowdsourcing or conducted research with human subjects...
- 320 (a) Did you include the full text of instructions given to participants and screenshots, if  
321 applicable? [N/A] Not applicable
- 322 (b) Did you describe any potential participant risks, with links to Institutional Review  
323 Board (IRB) approvals, if applicable? [N/A] Not applicable
- 324 (c) Did you include the estimated hourly wage paid to participants and the total amount  
325 spent on participant compensation? [N/A] Not applicable