

# Flash Cache: Reducing Bias in Radiance Cache Based Inverse Rendering


Benjamin Attal<sup>1</sup>, Dor Verbin<sup>2</sup>, Ben Mildenhall<sup>2</sup>, Peter Hedman<sup>2</sup>,  
Jonathan T. Barron<sup>2</sup>, Matthew O’Toole<sup>1</sup>, and Pratul P. Srinivasan<sup>2</sup>

<sup>1</sup>Carnegie Mellon University    <sup>2</sup>Google Research  
<https://benattal.github.io/flash-cache/>

**Abstract.** State-of-the-art techniques for 3D reconstruction are largely based on volumetric scene representations, which require sampling multiple points to compute the color arriving along a ray. Using these representations for more general inverse rendering — reconstructing geometry, materials, and lighting from observed images — is challenging because recursively path-tracing such volumetric representations is expensive. Recent works alleviate this issue through the use of radiance caches: data structures that store the steady-state, infinite-bounce radiance arriving at any point from any direction. However, these solutions rely on approximations that introduce bias into the renderings and, more importantly, into the gradients used for optimization. We present a method that avoids these approximations while remaining computationally efficient. In particular, we leverage two techniques to reduce variance for unbiased estimators of the rendering equation: (1) an occlusion-aware importance sampler for incoming illumination and (2) a fast cache architecture that can be used as a control variate for the radiance from a high-quality, but more expensive, volumetric cache. We show that by removing these biases our approach improves the generality of radiance cache based inverse rendering, as well as increasing quality in the presence of challenging light transport effects such as specular interreflections.

## 1 Introduction

Volume rendering has been remarkably effective for novel view synthesis and 3D reconstruction, as demonstrated by approaches such as Neural Radiance Fields (NeRFs) [28] and Gaussian Splatting [18]. However, it is difficult to use volume rendering for the more general inverse rendering problem of recovering materials and lighting in addition to geometry. The issue is largely computational: volume rendering techniques generally query the appearance and density of the volume at many sample points to compute the light arriving along each ray. Combining this with physically-accurate rendering models that simulate global illumination (instead of simple shading models that only consider direct lighting) dramatically increases computational complexity, because evaluating the rendering integral requires recursively casting additional scattered rays from each sample location.



**Fig. 1:** Results on two captured scenes from the Stanford ORB dataset [21]. Our model synthesizes realistic novel views of real-world objects that exhibit complicated specular materials, and recovers accurate material roughnesses (black represents a perfect mirror, white is perfectly rough) and albedo maps.

Recent methods [15, 48] have leveraged radiance caches to simulate global illumination tractably within an inverse rendering optimization pipeline. Instead of integrating over the domain of  $n$ -bounce light paths, radiance caches store a representation of the steady-state light arriving at any point in the volume from any direction. This is a powerful paradigm that removes the need for recursively evaluating the rendering integral by considering a single reflection of the incoming steady-state (as opposed to direct) light at any point on the scene geometry towards the direction of the viewer.

Existing works that combine radiance caches with volumetric inverse rendering fall into two categories. The first uses trained NeRF models as radiance caches, and volume renders the NeRF along rays from points within the scene to estimate incoming light. This type of radiance cache has high quality but tends to be computationally expensive to evaluate; methods that use NeRF as a radiance cache, such as TensoIR [15], only compute the reflection integral at a single point (the expected ray termination) along each camera ray. The second category relies on cheap radiance caches such as a small MLPs, as in InvRender [48] and NeRO [25], to evaluate incoming lighting with a single ray query. Though fast, these caches fail to capture high-frequency and near-field illumination. We observe that both of these existing approaches introduce bias into the pipeline: the first category renders a NeRF model at only a single evaluation along each camera ray, which results in a biased estimation of the volume rendering integral, and the second category approximates incoming light with a low-capacity model, which results in an inaccurate and biased estimation of the reflection integral.

We present a system for volumetric inverse rendering based on the basic goal of avoiding introducing bias into the rendering procedure as much as possible, while leveraging variance reduction techniques to keep the rendering model computationally tractable. This is made possible by the following key technical contributions:

- A strategy for efficiently estimating the light reflected by volumetric geometry along each camera ray by using a neural field to parameterize a spatially-varying mixture of von-Mises Fisher distributions for light importance sampling. This captures spatially-varying illumination effects such as light occlusion and interreflections and lets our system efficiently estimate the rendering equation with fewer queries of the radiance cache.
- An architecture for a lightweight high-capacity radiance cache that can query the incoming illumination from any direction at any point within the volume, along with a strategy for using this cache as a control variate to avoid introducing rendering bias. This lets us efficiently model high-frequency and near-field illumination without requiring the full volume rendering necessary when using NeRF as a radiance cache.

We demonstrate that our system recovers improved geometry and materials over prior work, and in particular enables more accurate recovery of material parameters in regions affected by near-field lighting effects such as specular interreflections. Furthermore, since our system does not use a simplified model of volumetric geometry, it preserves the advantages of NeRF-like volumetric representations for capturing thin structures. We validate the benefits of our system with evaluation on existing benchmarks and rendered scenes that exhibit challenging lighting effects.

## 2 Related Work

*Inverse rendering.* Inverse rendering aims to recover scene attributes like materials, lighting, and geometry, from a set of images [1, 8, 35, 45]. A common approach is to use a physically-based model of light transport [16, 34, 40] and differentiable rendering so that it is possible to leverage gradient-based optimization for decomposing the scene’s appearance into its constituent parts [7, 14, 22]. Recent advances in view synthesis demonstrate the benefit of volume rendering neural field representations for the task of reproducing realistic images of a captured scene from unseen views [2, 28, 41]. Neural field representations also improve inverse rendering methods, combining ideas from volume rendering to represent geometry and physically-based rendering to represent appearance [3, 12, 26, 31, 39].

