

# VMINer: Versatile Multi-view Inverse Rendering with Near- and Far-field Light Sources

Fan Fei<sup>1,2,4</sup> Jiajun Tang<sup>1,2,4</sup> Ping Tan<sup>3,4</sup> Boxin Shi<sup>1,2#</sup>

<sup>1</sup>National Key Laboratory for Multimedia Information Processing, School of Computer Science, Peking University

<sup>2</sup>National Engineering Research Center of Visual Technology, School of Computer Science, Peking University

<sup>3</sup>Hong Kong University of Science and Technology <sup>4</sup>Light Illusions

## Abstract

This paper introduces a versatile multi-view inverse rendering framework with near- and far-field light sources. Tackling the fundamental challenge of inherent ambiguity in inverse rendering, our framework adopts a lightweight yet inclusive lighting model for different near- and far-field lights, thus is able to make use of input images under varied lighting conditions available during capture. It leverages observations under each lighting to disentangle the intrinsic geometry and material from the external lighting, using both neural radiance field rendering and physically-based surface rendering on the 3D implicit fields. After training, the reconstructed scene is extracted to a textured triangle mesh for seamless integration into industrial rendering software for various applications. Quantitatively and qualitatively tested on synthetic and real-world scenes, our method shows superiority to state-of-the-art multi-view inverse rendering methods in both speed and quality.

## 1. Introduction

The reconstruction of 3D scenes from multi-view RGB imagery has experienced significant advancements following the development of Neural Radiance Fields (NeRF) [19], and both the speed and quality of the reconstruction have reached an unprecedented level [18, 20]. Despite these improvements, a common limitation is to represent only the radiance field, which is a complex product of the external lighting interacting with the intrinsic geometry and material of the scene [11]. The entangled nature of this representation generally hampers the ability to accurately render the scene under unseen lighting conditions, because the influence of the original lighting is embedded within the newly rendered scene. Recent approaches [24, 41, 46] have incorporated inverse rendering [22] to separate material properties and lighting effects, extending the application of the reconstruction beyond novel view synthesis to novel scenarios involving relighting and material editing [2]. Some

# Corresponding author. E-mail: shiboxin@pku.edu.cn.

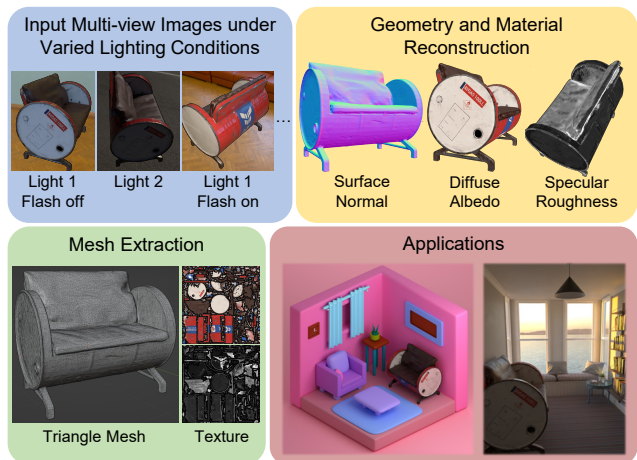


Figure 1. Given multi-view images under possibly varied lighting, our method leverages all present lighting conditions to reconstruct scene geometry and material disentangled with lighting. The trained fields are extracted to textured meshes for seamless integration into industrial renderers for various applications.

methods [2, 6, 9, 21, 32] have advanced to the extraction of detailed triangle meshes with material UV textures, which can serve as economical and lifelike 3D models for gaming and cinematography industries [9, 21], thereby marking a transformative step in digital asset creation.

Inverse rendering presents several fundamental challenges, one of them being its severe inherent ambiguity [41]. To combat this, multi-view inverse rendering methodologies usually reduce the ambiguity by imposing various constraints on each scene component. These methods can be categorized based on the assumption on the amount and types of lighting conditions present in the input images (Tab. 1). The majority of methods [9, 21, 24, 35, 36, 39, 41, 42] assumes the imagery to be captured under one fixed lighting condition (rows 1, 2). However, in such scenarios, only one case of scene appearances among all possible cases under different lighting is observed, causing severe **lighting-material ambiguity**, posing considerable difficulties to material estimation. In a bid to alleviate this ambiguity, recent techniques [2, 6, 10, 12, 23, 32] utilize input

Table 1. A summary of multi-view inverse rendering methods based on their assumptions and usages of input lighting conditions.

Methods	# of far-field light	# of near-field light	Material optimization	Lighting-material ambiguity	Capture workload
(1) [24, 39, 41, 42]	Single	None	Only from single	Strong	Low
(2) NeILF [35]	Single 5D incident light field				
(3) [2, 10, 12, 32] <sup>a</sup>	Multiple	None	Only from far-field	Moderate	High
(4) WildLight [6]	Single	Single	Only from near-field	Moderate	Low
(5) Ours	Single	Single	From near- and far-field	Weak	Low
	Multiple	Single or multiple		Almost none	High

<sup>a</sup> These methods also support input images under single far-field lighting. In that case, they belong to single far-field lighting methods (row 1).

images under varied lighting to discern materials from lighting (rows 3, 4). This demands more complex lighting models. Existing varied-lighting methods [2, 10, 12, 23, 32] typically model multiple, but only far-field, lighting conditions to maintain manageable unknown lighting parameters. This means throughout the capture, the lighting setups should be largely changed by either moving the objects to reconstruct or adapting far-field light sources, causing **increased workload** to the data collection process. Alternatively, near-field light sources, such as flashlights, can function as easily controllable light sources for disambiguation. WildLight [6] suggests the use of a camera-located flashlight in addition to the ambient lighting for inverse rendering. Nevertheless, it limits itself by using only the isolated appearance under the flashlight for material estimation, **not harnessing available observations** under the ambient lighting and degrading the quality of its estimated material (row 4).

Addressing the challenges and restrictions previously outlined, this paper introduces VMINer, a **Versatile Multi-view Inverse rendering framework with Near- and far-field light sources** (row 5). The distinctive feature of our framework is that it uses a lightweight yet inclusive lighting model for different far- and near-field light sources and efficiently leverages available observations under each lighting to disentangle lighting from other components, **making the most of whatever lighting conditions are at hand** to boost the practicality and quality of the reconstruction, as shown in Fig. 1. Although our method accepts the simplest setting of inputs under single far-field lighting, a more effective compromise is to employ an additional flashlight to enhance quality without overly complicating the capture process. The ideal scenario for the highest quality reconstruction would involve capturing the scene under varied far-field lighting and using a flashlight.

Nonetheless, implementing such a versatile framework is far from straightforward. The method has to render the scene under different lighting conditions, including spatially-varying and changing ones to represent possibly moving near-field lights. We propose to solve the problem by modeling **near-field lights as point sources** in 3D space with adjustable positions and intensities, which could be **anisotropic and not necessarily camera-aligned** (*e.g.*,

stationary desk lamps). This representation effectively approximates a broad range of common near-field lights and is easy for spatially-consistent editing and optimization. We use 3D implicit fields to model the scenes’ radiance, shape, and material and minimize the error of the neural or physically-based re-rendered scene appearance under each input lighting. To account for multiple lighting situations and moving near-field lights, we integrate specially tailored multi-layer perceptrons (MLPs) that additionally process light directions and embeddings. Post-training, the fields are converted into a textured triangle mesh, ready for **seamless integration** into industrial rendering software like Blender [1], facilitating a broad spectrum of applications. Our extensive experiments, covering both synthetic and real-world scenes, demonstrate that VMINer surpasses prior methods quantitatively and qualitatively.

