

MAIR: Multi-view Attention Inverse Rendering with 3D Spatially-Varying Lighting Estimation

JunYong Choi^{1,2} SeokYeong Lee^{1,2} Haesol Park¹ Seung-Won Jung² Ig-Jae Kim^{1,3,4} Junghyun Cho^{1,3,4}

¹Korea Institute of Science and Technology(KIST) ²Korea University

³AI-Robotics, KIST School, University of Science and Technology

⁴Yonsei-KIST Convergence Research Institute, Yonsei University

{happily, shapin94, haesol, drjay, jhcho}@kist.re.kr swjung83@korea.ac.kr

Abstract

We propose a scene-level inverse rendering framework that uses multi-view images to decompose the scene into geometry, a SVBRDF, and 3D spatially-varying lighting. Because multi-view images provide a variety of information about the scene, multi-view images in object-level inverse rendering have been taken for granted. However, owing to the absence of multi-view HDR synthetic dataset, scene-level inverse rendering has mainly been studied using single-view image. We were able to successfully perform scene-level inverse rendering using multi-view images by expanding OpenRooms dataset and designing efficient pipelines to handle multi-view images, and splitting spatially-varying lighting. Our experiments show that the proposed method not only achieves better performance than single-view-based methods, but also achieves robust performance on unseen real-world scene. Also, our sophisticated 3D spatially-varying lighting volume allows for photorealistic object insertion in any 3D location.

1. Introduction

Inverse rendering is a technology used to estimate material, lighting, and geometry from RGB color images. Decomposing a scene through inverse rendering enables various applications such as object insertion, relighting, and material editing in VR and AR. However, since inverse rendering is an ill-posed problem, previous studies have focused only on a part of the inverse rendering, such as intrinsic image decomposition [5, 6, 14, 21], shape from shading [15, 29, 46], and material estimation [8, 9, 22, 24, 32].

Recent advances in GPU-accelerated physically-based




Figure 1. The result of inverse rendering and floating chrome sphere insertion in the unseen real-world scene. Since the single-view-based method [19] relies only on contextual information, it has difficulty estimating the complex material, geometry of the real-world. Notice the spatially-consistent albedo of apples, elaborated normal, and realistic lighting reflected on the inserted object. In previous work [19], they have limitations on floating object insertion because they use per-pixel lighting.

rendering algorithms have made constructing large-scale photorealistic indoor high dynamic range (HDR) image dataset that include geometries, materials, and spatially-varying lighting [25]. The availability of such dataset and the recent success of deep learning technology have enabled seminal works on single-view-based inverse rendering [19, 20, 23, 40, 49, 50]. These methods have fundamental limitations in that they are prone to bias in training dataset despite having shown promising results. Specifically, single-view-based inverse rendering must refer to specular reflectance from the contextual information of the image, making it less reliable for predicting complex SVBRDF (spatially-varying bidirectional reflectance distribution function) in the real-world. Fig. 1 shows such an example, where decomposition has severely failed owing to the complexity of the real-world scene. In addition, depth-scale ambiguity makes it challenging to employ these methods to 3D applications such as object insertion.

In this paper, we introduce **MAIR**, the scene-level Multi-

This work was partly supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT)(No.2020-0-00457, 50%) and KIST Institutional Program(Project No.2E32301, 50%).

view Attention Inverse Rendering pipeline. MAIR exploits multiple RGB observations of the same scene point from multiple cameras and, more importantly, it utilizes multi-view stereo (MVS) depth as well as scene context to estimate SVBRDF. As a result, the method becomes less dependent on the implicit scene context, and shows better performance on unseen real-world images. However, the processing of multi-view images inherently requires a high computational cost to handle multiple observations with occlusions and brightness mismatches. To remedy this, we design a three-stage training pipeline for MAIR that can significantly increase training efficiency and reduce memory consumption. Spatially-varying lighting consists of direct lighting and indirect lighting. Indirect lighting affected by the surrounding environment makes inverse rendering difficult. Therefore, in Stage 1, we first estimate the direct lighting and geometry, which reflect the amount of light entering each point and in which direction the specular reflection appears. We estimate the material in Stage 2 using the estimates of the direct lighting, geometry, and multi-view color images. In Stage 3, we collect all the material, geometry, and direct lighting information and finally estimate 3D spatially-varying lighting, including indirect lighting. The MAIR pipeline is shown in Fig. 2. We created the OpenRooms Forward Facing (OpenRooms FF) dataset as an extension of OpenRooms [25] to train the proposed network.

Our contribution can be summarized as follows:

- As summarized in Tab. 1, we believe this is the first demonstration of using multi-view images to decompose the scene into geometry, complex material, and 3D spatially-varying lighting without test-time optimization. Also, we release OpenRooms FF dataset.
- We propose a framework that can efficiently train multi-view inverse rendering networks. Our framework increases the training efficiency by decomposing lighting and separating the scene components by stage.
- Our method achieves better inverse rendering performance than the existing single-view-based method, and realistic object insertion in real-world is possible by reproducing 3D lighting of the real-world.

2. Related Works

Inverse rendering. Research on inverse rendering has received significant attention in recent years owing to the development of deep learning technology. Yu *et al.* [42] performed outdoor inverse rendering with multi-view self-supervision, but their lighting is simple distant lighting. A pioneering work by Li *et al.* [19] conducted inverse rendering on a single image, and IRISformer [50] further improved the performance by replacing convolutional neural

Method	Input	Material	Lighting	Insertion
VSG [40]	single	diffuse	volume	any
Li <i>et al.</i> [19]	single	microfacet	per-pixel	surface
IRISformer [50]	single	microfacet	per-pixel	surface
PBE [20]	single	microfacet	parametric	any
SOLD [36]	single	✗	per-pixel	surface
Zhu <i>et al.</i> [49]	single	microfacet(metalic)	volume	any
lighthouse [35]	stereo	✗	volume	any
FreeView [31]	multi	glossy	irradiance	✗
Zhang <i>et al.</i> [48]	multi	microfacet	object	✗
PhotoScene [41]	any	microfacet	parametric	any
Intrinsic3D [26]	multi RGBD	diffuse	volume	any
Zhang <i>et al.</i> [44]	multi RGBD	diffuse	parametric	any
MAIR (Ours)	multi	microfacet	volume	any

Table 1. Compared to previous works, our work is the first demonstration to perform multi-view scene-level inverse rendering.

networks (CNNs) with a Transformer [43]. PhyIR [23] addressed this problem using panoramic image. However, the lighting representation in [19, 23, 50] is a 2D per-pixel environment map, which is insufficient for modeling 3D lighting. Li *et al.* [20] adopted a parametric 3D lighting representation; however, it fixed with two types of indoor light sources. The recently introduced Zhu *et al.* [49] demonstrated realistic 3D volumetric lighting based on ray tracing. But, since these prior works [19, 20, 23, 42, 49, 50] are all single-view-based, they inherently rely on the scene context, making these methods less reliable for unseen images. In contrast, the proposed method can estimate BRDF and geometry more accurately by utilizing the multi-view correspondences as additional cues for inverse rendering.