*Variance-reduction for inverse rendering.* Among the different paradigms for inverse rendering using neural fields, cache-based inverse rendering is the most closely related to ours. These methods rely on using a radiance cache for storing the incoming radiance at every point and along every direction. These radiance caches capture global illumination effects, replacing a prohibitively expensive recursive path tracing procedure with a much cheaper cache look-up. This technique has been used extensively in the past for computing indirect illumination

and global illumination effects [20, 42], and consequently used for real-time rendering applications [30, 37, 38].

In the context of inverse rendering, radiance caches are used for storing the hemisphere of incoming radiance at every point. This hemisphere is then integrated against the local BRDF to yield outgoing radiance, which is optimized to match the observed image pixels. NeRFactor [47] only considers direct illumination effects for decomposing lights and materials, but caches the visibility of each secondary ray into an MLP to model self occlusions. Neural Incident Light Field (NeILF) [43] and Neural Reflective Objects (NeRO) [25] and InvRender [48] parameterize the radiance cache using an MLP modeling the incident light field. However, they only supervise the outgoing radiance, which does not enforce consistency between the radiance cache and incoming radiance. An improved version of NeILF [46] solves this by penalizing differences between the radiance cache and the corresponding outgoing radiance. These methods all use inexpensive caches which tend to fail to capture high-frequency and near-field illumination. Most recently, concurrent work to ours from Ling *et al.* [23] uses a NeRF-based cache, along with an optimizable mixture of Gaussians model for importance sampling the 3D distribution of light sources in the scene. Our work also uses importance sampling to reduce Monte Carlo noise, but our distributions account for occlusions — a key part of variance reduction [40]. Along with importance sampling, we further reduce noise using control variates, two techniques that have been extensively explored in the context of forward rendering [4, 29, 33].

### 3 Preliminaries

In this work, we consider the problem of inverse rendering of geometry, materials, and lighting from a set of input images. We leverage a combination of volumetric rendering and physically-based Monte Carlo rendering in order to accomplish this task. In particular, we optimize the parameters  $\Phi$  of a scene representation using a regularization term  $\mathcal{L}_{reg}$  and a loss term such that the color value  $L_i(\mathbf{o}, \boldsymbol{\omega}_o; \Phi)$  incident at a pixel corresponding to a camera ray  $(\mathbf{o}, \boldsymbol{\omega}_o)$  — where  $\mathbf{o}$  is its origin and  $\boldsymbol{\omega}_o$  its direction — matches with the corresponding measured image pixel  $I(\mathbf{o}, \boldsymbol{\omega}_o)$ :

$$\min_{\Phi} \sum_{\mathbf{o}, \boldsymbol{\omega}_o} \|I(\mathbf{o}, \boldsymbol{\omega}_o) - L_i(\mathbf{o}, \boldsymbol{\omega}_o; \Phi)\|^2 + \mathcal{L}_{reg}(\Phi) \quad (1)$$

#### 3.1 Volume Rendering and Neural Radiance Fields

In the volume rendering setting, the incident radiance  $L_i$  of a ray  $(\mathbf{o}, \boldsymbol{\omega}_o)$  passing through a neural volume, is computed by integrating the differential outgoing radiance  $L_o$  at every point along that ray:

$$L_i(\mathbf{o}, \boldsymbol{\omega}_o) = \int_0^\infty L_o(\mathbf{x}(t), \boldsymbol{\omega}_o) T(\mathbf{o}, \mathbf{x}(t)) \sigma(\mathbf{x}(t)) dt, \quad (2)$$

where  $T(\mathbf{o}, \mathbf{x}(t)) = \exp(-\int_0^t \sigma(\mathbf{x}(s)) ds)$  denotes the transmittance from  $\mathbf{o}$  to  $\mathbf{x}(t)$ , and  $\mathbf{x}(t) = \mathbf{o} - t\boldsymbol{\omega}_o$  is a point along the ray parameterized by  $t$ . This integral is generally intractable to evaluate analytically, but it can be approximated via quadrature [27], by sampling density and radiance at a set of points  $\{\mathbf{x}(t_k)\}_{k=1}^N$ :

$$\hat{L}_i(\mathbf{o}, \boldsymbol{\omega}_o) = \sum_{k=1}^N w_k L_o(\mathbf{x}(t_k), \boldsymbol{\omega}_o), \quad (3)$$

$$w_k = \left(1 - e^{-\sigma(\mathbf{x}(t_k))(t_{k+1}-t_k)}\right) \exp\left(-\sum_{j=1}^{k-1} \sigma(\mathbf{x}(t_j))(t_{j+1}-t_j)\right).$$

In NeRF [28], the volume density  $\sigma$  and outgoing radiance  $L_o$  are parameterized using a multi-layer perceptron (MLP). Follow-up work [6, 9, 32] relies on optimizable grids or hash tables along with smaller neural networks. In both cases, a function maps every position  $\mathbf{x}$  and direction  $\boldsymbol{\omega}_o$  into the volume density  $\sigma(\mathbf{x})$  at that location and the outgoing radiance from it in the specified direction  $L_o(\mathbf{x}, \boldsymbol{\omega}_o)$ . Together these give the radiance along the ray using Equation 3.

### 3.2 Physics-Based Monte Carlo Light Transport

We use ideas from physically-based rendering for combining incoming radiance  $L_i$  with a BRDF  $f$  to produce outgoing radiance  $L_o$  using the rendering equation:

$$L_o(\mathbf{x}(t), \boldsymbol{\omega}_o) = \int_{\Omega} f(\mathbf{x}(t), \boldsymbol{\omega}_i, \boldsymbol{\omega}_o) L_i(\mathbf{x}, \boldsymbol{\omega}_i) (\mathbf{n} \cdot \boldsymbol{\omega}_i) d\boldsymbol{\omega}_i, \quad (4)$$

where  $\Omega = \{\boldsymbol{\omega}_i : \mathbf{n} \cdot \boldsymbol{\omega}_i > 0\}$  is the positive hemisphere with respect to normal  $\mathbf{n}$ .


