

Neural Microfacet Fields for Inverse Rendering

Alexander Mai
 UC San Diego
 atm008@ucsd.edu

Dor Verbin
 Google Research
 dorverbin@google.com

Falko Kuester
 UC San Diego
 fkuester@ucsd.edu

Sara Fridovich-Keil
 UC Berkeley
 sfk@eecs.berkeley.edu

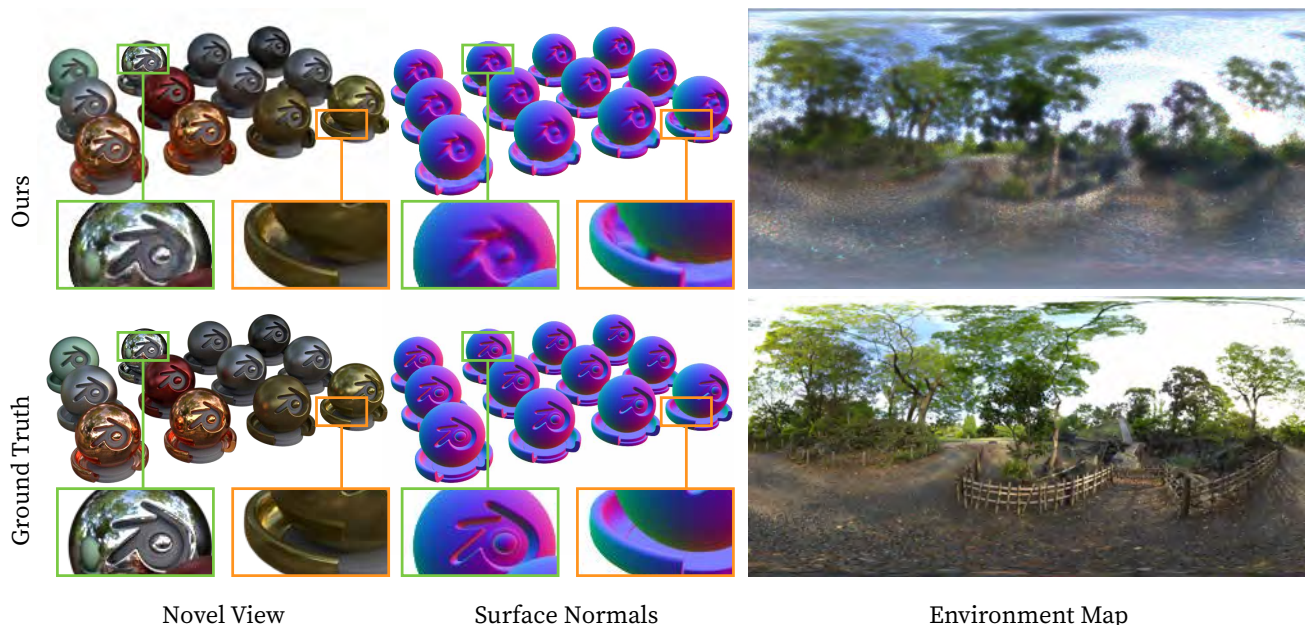


Figure 1: Our method (top) recovers materials, geometry, and illumination that closely resemble the ground truth (bottom), optimizing directly from calibrated images of a scene. Here we show results on the *materials* scene from NeRF [27]: a rendered novel view (left), surface normals (middle), and environment map illumination (right). Insets show high-fidelity reflections, including interreflections, as well as accurate geometry, even in concave regions.

Abstract

We present *Neural Microfacet Fields*, a method for recovering materials, geometry, and environment illumination from images of a scene. Our method uses a microfacet reflectance model within a volumetric setting by treating each sample along the ray as a (potentially non-opaque) surface. Using surface-based Monte Carlo rendering in a volumetric setting enables our method to perform inverse rendering efficiently by combining decades of research in surface-based light transport with recent advances in volume rendering for view synthesis. Our approach outperforms prior work in inverse rendering, capturing high fidelity geometry and high frequency illumination details; its novel view synthesis results are on par with state-of-the-art methods that do not recover illumination or materials.

1. Introduction

Simultaneous recovery of the light sources illuminating a scene and the materials and geometry of objects inside it, given a collection of images, is a fundamental problem in computer vision and graphics. This decomposition enables editing and downstream usage of a scene: rendering it from novel viewpoints, and arbitrarily changing the scene’s illumination, geometry, and material properties. This disentanglement is especially useful for creating 3D assets that can be inserted into other environments and realistically rendered under novel lighting conditions.

Recent methods for novel view synthesis based on neural radiance fields [27] have been highly successful at decomposing scenes into their geometry and appearance components, enabling rendering from new, unobserved view-

points. However, the geometry and appearance recovered are often of limited use in manipulating either materials or illumination, since they model each point as a direction-dependent emitter rather than as reflecting the incident illumination. To tackle the task of further decomposing appearance into illumination and materials, we return to a physical model of light-material interaction, which models a surface as a distribution of microfacets that *reflect* light rather than emitting it. By explicitly modeling this interaction during optimization, our method can recover both material properties and the scene’s illumination.

Our method uses a Monte Carlo rendering approach with a hybrid surface-volume representation, where the scene is parameterized as a 3D field of microfacets: the scene’s geometry is represented as a volume density, but its materials are parameterized using a spatially varying Bidirectional Reflectance Distribution Function (BRDF). The volumetric representation of geometry has been shown to be effective for optimization [27, 44], and treating each point in space as a microfaceted surface allows us to use ideas stemming from decades of prior work on material parameterization and efficient surface-based rendering. Despite its volumetric parameterization, we verify experimentally that our model shrinks into a surface around opaque objects, with all contributions to the color of a ray coming from the vicinity of its intersection with the object.