Lighting estimation. Lighting estimation has been studied not only as a sub-task of the inverse rendering but also an important research topic [10–12, 33]. In Lighthouse [35], 3D spatially-varying lighting was obtained by estimating the RGBA lighting volume. Wang *et al.* [40] estimated a more sophisticated 3D lighting volume by replacing RGB with a spherical Gaussian in lighting volume. However, because the methods in [35, 40] assume Lambertian reflectance, they cannot represent complex indirect lighting and have limitations in expressing HDR lighting due to weak-supervision with LDR dataset [18]. On the other hand, the proposed method can handle complex SVBRDF and HDR lighting well because we trained our model on the large indoor HDR dataset [25]. Recently, a study on spatially-varying lighting estimation in the outdoor scenes [36, 39] was introduced. Because they focus on outdoor street scenes, they cannot clearly reproduce the indirect lighting by the scene material.

Multi-view inverse rendering and neural rendering. Earlier works such as Intrinsic3D [26] and Zhang *et al.* [44] perform multi-view inverse rendering without deep learning, but require additional equipment to obtain RGB-D images. 3D geometry-based methods [2, 17, 28] require additional computation to generate mesh. PhotoScene [41], in particular, requires external CAD geometry, which should be manually aligned for each object. Moreover, all of these methods require test-time optimization. In contrast, we only

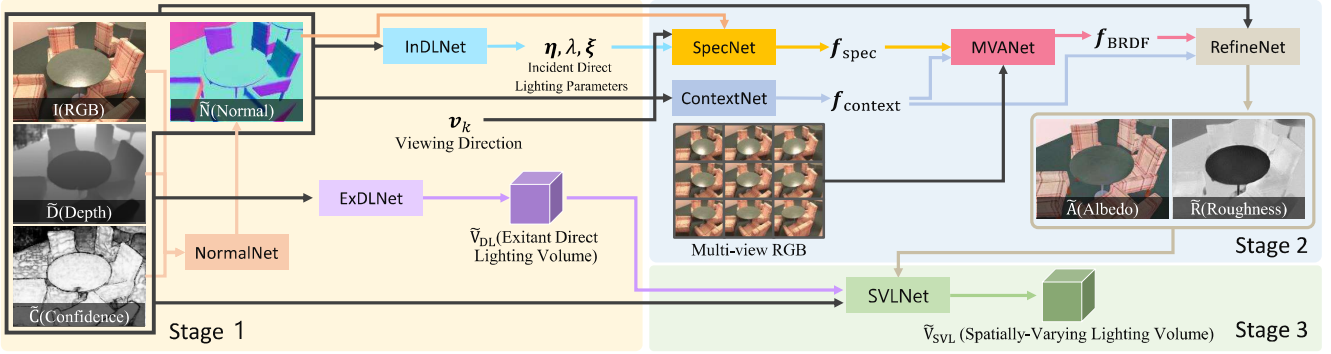


Figure 2. MAIR’s entire pipeline. Our method has reduced the difficulty of inverse rendering by splitting the scene components as small as possible, and progressively estimating the scene components.

need per-view depth maps and we do not require 3D geometry or test-time optimization, making it more computationally efficient and much easier to apply to more general scenes. Philip *et al.* [30, 31] demonstrated a successful relighting with multi-view images; however, these methods cannot be used for applications such as object insertion because they use trained neural renderers. The advent of NeRF [3, 4, 27] has led to a breakthrough in the field of neural rendering research. Several recent methods [7, 34, 37, 45, 47, 48] have successfully performed inverse rendering using multi-view images; however, they are difficult to apply to scene-level inverse rendering because they are only trained and tested on object-centric images.

3. Method

In this section, we describe the detailed architecture of the proposed network, MAIR. Let K denote the number of viewpoints; then, the inputs to the network are K triples, where each triple is composed of an RGB image with H, W size ($I \in \mathbb{R}^{3 \times H \times W}$), depth map ($\tilde{D} \in \mathbb{R}^{H \times W}$), and its confidence map ($\tilde{C} \in \mathbb{R}^{H \times W}$). \tilde{D} and \tilde{C} are obtained using a state-of-the-art MVS model [13]. We designed a three-stage structure that progressively estimates the normal, direct lighting, material, and spatially-varying lighting. The entire MAIR pipeline is summarized in Fig. 2.

3.1. Stage 1 - Target View Analysis Stage

Stage 1 of MAIR comprises three estimation networks: Normal map (NormalNet), **I**ncident **D**irect **L**ighting (InDLNet), and **E**xitant **D**irect **L**ighting (ExDLNet). Inspired by recent studies [19, 35, 40], we adopt spatially-varying spherical Gaussians (SVSGs) [19] and volumetric spherical Gaussian (VSG) [40] for the representation of incident lighting and exitant lighting, respectively.