Like the volume rendering integral in Equation 2, evaluating the rendering equation analytically is, in general, intractable. Most approaches therefore use Monte Carlo integration in order to approximate Equation 4. In particular, if we randomly sample  $M$  incoming directions from some distribution  $p$  over incoming directions,  $\boldsymbol{\omega}_i^{(1)}, \dots, \boldsymbol{\omega}_i^{(M)} \sim p(\boldsymbol{\omega}_i)$ , then we can construct an unbiased estimator:

$$\hat{L}_o(\mathbf{x}(t), \boldsymbol{\omega}_o) = \frac{1}{M} \sum_{j=1}^M f(\mathbf{x}(t), \boldsymbol{\omega}_i^{(j)}, \boldsymbol{\omega}_o) L_i(\mathbf{x}(t), \boldsymbol{\omega}_i^{(j)}) (\mathbf{n} \cdot \boldsymbol{\omega}_i^{(j)}) / p(\boldsymbol{\omega}_i^{(j)}), \quad (5)$$

where performing optimization using the estimator in Equation 5 is, in expectation, equivalent to optimizing using Equation 4 [19].

We can now plug the Monte Carlo-based estimator in Equation 5 into the volume rendering estimator of Equation 3 to arrive at the combined estimator:

$$\hat{L}_i(\mathbf{o}, \boldsymbol{\omega}_o) = \sum_{k=1}^N \frac{w_k}{M} \sum_{j=1}^M f(\mathbf{x}(t_k), \boldsymbol{\omega}_i^{(j,k)}, \boldsymbol{\omega}_o) \hat{L}_i(\mathbf{x}(t_k), \boldsymbol{\omega}_i^{(j,k)}) (\mathbf{n} \cdot \boldsymbol{\omega}_i^{(j,k)}) / p(\boldsymbol{\omega}_i^{(j,k)}) \quad (6)$$



**Fig. 2:** Our pipeline for physically-based inverse rendering combines volume rendering with physically-based Monte Carlo rendering to recover geometry, materials, and lighting. We first use an estimator (Equation 11) to approximate the volume rendering integral (Equation 2) with only a single sample point per ray. Given this sample point, we use a spatially-varying mixture-of-vMFs (Section 4.4) to importance-sample directions based on the distribution of incoming radiance. We then produce a low-variance estimate of the rendering equation (Equation 4) using a NeRF as a radiance cache, and a fast cache approximation of the NeRF (Equation 9) that is far less expensive to evaluate.

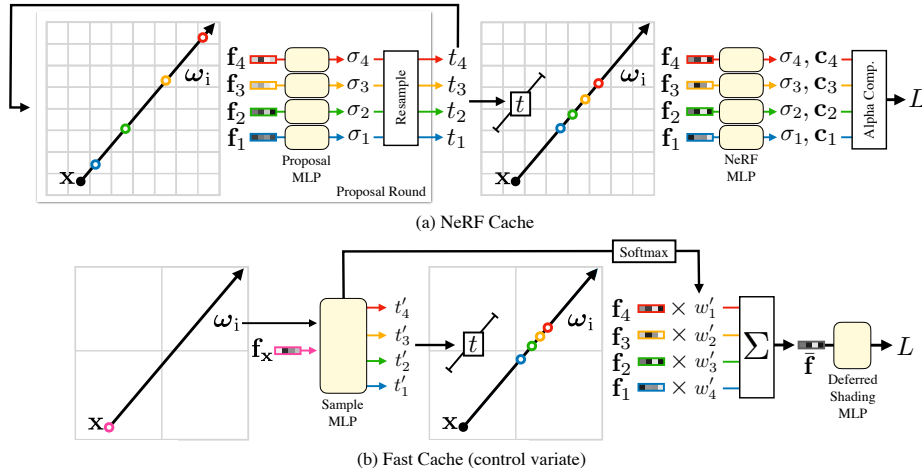
Unfortunately, Equation 6 requires recursive evaluation of both the volume rendering integral and the rendering equation, since  $\hat{L}_i$  is now defined with respect to itself. If one is not careful, the recursive nature of the rendering equation can quickly lead to a rapid increase in computational cost.

To deal with this issue, we return to the discussion in Section 3.1. It is possible to define a neural volume that *directly* predicts outgoing radiance, without requiring recursive evaluation. In other words, we can replace the incoming radiance  $\hat{L}_i$  with the output of a neural radiance field  $\hat{L}_i^{cache}$ , which relates the incoming and outgoing radiance according to Equation 3. It is worth remarking that for this to be useful, evaluating the radiance cache  $\hat{L}_i^{cache}$  for a given ray should be cheaper than performing recursive path tracing.

## 4 Method

Our inverse rendering method is based on a volumetric rendering approach to physically-based rendering, as described by Equation 6. However, naively using the rendering model of Equation 6 is extremely expensive for two main reasons. First, it requires volumetrically rendering the incoming radiance at every point along the ray. Second, for every such point we need to use Monte Carlo integration, which typically requires many samples of incoming radiance. Our work aims to resolve both issues without introducing bias in the process, which may adversely affect the recovery of materials and lighting.

In this section, we present the design of our system. We begin by describing the incoming radiance cache in Section 4.1, along with a novel “fast cache” architecture that is cheap to evaluate. Designing this cache as a control variate for incoming illumination allows us to use many incoming rays, which increases the accuracy of our method without requiring significant compute. In Section 4.2, we



**Fig. 3:** An illustration of our two radiance cache architectures. (a) The NeRF-based cache, which is accurate yet expensive to evaluate, uses proposal MLPs to predict sample point distances  $\{t_k\}$  through an iterative resampling procedure. These distances are then used to generate sample points at which we query density and outgoing radiance. These quantities are combined to approximate the volume rendering integral via quadrature 3. (b) Rather than predict sample point distances via proposal sampling, our fast cache directly outputs distances  $\{t'_k\}$  as well as weights  $\{w'_k\}$  given incoming ray parameters  $(\mathbf{x}, \boldsymbol{\omega}_i)$ . The distances are used to generate sample points at which we query features from an NGP grid. The resulting features are averaged together, and then fed through a deferred shading MLP which predicts the final color for the incoming ray. Though our illustration depicts four samples in both cases, the NeRF-based cache (a) uses 64 samples for the first two resampling rounds and 32 samples for the final one, while our fast cache only uses 8 samples overall.

discuss our physically-based material representation. Finally, in Sections 4.3 and 4.4, we present our efficient sampling schemes based on a simple unbiased estimator for the volumetric rendering integral, and a novel optimizable occlusion-aware spatially-varying importance sampling scheme.

