

# Differentiable Inverse Rendering with Interpretable Basis BRDFs

Hoon-Gyu Chung    Seokjun Choi    Seung-Hwan Baek  
 POSTECH

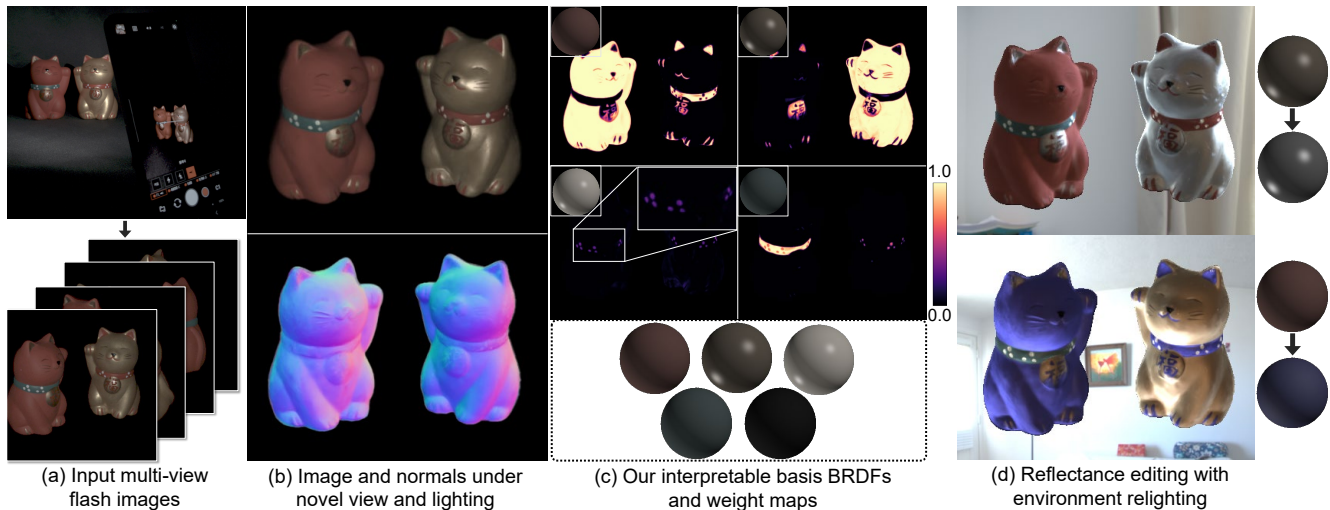


Figure 1. We propose a differentiable inverse rendering method with (a) multi-view flash photography inputs. Our analysis-by-synthesis method achieves not only (b) novel-view relighting and accurate geometry reconstruction, but also (c) interpretable basis BRDFs and their spatially-separated weights. This allows for (d) intuitive scene editing.

## Abstract

Inverse rendering seeks to reconstruct both geometry and spatially varying BRDFs (SVBRDFs) from captured images. To address the inherent ill-posedness of inverse rendering, basis BRDF representations are commonly used, modeling SVBRDFs as spatially varying blends of a set of basis BRDFs. However, existing methods often yield basis BRDFs that lack intuitive separation and have limited scalability to scenes of varying complexity. In this paper, we introduce a differentiable inverse rendering method that produces interpretable basis BRDFs. Our approach models a scene using 2D Gaussians, where the reflectance of each Gaussian is defined by a weighted blend of basis BRDFs. We efficiently render an image from the 2D Gaussians and basis BRDFs using differentiable rasterization and impose a specular-weighted rendering loss with the input flash photography images. During this analysis-by-synthesis optimization process of differentiable inverse rendering, we dynamically adjust the number of basis BRDFs to fit the target scene while encouraging sparsity in the basis weights. This ensures that the reflectance of each Gaussian is represented by only a few basis BRDFs. This approach enables

the reconstruction of accurate geometry and interpretable basis BRDFs that are spatially separated. Consequently, the resulting scene representation, comprising basis BRDFs and 2D Gaussians, supports physically-based novel-view relighting and intuitive scene editing.

## 1. Introduction

Inverse rendering aims to reconstruct geometry and reflectance from captured images, a fundamental problem in computer vision and computer graphics. Recent advances in differentiable rendering methods [18, 29] have facilitated analysis-by-synthesis differentiable inverse rendering. These methods enable the optimization of scene parameters by minimizing the discrepancy between captured and synthesized images [2, 8, 12, 45, 47]. However, achieving accurate inverse rendering remains challenging, especially when the available lighting and viewing angles are limited. A promising approach to mitigate this problem is to exploit the spatial coherence of SVBRDFs using basis BRDF representations [1, 19, 27, 28, 32]. Representing the BRDF at each point as a spatially varying blend of basis BRDFs allows gathering multiple points to accurately fit basis BRDFs

and their per-point weights.

Early methods often focus solely on estimating basis BRDFs and their weights for objects with known 3D geometry [20, 21, 50]. Moreover, they typically require excessive optimization time and suffer from low accuracy. Recent methods jointly optimize geometry and basis BRDFs [2, 8, 30]. These approaches often result in non-interpretable basis BRDFs which are spatially entangled: each scene point is represented by many basis BRDFs with high weights and those basis BRDFs are often non-interpretable. Thus, their results are impractical to use for downstream tasks such as scene editing. Additionally, they are often limited to using a fixed number of basis BRDFs, regardless of scene complexity.

In this paper, we propose a differentiable inverse rendering method that jointly estimates geometry and reflectance. We represent geometry as a set of 2D Gaussians [9, 14], each with shape parameters and basis BRDF weights. The reflectance of each Gaussian is modeled as a spatially varying blend of basis BRDFs. We dynamically adjust the number of basis BRDFs during the analysis-by-synthesis optimization, promoting sparsity in the basis BRDF weights. To enhance training stability, we employ a weighted photometric loss that focuses on potentially specular regions. We demonstrate that our method obtains not only accurate geometry but also scalable and interpretable basis BRDFs which sparsely represent SVBRDFs as shown in Figure 1.

In summary, we make the following contributions:

- A differentiable inverse rendering method that jointly estimates 2D Gaussians and basis BRDFs, obtaining interpretable basis BRDFs.
- A basis BRDF control method and sparsity regularizer that dynamically adjust the number of basis BRDFs during analysis-by-synthesis optimization, promoting interpretability through sparse blending.
- Extensive evaluation compared with existing methods, and demonstration of intuitive scene editing.

## 2. Related Work

**Inverse Rendering** Learning-based inverse rendering methods train neural networks such as CNNs [22, 23, 33, 34, 38, 43], transformers [51], and diffusion models [25, 31, 40] on datasets, enabling efficient inference by exploiting prior distributions in the training data. They often suffer from out-of-distribution inputs, resulting in physically inaccurate results.