Normal map estimation. Unlike single-view-based methods [19, 20, 40, 50], where normal information should be inferred from the scene context, the normal map (\tilde{N}) can be directly derived from the depth map. Thus, NormalNet

can show robust performance especially for real-world images, where the distribution of image contents and geometry largely differs from the training data. Still, use of other available information, including the RGB, depth gradient map ($\nabla \tilde{D} \in \mathbb{R}^{H \times W}$), and confidence map, can help NormalNet better handle unreliable depth predictions. Specifically, NormalNet is formulated as follows:

$$\tilde{N} = \text{NormalNet}(I, \tilde{D}, \nabla \tilde{D}, \tilde{C}), \tilde{N} \in \mathbb{R}^{3 \times H \times W}. \quad (1)$$

Incident direct lighting estimation. Given I , \tilde{D} , \tilde{C} , and \tilde{N} obtained using NormalNet, we estimate SVSGs as a lighting representation of incident direct lighting, which is proven to be effective in modeling environment map [19]. The proposed InDLNet is formulated as follows:

$$\{\xi_s\}, \{\lambda_s\}, \{\eta_s\} = \text{InDLNet}(I, \tilde{N}, \tilde{D}, \tilde{C}), \quad (2)$$

where $\xi_s \in \mathbb{R}^2$ is the direction vector outward from the center of the unit sphere, $\lambda_s \in \mathbb{R}$ is sharpness, and $\eta_s \in \mathbb{R}^3$ is intensity. The environment map is then parameterized with S_D SG lobes $\{\xi_s, \lambda_s, \eta_s\}_{s=1}^{S_D}$. For the s -th SG, its radiance $\mathcal{G}(\mathbf{l})$ in the direction $\mathbf{l} \in \mathbb{R}^2$ can be obtained as

$$\mathcal{G}(\mathbf{l}; \eta_s, \lambda_s, \xi_s) = \eta e^{\lambda_s(\mathbf{l} \cdot \xi_s - 1)}, \quad (3)$$

Using all S_D SG lobes, the incident radiance $\mathcal{R}_i(\mathbf{l})$ in the direction \mathbf{l} is expressed as

$$\mathcal{R}_i(\mathbf{l}) = \sum_{s=1}^{S_D} \mathcal{G}(\mathbf{l}; \eta_s, \lambda_s, \xi_s), \quad (4)$$

Li *et al.* [19] used $S_D = 12$ to represent complex spatially-varying lighting; however, we found $S_D = 3$ to be sufficient to model much simpler direct lighting. Also we used global intensity to make the SVSGs spatially coherent.

Exitant direct lighting estimation. Although effective, the above environment map alone is insufficient to model lighting in a 3D space. Thus, we adopt a voxel-based representation called VSG [40] to further model exitant direct lighting. ExDLNet estimates exitant direct lighting volume \tilde{V}_{DL} as

$$\tilde{V}_{DL} = \text{ExDLNet}(I, \tilde{N}, \tilde{D}, \tilde{C}), \tilde{V}_{DL} \in \mathbb{R}^{8 \times X \times Y \times Z}, \quad (5)$$

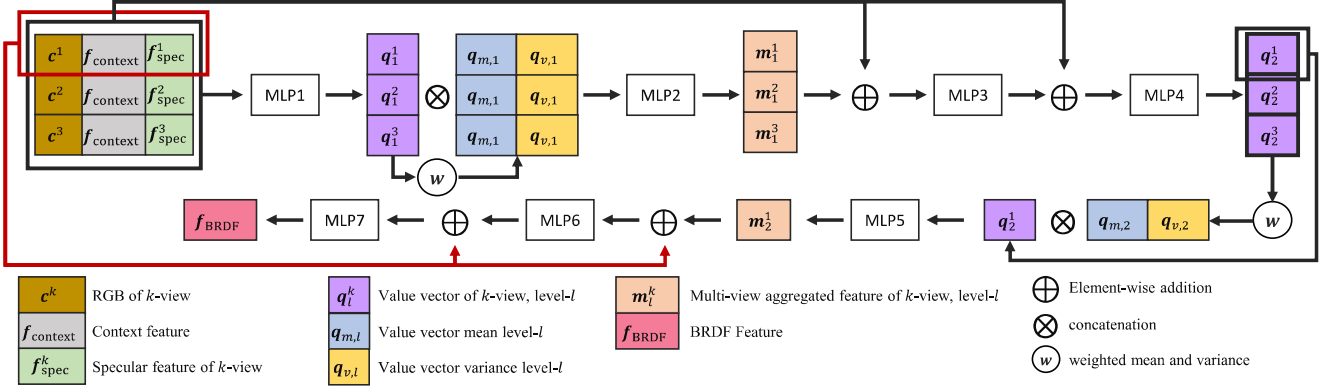


Figure 3. An illustration of MVANet when $K=3$. MVANet creates a value vector by encoding color, context feature, and specular feature, and uses multi-view weights as attention to create multi-view aggregated features. Since our goal is to obtain the BRDF of target-view(1-view), in level-2, only the value vector of target view is processed.

where X , Y , and Z are the sizes of the volume. Each voxel in \tilde{V}_{DL} contains opacity α and SG parameters (η, ξ, λ) . From VSG, alpha compositing in the direction \mathbf{l} allows us to calculate the incident radiance $\mathcal{R}_e(\mathbf{l})$ as follows:

$$\mathcal{R}_e(\mathbf{l}) = \sum_{n=1}^{N_R} \prod_{m=1}^{n-1} (1 - \alpha_m) \alpha_n \mathcal{G}(-\mathbf{l}; \eta_n, \lambda_n, \xi_n), \quad (6)$$

where N_R is the number of ray samples, and η_n, λ_n , and ξ_n are the SG parameters of the sample.

The intensity at which light converges at a point, *i.e.*, incident radiance, serves as guidance for the network to infer diffuse and specular reflections. Meanwhile, information on how light is present in 3D space, *i.e.*, exitant radiance, helps the network infer indirect lighting. Fig. 4 is an explanation of the incident/exitant direct lighting. Note that these two representations are not convertible and have different physical-meanings. Please see supplementary for a detailed description of direct lighting.




Figure 4. Explanation of the incident/exitant direct lighting. E means per-pixel direct lighting environment map.