To summarize, our method (1) combines aspects of volume-based and surface-based rendering for effective optimization, enabling reconstructing high-fidelity scene geometry, materials, and lighting from a set of calibrated images; (2) uses an optimizable microfacet material model rendered using Monte Carlo integration with multi-bounce raytracing, allowing for realistic interreflections on nonconvex objects; and (3) is efficient: it optimizes a scene from scratch in ~ 3 hours on a single NVIDIA GTX 3090.

2. Related work

Our work lies in the rich field of inverse rendering, in which the goal is to reconstruct the geometry, material properties, and illumination that gave rise to a set of observed images. This task is a severely underconstrained inverse problem, with challenges ranging from lack of differentiability [22] to the computational cost-variance tradeoff of the forward rendering process [17].

Recent progress in inverse rendering, and in particular in view synthesis, has been driven by modeling scenes as radiance fields [27, 3, 4], which can produce photorealistic models of a scene based on calibrated images.

Inverse rendering. Inverse rendering techniques can be categorized based on the combination of unknowns recovered and assumptions made. Common assumptions include far field illumination, isotropic BRDFs, and no in-

terreflections or self-occlusions. Early work by Ramamoorthi and Hanrahan [33] handled unknown lighting, texture, and BRDF by using spherical harmonic representations of both BRDF and lighting, which allowed recovering materials and low frequency lighting components. More recent methods used differentiable rendering of known geometry, first through differentiable rasterization [24, 23, 9] and later through differentiable ray tracing [22, 2, 31]. Later methods built on differentiable ray tracing, making use of Signed Distance Fields (SDFs) to also reconstruct geometry [46, 29, 17].

Following the success of NeRF [27], volumetric rendering has emerged as a useful tool for inverse rendering. Some methods based on volume rendering assume known lighting and only recover geometry and materials [37, 5, 1], while others solve for both lighting and materials, but assume known geometry [25, 49] or geometry without self-occlusions [7, 6]. Some methods simultaneously recover illumination, geometry, and materials, but assume that illumination comes from a single point light source [16, 45]. However, to the best of our knowledge none of these existing volumetric inverse rendering methods are able to capture high frequency lighting (and appearance of specular objects) from just the input images themselves.

An additional challenge is in handling multi-bounce illumination, or interreflections, in which light from a source bounces off multiple objects before reaching the camera. In this case, computational tradeoffs are unavoidable due to the exponential growth of rays with the number of such bounces. Park *et al.* [31] model interreflections assuming known geometry, but do not model materials, which is equivalent to treating all objects as perfect mirrors. Other methods use neural networks to cache visibility fields [37, 49] or radiance transfer fields [25, 16]. Our method handles interreflections by casting additional rays through the scene, using efficient Monte Carlo sampling.

Volumetric view synthesis. We base our representation of geometry on recent advances in volumetric view synthesis, following NeRF [27]. Specifically, we retain the idea of using differentiable volumetric rendering to model geometry, using a voxel-based representation of the underlying density field [14, 28, 8].

Of particular relevance to our work are prior methods such as [41, 15] that are specifically designed for high-fidelity appearance of glossy objects. In general, most radiance field models fail at rendering high-frequency appearance caused by reflections from shiny materials under natural illumination, instead rendering blurry appearance [47]. To enable our method to handle these highly specular materials, and to improve the normal vectors estimated by our method, which are key to the rendered appearance, we utilize regularizers from Ref-NeRF [41].

3. Preliminaries

Our method combines aspects of *volumetric* and *surface-based* rendering; we begin with a brief introduction to each before describing our method in Section 4.

3.1. Volume Rendering

The core idea in emission-absorption volume rendering is that light accumulates along rays, with “particles” along the ray both emitting and absorbing light. The color measured by a camera pixel corresponding to a ray with origin \mathbf{p}_c and direction $-\hat{\omega}_o$ is:

$$L(\mathbf{p}_c, \hat{\omega}_o) = \int_0^\infty T(t) \sigma(\vec{\mathbf{r}}(t)) L_o(\vec{\mathbf{r}}(t), \hat{\omega}_o) dt, \quad (1)$$

$$\text{where } T(t) = \exp\left(-\int_0^t \sigma(\vec{\mathbf{r}}(t')) dt'\right), \quad (2)$$

where $\vec{\mathbf{r}}(t) = \mathbf{p}_c - t\hat{\omega}_o$ is a camera ray, $\sigma(\mathbf{p})$ is the density at point \mathbf{p} in the volume, T denotes transmittance along the ray, and L_o is the outgoing radiance. This formula is often approximated numerically using quadrature, following [26]:

$$L(\mathbf{p}_c, \hat{\omega}_o) \approx \sum_{j=0}^{N-1} w_j L_o(\vec{\mathbf{r}}(t_j), \hat{\omega}_o), \quad (3)$$

$$\text{where } w_j = T_j (1 - \exp(-\sigma(\vec{\mathbf{r}}(t_j))(t_{j+1} - t_j))), \quad (4)$$

$$\text{and } T_j = \exp\left(-\sum_{k=0}^{j-1} \sigma(\vec{\mathbf{r}}(t_k))(t_{k+1} - t_k)\right). \quad (5)$$

In this volume rendering paradigm, multiple 3D points can contribute to the color of a ray, with nearer and denser points contributing most.