#### 4.1 Incoming Radiance Cache

Our incoming radiance cache is based on the architecture of Zip-NeRF [2], and our appearance model is based on Ref-NeRF [41]. In particular, we borrow Zip-NeRF’s usage of a multi-resolution hash grid as a neural graphics primitive (NGP) [32]. See Figure 3(a) for an illustration of the NeRF-based radiance cache, and see a complete description of it in the supplement.

In order to increase the efficiency of our method, we wish to use an accelerated cache for querying incoming radiance instead of relying on an expensive NeRF-based cache which requires evaluating many samples per incoming ray. We therefore design a “fast cache” and supervise it to match the incoming radi-

ance of the expensive, NeRF-based cache. See Figure 3 for a diagram comparing the two caches.

Our fast cache works by first passing the incident ray’s origin and direction ( $\mathbf{x}$  and  $\boldsymbol{\omega}_i$ ) through an optimizable function  $g_{\text{sample}}$ , implemented as a low-resolution hash-encoding followed by a small sample MLP. This sampling function returns a set of sample point distances along the ray and their corresponding weights:

$$\{t'_k, w'_k\}_{k=1}^S = g_{\text{sample}}(\mathbf{x}, \boldsymbol{\omega}_i). \quad (7)$$

Each distance  $t'_k$  corresponds to a new sample point  $\mathbf{x}'_k = \mathbf{x} + t'_k \boldsymbol{\omega}_i$  along the incident ray. Each of these samples is passed through a separate low-resolution NGP to yield a set of corresponding features  $\{\mathbf{f}_k\}_{k=1}^S$ , which are combined with their corresponding normalized (via a softmax) weights  $w'_k$  to yield a ray feature vector which encodes the appearance of that ray. This feature is then decoded into the incident radiance by using a small MLP, similar to the “deferred shading” approach of Hedman *et al.* [13]:

$$\hat{L}_i^{\text{fast}} = \text{MLP}_{\text{color}} \left( \sum_{k=1}^S w'_k \mathbf{f}_k \right). \quad (8)$$

Even for a small number of samples (we use  $S = 8$  for all experiments), we find that this design serves as an effective cache for incoming radiance that captures high-frequency near-field illumination. See Figure 3 for an overview of the NeRF-based and fast caches, as well as figure 5 for an example of the output of the fast cache. Our supplement contains more details about both models.

Instead of simply replacing the NeRF-based cache with our fast cache, we combine the accuracy of the NeRF-based cache with the efficiency of the fast cache by treating the fast cache as a control variate [17, 33] for incoming radiance. In order to do this, we sample two sets of incident directions using an importance sampling distribution  $p$  (to be introduced in Section 4.4):  $M'$  samples  $\{\boldsymbol{\omega}_i'^{(j)}\}_{j=1}^{M'}$ , and  $M \ll M'$  additional independent samples  $\{\boldsymbol{\omega}_i^{(j)}\}_{j=1}^M$ . We then use the  $M'$  first samples to evaluate the fast cache  $\hat{L}_i^{\text{fast}}(\mathbf{x}, \boldsymbol{\omega}_i)$  via Equation 8, and the other  $M$  samples to evaluate the error between the NeRF-based cache  $\hat{L}_i^{\text{NeRF}}(\mathbf{x}, \boldsymbol{\omega}_i)$  and the fast cache. We then add the two estimates to arrive at an unbiased estimator for the outgoing radiance:

$$\begin{aligned} \hat{L}_o^{\text{fast}} &= \frac{1}{M'} \sum_{j=1}^{M'} f(\mathbf{x}, \boldsymbol{\omega}_i'^{(j)}, \boldsymbol{\omega}_o) \hat{L}_i^{\text{fast}}(\mathbf{x}, \boldsymbol{\omega}_i'^{(j)}) (\mathbf{n} \cdot \boldsymbol{\omega}_i'^{(j)}) / p(\boldsymbol{\omega}_i'^{(j)}), \\ \Delta \hat{L}_o &= \frac{1}{M} \sum_{j=1}^M f(\mathbf{x}, \boldsymbol{\omega}_i^{(j)}, \boldsymbol{\omega}_o) \left( \hat{L}_i^{\text{NeRF}}(\mathbf{x}, \boldsymbol{\omega}_i^{(j)}) - \hat{L}_i^{\text{fast}}(\mathbf{x}, \boldsymbol{\omega}_i^{(j)}) \right) (\mathbf{n} \cdot \boldsymbol{\omega}_i^{(j)}) / p(\boldsymbol{\omega}_i^{(j)}), \\ \hat{L}_o &= \hat{L}_o^{\text{fast}} + \Delta \hat{L}_o. \end{aligned} \quad (9)$$

Note that this requires  $M$  evaluations of the NeRF-based cache and  $M + M'$  evaluations of the fast cache. Given sufficiently accurate cache (see Section 5 for



an example) that is indeed cheap to evaluate, this process results in the quality of using  $M'$  incoming rays from the accurate NeRF-based cache, by only requiring a fraction of the computation.

On top of the two caches, we also have an explicit model for far-field illumination based on an MLP,  $\hat{L}_1^{env}(\omega_i) = \text{MLP}_{env}(\omega_i)$ , which is multiplied by the transparency of the ray  $1 - \sum_{k=1}^N w_k$  and added to the radiance of the cache.