3.2. Stage 2 - Material Estimation Stage

For BRDF estimation, the specular radiance must be considered. To obtain a specular radiance feature \mathbf{f}_{spec} , we propose a SpecNet. The diffuse radiance κ and specular radiance γ of the microfacet BRDF model [16] are as follows:

$$\kappa = \frac{a}{\pi} \int_{\mathbf{l}} L(\mathbf{l}) \mathbf{n} \cdot \mathbf{l} d\mathbf{l}, \kappa \in \mathbb{R}^3 \quad (7)$$

$$\gamma = \int_{\mathbf{l}} L(\mathbf{l}) \mathcal{B}_s(\mathbf{v}, \mathbf{l}, \mathbf{n}, r) \mathbf{n} \cdot \mathbf{l} d\mathbf{l}, \gamma \in \mathbb{R}^3 \quad (8)$$

where \mathbf{a} is diffuse albedo, \mathbf{l} is lighting direction, $L(\mathbf{l})$ is lighting intensity, \mathcal{B}_s is specular BRDF, \mathbf{n} is normal, r is roughness, and \mathbf{v} is viewing direction, respectively. Since Eq. (8) is highly complicated, it is necessary to efficiently encode inputs to make it easier for the network to learn. To this intent, we first rewrite the arguments of \mathcal{B}_s as following:

$$[\mathcal{F}(\mathbf{v}, \mathbf{h}), (\mathbf{n} \cdot \mathbf{h})^2, \mathbf{n} \cdot \mathbf{l}, \mathbf{n} \cdot \mathbf{v}, r], \quad (9)$$

where \mathcal{F} is the Fresnel equation, and \mathbf{h} is the half vector. Then, we approximate the lighting of Eq. (8) with the SVSGs of Eq. (2). Since each SG lobe $\{\xi, \lambda, \eta\}$ can be thought of as an individual light source, ξ, η , and λ can be regarded as $\mathbf{l}, L(\mathbf{l})$, and a parameter to approximate the integral, respectively. Consequently, γ can be read as follows:

$$\gamma = \sum_{s=1}^{S_D} g(\mathcal{F}(\mathbf{v}, \mathbf{h}_s), (\mathbf{n} \cdot \mathbf{h}_s)^2, \mathbf{n} \cdot \xi_s, \mathbf{n} \cdot \mathbf{v}, \eta_s, \lambda_s, r), \quad (10)$$

where g is a newly defined function from our reparameterization. Using Eq. (10), we define SpecNet as follows:

$$\mathbf{f}_{\text{spec}}^k = \sum_{s=1}^{S_D} m_s \text{SpecNet}(\mathcal{F}(\mathbf{v}_k, \mathbf{h}_{s,k}), (\tilde{\mathbf{n}} \cdot \mathbf{h}_{s,k})^2, \tilde{\mathbf{n}} \cdot \xi_s, \tilde{\mathbf{n}} \cdot \mathbf{v}_k, \eta_s, \lambda_s), \quad (11)$$

$$m_s = \begin{cases} 1 & \text{if } \|\eta_s\|_1 \tilde{\mathbf{n}} \cdot \xi_s > 0, \\ 0 & \text{else,} \end{cases} \quad (12)$$

where k is k -th view. A binary indicator m_s is used to exclude SG lobes from γ if the intensity of the light source ($\|\eta_s\|_1$) is 0 or the dot product of the normal and light axis ($\tilde{\mathbf{n}} \cdot \xi_s$) is less than 0. Since SpecNet approximates Eq. (8) with this physically-motivated encoding, \mathbf{f}_{spec} can include feature for specular radiance information. In addition to SpecNet, we use ContextNet to obtain a context feature map $\mathbf{f}_{\text{context}} = \text{ContextNet}(\mathbf{I}, \tilde{\mathbf{D}}, \tilde{\mathbf{C}}, \tilde{\mathbf{N}})$ that contains

the local context of the scene. All views share $\mathbf{f}_{\text{context}}$ of the target view.

Next, a **Multi-View Aggregation network (MVANet)** is used to aggregate \mathbf{f}_{spec} , $\mathbf{f}_{\text{context}}$, and RGB across the pixels from all K views, which corresponds to the target view pixel considering MVS depths. However, some of these pixel values might have negative effect if they are from the wrong surfaces due to occlusion. To consider occlusion, the depth projection error in k -view, denoted as $e_k = \max(-\log(|\tilde{d}_k - z_k|), 0)$, is calculated. \tilde{d}_k is the depth at the pixel position obtained by projecting a point seen from the target view onto k -view, and z_k is the distance between the point and the camera center of k -view. The depth projection error $\mathbf{e} \in \mathbb{R}^K$ is obtained by aggregating e_k from all K views. We use multi-view weight $\mathbf{w} = \frac{\mathbf{e}}{\|\mathbf{e}\|_1}$, $\mathbf{w} \in \mathbb{R}^K$ as attention weights during the multi-view feature aggregation in MVANet. Our intuition for material estimation is to consider the mean and variance of RGB. MVANet first encodes the input for each view to produce a value vector \mathbf{q} , and produces a mean and variance of \mathbf{q} according to \mathbf{w} [38]. It is encoded again and produces a multi-view aggregated feature \mathbf{m} . Since \mathbf{m} is created from weighted means and variances, it has multi-view information considering occlusion. This process is repeated once again for the target view. See Fig. 3 for detailed structure of MVANet.

Since MVANet exploits only local features, long-range interactions within the image need to be further considered for inverse rendering [50]. Thus, we propose RefineNet for albedo ($\tilde{\mathbf{A}} \in \mathbb{R}^{3 \times H \times W}$), roughness ($\tilde{\mathbf{R}} \in \mathbb{R}^{H \times W}$) estimation using \mathbf{f}_{BRDF} from MVANet.

$$\tilde{\mathbf{A}}, \tilde{\mathbf{R}} = \text{RefineNet}(\mathbf{I}, \tilde{\mathbf{D}}, \tilde{\mathbf{C}}, \tilde{\mathbf{N}}, \mathbf{f}_{\text{BRDF}}, \mathbf{f}_{\text{context}}). \quad (13)$$

3.3. Stage 3 - Lighting Estimation Stage

In stage 3, **Spatially Varying Lighting Estimation Network (SVLNet)** infers 3D lighting with direct lighting, geometry, and material. To this end, we create a visible surface volume ($\mathbf{T} \in \mathbb{R}^{10 \times X \times Y \times Z}$). Although Wang *et al.* [40] used a similar representation, they used a Lambertian reflectance model, which cannot represent complex lighting due to specularly. In contrast, we initialize \mathbf{T} by reprojecting $\mathbf{I}, \tilde{\mathbf{N}}, \tilde{\mathbf{A}}, \tilde{\mathbf{R}}$, which can model specularly. For each voxel, let (u, v) and d denote the projected coordinate of the center point and the depth, respectively. Then, the local feature $\mathbf{t} \in \mathbb{R}^{10}$ for each voxel is initialized as follows.