The analysis-by-synthesis framework enables inverse rendering by inverting the forward rendering process. Recent advancements have been particularly sparked by differentiable rendering techniques using volumetric rendering [3–7, 10, 11, 17, 26, 36, 37, 39, 41, 44–48] and rasterization [2, 12, 16, 24, 35, 52]. Recent rasterization-based methods using isotropic points [8], 2D Gaussians [49], and

3D Gaussians [2, 12, 16, 24, 35, 52] exhibit both efficiency and accuracy. However, they suffer when insufficient light-view angular samples are provided. Incorporating basis BRDFs into the modern analysis-by-synthesis framework has shown promising performance, however the reconstructed basis BRDFs are often non-interpretable, limiting their usefulness in downstream tasks such as scene editing [2, 8]. In contrast, our novel method obtains 2D Gaussians with accurate geometry and interpretable basis BRDFs.

**Basis BRDF** Basis BRDFs have often been used as a representation of SVBRDFs, decomposing the SVBRDF into a blend of basis BRDFs. This approach is based on the spatial coherence of reflectance, assuming that the BRDFs of many pixels are similar in real-world scenes [27, 28]. Using basis BRDFs helps reduce the ill-posedness of inverse rendering because it enables gathering many light-angular observations for each basis BRDF reconstruction. Existing methods jointly optimize basis BRDFs and spatially varying weights for planar [19, 32] and non-planar objects [1, 13, 15, 30]. Recently, two methods [2, 8] demonstrated the effectiveness of basis BRDFs in modern differentiable analysis-by-synthesis frameworks. While they showed promising results, the estimated basis BRDFs are non-interpretable as shown in Figure 6, which limits their utility in downstream tasks such as scene editing. Moreover, generalizing their methods to varying degrees of scene complexity is challenging because the number of basis BRDFs is fixed. To overcome these problems, we are inspired by the decade-ago work of Zhou et al. [50], which proposed adjusting the number of basis BRDFs during optimization. Their method assumes known geometry, does not exploit a differentiable rendering pipeline, and suffers from low reconstruction accuracy and efficiency. We address these challenges by developing a differentiable inverse rendering method that jointly estimates 2D Gaussians and basis BRDFs to obtain interpretable basis BRDFs, where the number of basis BRDFs are automatically controlled.

## 3. Method

Our method consists of two parts. In Section 3.1, we focus on optimizing 2D Gaussians and basis BRDFs without concerning the interpretability of basis BRDFs. Next, we extend the method to obtaining interpretable basis BRDFs in Section 3.2. Figure 2 shows the process of our method over analysis-by-synthesis iterations. We obtain not only accurate geometry but also interpretable basis BRDFs whose number adapts to the scene and the basis BRDFs represent SVBRDFs with spatially-separated basis-BRDF weights.

### 3.1. Gaussian Inverse Rendering with Basis BRDFs

We introduce our method that jointly estimates geometry and a fixed number of basis BRDFs in a differentiable man-

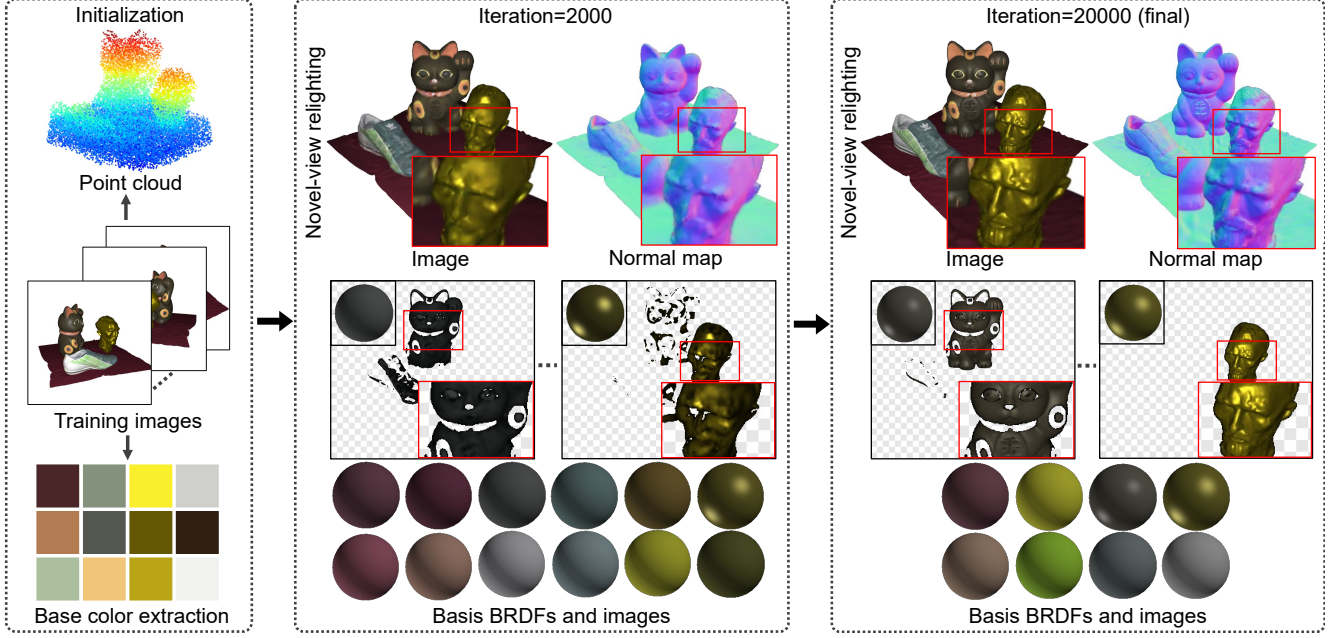


Figure 2. **The process of our analysis-by-synthesis iterations.** Given a set of multi-view photometric images, we initialize point cloud and extract base color for basis BRDFs. We jointly optimize 2D Gaussians and basis BRDFs by comparing the differentially-rendered images and the input images. Our method enables obtaining interpretable basis BRDFs with spatially-separated basis-BRDF weights and the number of basis BRDFs adapts to the scene.

ner from multi-view flash-photography images.

**Geometry** We use 2D Gaussians as geometric primitives [14], where each Gaussian  $g \in G$  is an elliptical disk, and  $G$  is the set of all Gaussians. Each Gaussian  $g$  is parameterized by its center location  $\mathbf{p} \in \mathbb{R}^{3 \times 1}$ , principal vectors  $\mathbf{t} \in \mathbb{R}^{3 \times 2}$ , and scaling factors  $\mathbf{s} \in \mathbb{R}^{2 \times 1}$ . The Gaussian surface normal  $\mathbf{n}$  is computed as the cross product of the principal vectors.