## 2. Related Work

**Preliminary.** The rendering equation [11] models the reflected radiance from any point on a surface as a result of an intricate surface integral. This integral comprises contributions from three fundamental scene components: lighting, material, and geometry. The equation is expressed as:

$$L_o(\mathbf{x}, \boldsymbol{\omega}_o) = \int_{\Omega} \underbrace{L_i(\mathbf{x}, \boldsymbol{\omega}_i)}_{\text{lighting}} \underbrace{f(\mathbf{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o)}_{\text{material}} \underbrace{(\boldsymbol{\omega}_i \cdot \mathbf{n})^+}_{\text{geometry}} d\boldsymbol{\omega}_i. \quad (1)$$

In this equation,  $L_o(\mathbf{x}, \boldsymbol{\omega}_o)$  represents the radiance reflected in the outgoing direction  $\boldsymbol{\omega}_o$  from a surface point  $\mathbf{x}$  in 3D space.  $\Omega$  denotes the unit hemisphere encompassing all incident directions  $\boldsymbol{\omega}_i$  with  $\boldsymbol{\omega}_i \cdot \mathbf{n} > 0$ .  $L_i(\mathbf{x}, \boldsymbol{\omega}_i)$  is the incident radiance at  $\mathbf{x}$  from direction  $\boldsymbol{\omega}_i$ , and  $f(\mathbf{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o)$  is the bidirectional reflectance distribution function (BRDF), denoting the ratio of light reflected along  $\boldsymbol{\omega}_o$  at  $\mathbf{x}$  from  $\boldsymbol{\omega}_i$ . The term  $(\boldsymbol{\omega}_i \cdot \mathbf{n})^+$ , where  $x^+ \triangleq \max(x, 0)$ , signifies the cosine weakening factor. The incident lighting  $L_i(\mathbf{x}, \boldsymbol{\omega}_i)$  may originate from other points within the scene, necessitating recursive evaluation of this integral.

**Multi-view Inverse Rendering.** Multi-view inverse rendering techniques [2–4, 6, 9, 10, 12, 21, 23, 24, 34–36, 38–43] generally employ implicit or explicit 3D fields to model the scene’s geometry and material, while also estimating

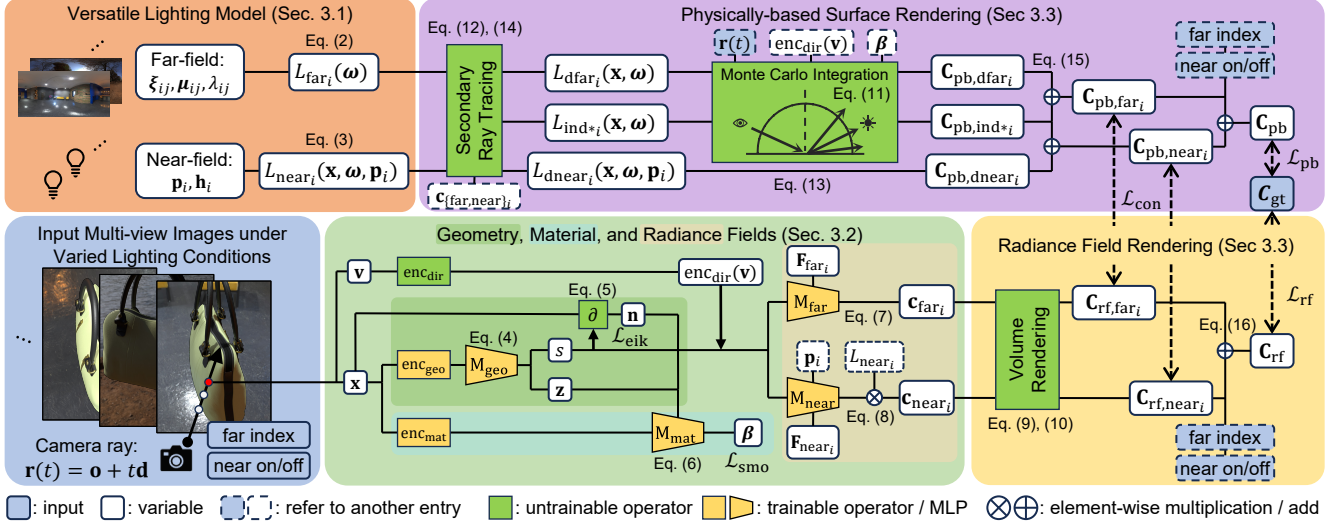


Figure 2. **Bottom-left:** VMIner takes as input multi-view RGB images with foreground masks under possibly varied far- and near-field lighting. The total number of lighting conditions and their types should be given. Each image is illuminated by one far-field lighting whose index is “far index”, and near-field lights whose on/off states are “near on/off”. **Top-left:** VMIner models each far-field lighting ( $\text{far}_i$ ) and near-field lighting ( $\text{near}_i$ ) separately as parametric models, from which incoming radiance is queried at any position  $\mathbf{x}$  and direction  $\omega$ . **Bottom-mid:** It utilizes 3D implicit representations for scene geometry, material, and radiance. Four MLPs process  $\mathbf{x}$  and view direction  $\mathbf{v}$  encoded by different encoders  $\text{enc}_{\{\text{geo}, \text{mat}, \text{dir}\}}$  with light embeddings  $\mathbf{F}_{\{\text{far}, \text{near}\}_i}$  to get the SDF  $s$ , the neural appearance descriptor  $\mathbf{z}$ , the BRDF parameters  $\beta$ , and the radiance  $\mathbf{c}_{\{\text{far}, \text{near}\}_i}$  under each lighting. **Bottom-right:** It then uses radiance field rendering to re-render the appearance  $\mathbf{C}_{\text{rf}, \{\text{near}, \text{far}\}_i}$  separately under each lighting condition, added up to  $\mathbf{C}_{\text{rf}}$  according to the per-image lighting condition. **Top-right:** Physically-based surface rendering uses Monte Carlo integration to evaluate both direct illumination  $L_{\text{d}\{\text{near}, \text{far}\}_i}$  with secondary visibility and indirect illumination  $L_{\text{ind}\{\text{near}, \text{far}\}_i}$ . The rendered appearances are aggregated to  $\mathbf{C}_{\text{pb}}$ .  $\mathbf{C}_{\{\text{rf}, \text{pb}\}}$  are compared with the observation  $\mathbf{C}_{\text{gt}}$  to train the scene model.  $\mathbf{C}_{\text{rf}, \{\text{near}, \text{far}\}_i}$  are also used as additional supervision signals for  $\mathbf{C}_{\text{pb}, \{\text{near}, \text{far}\}_i}$ .

lighting to recreate the scene’s appearance. They optimize each scene component from scratch by aligning the rendered with the observed appearance, while avoiding potential local minima, which, although do not faithfully represent the scene, can replicate the observed input. Addressing this inherently ambiguous nature of inverse rendering, these methods typically impose constraints on each scene component. For geometry and material, assumptions of known geometry [35, 36, 46] are made, along with leveraging the recovered shape from radiance fields [6, 10, 41], adopting Lambertian or uniform materials [23, 39], controlling BRDF parameter smoothness [10], or integrating a learned BRDF latent space [41]. Regarding lighting, many approaches presuppose known lighting conditions [24, 40] or fixed lighting shared across input images [35, 38, 41, 42]. More recent techniques [2, 6, 10, 12, 23, 32] have turned to using images with varied lighting conditions to mitigate the lighting-material ambiguity. They typically accept input images that can be grouped into sets, each illuminated by one distinct and fixed far-field lighting. WildLight [6] exceptionally employs a camera-located flashlight along with fixed ambient lighting. However, it does not model the ambient lighting and thus its material property estimation relies exclusively on flashlight observations, where the light and view direction always coincide [5], limiting its BRDF estimation ability. Our VMIner, in contrast, accommodates

input images under single or multiple far-field lighting conditions with near-field light sources and leverages each of them to enhance the quality of the estimated material.