$$\mathbf{t} = \left[\rho \mathbf{I}(u, v), \rho \tilde{\mathbf{N}}(u, v), \rho \tilde{\mathbf{A}}(u, v), \rho \tilde{\mathbf{R}}(u, v) \right], \quad (14)$$

where $\rho = e^{-\tilde{\mathbf{C}}(u, v)(d - \tilde{\mathbf{D}}(u, v))^2}$. Note that the confidence ($\tilde{\mathbf{C}}$) is used to reflect the accuracy of the depth. \mathbf{T} and $\tilde{\mathbf{V}}_{\text{DL}}$ are fed to SVLNet, producing outputs $\tilde{\mathbf{V}}_{\text{SVL}}$ representing 3D spatially-varying lighting volume, as follows:

$$\tilde{\mathbf{V}}_{\text{SVL}} = \text{SVLNet}(\tilde{\mathbf{V}}_{\text{DL}}, \mathbf{T}), \tilde{\mathbf{V}}_{\text{SVL}} \in \mathbb{R}^{8 \times X \times Y \times Z}. \quad (15)$$

In previous work [40], they used implicit global feature volume, which is unclear what information it contains, but we specified $\tilde{\mathbf{V}}_{\text{DL}}$ explicitly so that the network can learn the interaction of light source, material, and geometry.

4. Implementation Details

Dataset. OpenRooms [25] provides many HDR images and various ground truths for indoor scenes; however, the distribution of camera poses is too random to be used in multi-view applications. Therefore, we created OpenRooms FF. Among OpenRooms [25] images, 23,618 images were selected to be appropriate for multi-view rendering based on the camera position and minimum depth, and eight neighboring frames were rendered by translating the camera for each image. For a detailed description of OpenRooms FF, please see supplementary material.

Training and loss. We used nine images ($K=9$) in our experiments. We trained each stage separately since the ground truth required for each stage is available in OpenRooms FF. We trained lighting with only 2D per-pixel ground truth lighting, but our VSG can represent lighting well in any 3D space (See (b) of Fig. 8.). A brief description of the loss for each stage is provided in Tab. 2. For a detailed description of training, loss function, and network architecture, please see supplementary material.

Stage	Network	Loss
1	NormalNet	MSE(N) + L1 angular(N)
	InDLNet	si-MSE(Lighting)
	ExDLNet	si-MSE(Lighting)
2	All	si-MSE(A) + MSE(R)
3	All	si-MSE(Lighting) + MSE(Re-rendering)

Table 2. Training loss for MAIR. si- means scale invariant.

5. Experiments

Since no previous works, to the best of our best knowledge, addressed multi-view scene-level inverse rendering, we compare our method qualitatively and quantitatively with a single-view-based method [19]. We also conduct a qualitative performance evaluation on the real-world dataset [38]. Last, we demonstrate that our realistic 3D spatially-varying lighting through virtual object insertion.

5.1. Evaluation on Synthetic Data

For quantitative comparison, we trained Li *et al.* [19] on our OpenRooms FF. Tab. 3 shows the performance comparison for all test images in OpenRooms FF. Material and geometry errors were measured by MSE and lighting error was measured by log-space MSE with ground truth per-pixel environment map. Our method outperformed the previous method in geometry and material estimation. Moreover, we achieved better lighting performance even though the previous method directly estimates per-pixel lighting

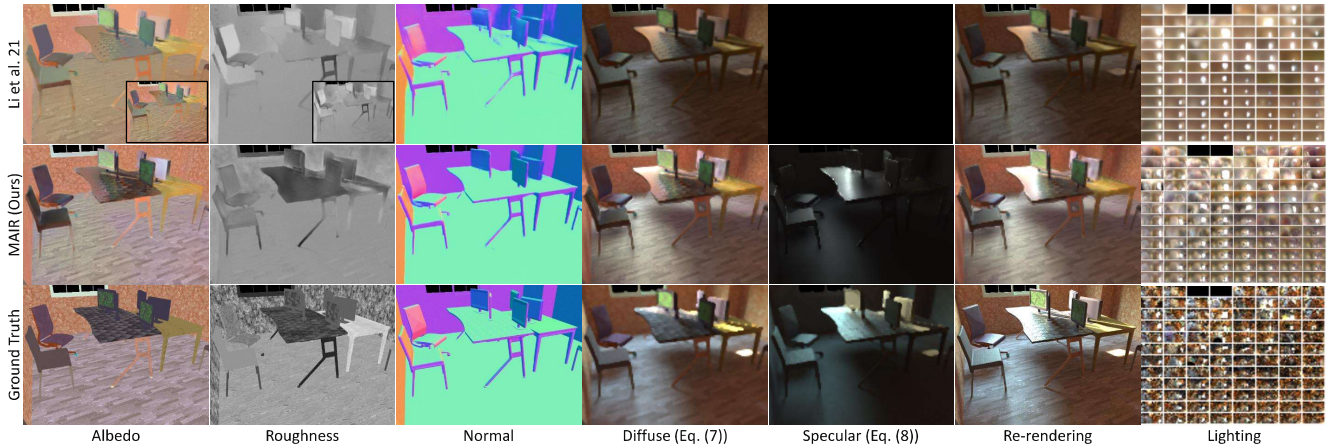


Figure 5. Material, geometry, lighting estimation on OpenRooms FF test data. The small insets in the first row are the material estimation processed by Bilateral Solvers (BS). Even for scenes where inverse rendering is difficult due to strong specular radiance, our method can obtain material, geometry, and lighting more accurately. Please see the lighting and material of the desk.