**Basis BRDFs** We use a basis BRDF representation, where we define  $N$  basis BRDFs. We model the  $i$ -th basis BRDF  $f_i$  with base color  $\mathbf{b}_i$ , roughness  $\sigma_i$ , and metallic parameter  $m_i$  using the simplified Disney BRDF model [42]:

$$f_i(\mathbf{i}, \mathbf{o}) = \frac{1 - m_i}{\pi} \mathbf{b}_i + \frac{D(\mathbf{h}; \sigma_i) F(\mathbf{o}, \mathbf{h}; \mathbf{b}_i, m_i) G(\mathbf{i}, \mathbf{o}, \mathbf{n}; \sigma_i)}{4(\mathbf{n} \cdot \mathbf{i})(\mathbf{n} \cdot \mathbf{o})}, \quad (1)$$

where  $\mathbf{i}$  and  $\mathbf{o}$  are the incident and outgoing directions, and  $\mathbf{h} = (\mathbf{i} + \mathbf{o})/2$  is the half-way vector.  $D$  is the normal distribution function,  $F$  is the Fresnel term, and  $G$  is the geometric attenuation. Details of these terms are provided in the Supplemental Document.

We assign blending weights  $w_i(g)$  for every  $i$ -th basis BRDF per each Gaussian  $g$ , and represent the BRDF  $f$  of

the Gaussian  $g$  as:

$$f(\mathbf{i}, \mathbf{o}; g) = \sum_{i=1}^N w_i(g) f_i(\mathbf{i}, \mathbf{o}). \quad (2)$$

**Scene Representations** In summary, our scene representation consists of per-Gaussian parameters  $G$  and basis BRDF parameters  $R$ :

$$g = \underbrace{\{\mathbf{p}, \mathbf{t}, \mathbf{s}\}}_{\text{geometric}} \cup \underbrace{\{\{w_i(g)\}_{i=1}^N, \alpha\}}_{\text{photometric}}, \text{ where } g \in G, \\ R = \{m_i, \mathbf{b}_i, \sigma_i\}_{i=1}^N. \quad (3)$$

$\alpha$  is the Gaussian opacity. To initialize the basis BRDF parameters  $R$ , we perform k-means clustering with a fixed number  $N$  on the input multi-view flash photography images  $I'$ . The base colors  $\{\mathbf{b}_i\}_{i=1}^N$  are initialized as the mean values of the clusters. The roughness  $\{\sigma_i\}_{i=1}^N$  and metallic parameters  $\{m_i\}_{i=1}^N$  are set to initial values:  $\sigma_i = 0.5$  and  $m_i = 0.0$ . We initialize the weights  $w_i(g)$  uniformly and set  $\alpha$  following Huang et al. [14].

**Differentiable Rendering** We render an image under the flash light, modeled as a point light source, using differentiable rasterization. Specifically, we compute the radiance  $L$  of each Gaussian  $g$  as:

$$L(\mathbf{i}, \mathbf{o}; g) = (\mathbf{n} \cdot \mathbf{i}) f(\mathbf{i}, \mathbf{o}; g) E(g), \quad (4)$$

where  $E(g)$  is the incident light intensity on Gaussian  $g$ .

After computing the radiance, we rasterize the Gaussians onto the image plane, sort them by depth, and accumulate radiance values using alpha blending [14] to produce the final pixel value of the rendered image  $I$ . We define this rendering as a function  $\text{render}(\cdot)$ :

$$I(u) = \text{render}(\{L_i\}_{i=1}^M) \\ = \sum_{i=1}^M L_i \alpha_i \mathcal{G}_i(\mathbf{r}(u)) \prod_{j=1}^{i-1} (1 - \alpha_j \mathcal{G}_j(\mathbf{r}(u))), \quad (5)$$

where  $u \in U$  is a pixel,  $M$  is the number of Gaussians projected onto pixel  $u$ .  $\{L_i\}_{i=1}^M$  and  $\{\alpha_i\}_{i=1}^M$  are the radiance and opacity values of the depth-sorted  $i$ -th Gaussian, respectively. The vector  $\mathbf{r}(u)$  is the ray coming from corresponding pixel  $u$ , and  $\mathcal{G}_i(\mathbf{r}(u))$  is the Gaussian-filtered distance between the center of  $i$ -th Gaussian and the intersection point between the ray  $\mathbf{r}(u)$  and the  $i$ -th Gaussian. Details are provided in the Supplemental Document.

**Optimization** We optimize the Gaussian parameters  $G$  and the basis BRDF parameters  $R$  by minimizing the loss function  $\mathcal{L}$  in an analysis-by-synthesis manner:

$$\mathcal{L} = \mathcal{L}_{\text{render}} + \lambda_{\text{geom}} \mathcal{L}_{\text{geom}} + \lambda_{\text{mask}} \mathcal{L}_{\text{mask}}, \quad (6)$$

where  $\mathcal{L}_{\text{render}}$  penalizes the difference between the rendered image  $I$  and the input image  $I'$  across all input views.  $\mathcal{L}_{\text{geom}}$  is a geometric regularization term for depth distortion and normal consistency [14], and  $\mathcal{L}_{\text{mask}}$  is the cross-entropy loss between the rendered mask and the ground-truth mask.  $\lambda_{\text{geom}}$  and  $\lambda_{\text{mask}}$  are the balancing weights. Details are provided in the Supplemental Document.

**Specular-weighted Rendering Loss** Multi-view flash photography typically generates fewer observations of specular-dominant pixels compared to diffuse-dominant pixels. To address this imbalance, we weight specular observations by using the potentially-specular weight map:

$$H(u) = 1 + \lambda_{\theta_h} (\cos \theta_h(u))^k, \quad (7)$$

where  $\theta_h$  is the rendered half-way angle map as  $\theta_h(u) = \text{render}(\{\arccos(\mathbf{n}_i \cdot \mathbf{h}_i)\}_{i=1}^M)$ .  $\mathbf{n}_i$  and  $\mathbf{h}_i$  are the normal and half-way vector of Gaussian  $i$ , respectively. We set the scalar  $\lambda_{\theta_h} = 5$  and  $k = 10$  to strongly weight specular regions.

We use the potentially-specular weight map  $H$  for calculating the final rendering loss:

$$\mathcal{L}_{\text{render}} = \frac{1}{|U|} \sum_{u \in U} H(u) ((1 - \lambda_s) \mathcal{L}_1(u) + \lambda_s \mathcal{L}_{\text{SSIM}}(u)), \quad (8)$$

where  $\mathcal{L}_1(u) = \|I(u) - I'(u)\|_1$ ,  $\mathcal{L}_{\text{SSIM}}(u)$  is the SSIM loss between  $I(u)$  and  $I'(u)$ , and  $\lambda_s$  is a balancing weight.