3.2. Surface Rendering

In surface rendering, and assuming fully-opaque surfaces, the color of a ray is determined solely by the light reflected by the first surface it encounters. Consider that the ray $\vec{\mathbf{r}}$ from camera position \mathbf{p}_c in direction $-\hat{\omega}_o$ intersects its first surface at a 3D position \mathbf{p} . The ray color is then:

$$L(\mathbf{p}_c, \hat{\omega}_o) = \int_{\mathbb{S}^2} f(\mathbf{p}, \hat{\omega}_o, \hat{\omega}_i) L_i(\mathbf{p}, \hat{\omega}_i) (\hat{\mathbf{n}}(\mathbf{p}) \cdot \hat{\omega}_i)^+ d\hat{\omega}_i, \quad (6)$$

where $\hat{\omega}_i$ is the direction of incident light, $\hat{\mathbf{n}}(\mathbf{p})$ is the surface normal at \mathbf{p} , f is the BRDF describing the material of the surface at \mathbf{p} , L_i is the incident radiance, and $(\hat{\mathbf{n}}(\mathbf{p}) \cdot \hat{\omega}_i)^+$ is a truncated cosine lobe (*i.e.* its negative values are clipped to zero) facing outward from the surface. Note that this equation is recursive: L_i inside the integral may be the outgoing radiance L_o coming from a different scene point. The integral in Equation 6 is also typically approximated by discrete (and often random) sampling, and is the subject of a rich body of work [11, 40].

4. Method

We present Neural Microfacet Fields to tackle the problem of inverse rendering by combining volume and surface rendering, as shown in Figure 2. Our method takes as input a collection of images (100 in our experiments) with known cameras, and outputs the volumetric density and normals, materials (BRDFs), and far-field illumination (environment map) of the scene. We assume that all light sources are infinitely far away from the scene, though light may interact locally with multiple bounces through the scene.

In this section, we describe our representation of a scene and the rendering pipeline we use to map this representation into pixel values. Section 4.1 introduces the main idea of our method, to build intuition before diving into the details. Section 4.2 describes our representation of the scene geometry, including density and normal vectors. Section 4.3 describes our representation of materials and how they reflect light via the BRDF. Section 4.4 introduces our parameterization of illumination, which is based on a far-field environment map equipped with an efficient integrator for faster evaluations of the rendering integral. Finally, Section 4.5 describes the way we combine these different components to render pixel values in the scene.

4.1. Main Idea

The key to our method is a novel combination of the volume rendering and surface rendering paradigms: we model a density field as in volume rendering, and we model outgoing radiance at every point in space using surface-based light transport (approximated using Monte Carlo ray sampling). Volume rendering with a density field lends itself well to optimization: initializing geometry as a semi-transparent cloud creates useful gradients (see Figure 3), and allows for changes in geometry and topology. Using surface-based rendering allows modeling the interaction of light and materials, and enables recovering these materials.

We combine these paradigms by modeling a *microfacet field*, in which each point in space is endowed with a volume density and a local micro-surface. Light accumulates along rays according to the volume rendering integral of Equation 1, but the outgoing light of each 3D point is determined by surface rendering as in Equation 6, using rays sampled according to its local micro-surface. This combination of volume-based and surface-based representation and rendering, shown in Figure 2, enables us to optimize through a severely underconstrained inverse problem, recovering geometry, materials, and illumination simultaneously.

4.2. Geometry Parameterization

We represent geometry using a low-rank tensor data structure based on TensorRF [8], with small modifications described in Appendix C. Our model stores both density




Figure 2: To render the color of a ray cast through the scene, we (a) evaluate density at each sample and compute each sample’s volume rendering quadrature weight w_i , then (b) query the material properties and surface normal (flipped if it does not face the camera) at each sample point, which are used to (c) compute the color of each sample by using Monte Carlo integration of the surface rendering integral, where the number of samples used is proportional to the quadrature weight w_i . This sample color is then accumulated along the ray using the quadrature weight to get the final ray color.

σ and a spatially-varying feature that is decoded into the material’s BRDF at every point in space. We initialize our model at low resolution and gradually upsample it during optimization (see Appendix C for details).

Similar to prior work [37, 41], we use the negative normalized gradient of the density field as a field of “volumetric normals.” However, like [20], we found that numerically computing spatial gradients of the density field using finite differences rather than using analytic gradients leads to normal vectors that we can use directly, without using features predicted by a separate MLP. Additionally, these numerical gradients can be efficiently computed using 2D and 1D convolution using TensorRF’s low-rank density decomposition (see Appendix C). These accurate normals are then used for rendering the appearance at a volumetric microfacet, as will be discussed in Sections 4.3 and 4.5.

Our volumetric normals are regularized using the orientation loss \mathcal{R}_o introduced by Ref-NeRF [41]:

$$\mathcal{R}_o = \sum_j w_j \max(0, -\hat{\mathbf{n}}(\mathbf{p}_j) \cdot \hat{\omega}_o)^2, \quad (7)$$

where $\hat{\omega}_o$ is the view direction facing towards the camera, and $\hat{\mathbf{n}}(\mathbf{p}_j)$ is the normal vector at the j th point along the ray. The orientation loss \mathcal{R}_o penalizes normals that face away from the camera yet contribute to the color of the ray (as quantified by weights w_j).

Because our volumetric normals are derived from the density field, this regularizer has a direct effect on the reconstructed geometry: it decreases the weight of backwards-facing normals by decreasing their density or increasing the density between them and the cameras, thereby promoting

hard surfaces and improving reconstruction. Note that unlike Ref-NeRF, we do not use “predicted normals” for surface rendering, as our Gaussian-smoothed derivative filter achieves similar effect.