Figure 6. Inverse rendering results for unseen real-world data. Li *et al.* [19] failed the geometry estimation and predicted the material almost exactly as the input image, while we can remove shadow and disentangle the material and lighting from the image.

but we obtain per-pixel lighting from the 3d lighting volume. For the re-rendering, the previous method showed better performance on average because Li *et al.* [19]’s cascade structure self-supervises re-rendering through the rendering layer. However, in scenes where specular radiance appears strongly, we found that our method robustly and accurately estimates the details of lighting and materials, thereby yielding more realistic result, as shown in Fig. 5.

MSE ($\times 10^{-2}$)	Li <i>et al.</i> [19]	MAIR (Ours)
Albedo ↓	0.569	0.368 (−0.201)
Normal ↓	2.71	1.36 (−1.35)
Roughness ↓	3.66	2.70 (−0.96)
Lighting ↓	13.74	12.04 (−1.70)
Re-rendering ↓	0.554	0.633 (+0.079)

Table 3. Quantitative comparison of material, geometry, and lighting in OpenRooms FF. Albedo and roughness of Li *et al.* [19] are processed with a Bilateral Solver (BS).

5.2. Evaluation on Real-world Data

We evaluate the performance of inverse rendering for the real-world scenes in the IBRNet dataset [38] in which the scene context is difficult to grasp. As shown in Fig. 6, the single-view-based method fails in geometry estimation due to the lack of the scene context, leading to the inaccurate disentanglement of the material and illumination. Our method, on the other hand, obtains a reasonable geometry, removes shadows, and predicts spatially-consistent albedo. More results can be found in the supplementary material.

5.3. Evaluation on Object Insertion

We further demonstrate the effectiveness of our robust inverse rendering on the object insertion task. In Fig. 7, 8 we provide a comparison with the single-view method [19] and the stereo method [35]. It should be noted that, we use the normal from MAIR for Lighthouse [35] results because Lighthouse [35] does not provide scene geometry. In Li *et*




Figure 7. Object insertion comparison for OpenRooms FF test scene. Our method can insert the sphere to match the scene shadow, reproduce color bleeding, and put the sphere in the air.



Figure 8. Object [1] and floating chrome sphere insertion comparison for unseen real-world data from IBRNet dataset [38]. Our method estimate lighting robustly for unseen images, and we can insert objects more realistically than previous methods.

al. [19]’s object insertion application, the user must specify the plane on which the object is located. Therefore, object cannot be placed on another object or floated in the air, and shadows are only cast on the plane. In our method, it is possible to insert objects freely into the 3D space without restrictions on the shadow. We also conducted a user study on the object insertion, in which respondents chose the most realistic images rendered by competing methods for 25 scenes. As shown in the Tab. 4, our method shows the best performance in both synthetic and real-world scenes.

Scene	Lighthouse [35]	Li <i>et al.</i> [19]	MAIR
OpenRooms FF	-	0.220	0.780
IBRNet [38]	0.120	0.180	0.700
IBRNet [38](chrome)	0.186	0.025	0.789

Table 4. User study results on virtual object insertion. The design and details are described in supplementary material.

Object insertion in indoor test dataset. In (a) of Fig. 7, the sphere on the floor is inserted to match the geometry and lighting of the scene, and the sphere in the air is inserted to express the appropriate shadows for the scene geometry. In (b) of Fig. 7, the sphere inserted on the floor represents the color bleeding, showing that our method is capable of indirect lighting estimation.

Object insertion in real-world unseen dataset. To show the generalization ability of the proposed model trained only on the synthetic indoor images, we tested object insertion on the real-world unseen dataset [38]. In (a) of Fig. 8, Li *et al.* [19] failed to separate lighting from the material, and as a result, the appearance of the object was colored with the material on the floor. Our method successfully separate material and lighting from the image, and realistic lighting is represented on the inserted object. Lighthouse [35] was more vulnerable to unseen real-world data because it did not consider scene geometry and materials, even though they use stereo images. Their objects have a similar appearance regardless of the scene. In (b) Fig. 8, we inserted a chrome sphere to validate the indirect lighting estimation accuracy of the proposed method. Because lighthouse [35] was trained with an indoor LDR dataset, it always tended to create an environment map with an indoor scene on the sphere, and failed to reproduce HDR lighting. Li *et al.* [19]’s pixel lighting was not sufficient to express indirect lighting of the scene reflected in chrome sphere. In contrast, our chrome sphere shows realistic lighting that reflects the surrounding environment. Please refer to the supplementary material for more results.

5.4. Ablation Study

Design of stage 2. Tab. 5 shows experimental results for network design choices in stage 2. It shows that MVANet requires a local context of ContextNet for material estimation. In addition, RefineNet is necessary to compensate for the weakness of the pixel-wise operation in MVANet. We also validate the effect of f_{spec} : “w/o f_{spec} ” is the result of the model trained without f_{spec} , and “w/o reparameterize” is the result of the model trained without physically-motivated encoding as follows:

$$f_{\text{spec}}^k = \sum_{s=1}^{S_D} \text{SpecNet}(v_k, \tilde{n}, \xi_s, \eta_s, \lambda_s). \quad (16)$$

Improvement in roughness estimation implies that our physically-motivated encoding considering the microfacet BRDF [16] model helps in estimating specular radiance.

MSE ($\times 10^{-2}$)	Albedo \downarrow	Roughness \downarrow
w/o ContextNet	0.498	4.064
w/o RefineNet	0.661	3.971
w/o f_{spec}	0.436	3.056
w/o reparameterize	0.428	3.085
Ours	0.423	2.739

Table 5. ablation studies in stage 2.

Attention with multi-view weight. We also validate the effect of multi-view weight(w) in MVANet. Fig. 9 shows the material estimation accuracy according to the number of views. When training without w , increasing the number of views seems ineffective, possibly due to the negative effect of noises in multi-view features introduced by occlusion. On the other hand, MVANet can selectively focus on the information required for material estimation with w , making the accuracy increase with the number of views.




Figure 9. material evaluation depending on the number of views.

The necessity of exitant direct lighting. Fig. 10 shows that \tilde{V}_{DL} is helpful for indirect lighting estimation. The result of training with a global lighting feature as in [40], instead of \tilde{V}_{DL} , is also compared (with GLF). The \tilde{V}_{DL} helps network to converge quickly, which indicates that \tilde{V}_{DL} helps the network infers indirect lighting.