**Lighting Models.** Spatially-uniform (SU) lighting models rely on the assumption that the scene’s lighting originates from a distant source, modeled as spherical Gaussians (SGs) [2, 10], spherical harmonics (SH) [12, 23], or MLPs [32]. In contrast, spatially-varying (SV) lighting models acknowledge the presence of near-field light sources, causing different locations to receive different incident lighting. They include parametric 3D lights [7, 17, 36, 37], outgoing light fields such as volumetric SGs [31] and neural out-of-view lighting volumes [45], and incident light fields like SV environment maps [47], SVSGs [16], and neural incident light fields (NeILF) [35, 38]. Our VMIner integrates distant lights as SGs and near-field lights as point lights with adjustable position and intensity. It is less complex compared to other SV lighting models, facilitating spatially consistent editing and optimization, while being able to approximate common lighting setups in real life.

### 3. Proposed Method

As illustrated in Fig. 2, VMIner reconstructs the scene lighting (Sec. 3.1) and 3D fields (Sec. 3.2) from input images using differentiable rendering (Sec. 3.3). The training scheme and loss functions are described in Sec. 3.4.

### 3.1. Versatile Lighting Model

This subsection models the lighting term in Eq. (1) and corresponds to the top-left part of Fig. 2.

VMiner harnesses the diversity of lighting conditions present in the input RGB images to disentangle lighting effects. The lighting model plays a crucial role here: 1) Since lighting conditions are unknown initially, they must be optimized alongside other scene parameters using physically-based rendering during training. 2) Different lighting setups, each compatible with the model, are necessary during the capture phase to achieve this goal. Given these prerequisites, it is essential to design a lighting model that accommodates common far-field and near-field light sources, while being lightweight enough for efficient and robust optimization and physically-based rendering. With these considerations, we propose a versatile lighting model that is both inclusive and lightweight. This representation models far-field and near-field lighting separately as below.

**Far-field Lighting.** In our model, each far-field lighting condition  $i$ , where  $i$  belongs to the set  $\mathbb{N}_{N_{\text{far}}} \triangleq \{1, \dots, N_{\text{far}}\}$  and  $N_{\text{far}}$  denotes the number of far-field lighting, is represented using 128-lobe SGs [28] commonly adopted in existing far-field lighting techniques [33, 39, 42]. For each lobe  $j$  of lighting  $i$ , there are six parameters: lobe axis  $\xi_{ij} \in \mathbb{S}^2$ , lobe sharpness  $\lambda_{ij} \in \mathbb{R}_+$ , and lobe RGB amplitude  $\mu_{ij} \in \mathbb{R}_+^3$ . The incident radiance from far-field lighting  $i$  along direction  $\omega$  at any 3D position  $\mathbf{x}$  is calculated using the formula (ignoring visibility for now):

$$L_{\text{far}_i}(\omega) = \sum_{j=1}^{128} c_{ij} \mu_{ij} e^{\lambda_{ij}(\omega \cdot \xi_{ij} - 1)}, \quad (2)$$

where  $c_{ij} = \lambda_{ij} / (2\pi(1 - e^{-2\lambda_{ij}}))$  acts as the normalization factor related to roughness. Notably,  $\mathbf{x}$  does not appear as an input to  $L_{\text{far}_i}$ , reflecting the nature of distant lighting.

**Near-field Lighting.** Each near-field lighting  $i$ , included in the set  $\mathbb{N}_{N_{\text{near}}}$  with  $N_{\text{near}}$  representing the number of near-field light sources, is modeled as a point light. These point lights can have moving positions  $\mathbf{p}_i \in \mathbb{R}^3$  and exhibit anisotropic radiation characterized by  $l^{\text{th}}$ -order ( $l$  can be 0, 1, or 2) SH coefficients  $\mathbf{h}_i \in \mathbb{R}^{3 \times (l+1)^2}$ . Our method accommodates two types of near-field lighting: camera-located lights, positioned at the camera ray origin for each image, and stationary lights, which remain fixed across all images. We observe that in neural radiance field rendering, the radiance under a stationary near-field light, especially when the light is active in all images, can be challenging to distinguish from radiance under far-field lighting. Also, the radiance from these two sources aids in material estimation in a similar way. Therefore, in practical applications, we favor using a moving flashlight as the near-field light to provide unique information for material estimation. The incident radiance from near-field lighting  $i$  at a 3D po-

sition  $\mathbf{x}$  is computed as (also ignoring visibility for now):

$$L_{\text{near}_i}(\mathbf{x}, \omega, \mathbf{p}_i) = \begin{cases} \frac{\text{SH}(\omega; \mathbf{h}_i)}{\|\mathbf{p}_i - \mathbf{x}\|_2^2} & \text{if } \omega = \frac{\mathbf{p}_i - \mathbf{x}}{\|\mathbf{p}_i - \mathbf{x}\|_2}, \\ 0 & \text{otherwise} \end{cases}, \quad (3)$$

where  $\text{SH}(\omega; \mathbf{h})$  calculates the SH at direction  $\omega$  with coefficients  $\mathbf{h}$ , and  $1/\|\mathbf{p}_i - \mathbf{x}\|_2^2$  signifies the inverse-square lighting attenuation for point lights. Here we include  $\mathbf{p}_i$  as an input because it may be set differently across images. It is important to note that with respect to the incident direction  $\omega$ ,  $L_{\text{far}}$  is a continuous function, while  $L_{\text{near}}$  is a discrete function, being non-zero only in a single direction. As a result, to render appearances under direct lighting, Monte Carlo integration is essential for far-field lighting, while a simple multiplication suffices for near-field lighting. Further details about this process are provided in Sec. 3.3.

### 3.2. Geometry, Material, and Radiance Fields

This subsection models the material and geometry terms in Eq. (1) and corresponds to the bottom-mid part of Fig. 2.