4.3. Material Representation

We write our spatially varying BRDF model f as a combination of diffuse and specular components:

$$f(\hat{\omega}_o, \hat{\omega}_i) = \frac{\rho}{\pi} (1 - F_r(\hat{\mathbf{h}})) + F_r(\hat{\mathbf{h}}) f_s(\hat{\omega}_o, \hat{\omega}_i), \quad (8)$$

where ρ is the RGB albedo, $\hat{\mathbf{h}} = \frac{\hat{\omega}_o + \hat{\omega}_i}{\|\hat{\omega}_o + \hat{\omega}_i\|}$ is the half vector, $F_r(\hat{\mathbf{h}})$ is the Fresnel term, $f_s(\hat{\omega}_o, \hat{\omega}_i)$ is the specular component of the BRDF for outgoing view direction $\hat{\omega}_o$ and incident light direction $\hat{\omega}_i$. The spatial dependence of these terms on the point \mathbf{p} is omitted for brevity. We use the Schlick approximation [35] for the Fresnel term:

$$F_r(\hat{\mathbf{h}}) = F_0(\mathbf{p}) + (1 - F_0(\mathbf{p}))(1 - \hat{\mathbf{h}} \cdot \hat{\omega}_o)^5, \quad (9)$$

where $F_0(\mathbf{p}) \in [0, 1]^3$ is the spatially varying reflectance at the normal incidence at the point \mathbf{p} , and we base our specular BRDF on the Cook-Torrance BRDF [38], using:

$$f_s(\hat{\omega}_o, \hat{\omega}_i) = \frac{D(\hat{\mathbf{h}}; \alpha, \hat{\mathbf{n}}, \hat{\omega}_o) G_1(\hat{\omega}_o, \hat{\mathbf{h}}) g(\hat{\omega}_i, \hat{\omega}_o)}{4(\hat{\mathbf{n}} \cdot \hat{\omega}_o)(\hat{\mathbf{n}} \cdot \hat{\omega}_i)}, \quad (10)$$

where D is a Trowbridge-Reitz distribution [39] (popularized by the GGX BRDF model [42]), G_1 is the Smith shadow masking function for the Trowbridge-Reitz distribution, and g is a shallow multilayer perceptron (MLP)




Figure 3: **Snapshots of the toaster scene during optimization.** The second row shows a cross section of the weights along each ray taken along the dotted line. Early in training the object geometry is cloudy and the environment map is uniform, but as training proceeds the object develops a sharp surface and the environment map converges.

with a sigmoid nonlinearity at its output. The distribution D models the roughness α of the material, and it is used for importance sampling, as described in Section 4.5). The MLP g captures other material properties not included in its explicit components.

The parameters for each of the diffuse and specular BRDF components are stored as features \mathbf{x} in the TensorRF representation (alongside density σ), allowing them to vary in space. We compute the roughness α , albedo ρ and reflectance at the normal incidence F_0 by applying a single linear layer with sigmoid activation to the spatially-localized features \mathbf{x} . Details about the architecture of the MLP g and its input encoding can be found in Appendix B.

We can approximately evaluate the rendering equation integral in Equation 6 more efficiently by assuming all microfacets at a point \mathbf{p} have the same irradiance $E(\mathbf{p})$:

$$\int_{\mathbb{S}^2} f(\mathbf{p}, \hat{\omega}_o, \hat{\omega}_i) L_i(\mathbf{p}, \hat{\omega}_i) (\hat{\mathbf{n}}(\mathbf{p}) \cdot \hat{\omega}_i)^+ d\hat{\omega}_i \quad (11)$$

$$\approx \int_{\mathbb{S}^2} F_r(\hat{\mathbf{h}}) f_s(\mathbf{p}, \hat{\omega}_o, \hat{\omega}_i) L_i(\mathbf{p}, \hat{\omega}_i) (\hat{\mathbf{n}}(\mathbf{p}) \cdot \hat{\omega}_i)^+ d\hat{\omega}_i + \frac{\rho(\mathbf{p})}{\pi} E(\mathbf{p}) \int_{\mathbb{S}^2} (1 - F_r(\hat{\mathbf{h}})) d\hat{\omega}_i. \quad (12)$$

The irradiance $E(\mathbf{p})$, defined as:

$$E(\mathbf{p}) = \int_{\mathbb{S}^2} L_i(\mathbf{p}, \hat{\omega}_i) (\hat{\mathbf{n}}(\mathbf{p}) \cdot \hat{\omega}_i)^+ d\hat{\omega}_i, \quad (13)$$

can then be easily evaluated using an irradiance environment map approximated by low-degree spherical harmonics, as done by Ramamoorthi and Hanrahan [32]. At every optimization step, we obtain the current irradiance environment map by integrating the environment map with spherical harmonic functions up to degree 2, and combining the result with the coefficients of a clamped cosine lobe pointing in direction $\hat{\mathbf{n}}(\mathbf{p})$ to obtain the irradiance $E(\mathbf{p})$ [32]. Equation 12 can then be importance sampled according to D and integrated using Monte Carlo sampling of incoming light.