6. Discussion

Comparisons with 3D geometry-based methods. Fig. 11 shows the comparison results with PhotoScene [41] and MVIR [17] in 2 views. The inference time is 2s for ours, 9m 52s for MVIR [17] and 10m 40s for PhotoScene [41] on




Figure 10. lighting evaluation depending on the input source.

RTX 2080 Ti. MVIR [17] fails to generate geometry, showing severe artifacts. The PhotoScene [41] shows a complete scene reconstruction thanks to the CAD geometry and material graph, but the synthesized images largely differ from the original scene.

Inter-view consistency. Unlike 3D geometry-based methods [17, 41], which leverage global 3D geometry to achieve consistency, Our MAIR operates in pixel-space and therefore does not explicitly guarantee inter-view consistency. Still, we did not observe significant inconsistencies during our experiment, presumably due to the role of MVANet, which can be seen in Fig. 11. For object insertion, we used lighting of the center-view.




Figure 11. Comparisons with PhotoScene [41] and MVIR [17].

Limitation. One possible limitation comes from the cascaded nature of the pipeline. If depth estimation fails due to, for example, presence of dynamic objects or large textureless region, our MAIR will not work properly (See Fig. 12.). Another possible limitation comes from the VSG representation. Although VSG can express 3D lighting effectively, it cannot be applied to applications such as light source editing because it is non-parametric.



Figure 12. Failure case when depth prediction fails.

Conclusion. We presented the first practical multi-view scene-level inverse rendering method by creating a multi-view HDR synthetic dataset. Compared to the single-view based methods, we found that our method is more robust for unseen real-world scenes, providing high-quality virtual object insertion results. We believe that our work can elevate image-based rendering and physically-based rendering together, so realizing a higher level of inverse rendering and scene reconstruction.

References

- [1] The stanford 3d scanning repository. <https://graphics.stanford.edu/data/3Dscanrep/>. 7
- [2] Dejan Azinovic, Tzu-Mao Li, Anton Kaplanyan, and Matthias Nießner. Inverse path tracing for joint material and lighting estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2447–2456, 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5855–5864, 2021. 3
- [4] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5470–5479, 2022. 3
- [5] Harry Barrow, J Tenenbaum, A Hanson, and E Riseman. Recovering intrinsic scene characteristics. *Comput. vis. syst.*, 2(3-26):2, 1978. 1
- [6] Sean Bell, Kavita Bala, and Noah Snavely. Intrinsic images in the wild. *ACM Transactions on Graphics (TOG)*, 33(4):1–12, 2014. 1
- [7] Mark Boss, Raphael Braun, Varun Jampani, Jonathan T Barron, Ce Liu, and Hendrik Lensch. Nerf: Neural reflectance decomposition from image collections. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12684–12694, 2021. 3
- [8] Valentin Deschaintre, Miika Aittala, Fredo Durand, George Drettakis, and Adrien Bousseau. Single-image svbrdf capture with a rendering-aware deep network. *ACM Transactions on Graphics (ToG)*, 37(4):1–15, 2018. 1
- [9] Valentin Deschaintre, Miika Aittala, Fredo Durand, George Drettakis, and Adrien Bousseau. Single-image svbrdf capture with a rendering-aware deep network. *ACM Transactions on Graphics (ToG)*, 37(4):1–15, 2018. 1
- [10] Marc-André Gardner, Yannick Hold-Geoffroy, Kalyan Sunkavalli, Christian Gagné, and Jean-François Lalonde. Deep parametric indoor lighting estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7175–7183, 2019. 2
- [11] Marc-André Gardner, Kalyan Sunkavalli, Ersin Yumer, Xiaohui Shen, Emiliano Gambaretto, Christian Gagné, and Jean-François Lalonde. Learning to predict indoor illumination from a single image. *arXiv preprint arXiv:1704.00090*, 2017. 2
- [12] Mathieu Garon, Kalyan Sunkavalli, Sunil Hadap, Nathan Carr, and Jean-François Lalonde. Fast spatially-varying indoor lighting estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6908–6917, 2019. 2
- [13] Khang Truong Giang, Soohwan Song, and Sungho Jo. Curvature-guided dynamic scale networks for multi-view stereo. *arXiv preprint arXiv:2112.05999*, 2021. 3
- [14] Roger Grosse, Micah K Johnson, Edward H Adelson, and William T Freeman. Ground truth dataset and baseline evaluations for intrinsic image algorithms. In *2009 IEEE 12th International Conference on Computer Vision*, pages 2335–2342. IEEE, 2009. 1
- [15] Berthold KP Horn and Michael J Brooks. *Shape from shading*. MIT press, 1989. 1
- [16] Brian Karis and Epic Games. Real shading in unreal engine 4. *Proc. Physically Based Shading Theory Practice*, 4(3):1, 2013. 4, 8
- [17] Kichang Kim, Akihiko Torii, and Masatoshi Okutomi. Multi-view inverse rendering under arbitrary illumination and albedo. In *European conference on computer vision*, pages 750–767. Springer, 2016. 2, 8
- [18] Wenbin Li, Sajad Saeedi, John McCormac, Ronald Clark, Dimos Tzoumanikas, Qing Ye, Yuzhong Huang, Rui Tang, and Stefan Leutenegger. InteriorNet: Mega-scale multi-sensor photo-realistic indoor scenes dataset. *arXiv preprint arXiv:1809.00716*, 2018. 2
- [19] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2475–2484, 2020. 1, 2, 3, 5, 6, 7
- [20] Zhengqin Li, Jia Shi, Sai Bi, Rui Zhu, Kalyan Sunkavalli, Miloš Hašan, Zexiang Xu, Ravi Ramamoorthi, and Manmohan Chandraker. Physically-based editing of indoor scene lighting from a single image. *arXiv preprint arXiv:2205.09343*, 2022. 1, 2, 3
- [21] Zhengqi Li and Noah Snavely. Cgintrinsics: Better intrinsic image decomposition through physically-based rendering. In *Proceedings of the European conference on computer vision (ECCV)*, pages 371–387, 2018. 1
- [22] Zhengqin Li, Kalyan Sunkavalli, and Manmohan Chandraker. Materials for masses: Svbrdf acquisition with a single mobile phone image. In *Proceedings of the European conference on computer vision (ECCV)*, pages 72–87, 2018. 1
- [23] Zhen Li, Lingli Wang, Xiang Huang, Cihui Pan, and Jiaqi Yang. Phyr: Physics-based inverse rendering for panoramic indoor images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12713–12723, 2022. 1, 2
- [24] Zhengqin Li, Zexiang Xu, Ravi Ramamoorthi, Kalyan Sunkavalli, and Manmohan Chandraker. Learning to reconstruct shape and spatially-varying reflectance from a single image. *ACM Transactions on Graphics (TOG)*, 37(6):1–11, 2018. 1
- [25] Zhengqin Li, Ting-Wei Yu, Shen Sang, Sarah Wang, Meng Song, Yuhan Liu, Yu-Ying Yeh, Rui Zhu, Nitesh Gundavarapu, Jia Shi, et al. Openrooms: An open framework for photorealistic indoor scene datasets. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7190–7199, 2021. 1, 2, 5
- [26] Robert Maier, Kihwan Kim, Daniel Cremers, Jan Kautz, and Matthias Nießner. Intrinsic3D: High-quality 3D reconstruction by joint appearance and geometry optimization with