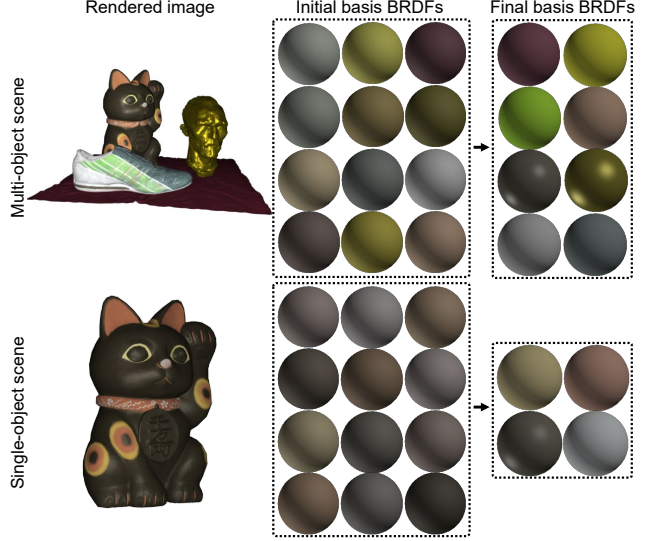


Figure 3. **Basis BRDFs with varying scene complexity.** We adjust the number of basis BRDFs during optimization to adapt to the scene complexity. We initialize the same number of basis BRDFs for both scenes in this example.

### 3.2. Interpretable Basis BRDFs

While our method in Section 3.1 offers accurate inverse rendering, the reconstructed basis BRDFs lack of interpretability. Here, we extend the method to obtain interpretable basis BRDFs whose number adapts to the scene as shown in Figure 3. We develop basis BRDF control method of merge and removal in addition to impose sparsity on basis BRDF weights.

**Sparsity of Basis BRDF Weights** For obtaining interpretable basis BRDF, each scene point needs to be represented with only few basis BRDFs with high weights. This enables spatially separating SVBRDFs with sparse basis BRDFs. We impose sparsity on the basis BRDF weights  $\{w_i(g)\}_{i=1}^N$  for each Gaussian  $g$ , making each Gaussian represented by only a few basis BRDFs.

First, we apply a softmax function with a low temperature  $T = 0.0125$  to the weights, promoting sparse per-Gaussian weights:

$$w_i(g) \leftarrow \frac{\exp(w_i(g)/T)}{\sum_{i'=1}^N \exp(w_{i'}(g)/T)}. \quad (9)$$

Second, we apply an entropy-based sparsity regularizer on the Gaussian weight  $w_i$  and rendered weight image  $W_i$  for

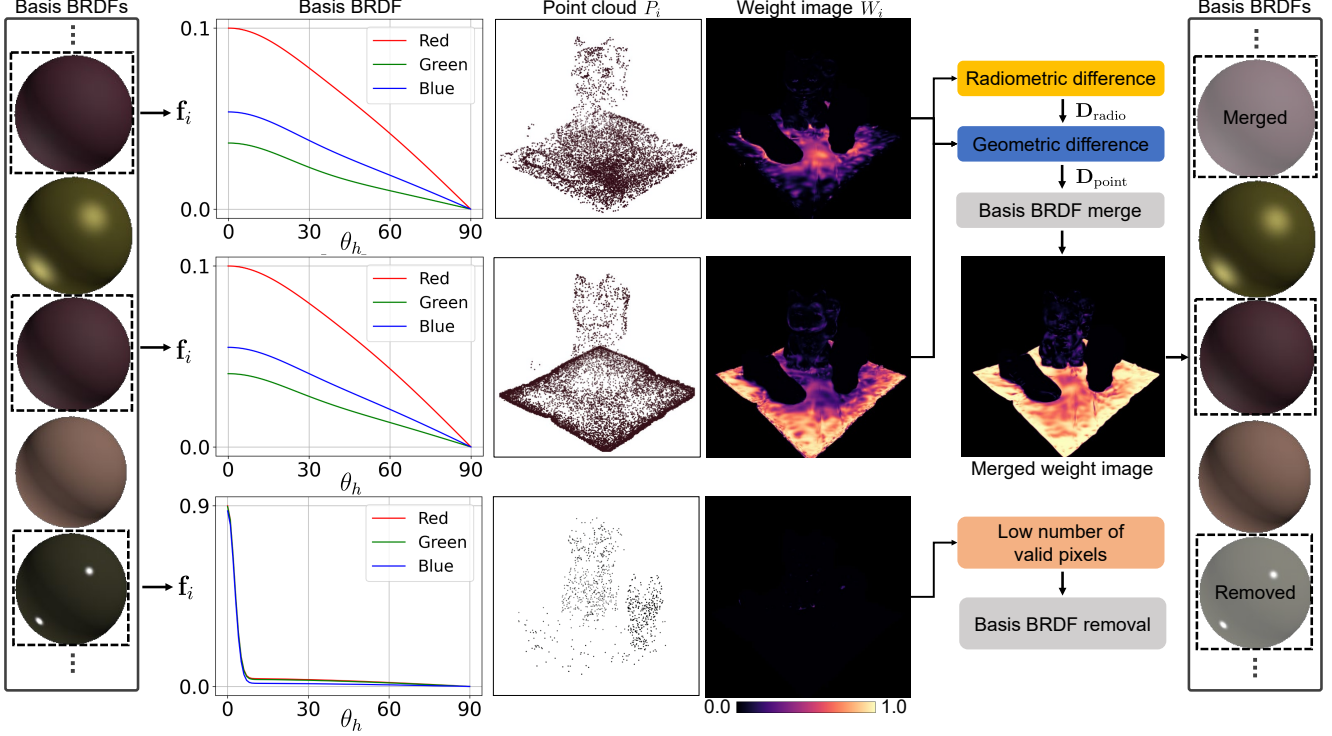


Figure 4. **Basis BRDF control.** During the optimization, we compute the values of each basis BRDF for sampled half-way angles  $\theta_h$  from which radiometric difference is obtained. We compute the geometric difference between point clouds of basis BRDFs. If two basis BRDFs are radiometrically and geometrically similar, we merge them. If the rendered weight map  $W_i$  has few valid pixels, we remove the basis BRDF.

each  $i$ -th basis BRDF:

$$\mathcal{L}_{\text{sparse}} = -\frac{1}{|G|} \sum_{g \in G} \sum_{i=1}^N w_i(g) \log w_i(g) - \frac{1}{|U|} \sum_{u \in U} \sum_{i=1}^N W_i(u) \log W_i(u), \quad (10)$$

$$W_i = \text{render}(\{w_i(g)\}_{g \in G}). \quad (11)$$

**Basis BRDF Merge** Basis BRDFs often become similar during optimization, making the resulting basis BRDFs non-intuitive with duplicates. We address this by merging similar basis BRDFs during analysis-by-synthesis optimization, as shown in Figure 4. We compute the values of each basis BRDF  $f_i$  at uniformly-sampled half-way angles  $\theta_h$  for the RGB channels, resulting in a matrix  $\mathbf{f}_i \in \mathbb{R}^{C \times 3}$ .  $C$  is the number of half-way angles. We then compute the radiometric difference between pairs of basis BRDFs, represented as a matrix  $\mathbf{D}_{\text{radio}} \in \mathbb{R}^{N \times N}$ , which is defined as:

$$\mathbf{D}_{\text{radio}}(i, j) = \frac{1}{C} \|\mathbf{f}_i - \mathbf{f}_j\|_2. \quad (12)$$

We also compute the geometric difference between the every pair of basis BRDFs, represented as a matrix  $\mathbf{D}_{\text{point}} \in$

$\mathbb{R}^{N \times N}$  using the Chamfer distance between point clouds represented by each basis BRDF:

$$\mathbf{D}_{\text{point}}(i, j) = \frac{1}{|P_i|} \sum_{\mathbf{p} \in P_i} \min_{\mathbf{q} \in P_j} \|\mathbf{p} - \mathbf{q}\|_2^2 + \frac{1}{|P_j|} \sum_{\mathbf{q} \in P_j} \min_{\mathbf{p} \in P_i} \|\mathbf{q} - \mathbf{p}\|_2^2, \quad (13)$$

where  $P_i$  and  $P_j$  are the point clouds for the  $i$ -th and  $j$ -th basis BRDFs, consisting of the center locations of Gaussians where the  $i$ -th or  $j$ -th basis BRDF has the highest weight.