## 4.2 Material Representation

Like Ref-NeRF [41], we extract analytic normals using the negative gradient of the density field, and leverage “predicted normals” output by the NGP, which are tied to the analytic normals using an L2 loss. Our material model is based on the Disney-GGX BRDF [5], which is combined with the predicted normals to obtain the radiance value as in Equation 4. The spatially-varying parameters of the BRDF model are output by a separate NGP,  $g_{material}$ :

$$\left(m(\mathbf{x}), r(\mathbf{x}), \mathbf{a}(\mathbf{x})\right) = g_{material}(\mathbf{x}), \quad (10)$$

where  $m$  is metalness,  $r$  is roughness, and  $\mathbf{a}$  is albedo.

## 4.3 Efficient Gradients for Volume Rendering

Evaluating the volume rendering integral using the quadrature technique in Equation 3 means querying the appearance and geometry of the volume at  $N$  sample point locations per ray, each of which requires Monte Carlo integration. However, this can be accelerated considerably by replacing the sum with a Monte Carlo estimate along the ray using  $K \ll N$  samples [11]. We draw  $K$  indices from a categorical distribution  $j_1, \dots, j_K \sim \text{Cat}(w_1, \dots, w_N)$ , and replace Equation 3 with:


$$\hat{L}_i(\mathbf{o}, \omega_o) = \frac{1}{K} \sum_{k=1}^K L_o(\mathbf{x}(t_{j_k}), \omega_o). \quad (11)$$

The estimator of Equation 11 is unbiased, *i.e.*, its expected value equals the quadrature value of Equation 3. While this still requires querying volume density for  $N$  points, we only need to evaluate the costly physically-based appearance at  $K$  points. Furthermore, if the render weight distribution has a single peak (often the case in practice), then this estimator has low variance, even for  $K = 1$ .

## 4.4 Low-Variance Gradients for Monte Carlo Rendering

The unbiased estimator in Equation 11 reduces the number of secondary ray samples from  $N \cdot M$  to  $K \cdot M$ . However, a large number of secondary rays  $K \cdot M$  may be necessary for convergence, but may be prohibitively expensive. We reduce the number of samples by leveraging occlusion-aware importance sampler.

**Fig. 4:** An illustration of our occlusion-aware importance sampler. We use an NGP to predict the parameters of a mixture of vMFs distribution for each point in space (Equation 12), which can produce different lobes for different light sources. Because the parameters of these distributions are *spatially-varying*, if a light source is occluded from the perspective of one surface point, then the model can “turn off” the lobe for that light source.



Importance sampling for incoming radiance is common technique in physically-based forward rendering and inverse rendering. Here, we aim to build a lightweight importance sampler that is capable of approximating the true distribution of incoming radiance, including occlusions, from both direct and indirect light sources. For this purpose, we use an NGP to map a point  $\mathbf{x}$  into a set of 5L scalars representing the parameters  $\{\boldsymbol{\mu}'_\ell, \kappa_\ell, \lambda_\ell\}_{\ell=1}^L$  which parameterize a mixture of von Mises-Fisher (vMF) distributions as:

$$q(\boldsymbol{\omega}; \mathbf{x}) = \frac{1}{Z} \sum_{\ell=1}^L \lambda_\ell(\mathbf{x}) \text{vMF}(\boldsymbol{\omega}; \boldsymbol{\mu}_\ell(\mathbf{x}), \kappa_\ell(\mathbf{x})), \text{ with } \boldsymbol{\mu}_\ell(\mathbf{x}) = \frac{\boldsymbol{\mu}'_\ell(\mathbf{x}) - \mathbf{x}}{\|\boldsymbol{\mu}'_\ell(\mathbf{x}) - \mathbf{x}\|}, \quad (12)$$

where  $Z = \sum_{\ell=1}^L \lambda_\ell(\mathbf{x})$  is the partition function used for normalizing  $q$ , and the coefficients  $\{\lambda_\ell\}_{\ell=1}^L$  and  $\{\kappa_\ell\}_{\ell=1}^L$  are nonnegative numbers representing the mixture weights and vMF concentration parameters respectively.

The parameterization of the means  $\{\boldsymbol{\mu}'_\ell\}_{\ell=1}^L$  in Equation 12 is designed so that a single vMF lobe can be assigned to the location of single (direct or indirect) light source in 3D, which is then projected onto the sphere.


We optimize the parameters describing the distribution  $q$  at every location to match the true distribution of incoming radiance, taking occlusions into account. Concretely, we use the same importance sampling scheme used for evaluating the rendering integral in Equation 5, to minimize the loss:

$$\mathcal{L}_{\text{vMF}} = \frac{1}{M} \sum_{j=1}^M \left\| Z \cdot q(\boldsymbol{\omega}_i^{(j)}; \mathbf{x}) - \bar{L}(\mathbf{x}, \boldsymbol{\omega}_i^{(j)}) \right\|^2 / p(\boldsymbol{\omega}_i^{(j)}), \quad (13)$$

where  $p$  is a distribution used to generate the samples  $\{\boldsymbol{\omega}_i^{(j)}\}_{j=1}^M$  (as well as the  $M'$  cheap cache samples used in Section 4.1), which is based the mixture distribution  $q$  and a standard material-based importance sampling distribution (see supplement for more details), and  $\bar{L}(\mathbf{x}, \boldsymbol{\omega}_i) = \|\hat{L}_i^{\text{NeRF}}(\mathbf{x}, \boldsymbol{\omega}_i)\|_2$  is the L2 norm of the incident radiance, taken across channels.

## 5 Experiments

Below, we provide implementation details of our approach, and describe how we validate it on both synthetic and real inverse rendering benchmarks.



**Fig. 5:** (*left*) We choose two points on the hotdog for which to visualize distributions of incoming radiance. (*right*) We show (a) incoming radiance from the NeRF cache, which is high quality but slow to evaluate, (b) incoming radiance from the fast cache, which is slightly lower quality but fast to evaluate, and (c) the probability density of the spatially-varying mixture of vMFs, which is able to capture occlusions in the incoming radiance distribution.