We sample half vectors $\hat{\mathbf{h}}$ from the distribution of visible normals $D_{\hat{\omega}_o}$ [18], which is defined as:

$$D_{\hat{\omega}_o}(\hat{\mathbf{h}}) = \frac{G_1(\hat{\omega}_o, \hat{\mathbf{h}}) |\hat{\omega}_o \cdot \hat{\mathbf{h}}| D(\hat{\mathbf{h}})}{|\hat{\omega}_o \cdot \hat{\mathbf{n}}|}. \quad (14)$$

However, to perform Monte Carlo integration of the rendering equation, we need to convert from half vector space to the space of incoming light, which requires multiplying by the determinant of the Jacobian of the reflection equation $\hat{\omega}_i = 2(\hat{\omega}_o \cdot \hat{\mathbf{h}})\hat{\mathbf{h}} - \hat{\omega}_o$, which is $4(\hat{\omega}_o \cdot \hat{\mathbf{h}})$ [42]. Multiplying by the $\hat{\omega}_i \cdot \hat{\mathbf{n}}$ term from Equation 6 as well as the

Jacobian and Equation 14 results in the following Monte Carlo estimate:

$$L(\mathbf{p}, \hat{\omega}_o) \approx \frac{1}{N} \sum_{n=1}^N F_r(\hat{\mathbf{h}}^n) g(\mathbf{x}, \hat{\omega}_o, \hat{\omega}_i^n) L_i(\mathbf{p}, \hat{\omega}_i^n) + (1 - F_r(\hat{\mathbf{h}}^n)) \frac{\rho(\mathbf{p})}{\pi} E(\mathbf{p}), \quad (15)$$

where $\hat{\mathbf{h}}^n \sim D_{\hat{\omega}_o}(\cdot; \alpha(\mathbf{p}), \hat{\mathbf{n}}(\mathbf{p}), \hat{\omega}_o)$,
and $\hat{\omega}_i^n = 2(\hat{\omega}_o \cdot \hat{\mathbf{h}}^n) \hat{\mathbf{h}}^n - \hat{\omega}_o$.

4.4. Illumination

We model far field illumination using an environment map, represented using an equirectangular image with dimensions $H \times W \times 3$, with $H = 512$ and $W = 1024$. We map the optimizable parameters in our environment map parameterization into high dynamic range RGB values by applying an elementwise exponential function.

We use the term *primary* to denote a ray originating at the camera, and *secondary* to denote a ray bounced from a surface to evaluate its reflected light (whether that light arrives directly from the environment map, or from another scene point).

To minimize sampling noise, instead of using a single environment map element per secondary ray, we use the mean value over an axis-aligned rectangle in spherical coordinates, with the solid angle covered by the rectangle adjusted to match the sampling distribution D at that point. Concretely, to query the environment map at a given incident light direction, we first compute its corresponding spherical coordinates (θ, ϕ) , where θ and ϕ are the polar and azimuthal angles respectively. We then compute the mean value of the environment map over a (spherical) rectangle centered at (θ, ϕ) , whose size $\Delta\theta \times \Delta\phi$ we constrain to have aspect ratio $\frac{\Delta\theta}{\Delta\phi} = \sin\theta$. To choose the solid angle of the rectangle, $\Delta\theta \cdot \Delta\phi$, we modify the method from [10], which is based on Nyquist’s sampling theorem (see Appendix A). We compute these mean values efficiently using integral images, also known as summed-area tables [13].

4.5. Rendering

In this section, we describe how the model components representing geometry (Section 4.2), materials (Section 4.3), and illumination (Section 4.4) are combined to render the color of a pixel. For each primary ray, we choose a set of sample points following the rejection-sampling strategy of [21, 28] to prioritize points near object surfaces. We query our geometry representation [8] for the density σ_j at each point, and use Equation 4 to estimate the contribution weight w_j of each point to the final ray color.

To compute the color for each of these points, we compute the irradiance from the environment map and apply

Equation 15 to obtain $L(\mathbf{p}_j, \hat{\omega}_o)$, as described in Section 4.3.

For each primary ray sample with weight w_j , we allocate $N = \lfloor w_j M \rfloor$ secondary rays, where $M = 128$ is an upper bound on the total number of secondary rays for each primary ray (since $\sum_j w_j \leq 1$). The secondary rays are sampled according to the Trowbridge-Reitz distribution using the normal vector $\hat{\mathbf{n}}(\mathbf{p}_j)$ and roughness value α_j at the current sample.

When computing the incoming light $L_i(\mathbf{p}, \hat{\omega}_i^n)$, we save memory by randomly selecting a fixed number R of secondary rays to interreflect through the scene while others index straight into the environment map. We importance sample these R secondary rays according to the largest channel of the weighted RGB multiplier $w_j \cdot g(\mathbf{x}_j, \hat{\omega}_i, \hat{\omega}_o)$. In practice, we find that adding a small amount of random noise to the weighted RGB multiplier before choosing the R largest values improves performance by slightly increasing the variation of the selected rays.

The remaining $N - R$ secondary rays with lower contribution are rendered more cheaply by evaluating the environment map directly rather than considering further interactions with the scene, as described in Section 4.4.

This combined sample color value $L(\mathbf{p}_j, \hat{\omega}_o)$ is then weighted by w_j , and the resulting colors are summed along the primary ray samples to produce pixel values. Finally, the resulting pixel values are tonemapped to sRGB color and clipped to $[0, 1]$.

Dynamic batching. We apply the dynamic batch size strategy from NeRFacc [21] to the TensorRF [8] sampler, which controls the number of samples per batch using the number of primary rays per batch. Since the number of secondary rays scales with the number of primary rays, we bound the maximal number of primary rays to avoid casting too many secondary rays. We use the same method to control R , the number of secondary bounces that retrace through the scene. During test time, we shuffle the image to match the training distribution, then unshuffle the image to get the result.

We train using photometric loss and the normal penalty loss of Equation 7. Further details of our optimization and sampling methods can be found in Appendix C.

5. Experiments

We evaluate our method using two synthetic datasets: the *Blender* dataset from NeRF [27], and the *Shiny Blender* dataset from Ref-NeRF [41]. Both datasets contain objects rendered against a white background; the *Shiny Blender* dataset focuses on shiny materials with high-frequency reflections whereas the *Blender* dataset contains a mixture of specular, glossy, and Lambertian materials.