**Geometry.** Multi-view reconstruction methodologies generally hinge on two geometry representations: volume density [2, 10, 19, 41] and the signed distance function (SDF) [33, 39, 42, 44] fields. We choose the SDF field for our geometry representation due to its clearly defined surface at the zero-level isosurface, which simplifies and enhances the post-training mesh extraction process. Our approach utilizes implicit fields as MLPs over explicit structures like voxel grids for their compactness and flexibility. For each 3D position  $\mathbf{x}$ , a multi-resolution hash grid [20] is first employed for positional encoding, yielding a feature vector  $\text{enc}_{\text{geo}}(\mathbf{x}) \in \mathbb{R}^{16}$ . The geometry MLP  $M_{\text{geo}}$  then predicts the SDF value at  $\mathbf{x}$ :

$$s(\mathbf{x}), \mathbf{z}(\mathbf{x}) = M_{\text{geo}}(\mathbf{x}, \text{enc}_{\text{geo}}(\mathbf{x})), \quad (4)$$

where  $s \in \mathbb{R}$  represents the signed distance (positive outside the surface, negative inside), and  $\mathbf{z} \in \mathbb{R}^{13}$  is a descriptor of local appearance. To maintain differentiability in rendering and to aid in geometry optimization, during training we apply NeuS [29] techniques to transform SDF values along rays into volume densities. The surface normal  $\mathbf{n} \in \mathbb{S}^2$  is derived as the gradient of the SDF  $s$  with respect to  $\mathbf{x}$ :

$$\mathbf{n}(\mathbf{x}) = \frac{\partial s}{\partial \mathbf{x}} \bigg/ \left\| \frac{\partial s}{\partial \mathbf{x}} \right\|_2. \quad (5)$$

$M_{\text{geo}}$  only approximates a strict SDF field that has  $\left\| \frac{\partial s}{\partial \mathbf{x}} \right\|_2 = 1$ , so we normalize  $\mathbf{n}$  to ensure it is a unit vector.

**Material.** For accurate scene material recovery using physically-based re-rendering, it is imperative to model the material in such a way that the renderer can query the BRDF at any point on the surface. Considering that material properties are independent of lighting and view directions, we represent the scene’s material through an implicit 3D field, much like the SDF field utilized for geometry. To achieve

this, we employ another multi-resolution hash grid  $\text{enc}_{\text{mat}}$  for positional encoding. We adopt the simplified GGX BRDF model [27] with a fixed fresnel parameter to help alleviate ambiguity. The material MLP,  $M_{\text{mat}}$ , is tasked with predicting the SVBRDF parameters  $\beta(\mathbf{x})$ :

$$\beta(\mathbf{x}) = M_{\text{mat}}(\mathbf{x}, \mathbf{z}, \mathbf{n}, \text{enc}_{\text{mat}}(\mathbf{x})). \quad (6)$$

**Radiance.** Although our method does not strictly require a radiance field for rendering scene appearance – given that appearance could be rendered solely using physically-based surface rendering (PBR) – we find that incorporating a radiance field at the start of training significantly eases the optimization of scene geometry and appearance. PBR, while physically accurate, tends to introduce instability and slow down the training process due to its inherent ambiguity and computational intensity. Moreover, the radiance field-rendered results can serve as supplementary supervision, aiding PBR in better separating scene appearance under each lighting condition and thus further diminishing lighting-material ambiguity. Hence, we deploy radiance MLPs to predict view-dependent and lighting-dependent radiance at each position. For each lighting condition  $i$ , a light embedding  $\mathbf{F}_{\{\text{far}, \text{near}\}_i} \in \mathbb{R}^{16}$  is utilized. The view direction  $\mathbf{v} \in \mathbb{S}^2$  is processed through directional encoding that projects it onto the coefficients of the 3<sup>rd</sup>-order SH basis, yielding  $\text{enc}_{\text{dir}}(\mathbf{v}) \in \mathbb{R}^{4 \times 16}$  [20]. The far-field radiance MLP  $M_{\text{far}}$  computes the outgoing radiance  $\mathbf{c}_{\text{far}_i}$  from position  $\mathbf{x}$  along view direction  $\mathbf{v}$  under far-field lighting  $i$ :

$$\mathbf{c}_{\text{far}_i}(\mathbf{x}, \mathbf{v}) = M_{\text{far}}(\mathbf{x}, \mathbf{z}, \mathbf{n}, \text{enc}_{\text{dir}}(\mathbf{v}), \mathbf{F}_{\text{far}_i}), \quad (7)$$

where we incorporate the surface normal  $\mathbf{n}$  as an additional MLP input, following Instant-NSR [44] and WildLight [6], as this has shown to aid shape recovery.

The near-field radiance MLP differs slightly, as near-field lighting can move in different images, thus the current light position  $\mathbf{p}_i$  is a necessary input. Additionally, the radiance under near-field point light takes the form of the rendering equation Eq. (1) without the integral, allowing the incoming radiance and cosine term to be explicitly included. We follow ReNe [25] to input the relative direction of the point light  $\boldsymbol{\omega}_i = \frac{\mathbf{p}_i - \mathbf{x}}{\|\mathbf{p}_i - \mathbf{x}\|_2}$  instead of the absolute position  $\mathbf{p}_i$  into the MLP. Therefore, the outgoing neural radiance  $\mathbf{c}_{\text{near}_i}$  under near-field lighting  $i$  is calculated as:

$$\begin{aligned} \mathbf{c}_{\text{near}_i}(\mathbf{x}, \mathbf{v}, \mathbf{p}_i) &= M_{\text{near}}(\mathbf{x}, \mathbf{z}, \mathbf{n}, \text{enc}_{\text{dir}}(\mathbf{v}), \mathbf{F}_{\text{near}_i}, \boldsymbol{\omega}_i) \\ &\quad \otimes \frac{L_{\text{near}_i}(\mathbf{x}, \boldsymbol{\omega}_i)}{\|\mathbf{p}_i - \mathbf{x}\|_2^2} (\boldsymbol{\omega}_i \cdot \mathbf{n})^+, \end{aligned} \quad (8)$$

where  $\otimes$  denotes element-wise multiplication. On the right side, the second term denotes the incident radiance attenuated by square distance, the third term denotes the cosine weakening factor, and  $M_{\text{near}}$  accounts for other factors like reflectance and visibility. Owing to the instability of these terms at the training’s outset, we use an annealing strategy,

as recommended in prior works [44], gradually replacing default values (e.g., 1 for incident radiance) with those optimized during training, to ease through the process.

### 3.3. Differentiable Rendering

To simplify the notation, we omit  $\mathbf{p}_i$  and the lighting index  $i$  in the equations within this subsection. Our approach optimizes the lighting alongside the implicit fields primarily by reducing the discrepancy between the observed and differentially re-rendered appearances. The re-rendering is accomplished using two types of renderer: a neural radiance field renderer (Fig. 2 bottom-right) and a physically-based surface renderer (Fig. 2 top-right; Eq. (1)). We handle each lighting condition separately in rendering, and then aggregate the results based on per-image lighting condition.

**Volume Rendering for SDF Field.** Consider a camera ray  $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$  ( $t > 0$ ), with  $\mathbf{o} \in \mathbb{R}^3$  as its origin and  $\mathbf{d} \in \mathbb{S}^2$  as its direction. We incorporate importance sampling using an occupancy grid [14] to get  $N \leq 1024$  points along this ray, denoted as  $\mathbf{r}(t_i), i \in \mathbb{N}_N$ . Utilizing techniques from NeuS [29], we transform the signed distances  $s(\mathbf{r}(t_i))$  into discrete opacity values as follows:

$$\alpha_i = \left( \frac{\Phi_b(s(\mathbf{r}(t_i))) - \Phi_b(s(\mathbf{r}(t_{i+1})))}{\Phi_b(s(\mathbf{r}(t_i)))} \right)^+, \quad (9)$$

where  $\Phi_b(s) = 1/(1 + e^{-bs})$  is the cumulative opacity distribution function. The parameter  $b$  is trainable and tends to increase during training, focusing opacity more narrowly around the surface where  $s = 0$ . The accumulated transmittance is then calculated as  $T_i = \prod_{j=1}^{i-1} (1 - \alpha_j)$ .