If the radiometric difference  $\mathbf{D}_{\text{radio}}(i, j)$  is below a threshold  $\tau_{\text{merge}}$  and  $\mathbf{D}_{\text{point}}(i, j)$  is minimal among all  $\forall j$ , we merge the two basis BRDFs by deleting the basis BRDF with fewer associated Gaussians and reassigning the per-Gaussian basis BRDF weights as

$$w_j(g) \leftarrow w_i(g) + w_j(g), \quad \forall g \in G. \quad (14)$$

**Basis BRDF Removal** To obtain meaningful basis BRDFs only, we remove  $i$ -th basis BRDF if it contributes to only a few pixels during optimization as shown in Figure 4. We evaluate this using the rendered weight image  $W_i$

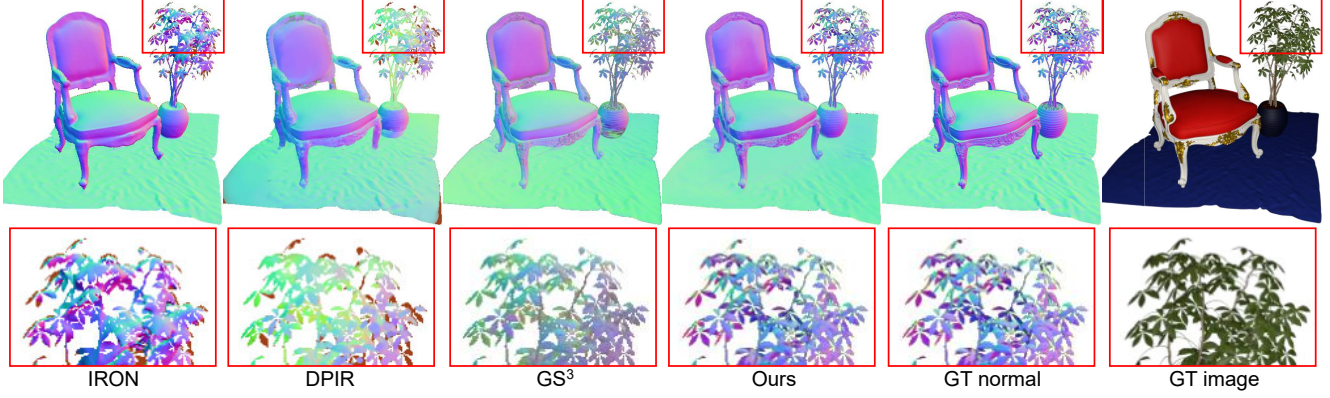


Figure 5. **Normal reconstruction.** Our method successfully recovers the detailed surface normal, outperforming other state-of-the-art inverse rendering methods: IRON [45], DPIR [8], GS<sup>3</sup> [2].

from Equation (11) as:

$$\frac{|\{u \in U | W_i(u) > \tau_{\text{removal-weight}}\}|}{|U|} < \tau_{\text{removal-number}}, \quad (15)$$

where  $\tau_{\text{removal-weight}} = 0.1$  is the threshold determining whether the  $i$ -th basis BRDF significantly contributes to that pixel  $u$ , and  $\tau_{\text{removal-number}} = 0.005$  is the threshold for the normalized number of such pixels.

**Scheduling** For a warm start, we first run our inverse rendering described in Section 3.1. After a predefined number of iterations, 6000 in our experiments, we perform basis BRDF merge and removal at predefined intervals 500 during iterations to refine the basis BRDFs. We incorporate the sparsity loss  $\mathcal{L}_{\text{sparse}}$  and optimize the scene parameters by minimizing:

$$\mathcal{L} = \mathcal{L}_{\text{render}} + \lambda_{\text{geom}} \mathcal{L}_{\text{geom}} + \lambda_{\text{mask}} \mathcal{L}_{\text{mask}} + \lambda_{\text{sparse}} \mathcal{L}_{\text{sparse}}, \quad (16)$$

where  $\lambda_{\text{sparse}}$  is the balancing weight. Refer to the Supplemental Document for the details on the optimization.

## 4. Results

**Datasets** We evaluate our approach on both synthetic and real-world photometric dataset. For the synthetic dataset, we render four complex scenes containing multiple objects, following the co-located flash photography. Each scene includes multiple objects. The dataset consists of 200 training images and 100 test images per scene. For the real-world photometric dataset, we capture one multi-object scene using a mobile phone with flash at 200/40 training/testing images. We conduct geometric and radiometric calibrations of the light and the camera.

### 4.1. Validation and Comparison

**Geometry** We assess our method in comparison with state-of-the-art analysis-by-synthesis inverse rendering methods: IRON [45], DPIR [8] and GS<sup>3</sup> [2]. Figure 5

	IRON	DPIR	GS <sup>3</sup>	Ours
MAE ↓	15.71	17.04	15.28	<b>9.81</b>
Train ↓	12h	2h	1h	<b>0.5h</b>

Table 1. **Normal reconstruction and training time.** Our method outperforms IRON [45], DPIR [8], and GS<sup>3</sup> [2] in quantitative numbers evaluated on the synthetic dataset.

and Table 1 show the results for geometry reconstruction, demonstrating the highest accuracy in normal reconstruction with mean angular error (MAE). IRON and DPIR employ neural signed distance functions to recover smooth surface normal, not only resulting in loss of details but also requiring extended training times for optimization as shown in Table 1. GS<sup>3</sup>, which is based on 3D Gaussian splatting with anisotropic spherical Gaussians, exhibits limited accuracy in shape reconstruction. Our method accurately models surface geometry and reflectance using 2D Gaussians and interpretable basis BRDFs, achieving the highest accuracy of normal reconstruction, especially for thin and convex objects. Moreover, we outperform the others in terms of training speed.