We evaluate standard metrics PSNR, SSIM [43], and LPIPS [48] on the novel view synthesis task for each dataset. To quantify the quality of our reconstructed geometry, we also evaluate the Mean Angular Error (MAE°) of the normal vectors. MAE° is evaluated over the same test set of novel views; for each view we take the dot product between the ground truth and predicted normals, multiply the arccos of this angle by the ground truth opacity of the pixel, and take the mean over all pixels. This gives us an error of 90° if the predicted normals are missing (*i.e.* if the pixel is mistakenly predicted as transparent). We summarize these quantitative metrics in Table 1, with per-scene results in Appendix D.

We provide qualitative comparisons of our reconstructed environment maps with prior inverse rendering approaches in Figures 5 and 6, and additional results in Appendix D. We also demonstrate two applications of our scene decomposition in Figure 4. For Figure 4 (a), we render the geometry and spatially varying BRDF recovered from the *materials* scene with the environment map optimized from the *helmet* scene, showing convincing new reflections while retaining the original object shape and material properties. For Figure 4 (b), we take this a step further by training the *toaster* and *car* scenes with the same neural material decoder, which enabled us to compose the two scenes under the environment map recovered from the *toaster* scene.

To quantitatively measure the quality of our method’s disentanglement, we evaluate it similarly to NeRFactor [49]. Since our method is specifically designed to handle more specular objects, we modify the Shiny Blender dataset from Ref-NeRF [41] by rendering the images in HDR under both the original lighting and an unseen lighting condition. We then optimize each method on each of the scenes with original lighting, then render the estimated geometry and materials using the unseen lighting condition. For computing quality metrics, we find the (ambiguous) per-channel scaling factor by minimizing the mean squared error. The scaled image is then evaluated against the ground truth relit image using PSNR, SSIM, and LPIPS.

The results are presented in Table 2, showing that our method is able to relight shiny objects with significantly increased accuracy when compared to state of the art inverse rendering methods such as NVDiffRec and NVDiffRecMC.

Table 1 also shows quantitative ablation studies on our model. If we do not use integral images for the far field illumination (“no integral image”), the sparse gradient on the environment map prevents the model from learning to use it to explain reflections. Using the derivative of the linear interpolation of the density instead of smoothed numerical derivatives (“analytical grad”), results in more holes in the geometry, limiting performance. If we do not utilize multiple ray bounces for interreflections (“single bounce”), all metrics are slightly worse, with most of the errors arising in




Figure 4: **Rendering with different illumination.** (a) shows the geometry and BRDF (shown integrated against uniform white lighting) recovered from the *materials* scene, rendered with the environment map recovered from the *helmet* scene. (b) shows the geometry and BRDF recovered from the *toaster* and *car* scene composed under the environment map recovered from the *toaster* scene.

regions with strong interreflections. Finally, replacing the neural network with the identity function (“no neural”), results in slightly worse performance, especially in regions with strong interreflections.

6. Discussion and Limitations

We introduced a novel and successful approach for the task of inverse rendering, using calibrated images alone to decompose a scene into its geometry, far-field illumination, and material properties. Our approach uses a combination of volumetric and surface-based rendering, in which we en-




Figure 5: **Results on the *car* scene**, compared to NVDiffRec [29] and NVDiffRecMC [17].




Figure 6: **Results on the *helmet* scene**, compared to NVDiffRec [29] and NVDiffRecMC [17].

dow each point in space with both a density and a local microsurface, so that it can both occlude and reflect light from its environment. We verified experimentally that our method, which enjoys both the optimization landscape of

volume rendering, as well as the richness and efficiency of surface-based Monte Carlo rendering, provides superior results relative to prior work.

However, our method is not without limitations. First,

	Blender				Shiny Blender			
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	MAE $^\circ$ \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	MAE $^\circ$ \downarrow
PhySG ¹	18.54	.847	.182	29.17	26.21	.921	.121	8.46
NVDiffRec ¹	28.79	.939	.068	7.80	29.90	.945	.114	6.08
NVDiffRecMC ¹	25.81	.904	.111	7.09	28.20	.902	.175	4.93
Ref-NeRF ²	33.99	.966	.038	23.22	35.96	.967	.059	18.38
Ours, no integral image	28.47	.920	.069	7.96	27.33	.869	.170	14.10
Ours, analytical derivative	28.94	.926	.064	6.77	29.22	.900	.140	7.43
Ours, single bounce	30.68	.944	.045	5.73	34.39	.962	.053	2.81
Ours, no neural	29.40	.933	.057	6.11	33.00	.955	.063	3.40
Ours	30.71	.940	.053	5.70	34.56	.963	.053	2.78

¹ requires object masks during training. ² view synthesis method, not inverse rendering. Red is best, followed by orange, then yellow.

Table 1: **Results on the *Blender* dataset from NeRF [27], and *Shiny Blender* dataset from Ref-NeRF [41].** We compute PSNR, SSIM, and LPIPS on the novel view synthesis task, and MAE $^\circ$ on the normals. Our method outperforms all prior methods except Ref-NeRF on view synthesis, and produces the most accurate geometric normals.