**Neural Radiance Field Rendering.** The color  $\mathbf{C}_{\text{rf}}$  of a camera ray under either far- or near-field lighting conditions, as rendered through the neural radiance field, is determined by accumulating the radiance along the ray’s path:

$$\mathbf{C}_{\text{rf}, \{\text{far}, \text{near}\}}(\mathbf{o}, \mathbf{d}) = \sum_{i=1}^N T_i \alpha_i \mathbf{c}_{\{\text{far}, \text{near}\}}(\mathbf{r}(t_i), -\mathbf{d}). \quad (10)$$

This formula sums the radiance from each sampled point to yield the color of the ray under the specified lighting.

**Physically-based Rendering.** Our method employs differentiable surface rendering, modeling the observed radiance as reflected from a single surface point. We compute the depth  $t$  of this surface point similarly to Eq. (10):  $t = \sum_{i=1}^N T_i \alpha_i t_i$ . Accumulation along the ray also applies to the BRDF parameters  $\beta$  and the surface normal  $\mathbf{n}$ . For each input lighting, we render the appearance that models secondary visibility and indirect illumination. The rendering method under direct illumination differs for far- and near-field lighting. For direct far-field lighting  $L_{\text{dfar}}$ , we employ Monte Carlo integration to evaluate the rendering equation (Eq. (1)) as the incident light comes from all

directions:

$$\begin{aligned} \mathbf{C}_{\text{pb,dfar}}(\mathbf{o}, \mathbf{d}) &= L_{\text{o}}(\mathbf{r}(t), -\mathbf{d}) \\ &= \frac{1}{S} \sum_{s=1}^S \frac{L_{\text{dfar}}(\mathbf{r}(t), \boldsymbol{\omega}_s) f(\mathbf{r}(t), \boldsymbol{\omega}_s, -\mathbf{d})(\boldsymbol{\omega}_s \cdot \mathbf{n})^+}{p(\boldsymbol{\omega}_s)}, \end{aligned} \quad (11)$$

where  $S$  is the number of sampled directions,  $\boldsymbol{\omega}_s$  a sampled incident direction, and  $p(\boldsymbol{\omega}_s)$  the probability density function (PDF) for the sampled direction. We use multiple importance sampling [26] to combine evaluations from different strategies: BRDF importance sampling [13], SG lighting importance sampling [33], and cosine importance sampling. we follow differentiable rendering works [15] to use a small  $S = 20$ , as stochastic gradient descent handles noisy gradients. The incident radiance, now taking the visibility of light sources into account, is given by:

$$L_{\text{dfar}}(\mathbf{r}(t), \boldsymbol{\omega}_s) = L_{\text{far}}(\boldsymbol{\omega}_s) V(\mathbf{r}(t), \boldsymbol{\omega}_s), \quad (12)$$

with  $V(\mathbf{r}(t), \boldsymbol{\omega}_s) \in [0, 1]$  representing visibility, or the inverse of the accumulated opacity along a secondary ray  $\mathbf{r}_{\text{sec}}(t') = \mathbf{r}(t) + t'\boldsymbol{\omega}_s$ , ( $t' > 0$ ), computed in a manner akin to Eq. (10). For direct near-field lighting  $L_{\text{dnear}}$ , the integral in Eq. (1) simplifies to a multiplication, as the light originates from a single direction:

$$\mathbf{C}_{\text{pb,dnear}}(\mathbf{o}, \mathbf{d}) = L_{\text{dnear}}(\mathbf{r}(t), \boldsymbol{\omega}) f(\mathbf{r}(t), \boldsymbol{\omega}, -\mathbf{d})(\boldsymbol{\omega} \cdot \mathbf{n})^+, \quad (13)$$

where  $\boldsymbol{\omega} = (\mathbf{p} - \mathbf{r}(t)) / \|\mathbf{p} - \mathbf{r}(t)\|_2$  is the direction of incident light, and  $L_{\text{dnear}}(\mathbf{r}(t), \boldsymbol{\omega})$  is the potentially occluded incident radiance, computed similarly to Eq. (12).

Indirect illumination  $L_{\text{ind}}$ , relevant to both far-field and near-field lighting, considers radiance reflected from the scene itself, potentially from all directions, thus necessitating Monte Carlo integration. The neural radiance field  $\mathbf{C}_{\text{rf}}$  substitutes multi-bounce path tracing for indirect illumination, given its encapsulation of the scene radiance under infinite lighting bounces. The radiance  $\mathbf{C}_{\text{pb,ind}\{\text{far, near}\}}$  from indirect illumination is evaluated similarly to Eq. (12), but with the  $L_{\text{dfar}}$  replaced by the indirect lighting  $L_{\text{ind}\{\text{far, near}\}}$ :

$$L_{\text{ind}*}(\mathbf{r}(t), \boldsymbol{\omega}_s) = \mathbf{c}_*(\mathbf{x}', -\boldsymbol{\omega}_s)(1 - V(\mathbf{r}(t), \boldsymbol{\omega}_s)), \quad (14)$$

where  $\mathbf{x}'$  is the intersection of the secondary ray  $\mathbf{r}_{\text{sec}}(t')$  with the scene geometry. The radiance from indirect illumination is also computed for each lighting condition separately, then combined with direct illumination to yield the complete PBR radiance for the camera ray under a specific lighting condition:

$$\mathbf{C}_{\text{pb,}\{\text{far, near}\}} = \mathbf{C}_{\text{pb,d}\{\text{far, near}\}} + \mathbf{C}_{\text{pb,ind}\{\text{far, near}\}}. \quad (15)$$

### 3.4. Training

**Training Schemes.** Our training process is divided into two sequential stages. In the first stage, the focus is on training the geometry and radiance field without employ-

ing PBR. The objectives are: 1) to recover the scene’s geometry, 2) to distinguish the appearance under each lighting condition, a necessity due to the potential presence of multiple light sources in one image, and 3) to utilize the reconstructed radiance field for indirect illumination. In the second stage, we train the material field and scene lighting using PBR, aiming to estimate material properties.

**Loss Functions.** We average the loss functions across batches of camera rays. During the first stage, our loss function compares the total re-rendered radiance under all lighting conditions present in an image – including a far-field light  $i$  and any active near-field lights – with the ground-truth (GT) color  $\mathbf{C}_{\text{gt}}$  from the input observation. This comparison trains the geometry and radiance field:

$$\mathcal{L}_{\text{rf}} = \|\mathbf{C}_{\text{rf, far}_i} + \sum_{\text{near } i \text{ is on}} \mathbf{C}_{\text{rf, near}_i} - \mathbf{C}_{\text{gt}}\|_2^2. \quad (16)$$

The Eikonal loss [8] is applied to the gradients of the SDF values  $s$  for geometric regularization:

$$\mathcal{L}_{\text{eik}} = \sum_{j=1}^N T_j \alpha_j \left( \left\| \frac{\partial s(\mathbf{r}(t_j))}{\partial \mathbf{r}(t_j)} \right\|_2 - 1 \right)^2, \quad (17)$$