**Fig. 6:** We demonstrate the impact of both of our variance reduction strategies by showcasing a rendering using 16spp from the *Hotdog* scene in the TensorIR-synthetic dataset. (Left) We use both the fast cache and vMF importance sampler, (Middle) we use only vMF importance sampler, (Right) we use neither strategy.

## 5.1 Implementation Details

*Model Parameters.* Our NGP parameters for volumetric density closely follow those of Zip-NeRF [2]. We do proposal resampling with two additional NGPs with maximum resolutions of 1024 and 512, both of which, as well as the final density NGP, use of 2-layer 64-hidden-unit MLPs before emitting density. The material NGP as well as another NGP used for diffuse appearance in the cache (see supplement for more details) each has a maximum grid resolution of 2048 and a 2-layer 64-hidden-unit MLP. For the fast cache discussed in Section 4.1, we use an NGP of maximum resolution 256 with a 4-layer 128-hidden-unit MLP to predict  $S = 8$  sample-point distances, and a maximum NGP of resolution 256 to predict features. For the incoming radiance importance sampler in Section 4.4, we use 128 vMF lobes, NGP with a maximum resolution of 2048, and a 2-layer 64-hidden-unit MLP. The hash maps of all NGPs have size  $2^{19}$ .

*Regularization.* We implement a smoothness regularizer  $\mathcal{L}_{BRDF}$  for BRDF parameters. In particular, we use a variant of TensorIR’s smoothness regularizer, with an  $L1$  rather than an  $L2$  loss (see the supplement for more details). We also enforce a loss  $\mathcal{L}_{consistency}$  that encourages the specular color and diffuse color from the NeRF cache to match the specular and diffuse lobes from physically-



**Fig. 7:** A visualization of ground-truth albedo and relighting for the *hotdog* (left) and *ficus* (right, cropped) scenes in the Blender dataset [28], alongside our and TensoIR’s results. Our method better decouples strong indirect illumination from the albedo (see the lip of the plate) while also reproducing more accurate highlights.

based rendering. This *indirectly* provides gradients through the physically-based model for volume rendering. We further apply the anti-aliased inter-level loss  $\mathcal{L}_{interlevel}$  from Zip-NeRF to the proposal grids, and  $L1$  regularization to the parameters of all density grids with a loss  $\mathcal{L}_{density}$ .


*Additional Details.* We perform two-stage optimization by (1) first optimizing the NeRF cache and an initial estimate of geometry to match a scene’s input color images, and (2) optimizing geometry, materials, fast cache, and importance sampler through Equation 1. For the second stage, we use an  $L2$  loss with the gradient trick [10], in order to ensure that our gradients for inverse rendering are unbiased. We leverage the same secondary rays that are used in rendering the physically-based model to supervise the fast cache and importance sampler. We *do not* apply full stop-gradients to any part of the pipeline, so that in the second stage geometry optimization can benefit from reflection and shading cues using the more accurate physically-based model of appearance. Altogether, we have the following:

$$\mathcal{L}_{reg} = \mathcal{L}_{normals} + \mathcal{L}_{BRDF} + \mathcal{L}_{consistency} + \mathcal{L}_{interlevel} + \mathcal{L}_{density} \quad (14)$$

All training and evaluation is carried out using a single NVIDIA A100 GPU.

## 5.2 Results on Synthetic Data

We first evaluate our approach on the TensoIR-synthetic dataset, which contains four scenes: *armadillo*, *ficus*, *hotdog*, and *lego*. Each scene consists of a set of 100 training views and 200 test views containing ground truth normals and albedo, as well as ground truth color under several different lighting conditions. We evaluate (a) the quality of our recovered normals, (b) the quality of our recovered albedo maps, (c) novel view synthesis quality for the *physically-based model*, and (d) relighting quality. Our method outperforms baselines in terms of albedo quality and relighting, while performing comparably to the best method for novel view synthesis. See Table 1, Figure 7, and the supplement for visualizations.



**Fig. 8:** Real results from Open Illumination [24]. Our normals are more accurate than those from TensorIR [15], and our albedo correctly removes specularities. Because of scale and color ambiguities, albedo maps are all visualized after using Photoshop’s auto-tone, auto-contrast, and auto-color features.

Method	Normal				Albedo			Novel View Synthesis			Relighting		
	MAE ↓	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓
NeRFactor	6.314	25.125	0.940	0.109	24.679	0.922	0.120	23.383	0.908	0.131			
InvRender	5.074	27.341	0.933	0.100	27.367	0.934	0.089	23.973	0.901	0.101			
TensorIR	4.100	29.275	0.950	0.085	<b>35.088</b>	<b>0.976</b>	<b>0.040</b>	28.580	<b>0.944</b>	0.081			
Ours	<b>3.355</b>	<b>30.274</b>	<b>0.955</b>	<b>0.085</b>	34.908	0.965	0.064	<b>29.724</b>	<b>0.944</b>	<b>0.080</b>			

**Table 1: TensorIR-Synthetic Dataset [15] Results.** Above, we present results of our method and 3 baselines on the TensorIR-synthetic dataset. We outperform all methods in terms of normal MAE, albedo quality, and relighting quality, while performing comparably to TensorIR in terms of novel view synthesis quality.

### 5.3 Results on Real Data

We further evaluate on 10 real scenes from the Open Illumination dataset. Each scene has 38 training views and 10 test views under 13 different lighting conditions. We split this dataset into two parts: (a) one that contains mostly diffuse objects, and (b) one that contains mostly glossy/specular objects. We evaluate novel view synthesis quality for our physically-based model, as well as relighting quality. We tend to outperform TensorIR on scenes with specular materials (visualized in Figure 8), while TensorIR performs better for diffuse materials.

### 5.4 Ablations

We ablate various design decisions and parts of our pipeline on the TensorIR synthetic dataset. In particular, we show that all of the following improve result quality: (a) volume rendering vs. median depth sampling, (b) vMF-importance sampling, and (c) our fast cache as a control variate for incoming illumination.