PSNR \uparrow	toaster	coffee	helmet	ball	teapot	car	mean
NVDiffRec ¹	12.56	20.30	17.73	13.21	33.00	19.84	19.44
NVDiffRecMC ¹	14.55	23.39	20.44	13.23	32.51	20.09	20.70
Ours	20.05	24.21	25.69	23.50	34.37	23.19	25.17
LPIPS \downarrow	toaster	coffee	helmet	ball	teapot	car	mean
NVDiffRec ¹	.378	.208	.229	.467	.027	.113	.237
NVDiffRecMC ¹	.290	.177	.217	.423	.024	.105	.206
Ours	.170	.142	.129	.191	.017	.063	.119
SSIM \uparrow	toaster	coffee	helmet	ball	teapot	car	mean
NVDiffRec ¹	.624	.870	.824	.626	.985	.853	.797
NVDiffRecMC ¹	.741	.913	.861	.761	.986	.869	.855
Ours	.864	.910	.918	.897	.987	.917	.916

¹ requires object masks during training. Red is best, followed by orange, then yellow.

Table 2: **Relighting on the *Shiny Blender* dataset from Ref-NeRF [35].**

although it can handle non-convex geometry, it assumes far field illumination and thus performs poorly when this assumption is not satisfied. This issue is most clear in the *coffee* scene in the Shiny Blender dataset, which has near-field light sources. It also does not handle interreflections very well, since the number of secondary bounces is limited, and due to our acceleration scheme of often directly querying the environment map, as explained in Section 4.5. Our model also does not handle refractive media, which is most clear in the *drums* and *ship* scenes in the Blender dataset. Our diffuse lighting model also assumes far field illumination, and thus fails to fully isolate shadows from the albedo, most obvious in the *lego* scene in the Blender dataset. These scenes are visualized in Appendix D. Another limitation of our model’s BRDF parameterization is that it struggles to represent anisotropic materials. Finally, our method exhibits some speckle noise in its renderings, particularly near bright spots, which may be alleviated by using a denoiser as was used in [17]. We believe these limitations would make for interesting future work, as well as applying our method to other field representations and larger scenes captured in the wild.

7. Acknowledgements

We would like to thank Tzu-Mao Li for his advice, especially regarding the BRDF formulation, as well as his La Jolla renderer, which we used as a reference. AM is supported by ALERTCalifornia, developing technology to stay ahead of disasters, and the National Science Foundation under award #CNS-1338192, MRI: Development of Advanced Visualization Instrumentation for the Collaborative Exploration of Big Data, and Kinsella Expedition Fund. This material is partially based upon work supported by the National Science Foundation under Award No. 2303178 to SFK.

References

- [1] Meghna Asthana, William Smith, and Patrik Huber. Neural apparent brdf fields for multiview photometric stereo. In *Proceedings of the 19th ACM SIGGRAPH European Conference on Visual Media Production*, pages 1–10, 2022. 2
- [2] Dejan Azinović, Tzu-Mao Li, Anton Kaplanyan, and Matthias Nießner. Inverse path tracing for joint material and lighting estimation. In *Proc. Computer Vision and Pattern Recognition (CVPR), IEEE*, 2019. 2
- [3] 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 *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*, pages 5835–5844. IEEE, 2021. 2
- [4] Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 5460–5469. IEEE, 2022. 2
- [5] Sai Bi, Zexiang Xu, Pratul Srinivasan, Ben Mildenhall, Kalyan Sunkavalli, Miloš Hašan, Yannick Hold-Geoffroy, David Kriegman, and Ravi Ramamoorthi. Neural re-