A silhouette loss using GT alpha  $\alpha_{\text{gt}}$  from the foreground mask and a normal smoothness loss [41] aid shape recovery:

$$\mathcal{L}_{\text{sil}} = ((1 - T_{N+1}) - \alpha_{\text{gt}})^2. \quad (18)$$

$$\mathcal{L}_{\text{ns}} = \sum_{j=1}^N T_j \alpha_j (\mathbf{n}(\mathbf{r}(t_j)) - \mathbf{n}(\mathbf{r}(t_j) + \boldsymbol{\epsilon}))^2, \quad (19)$$

where  $\boldsymbol{\epsilon} \in \mathbb{R}^3$  is a small random perturbation. The total loss for stage one is  $\mathcal{L}_{\text{rf}} + \lambda_{\text{eik}} \mathcal{L}_{\text{eik}} + \lambda_{\text{sil}} \mathcal{L}_{\text{sil}} + \lambda_{\text{ns}} \mathcal{L}_{\text{ns}}$ , with  $\lambda_*$  representing the loss weights. In stage two, PBR colors replace radiance field-rendered colors for comparison with  $\mathbf{C}_{\text{gt}}$ , and a novel self-consistency loss is introduced between neural and PBR radiance under each lighting for additional supervision. The direct supervision loss  $\mathcal{L}_{\text{pb}}$  mirrors  $\mathcal{L}_{\text{rf}}$ , with  $\mathbf{C}_{\text{rf,}\{\text{far, near}\}_i}$  replaced by  $\mathbf{C}_{\text{pb,}\{\text{far, near}\}_i}$ . The per-lighting self-consistency loss  $\mathcal{L}_{\text{con}}$  is defined as:

$$\mathcal{L}_{\text{con}} = \sum_{i=1}^{N_{\{\text{far, near}\}}} \|\mathbf{C}_{\text{rf,}\{\text{far, near}\}_i} - \mathbf{C}_{\text{pb,}\{\text{far, near}\}_i}\|_2^2. \quad (20)$$

This loss is crucial for discerning contributions from each lighting in images with multiple active light sources, thereby reducing ambiguity and improving material estimation. A material smoothness loss  $\mathcal{L}_{\text{ms}}$  is used similar to  $\mathcal{L}_{\text{ns}}$ . The total loss for stage two is  $\mathcal{L}_{\text{pb}} + \lambda_{\text{con}} \mathcal{L}_{\text{con}} + \lambda_{\text{ms}} \mathcal{L}_{\text{ms}}$ .

Our model is trained on an NVIDIA RTX 3090 GPU for a total of 40,000 steps. The first 30,000 steps of stage one take about 20 minutes, while the subsequent 10,000 steps of stage two require about 40 minutes. Post-training, we follow a procedure similar to WildLight [6] to extract the fields into textured meshes, which can be easily integrated into industry-standard rendering software like Blender [1],

Table 2. Quantitative comparison results with state-of-the-art methods averaged on 6 synthetic scenes. We show results of surface normal, diffuse albedo, view synthesis RGB, free-viewpoint (FV) relit RGB, the specular reflection part of FV relit RGB, and training time on a single RTX 3090 GPU. We mark the **best** and the second best results in each column.  $\uparrow$  ( $\downarrow$ ) means bigger (smaller) is better.

Method	Input lighting conditions	Normal	Albedo		View synthesis		FV relit		FV relit (spec)		Time
		MAnGE $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	PSNR $\uparrow$	SSIM $\uparrow$	PSNR $\uparrow$	SSIM $\uparrow$	PSNR $\uparrow$	SSIM $\uparrow$	
(1) TensoIR [10]		17.66	26.48	0.921	29.48	0.912	28.18	0.901	28.30	0.861	300 mins
(2) NVDiffRecMC [9]	Single far-field	16.24	26.52	0.915	27.13	0.913	26.61	0.901	25.57	0.836	150 mins
(3) Ours		12.39	24.50	0.882	28.20	0.934	27.46	0.921	27.56	0.871	45 mins
(4) WildLight [6]	Single far-field + Flashlight	11.49	28.86	0.940	29.87	0.929	30.44	0.930	27.71	0.863	1440 mins
(5) Ours		<u>10.89</u>	<u>31.62</u>	<u>0.953</u>	<u>32.09</u>	<b>0.953</b>	<u>32.00</u>	<u>0.953</u>	<u>30.75</u>	<u>0.906</u>	60 mins
(6) TensoIR [10]	Two far-field	16.24	27.18	0.929	29.68	0.912	28.66	0.902	28.46	0.863	300 mins
(7) Ours		11.70	26.07	0.902	29.59	0.942	29.27	0.934	29.09	0.890	45 mins
(8) Ours	Two far + Single near	<b>10.79</b>	<b>32.04</b>	<b>0.957</b>	<b>32.10</b>	<u>0.950</u>	<b>32.38</b>	<b>0.954</b>	<b>31.40</b>	<b>0.910</b>	60 mins

enabling fast and high-quality rendering suitable for various applications. Check the supplement for more details.

## 4. Experiments

Our experiments involve a comprehensive comparison with state-of-the-art methods (Sec. 4.2), ablation studies (Sec. 4.3), and evaluations on both synthetic datasets (Sec. 4.1) and real-world images (Sec. 4.4). The supplement shows further data creation details and more experiments.

### 4.1. Synthetic Datasets

We collect 6 synthetic scenes comprising a variety of scene geometries and materials. For training, each scene is rendered under four distinct lighting setups: 1) single far-field light, 2) single far-field light with a camera-located flashlight, 3) two far-field lights, and 4) two far-field lights with a near-field light source. We render 100 training images per setting using random viewpoints and per-image lighting conditions. The far-field lighting from the first setting is used in subsequent settings, with the last two sharing their far-field lighting. For testing, 200 images under one common far-field lighting are rendered to assess novel view synthesis under seen lighting, and another 200 under unseen far-field lighting to evaluate free-viewpoint (FV) relighting, which largely depends on the geometry and material estimation quality. This methodology provides a relatively fair comparison across methods with different abilities regarding supported input lighting conditions, trained under similar conditions and tested all using the same sets.

### 4.2. Comparison with State-of-the-art Methods

We compare our method with leading multi-view inverse rendering methods under each input lighting scenario. For single far-field lighting inputs (Tab. 1 row 1), we compare with NVDiffRecMC [9] and TensoIR [10] (single-lighting input). For multiple far-field lighting scenarios (Tab. 1 row 3), we compare with TensoIR [10]. For setups involving a far-field light and a flashlight (Tab. 1 row 4), we compare

with WildLight [6]. We use Blender [1] for high-quality relit results for our method, WildLight [6], and NVDiffRecMC [9], importing trained SDF and material fields as textured meshes. Since TensoIR [10] employs density-based geometry and cannot extract high-quality meshes, we utilize its differentiable renderer for relighting.

Our quantitative comparison employs mean angular error (MAnGE), peak signal-to-noise ratio (PSNR), and structural similarity (SSIM) [30] metrics for surface normal, diffuse albedo, view synthesis RGB, free-viewpoint (FV) relit RGB, and specular reflection of FV relit RGB. We follow NeRFactor [41] in assuming albedo and lighting brightness as scale-invariant, necessitating metric comparisons using corresponding scales. For visualization and metrics, we scale each RGB channel of albedo and relit images from all methods by a global scale to minimize mean squared error against ground truth.