## 6 Conclusion

In this paper, we have outlined a system for volumetric inverse rendering of geometry, materials, and lighting, where we reduce bias in two ways: (a) by approximating the volume rendering integral (Equation 2), and (b) by using Monte Carlo sampling to evaluate the rendering integral (Equation 5). In particular,

**Table 2:** A comparison to TensorIR [15] on the Open Illumination dataset [24] split into *diffuse* and *specular* scenes.

Method	PSNR $\uparrow$ ( <i>diffuse</i> )		PSNR $\uparrow$ ( <i>specular</i> )	
	NVS	Relighting	NVS	Relighting
TensorIR	<b>32.043</b>	<b>30.932</b>	27.115	26.955
Ours	31.781	29.679	<b>27.180</b>	<b>27.810</b>

**Table 3: Ablations** We progressively ablate (a) our control variate scheme (Equation 9), (b) our vMF importance sampler (Section 4.4), and (c) our estimator (Section 11), replacing it with median depth sampling.

Method	NVS PSNR $\uparrow$	Albedo PSNR $\uparrow$	MAE $\downarrow$
Ours	<b>34.915</b>	<b>30.345</b>	3.355
(a) Ours, no variate	34.629	30.329	<b>3.352</b>
(b) Ours, no variate, no sample	34.653	30.237	<b>3.352</b>
(c) Ours, no variate, no sample, no estimator	34.249	29.838	3.491

we accelerate volume rendering while providing correct *expected* gradients for optimization, and we accelerate the evaluation of the rendering equation by combining the high-quality NeRF-based radiance cache with an occlusion-aware importance sampler and a fast-cache-based control variate. This is in contrast to existing approaches which either use a single fixed intersected surface per ray, or use low-capacity and blurry representations of incoming radiance; both of which introduce bias. We show that our system achieves better quality material reconstruction and relighting than baselines on standard benchmarks containing both real and synthetic data, and justify our design decisions with ablations.

*Limitations & Future Work:* Despite our method’s reduced bias, the quadrature approximation to the volume rendering integral (Equation 3) is still a source of bias. Our proposed approach can potentially be combined with unbiased volume rendering methods such as differential ratio tracking [Nimier-David et al. 2022].

Our fast cache sometimes struggles to capture particularly thin structural details in near-field incoming illumination. While our control variate scheme prevents this from introducing bias into the optimization process, it does increase gradient variance. In principle, any implementation of a fast cache may be used, such as a high-quality baked representation of a neural radiance field [13, 36, 44]. We leave investigation of alternative fast cache architectures up to future work.

Note that we do not claim that our method is more efficient than existing approaches; rather, that we provide efficient tools to accelerate radiance cache queries and reduce rendering variance, allowing our less-biased & more physically accurate system to run tractably on commodity GPUs. Our method could potentially be combined with other tools (e.g. denoisers) to further reduce secondary ray sample counts, accelerate inference, and increase convergence speed.

*Acknowledgements* This work was carried out while Benjamin was an intern at Google Research. Authors thank Rick Szeliski, Aleksander Holynski, and Janne Kontkanen for fruitful discussions, as well as Isabella Liu for help with TensorIR comparisons. Benjamin Attal is supported by a Meta Research PhD Fellowship. Matthew O’Toole acknowledges support from NSF CAREER 2238485 and a Google gift.

## References

1. Barron, J.T., Malik, J.: Intrinsic Scene Properties from a Single RGB-D Image. CVPR (2013)
2. Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P.: Zip-NeRF: Anti-Aliased Grid-Based Neural Radiance Fields. ICCV (2023)
3. Bi, S., Xu, Z., Srinivasan, P., Mildenhall, B., Sunkavalli, K., Hašan, M., Hold-Geoffroy, Y., Kriegman, D., Ramamoorthi, R.: Neural reflectance fields for appearance acquisition (2020), arXiv:2008.03824
4. Bitterli, B., Wyman, C., Pharr, M., Shirley, P., Lefohn, A., Jarosz, W.: Spatiotemporal reservoir resampling for real-time ray tracing with dynamic direct lighting. ACM Trans. Graph. (2020)
5. Burley, B., Studios, W.D.A.: Physically-Based Shading at Disney. ACM Trans. Graph. (2012)
6. Chen, A., Xu, Z., Geiger, A., Yu, J., Su, H.: TensorRF: Tensorial radiance fields. ECCV (2022)
7. Chen, W., Ling, H., Gao, J., Smith, E., Lehtinen, J., Jacobson, A., Fidler, S.: Learning to Predict 3D Objects with an Interpolation-based Differentiable Renderer. NeurIPS (2019)
8. Debevec, P.: Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. ACM Trans. Graph. (1998)
9. Fridovich-Keil, S., Yu, A., Tancik, M., Chen, Q., Recht, B., Kanazawa, A.: Plenoxels: Radiance fields without neural networks. CVPR (2022)
10. Gkioulekas, I., Zhao, S., Bala, K., Zickler, T., Levin, A.: Inverse volume rendering with material dictionaries. ACM Trans. Graph. (2013)
11. Gupta, K., Hasan, M., Xu, Z., Luan, F., Sunkavalli, K., Sun, X., Chandraker, M., Bi, S.: MCNeRF: Monte Carlo rendering and denoising for real-time NeRFs. ACM Trans. Graph. (2023)
12. Hasselgren, J., Hofmann, N., Munkberg, J.: Shape, Light, and Material Decomposition from Images using Monte Carlo Rendering and Denoising. NeurIPS (2022)
13. Hedman, P., Srinivasan, P.P., Mildenhall, B., Barron, J.T., Debevec, P.: Baking neural radiance fields for real-time view synthesis. ICCV (2021)
14. Jakob, W., Speierer, S., Roussel, N., Vicini, D.: Dr.Jit: A Just-In-Time Compiler for Differentiable Rendering. ACM Trans. Graph. (2022)
15. Jin, H., Liu, I., Xu, P., Zhang, X., Han, S., Bi, S., Zhou, X., Xu, Z., Su, H.: TensorIR: Tensorial Inverse Rendering. CVPR (2023)
16. Kajiya, J.T.: The rendering equation. ACM Trans. Graph. (1986)
17. Kalos, M.H., Whitlock, P.A.: Monte Carlo Methods. John Wiley & Sons (2009)
18. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3D Gaussian Splatting for Real-Time Radiance Field Rendering. ACM Transactions on Graphics (2023)