**Basis BRDF** Figure 6 shows the estimated basis BRDFs and weight maps, comparing modern inverse-rendering methods using basis BRDFs: GS<sup>3</sup>, DPIR, our method without interpretable basis BRDFs (Section 3.1), and our method with interpretable basis BRDFs (Section 3.2). Both DPIR and GS<sup>3</sup> use basis BRDFs to leverage spatial coherence of specular reflectance. These methods yield non-interpretable basis BRDFs, as evident from the rendered weight maps and basis BRDF renderings. Our method introduced in Section 3.1 improves the interpretability of the basis BRDFs by jointly optimizing base colors. However, the resulting basis BRDFs are still often duplicated and non-intuitively classified. Moreover, all the methods have a fixed number of basis BRDFs, limited to modeling different complexity of scenes. In contrast, our full method (Section 3.2) obtains interpretable basis BRDFs whose number fits the target scene, as shown in Figure 3.

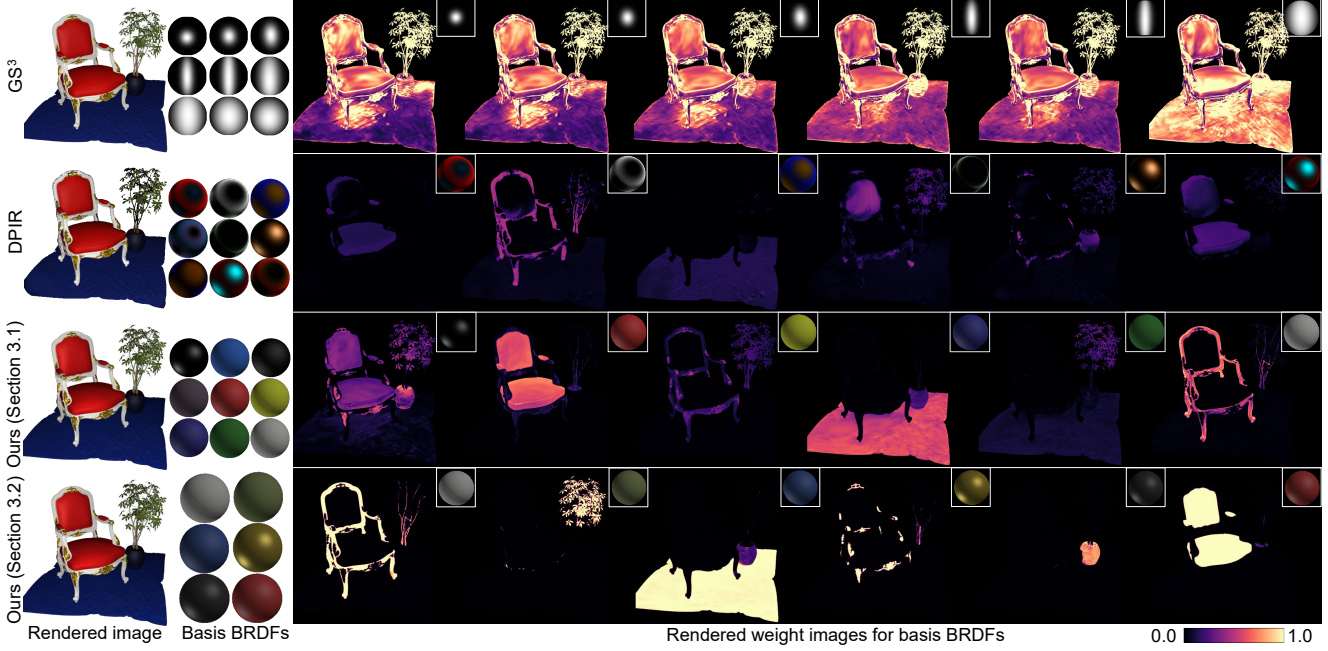


Figure 6. **Basis BRDFs and weight maps.** Our inverse-rendering methods (Section 3.1 and Section 3.2) obtain basis BRDFs that are explainable with spatially separated weight maps and intuitive basis BRDFs, compared to  $GS^3$  [2] and DPIR [8].

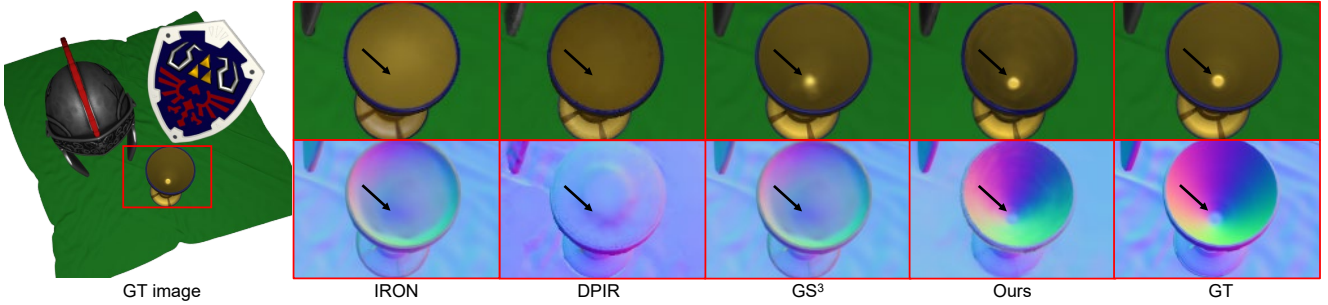


Figure 7. **Novel-view relighting.** Our method enables high quality novel-view relighting with the accurate surface normal. We achieve physically-valid relighting with accurate normal and basis BRDFs.

## 4.2. Applications

**Novel-view Relighting** Figure 7 shows novel-view relighting and estimated normals, compared with IRON, DPIR,  $GS^3$ . They suffer from inaccurate estimation of geometry. Despite inaccurate normals,  $GS^3$  manages to render high-quality relighting results. This is because neural networks compensate for such imperfection by learning residuals, as relighting is the main goal of  $GS^3$ . In contrast, our method reconstructs accurate geometry and interpretable basis BRDFs that enable not only novel-view relighting but also downstream tasks such as intuitive scene editing. For more results, refer to the Supplemental Document.

**Scene Editing** Our inverse rendering method enables selective editing based on basis BRDFs and weight images. Figure 8 shows that our interpretable basis BRDFs facilitate reflectance editing. Interpretable basis BRDFs enable prac-

tical material editing by adjusting desired BRDF parameters of base color, roughness, metallic based on their sparse weight maps. We then obtain a single object from the entire scene by selecting basis BRDFs that are utilized to represent the corresponding object. Next, we delete other basis BRDFs and Gaussians that have the highest weight for other basis BRDFs. We also filter out Gaussians based on Gaussian positions. This allows for extracting object mesh [14] and relighting it under an environment map. Figure 9 shows additional material-editing results on a real-world scene.

## 4.3. Ablation Study

**Sparsity Regularizer  $\mathcal{L}_{\text{sparse}}$**  Figure 10 shows that the sparsity regularizer and the basis BRDF control enable obtaining interpretable basis BRDFs.