The quantitative results are detailed in Tab. 2. We observe that when only input images under single far-field lighting are available (rows 1-3), our method gives comparable results as TensoIR [10] and NVDiffRecMC [9]. Adding a flashlight (rows 4, 5) significantly helps in geometry and material estimation (comparing rows 3, 5), where the greatest improvement is made regarding the diffuse albedo. Also, our method clearly surpasses WildLight [6] in both quality and speed in this setting. Leveraging another far-field lighting condition (rows 6, 7) brings similar effects as using a flashlight, but generally with a lower degree of improvement. Further adding a flashlight (row 8) gives the best results among all settings and methods, but the gain is marginal compared to the setting of single far-field lighting with a flashlight (comparing rows 5, 8).

Fig. 3 shows the qualitative comparison on two synthetic scenes: LEGO (left) with detailed geometry, and TROOPER (right) with highly-reflective SV materials. The results show that our method predicts more detailed and faithful shape, diffuse albedo, and specular parameters that could produce accurate and realistic relit results under unseen




Figure 3. Comparison with state-of-the-art methods. We test NVDiffRecMC [9] (single far-field lighting in input), TensorIR [10] (two far-field lighting in input), and WildLight [6] and our method (single far-field lighting with a flashlight in input) on two synthetic scenes. The intensity of the specular reflection is multiplied by  $3\times$  for the visualization purpose. Please zoom in for details.

Table 3. Ablation study on our method using inputs under single far-field lighting with a flashlight.

Model	Normal	Albedo	FV Relit	Relit (spec)
	MAnGE $\downarrow$	PSNR $\uparrow$	PSNR $\uparrow$	PSNR $\uparrow$
(1) w/o $\mathcal{L}_{con}$	8.452	27.70	29.85	29.08
(2) Mod. $c_{near}$	8.572	27.79	30.03	28.87
(3) w/o $enc_{mat}$	<b>8.263</b>	27.79	30.07	28.92
(4) Full model	8.298	<b>28.05</b>	<b>30.19</b>	<b>29.36</b>

lighting. In contrast, WildLight [6] tends to overly smooth geometry and material. TensorIR [10] can not handle highly-reflective surfaces, and lighting is baked into its predicted diffuse albedo for TROOPER.

### 4.3. Ablation Studies

In Tab. 3, we quantitatively evaluate ablation models on a subset of our synthetic scenes: 1) the model excluding the self-consistency loss  $\mathcal{L}_{con}$ , 2) the model where the output of the near-field radiance MLP is not explicitly multiplied by the cosine weakening factor and incident radiance (Eq. (8)), 3) the model with  $enc_{mat}(\mathbf{x})$  replaced by  $enc_{geo}(\mathbf{x})$ . The results show that all the above techniques enhance the geometry and materials estimated by our method.

### 4.4. Real-world Results

Fig. 4 shows the results of our method on two real-world scenes captured under single far-field lighting and a flashlight. Our method faithfully recovers the geometry and material of the objects and produces realistic relit results.

## 5. Conclusion

We introduce a versatile multi-view inverse rendering framework, distinguished by its ability to leverage input im-




Figure 4. The results of our method on two real-world scenes, using multi-view images under one ambient lighting and a flashlight. The specular reflection intensity is increased for visualization.

ages under varied far- and near-field light sources available in capture for better geometry and material estimation.

**Limitations.** VMINer does not model the background, necessitating a foreground mask for each image. It also does not consider unknown tone-mapping curves applied during the image signal processing (ISP) stage, using a fixed curve with  $\gamma = 2.2$  to the computed linear radiance. It depends on distinguishing the contribution from each lighting, which implies that, for optimal results, near-field light sources should be switched on/off during capture.

## Acknowledgments

This work is supported by the National Natural Science Foundation of China (Grand No. 62136001, 62088102).



## References

- [1] Blender Foundation. The Blender project - free and open 3D creation software. <https://www.blender.org>. Accessed: 2023-11-07. 2, 6, 7
- [2] Mark Boss, Raphael Braun, Varun Jampani, Jonathan T. Barron, Ce Liu, and Hendrik P. A. Lensch. NeRD: Neural reflectance decomposition from image collections. In *Proc. of IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 1, 2, 3, 4
- [3] Mark Boss, Varun Jampani, Raphael Braun, Ce Liu, Jonathan T. Barron, and Hendrik P. A. Lensch. Neural-PIL: Neural pre-integrated lighting for reflectance decomposition. In *Proc. of Neural Information Processing Systems (NeurIPS)*, 2021.
- [4] Mark Boss, Andreas Engelhardt, Abhishek Kar, Yuanzhen Li, Deqing Sun, Jonathan T. Barron, Hendrik P. A. Lensch, and Varun Jampani. SAMURAI: shape and material from unconstrained real-world arbitrary image collections. In *Proc. of Neural Information Processing Systems (NeurIPS)*, 2022. 2
- [5] Brent Burley. Physically-based shading at disney. In *SIGGRAPH 2012 Courses*, 2012. 3
- [6] Ziang Cheng, Junxuan Li, and Hongdong Li. WildLight: In-the-wild inverse rendering with a flashlight. In *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 1, 2, 3, 5, 6, 7, 8
- [7] Marc-André Gardner, Yannick Hold-Geoffroy, Kalyan Sunkavalli, Christian Gagné, and Jean-François Lalonde. Deep parametric indoor lighting estimation. In *Proc. of IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019. 3
- [8] Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization for learning shapes. In *Proc. of International Conference on Machine Learning (ICML)*, 2020. 6
- [9] Jon Hasselgren, Nikolai Hofmann, and Jacob Munkberg. Shape, light, and material decomposition from images using Monte Carlo rendering and denoising. In *Proc. of Neural Information Processing Systems (NeurIPS)*, 2022. 1, 2, 7, 8
- [10] Haiyan Jin, Isabella Liu, Peijia Xu, Xiaoshuai Zhang, Songfang Han, Sai Bi, Xiaowei Zhou, Zexiang Xu, and Hao Su. TensoIR: Tensorial inverse rendering. In *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 1, 2, 3, 4, 7, 8
- [11] James T. Kajiya. The rendering equation. In *Proc. of the ACM SIGGRAPH Conference and Exhibition On Computer Graphics and Interactive Techniques (SIGGRAPH)*, 1986. 1, 2
- [12] 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):56:1–56:12, 2022. 1, 2, 3
- [13] Eric P. Lafortune and Yves D. Willems. *Using the modified phong reflectance model for physically based rendering*. Katholieke Universiteit Leuven. Departement Computerwetenschappen, 1994. 6
- [14] Ruilong Li, Hang Gao, Matthew Tancik, and Angjoo Kanazawa. NerfAcc: Efficient sampling accelerates NeRFs. *arXiv preprint arXiv:2305.04966*, 2023. 5
- [15] Tzu-Mao Li, Miika Aittala, Frédo Durand, and Jaakko Lehtinen. Differentiable Monte Carlo ray tracing through edge sampling. *ACM Transactions on Graphics (TOG)*, 37(6):222, 2018. 6
- [16] Zhengqin Li, Mohammad Shafiei, Ravi Ramamoorthi, Kalyan Sunkavalli, and Manmohan Chandraker. Inverse rendering for complex indoor scenes: Shape, spatially-varying lighting and SVBRDF from a single image. In *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 3
- [17] Zhengqin Li, Jia Shi, Sai Bi, Rui Zhu, Kalyan Sunkavalli, Milos Hasan, Zexiang Xu, Ravi Ramamoorthi, and Manmohan Chandraker. Physically-based editing of indoor scene lighting from a single image. In *Proc. of European Conference on Computer Vision (ECCV)*, 2022. 3
- [18] Zhaoshuo Li, Thomas Müller, Alex Evans, Russell H. Taylor, Mathias Unberath, Ming-Yu Liu, and Chen-Hsuan Lin. Neuralangelo: High-fidelity neural surface reconstruction. In *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 1
- [19] 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. In *Proc. of European Conference on Computer Vision (ECCV)*, 2020. 1, 4
- [20] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (TOG)*, 41(4):102:1–102:15, 2022. 1, 4, 5
- [21] Jacob Munkberg, Wenzheng Chen, Jon Hasselgren, Alex Evans, Tianchang Shen, Thomas Müller, Jun Gao, and Sanja Fidler. Extracting triangular 3D models, materials, and lighting from images. In *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1, 2
- [22] Ravi Ramamoorthi and Pat Hanrahan. A signal-processing framework for inverse rendering. In *Proc. of the ACM SIGGRAPH Conference and Exhibition On Computer Graphics and Interactive Techniques (SIGGRAPH)*, 2001. 1
- [23] Viktor Rudnev, Mohamed Elgharib, William A. P. Smith, Lingjie Liu, Vladislav Golyanik, and Christian Theobalt. NeRF for outdoor scene relighting. In *Proc. of European Conference on Computer Vision (ECCV)*, 2022. 1, 2, 3
- [24] 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 *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 1, 2, 3
- [25] Marco Toschi, Riccardo De Matteo, Riccardo Spezialetti, Daniele De Gregorio, Luigi Di Stefano, and Samuele Salti. ReLight My NeRF: A dataset for novel view synthesis and relighting of real world objects. In *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 5