19. Kloek, T., Van Dijk, H.K.: Bayesian estimates of equation system parameters: an application of integration by Monte Carlo. *Econometrica: Journal of the Econometric Society* (1978)
20. Krivánek, J., Gautron, P.: Practical global illumination with irradiance caching. Springer Nature (2022)
21. Kuang, Z., Zhang, Y., Yu, H.X., Agarwala, S., Wu, E., Wu, J., et al.: Stanford-ORB: A Real-World 3D Object Inverse Rendering Benchmark. *NeurIPS Datasets & Benchmarks Track* (2023)
22. Li, T.M., Aittala, M., Durand, F., Lehtinen, J.: Differentiable Monte Carlo Ray Tracing through Edge Sampling. *ACM Trans. Graph.* (2018)
23. Ling, J., Yu, R., Xu, F., Du, C., Zhao, S.: NeRF as Non-Distant Environment Emitter in Physics-based Inverse Rendering. *arXiv:2402.04829* (2024)
24. Liu, I., Chen, L., Fu, Z., Wu, L., Jin, H., Li, Z., Wong, C.M.R., Xu, Y., Ramamoorthi, R., Xu, Z., et al.: OpenIllumination: A Multi-Illumination Dataset for Inverse Rendering Evaluation on Real Objects. *NeurIPS* (2024)
25. Liu, Y., Wang, P., Lin, C., Long, X., Wang, J., Liu, L., Komura, T., Wang, W.: NeRO: Neural Geometry and BRDF Reconstruction of Reflective Objects from Multiview Images. *ACM Trans. Graph.* (2023)
26. Mai, A., Verbin, D., Kuester, F., Fridovich-Keil, S.: Neural microfacet fields for inverse rendering. *ICCV* (2023)
27. Max, N.: Optical models for direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics* (1995)
28. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: NeRF: Representing scenes as neural radiance fields for view synthesis. *ECCV* (2020)
29. Müller, T., McWilliams, B., Rousselle, F., Gross, M., Novák, J.: Neural importance sampling. *ACM Trans. Graph.* (2019)
30. Müller, T., Rousselle, F., Novák, J., Keller, A.: Real-time neural radiance caching for path tracing. *ACM Trans. Graph.* (2021)
31. Munkberg, J., Hasselgren, J., Shen, T., Gao, J., Chen, W., Evans, A., Müller, T., Fidler, S.: Extracting Triangular 3D Models, Materials, and Lighting From Images. *CVPR* (2022)
32. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.* (2022)
33. Nicolet, B., Rousselle, F., Novak, J., Keller, A., Jakob, W., Müller, T.: Recursive control variates for inverse rendering. *ACM Trans. Graph.* (2023)
34. Pharr, M., Jakob, W., Humphreys, G.: *Physically based rendering: From theory to implementation*. MIT Press (2023)
35. Ramamoorthi, R., Hanrahan, P.: A signal-processing framework for inverse rendering. *ACM Trans. Graph.* (2001)
36. Reiser, C., Szeliski, R., Verbin, D., Srinivasan, P., Mildenhall, B., Geiger, A., Barron, J., Hedman, P.: MERF: Memory-Efficient Radiance Fields for Real-time View Synthesis in Unbounded Scenes. *ACM Trans. Graph.* (2023)
37. Scherzer, D., Nguyen, C.H., Ritschel, T., Seidel, H.P.: Pre-convolved radiance caching. *Computer Graphics Forum* (2012)
38. Silvennoinen, A., Lehtinen, J.: Real-time global illumination by precomputed local reconstruction from sparse radiance probes. *ACM Transactions on Graphics (TOG)* (2017)
39. Srinivasan, P.P., Deng, B., Zhang, X., Tancik, M., Mildenhall, B., Barron, J.T.: NeRV: Neural reflectance and visibility fields for relighting and view synthesis. *CVPR* (2021)



40. Veach, E.: Robust Monte Carlo methods for light transport simulation. Stanford University (1998)
41. Verbin, D., Hedman, P., Mildenhall, B., Zickler, T., Barron, J.T., Srinivasan, P.P.: Ref-NeRF: Structured view-dependent appearance for neural radiance fields. CVPR (2022)
42. Ward, G.J., Rubinstein, F.M., Clear, R.D.: A ray tracing solution for diffuse inter-reflection. ACM Trans. Graph. (1988)
43. Yao, Y., Zhang, J., Liu, J., Qu, Y., Fang, T., McKinnon, D., Tsin, Y., Quan, L.: NeiLF: Neural incident light field for physically-based material estimation. European Conference on Computer Vision (2022)
44. Yariv, L., Hedman, P., Reiser, C., Verbin, D., Srinivasan, P.P., Szeliski, R., Barron, J.T., Mildenhall, B.: BakedSDF: Meshing Neural SDFs for Real-Time View Synthesis. SIGGRAPH (2023)
45. Yu, Y., Debevec, P., Malik, J., Hawkins, T.: Inverse global illumination: Recovering reflectance models of real scenes from photographs. SIGGRAPH (1999)
46. Zhang, J., Yao, Y., Li, S., Liu, J., Fang, T., McKinnon, D., Tsin, Y., Quan, L.: NeILF++: Inter-reflectable light fields for geometry and material estimation. ICCV (2023)
47. Zhang, X., Srinivasan, P.P., Deng, B., Debevec, P., Freeman, W.T., Barron, J.T.: NeRFactor: Neural factorization of shape and reflectance under an unknown illumination. ACM Trans. Graph. (2021)
48. Zhang, Y., Sun, J., He, X., Fu, H., Jia, R., Zhou, X.: Modeling indirect illumination for inverse rendering. CVPR (2022)