**Specular-weighted Rendering Loss** We test the impact of the specular-weighted rendering loss that prioritizes

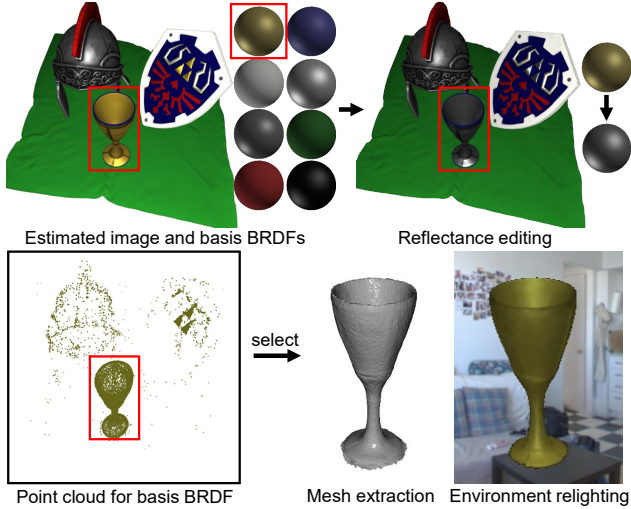


Figure 8. **Intuitive scene editing.** We modify our basis BRDF parameters to edit the reflectance. We also selectively extract mesh and relight it under the environment map.

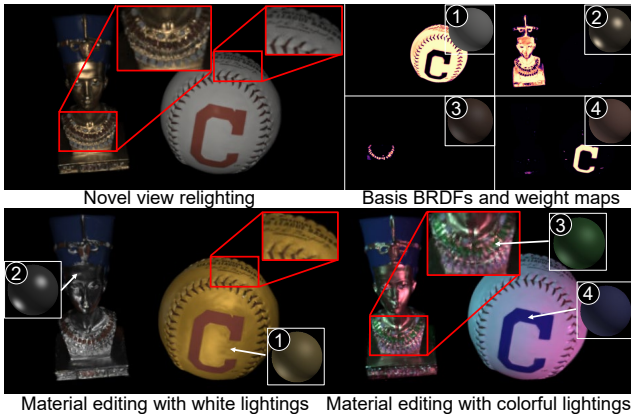


Figure 9. **Material editing and relighting on real-world objects.** We obtain high-quality novel-view relighting for a real-world scene, where interpretable basis BRDFs are reconstructed with their sparse weight maps. We edit basis BRDFs and relight the scene under novel illumination conditions.

potentially-specular areas using the rendered half-way angle map. Figure 11 shows that specular-weighted rendering loss produces a more stable reconstruction compared to not using the weighting.

## 5. Conclusion

In this paper, we have introduced a differentiable inverse rendering method that jointly estimates 2D Gaussians and basis BRDFs. Our method obtains interpretable basis BRDFs that can represent SVBRDFs in a spatially-separable manner, enabling accurate geometry and SVBRDF reconstruction. Also, the number of basis BRDFs automatically adjusts to a target scene. We demonstrate the effectiveness of our method compared to state-of-the-

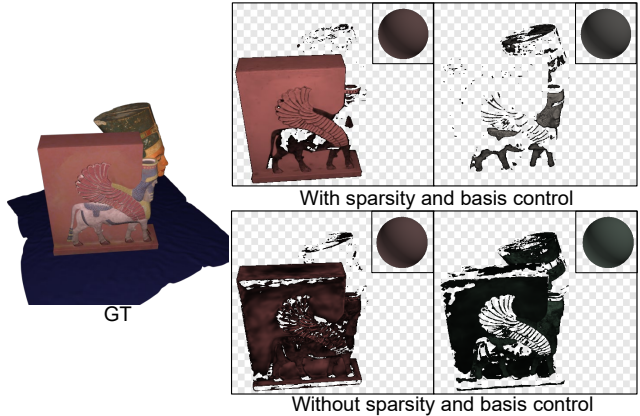


Figure 10. **Impact of sparsity regularizer and basis BRDF control.** Without sparsity loss and basis BRDF control of merge and removal, the resulting basis BRDFs are blended each other as shown in the corresponding images, making them non-interpretable for downstream tasks.

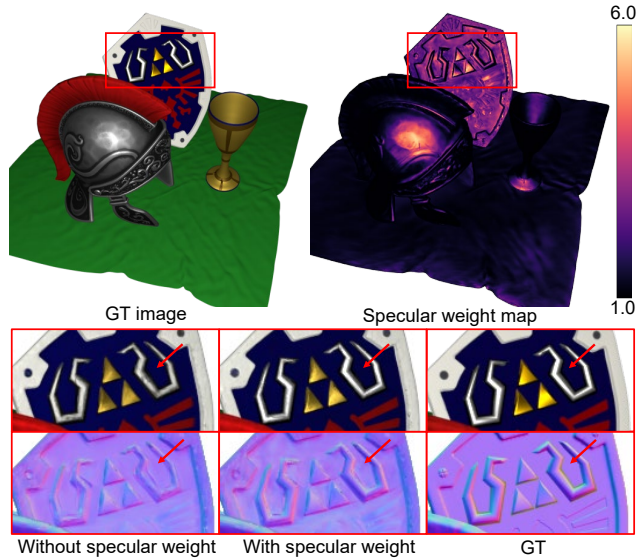


Figure 11. **Impact of specular-weighted loss.** Using specular-weighted rendering loss improves geometry and reflectance reconstruction quality by recovering high frequency features from specularly, which is emphasized by the specular weight map.

art approaches. Using the interpretable basis BRDFs, our method also facilitates downstream tasks such as intuitive scene editing.

**Limitations** Our basis BRDF representation assumes opaque objects, leaving modeling non-surface objects as a future work. Combining semantic information from a large vision model is also an interesting future work.

**Acknowledgments** Seung-Hwan Baek was supported by Korea NRF (RS-2023-00211658, RS-2024-00438532, RS-2025-00518643), and Korea IITP MSIT (RS-2023-00215700).