- [26] Eric Veach and Leonidas J. Guibas. Optimally combining sampling techniques for Monte Carlo rendering. In *Proc. of the ACM SIGGRAPH Conference and Exhibition On Computer Graphics and Interactive Techniques (SIGGRAPH)*, 1995. 6
- [27] Bruce Walter, Stephen R. Marschner, Hongsong Li, and Kenneth E. Torrance. Microfacet models for refraction through rough surfaces. In *Proc. of the Eurographics Symposium on Rendering Techniques*, 2007. 5
- [28] Jiaping Wang, Peiran Ren, Minmin Gong, John M. Snyder, and Baining Guo. All-frequency rendering of dynamic, spatially-varying reflectance. *ACM Transactions on Graphics (TOG)*, 28(5):133, 2009. 4
- [29] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. NeuS: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. In *Proc. of Neural Information Processing Systems (NeurIPS)*, 2021. 4, 5
- [30] 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 (TIP)*, 13(4):600–612, 2004. 7
- [31] Zian Wang, Jonah Philion, Sanja Fidler, and Jan Kautz. Learning indoor inverse rendering with 3D spatially-varying lighting. In *Proc. of IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 3
- [32] Zian Wang, Tianchang Shen, Jun Gao, Shengyu Huang, Jacob Munkberg, Jon Hasselgren, Zan Gojcic, Wenzheng Chen, and Sanja Fidler. Neural fields meet explicit geometric representations for inverse rendering of urban scenes. In *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 1, 2, 3
- [33] Haoqian Wu, Zhipeng Hu, Lincheng Li, Yongqiang Zhang, Changjie Fan, and Xin Yu. NeFII: Inverse rendering for reflectance decomposition with near-field indirect illumination. In *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 4, 6
- [34] Wenqi Yang, Guanying Chen, Chaofeng Chen, Zhenfang Chen, and Kwan-Yee K. Wong. PS-NeRF: Neural inverse rendering for multi-view photometric stereo. In *Proc. of European Conference on Computer Vision (ECCV)*, 2022. 2
- [35] Yao Yao, Jingyang Zhang, Jingbo Liu, Yihang Qu, Tian Fang, David McKinnon, Yanghai Tsin, and Long Quan. NeILF: Neural incident light field for physically-based material estimation. In *Proc. of European Conference on Computer Vision (ECCV)*, 2022. 1, 2, 3
- [36] Bohan Yu, Siqi Yang, Xuanning Cui, Siyan Dong, Baoquan Chen, and Boxin Shi. MILO: multi-bounce inverse rendering for indoor scene with light-emitting objects. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 45(8):10129–10142, 2023. 1, 2, 3
- [37] Chong Zeng, Guojun Chen, Yue Dong, Pieter Peers, Hongzhi Wu, and Xin Tong. Relighting neural radiance fields with shadow and highlight hints. In *Proc. of the ACM SIGGRAPH Conference and Exhibition On Computer Graphics and Interactive Techniques (SIGGRAPH)*, 2023. 3
- [38] Jingyang Zhang, Yao Yao, Shiwei Li, Jingbo Liu, Tian Fang, David McKinnon, Yanghai Tsin, and Long Quan. NeILF++: Inter-reflectable light fields for geometry and material estimation. In *Proc. of IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. 2, 3
- [39] 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 *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 1, 2, 3, 4
- [40] Kai Zhang, Fujun Luan, Zhengqi Li, and Noah Snavely. IRON: inverse rendering by optimizing neural SDFs and materials from photometric images. In *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 3
- [41] Xiuming Zhang, Pratul P. Srinivasan, Boyang Deng, Paul E. Debevec, William T. Freeman, and Jonathan T. Barron. NeRFactor: neural factorization of shape and reflectance under an unknown illumination. *ACM Transactions on Graphics (TOG)*, 2021. 1, 2, 3, 4, 6, 7
- [42] Yuanqing Zhang, Jiaming Sun, Xingyi He, Huan Fu, Rongfei Jia, and Xiaowei Zhou. Modeling indirect illumination for inverse rendering. In *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1, 2, 3, 4
- [43] Youjia Zhang, Teng Xu, Junqing Yu, Yuteng Ye, Yanqing Jing, Junle Wang, Jingyi Yu, and Wei Yang. NeMF: Inverse volume rendering with neural microflake field. In *Proc. of IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. 2
- [44] Fuqiang Zhao, Yuheng Jiang, Kaixin Yao, Jiakai Zhang, Liao Wang, Haizhao Dai, Yuhui Zhong, Yingliang Zhang, Minye Wu, Lan Xu, and Jingyi Yu. Human performance modeling and rendering via neural animated mesh. *ACM Transactions on Graphics (TOG)*, 41(6):235:1–235:17, 2022. 4, 5
- [45] Jingsen Zhu, Fujun Luan, Yuchi Huo, Zihao Lin, Zhihua Zhong, Dianbing Xi, Rui Wang, Hujun Bao, Jiayang Zheng, and Rui Tang. Learning-based inverse rendering of complex indoor scenes with differentiable Monte Carlo raytracing. In *Proc. of the ACM SIGGRAPH Conference and Exhibition on Computer Graphics and Interactive Techniques in Asia (SIGGRAPH Asia)*, 2022. 3
- [46] Jingsen Zhu, Yuchi Huo, Qi Ye, Fujun Luan, Jifan Li, Dianbing Xi, Lisha Wang, Rui Tang, Wei Hua, Hujun Bao, and Rui Wang.  $l^2$ -SDF: Intrinsic indoor scene reconstruction and editing via raytracing in neural SDFs. In *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 1, 3
- [47] Yongjie Zhu, Yinda Zhang, Si Li, and Boxin Shi. Spatially-varying outdoor lighting estimation from intrinsics. In *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 3