- spatially-varying lighting. In *International Conference on Computer Vision (ICCV)*, 2017. 2
- [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. 3
- [28] Merlin Nimier-David, Zhao Dong, Wenzel Jakob, and Anton Kaplanyan. Material and Lighting Reconstruction for Complex Indoor Scenes with Texture-space Differentiable Rendering. In Adrien Bousseau and Morgan McGuire, editors, *Eurographics Symposium on Rendering - DL-only Track*. The Eurographics Association, 2021. 2
- [29] Geoffrey Oxholm and Ko Nishino. Shape and reflectance from natural illumination. In *European Conference on Computer Vision*, pages 528–541. Springer, 2012. 1
- [30] Julien Philip, Michaël Gharbi, Tinghui Zhou, Alexei A Efros, and George Drettakis. Multi-view relighting using a geometry-aware network. *ACM Trans. Graph.*, 38(4):78–1, 2019. 3
- [31] Julien Philip, Sébastien Morgenthaler, Michaël Gharbi, and George Drettakis. Free-viewpoint indoor neural relighting from multi-view stereo. *ACM Transactions on Graphics (TOG)*, 40(5):1–18, 2021. 2, 3
- [32] Shen Sang and Manmohan Chandraker. Single-shot neural relighting and svbrdf estimation. In *European Conference on Computer Vision*, pages 85–101. Springer, 2020. 1
- [33] Shuran Song and Thomas Funkhouser. Neural illumination: Lighting prediction for indoor environments. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6918–6926, 2019. 2
- [34] 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. 3
- [35] Pratul P Srinivasan, Ben Mildenhall, Matthew Tancik, Jonathan T Barron, Richard Tucker, and Noah Snavely. Lighthouse: Predicting lighting volumes for spatially-coherent illumination. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8080–8089, 2020. 2, 3, 6, 7
- [36] Jiajun Tang, Yongjie Zhu, Haoyu Wang, Jun-Hoong Chan, Si Li, and Boxin Shi. Estimating spatially-varying lighting in urban scenes with disentangled representation. In *ECCV*, 2022. 2
- [37] 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. 3
- [38] Qianqian Wang, Zhicheng Wang, Kyle Genova, Pratul P Srinivasan, Howard Zhou, Jonathan T Barron, Ricardo Martin-Brualla, Noah Snavely, and Thomas Funkhouser. Ibrnet: Learning multi-view image-based rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4690–4699, 2021. 5, 6, 7
- [39] Zian Wang, Wenzheng Chen, David Acuna, Jan Kautz, and Sanja Fidler. Neural light field estimation for street scenes with differentiable virtual object insertion. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. 2
- [40] Zian Wang, Jonah Philion, Sanja Fidler, and Jan Kautz. Learning indoor inverse rendering with 3d spatially-varying lighting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12538–12547, 2021. 1, 2, 3, 5, 8
- [41] Yu-Ying Yeh, Zhengqin Li, Yannick Hold-Geoffroy, Rui Zhu, Zexiang Xu, Miloš Hašan, Kalyan Sunkavalli, and Manmohan Chandraker. Photoscene: Photorealistic material and lighting transfer for indoor scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18562–18571, June 2022. 2, 8
- [42] Ye Yu and William A. P. Smith. Outdoor inverse rendering from a single image using multiview self-supervision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. to appear. 2
- [43] Li Yuan, Yunpeng Chen, Tao Wang, Weihao Yu, Yujun Shi, Zi-Hang Jiang, Francis EH Tay, Jiashi Feng, and Shuicheng Yan. Tokens-to-token vit: Training vision transformers from scratch on imagenet. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 558–567, 2021. 2
- [44] Edward Zhang, Michael F. Cohen, and Brian Curless. Emptying, refurbishing, and relighting indoor spaces. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia 2016)*, 35(6), 2016. 2
- [45] 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. 3
- [46] Ruo Zhang, Ping-Sing Tsai, James Edwin Cryer, and Mubarak Shah. Shape-from-shading: a survey. *IEEE transactions on pattern analysis and machine intelligence*, 21(8):690–706, 1999. 1
- [47] 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. 3
- [48] Yuanqing Zhang, Jiaming Sun, Xingyi He, Huan Fu, Rongfei Jia, and Xiaowei Zhou. Modeling indirect illumination for inverse rendering. In *CVPR*, 2022. 2, 3
- [49] Jingsen Zhu, Fujun Luan, Yuchi Huo, Zihao Lin, Zhihua Zhong, Dianbing Xi, Rui Wang, Hujun Bao, Jiaxiang Zheng, and Rui Tang. Learning-based inverse rendering of complex indoor scenes with differentiable monte carlo raytracing. In *SIGGRAPH Asia 2022 Conference Papers*. ACM, 2022. 1, 2
- [50] Rui Zhu, Zhengqin Li, Janarбек Matai, Fatih Porikli, and Manmohan Chandraker. Irisformer: Dense vision transformers for single-image inverse rendering in indoor scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2822–2831, 2022. 1, 2, 3, 5