- reflectance fields for appearance acquisition. *arXiv preprint arXiv:2008.03824*, 2020. 2
- [6] Mark Boss, Raphael Braun, Varun Jampani, Jonathan T Barron, Ce Liu, and Hendrik Lensch. Nerf: Neural radiance fields from image collections. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12684–12694, 2021. 2
- [7] Mark Boss, Varun Jampani, Raphael Braun, Ce Liu, Jonathan Barron, and Hendrik Lensch. Neural-pil: Neural pre-integrated lighting for radiance decomposition. *Advances in Neural Information Processing Systems*, 34:10691–10704, 2021. 2
- [8] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. In *European Conference on Computer Vision*, pages 333–350. Springer, 2022. 2, 3, 6
- [9] Wenzheng Chen, Huan Ling, Jun Gao, Edward Smith, Jaakko Lehtinen, Alec Jacobson, and Sanja Fidler. Learning to predict 3d objects with an interpolation-based differentiable renderer. *Advances in neural information processing systems*, 32, 2019. 2
- [10] Mark Colbert and Jaroslav Krivanek. Gpu-based importance sampling. *GPU Gems*, 3:459–476, 2007. 6, 12
- [11] Robert L Cook and Kenneth E. Torrance. A reflectance model for computer graphics. *ACM Transactions on Graphics (ToG)*, 1(1):7–24, 1982. 3
- [12] Roy Cranley and Thomas NL Patterson. Randomization of number theoretic methods for multiple integration. *SIAM Journal on Numerical Analysis*, 13(6):904–914, 1976. 12
- [13] Franklin C Crow. Summed-area tables for texture mapping. In *Proceedings of the 11th annual conference on Computer graphics and interactive techniques*, pages 207–212, 1984. 6
- [14] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 5491–5500. IEEE, 2022. 2
- [15] Wenhao Ge, Tao Hu, Haoyu Zhao, Shu Liu, and Ying-Cong Chen. Ref-NeuS: Ambiguity-reduced neural implicit surface learning for multi-view reconstruction with reflection, 2023. 2
- [16] Michelle Guo, Alireza Fathi, Jiajun Wu, and Thomas Funkhouser. Object-centric neural scene rendering. *arXiv preprint arXiv:2012.08503*, 2020. 2
- [17] Jon Hasselgren, Nikolai Hofmann, and Jacob Munkberg. Shape, Light, and Material Decomposition from Images using Monte Carlo Rendering and Denoising. *arXiv:2206.03380*, 2022. 2, 8, 9, 15, 16, 17, 18, 19, 20, 21
- [18] Eric Heitz. Sampling the ggx distribution of visible normals. *Journal of Computer Graphics Techniques (JCGT)*, 7(4):1–13, 2018. 5
- [19] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 12
- [20] Zhengfei Kuang, Kyle Olszewski, Menglei Chai, Zeng Huang, Panos Achlioptas, and Sergey Tulyakov. NeROIC: Neural rendering of objects from online image collections. *ACM Transactions on Graphics (TOG)*, 41(4):1–12, 2022. 4
- [21] Ruilong Li, Matthew Tancik, and Angjoo Kanazawa. Nerfacc: A general nerf acceleration toolbox. *arXiv preprint arXiv:2210.04847*, 2022. 6
- [22] Tzu-Mao Li, Miika Aittala, Frédo Durand, and Jaakko Lehtinen. Differentiable monte carlo ray tracing through edge sampling. *ACM Trans. Graph.*, 37(6), dec 2018. 2
- [23] Shichen Liu, Tianye Li, Weikai Chen, and Hao Li. Soft rasterizer: A differentiable renderer for image-based 3d reasoning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019. 2
- [24] Matthew M Loper and Michael J Black. Opendr: An approximate differentiable renderer. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part VII 13*, pages 154–169. Springer, 2014. 2
- [25] Linjie Lyu, Ayush Tewari, Thomas Leimkühler, Marc Habermann, and Christian Theobalt. Neural radiance transfer fields for relightable novel-view synthesis with global illumination. *arXiv preprint arXiv:2207.13607*, 2022. 2
- [26] N. Max. Optical models for direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 1(2):99–108, 1995. 3
- [27] 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. 1, 2, 6, 9, 14
- [28] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.*, 41(4):102:1–102:15, July 2022. 2, 6
- [29] Jacob Munkberg, Jon Hasselgren, Tianchang Shen, Jun Gao, Wenzheng Chen, Alex Evans, Thomas Mueller, and Sanja Fidler. Extracting Triangular 3D Models, Materials, and Lighting From Images. *arXiv:2111.12503*, 2021. 2, 8, 15, 16, 17, 18, 19, 20, 21
- [30] Art B Owen. Randomly permuted (t, m, s)-nets and (t, s)-sequences. In *Monte Carlo and Quasi-Monte Carlo Methods in Scientific Computing: Proceedings of a conference at the University of Nevada, Las Vegas, Nevada, USA, June 23–25, 1994*, pages 299–317. Springer, 1995. 12
- [31] Jeong Joon Park, Aleksander Holynski, and Steven M Seitz. Seeing the world in a bag of chips. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1417–1427, 2020. 2
- [32] Ravi Ramamoorthi and Pat Hanrahan. An efficient representation for irradiance environment maps. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 497–500, 2001. 5
- [33] Ravi Ramamoorthi and Pat Hanrahan. A signal-processing framework for inverse rendering. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 117–128, 2001. 2

- [34] Szymon M Rusinkiewicz. A new change of variables for efficient brdf representation. In *Eurographics Workshop on Rendering Techniques*, pages 11–22. Springer, 1998. 12
- [35] Christophe Schlick. An inexpensive brdf model for physically-based rendering. In *Computer graphics forum*, volume 13, pages 233–246. Wiley Online Library, 1994. 4
- [36] Ilya Meerovich Sobol. On the distribution of points in a cube and the approximate calculation of integrals. *Journal of Computational Mathematics and Mathematical Physics*, 7(4):784–802, 1967. 12
- [37] Pratul P Srinivasan, Boyang Deng, Xiuming Zhang, Matthew Tancik, Ben Mildenhall, and Jonathan T Barron. Nerv: Neural reflectance and visibility fields for relighting and view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7495–7504, 2021. 2, 4
- [38] Kenneth E Torrance and Ephraim M Sparrow. Theory for off-specular reflection from roughened surfaces. *Josa*, 57(9):1105–1114, 1967. 4
- [39] TS Trowbridge and Karl P Reitz. Average irregularity representation of a rough surface for ray reflection. *JOSA*, 65(5):531–536, 1975. 4
- [40] Eric Veach. *Robust Monte Carlo methods for light transport simulation*. Stanford University, 1998. 3
- [41] 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. 2, 4, 6, 7, 9, 12, 14
- [42] Bruce Walter, Stephen R Marschner, Hongsong Li, and Kenneth E Torrance. Microfacet models for refraction through rough surfaces. In *Proceedings of the 18th Eurographics conference on Rendering Techniques*, pages 195–206, 2007. 4, 5
- [43] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 7
- [44] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan, editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 4805–4815, 2021. 2
- [45] Kai Zhang, Fujun Luan, Zhengqi Li, and Noah Snavely. Iron: Inverse rendering by optimizing neural sdfs and materials from photometric images. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2022. 2
- [46] Kai Zhang, Fujun Luan, Qianqian Wang, Kavita Bala, and Noah Snavely. Physg: Inverse rendering with spherical gaussians for physics-based material editing and relighting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5453–5462, 2021. 2
- [47] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. Nerf++: Analyzing and improving neural radiance fields. *arXiv preprint arXiv:2010.07492*, 2020. 2
- [48] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. 7
- [49] Xiuming Zhang, Pratul P Srinivasan, Boyang Deng, Paul Debevec, William T Freeman, and Jonathan T Barron. Nerfactor: Neural factorization of shape and reflectance under an unknown illumination. *ACM Transactions on Graphics (TOG)*, 40(6):1–18, 2021. 2, 7