## References

- [1] Neil Alldrin, Todd Zickler, and David Kriegman. Photometric stereo with non-parametric and spatially-varying reflectance. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2008. 1, 2
- [2] Zoubin Bi, Yixin Zeng, Chong Zeng, Fan Pei, Xiang Feng, Kun Zhou, and Hongzhi Wu. Gs<sup>3</sup>: Efficient relighting with triple gaussian splatting. In *SIGGRAPH Asia 2024 Conference Papers*, 2024. 1, 2, 6, 7
- [3] Mark Boss, Raphael Braun, Varun Jampani, Jonathan T Barron, Ce Liu, and Hendrik Lensch. Nerd: Neural reflectance decomposition from image collections. In *Int. Conf. Comput. Vis.*, pages 12684–12694, 2021. 2
- [4] Mark Boss, Andreas Engelhardt, Abhishek Kar, Yuanzhen Li, Deqing Sun, Jonathan Barron, Hendrik Lensch, and Varun Jampani. Samurai: Shape and material from unconstrained real-world arbitrary image collections. *Advances in Neural Information Processing Systems*, 35:26389–26403, 2022.
- [5] Mark Boss, Varun Jampani, Raphael Braun, Ce Liu, Jonathan Barron, and Hendrik Lensch. Neural-pil: Neural pre-integrated lighting for reflectance decomposition. *Adv. Neural Inform. Process. Syst.*, 34:10691–10704, 2021.
- [6] Baptiste Brument, Robin Bruneau, Yvain Quéau, Jean Mélou, François Bernard Lauze, Jean-Denis Durou, and Lilian Calvet. Rnb-neus: Reflectance and normal-based multi-view 3d reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5230–5239, 2024.
- [7] Ziang Cheng, Junxuan Li, and Hongdong Li. Wildlight: In-the-wild inverse rendering with a flashlight. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4305–4314, 2023. 2
- [8] Hoon-Gyu Chung, Seokjun Choi, and Seung-Hwan Baek. Differentiable point-based inverse rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4399–4409, 2024. 1, 2, 6, 7
- [9] Pinxuan Dai, Jiamin Xu, Wenxiang Xie, Xinguo Liu, Huamin Wang, and Weiwei Xu. High-quality surface reconstruction using gaussian surfels. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–11, 2024. 2
- [10] Andreas Engelhardt, Amit Raj, Mark Boss, Yunzhi Zhang, Abhishek Kar, Yuanzhen Li, Deqing Sun, Ricardo Martin Brualla, Jonathan T Barron, Hendrik Lensch, et al. Shinobi: Shape and illumination using neural object decomposition via brdf optimization in-the-wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19636–19646, 2024. 2
- [11] Fan Fei, Jiajun Tang, Ping Tan, and Boxin Shi. Vminer: Versatile multi-view inverse rendering with near-and far-field light sources. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11800–11809, 2024. 2
- [12] Jian Gao, Chun Gu, Youtian Lin, Hao Zhu, Xun Cao, Li Zhang, and Yao Yao. Relightable 3d gaussian: Real-time point cloud relighting with brdf decomposition and ray tracing. *arXiv preprint arXiv:2311.16043*, 2023. 1, 2
- [13] Dan B Goldman, Brian Curless, Aaron Hertzmann, and Steven M Seitz. Shape and spatially-varying brdfs from photometric stereo. *IEEE Trans. Pattern Anal. Mach. Intell.*, 32(6):1060–1071, 2009. 2
- [14] Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2d gaussian splatting for geometrically accurate radiance fields. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–11, 2024. 2, 3, 4, 7
- [15] Zhuo Hui and Aswin C Sankaranarayanan. Shape and spatially-varying reflectance estimation from virtual exemplars. *IEEE Trans. Pattern Anal. Mach. Intell.*, 39(10):2060–2073, 2016. 2
- [16] Yingwenqi Jiang, Jiadong Tu, Yuan Liu, Xifeng Gao, Xiaoxiao Long, Wenping Wang, and Yuxin Ma. Gaussian-shader: 3d gaussian splatting with shading functions for reflective surfaces. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5322–5332, 2024. 2
- [17] Haian Jin, Isabella Liu, Peijia Xu, Xiaoshuai Zhang, Songfang Han, Sai Bi, Xiaowei Zhou, Zexiang Xu, and Hao Su. Tensor: Tensorial inverse rendering. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 165–174, 2023. 2
- [18] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):1–14, 2023. 1
- [19] Jason Lawrence, Aner Ben-Artzi, Christopher DeCoro, Wojciech Matusik, Hanspeter Pfister, Ravi Ramamoorthi, and Szymon Rusinkiewicz. Inverse shade trees for non-parametric material representation and editing. *ACM Trans. Graph.*, 25(3):735–745, 2006. 1, 2
- [20] Hendrik PA Lensch, Jan Kautz, Michael Goesele, Wolfgang Heidrich, and Hans-Peter Seidel. Image-based reconstruction of spatially varying materials. In *Rendering Techniques 2001: Proceedings of the Eurographics Workshop in London, United Kingdom, June 25–27, 2001 12*, pages 103–114. Springer, 2001. 2
- [21] Hendrik PA Lensch, Jan Kautz, Michael Goesele, Wolfgang Heidrich, and Hans-Peter Seidel. Image-based reconstruction of spatial appearance and geometric detail. *ACM Transactions on Graphics (TOG)*, 22(2):234–257, 2003. 2
- [22] 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 *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2475–2484, 2020. 2
- [23] Zhengqin Li, Zexiang Xu, Ravi Ramamoorthi, Kalyan Sunkavalli, and Manmohan Chandraker. Learning to reconstruct shape and spatially-varying reflectance from a single image. *ACM Trans. Graph.*, 37(6):1–11, 2018. 2
- [24] Zhihao Liang, Qi Zhang, Ying Feng, Ying Shan, and Kui Jia. Gs-ir: 3d gaussian splatting for inverse rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21644–21653, 2024. 2
- [25] Yehonathan Litman, Or Patashnik, Kangle Deng, Aviral Agrawal, Rushikesh Zawat, Fernando De la Torre, and Shubham Tulsiani. Materialfusion: Enhancing inverse rendering with material diffusion priors. *ArXiv*, 2024. 2
- [26] Yuan Liu, Peng Wang, Cheng Lin, Xiaoxiao Long, Jiepeng Wang, Lingjie Liu, Taku Komura, and Wenping Wang. Nero:

- Neural geometry and brdf reconstruction of reflective objects from multiview images. *arXiv preprint arXiv:2305.17398*, 2023. 2
- [27] Wojciech Matusik. *A data-driven reflectance model*. PhD thesis, Massachusetts Institute of Technology, 2003. 1, 2
- [28] Wojciech Matusik, Hanspeter Pfister, Matthew Brand, and Leonard McMillan. Efficient isotropic brdf measurement. 2003. 1, 2
- [29] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. 1
- [30] Giljoo Nam, Joo Ho Lee, Diego Gutierrez, and Min H Kim. Practical svbrdf acquisition of 3d objects with unstructured flash photography. *ACM Trans. Graph.*, 37(6):1–12, 2018. 2
- [31] Yohan Poirier-Ginter, Alban Gauthier, Julien Phillip, J-F Lalonde, and George Drettakis. A diffusion approach to radiance field relighting using multi-illumination synthesis. In *Computer Graphics Forum*, volume 43, page e15147. Wiley Online Library, 2024. 2
- [32] Peiran Ren, Jiaping Wang, John Snyder, Xin Tong, and Bain-ing Guo. Pocket reflectometry. *ACM Transactions on Graphics (TOG)*, 30(4):1–10, 2011. 1, 2
- [33] Shen Sang and Manmohan Chandraker. Single-shot neural relighting and svbrdf estimation. In *Eur. Conf. Comput. Vis.*, pages 85–101. Springer, 2020. 2
- [34] Soumyadip Sengupta, Jinwei Gu, Kihwan Kim, Guilin Liu, David W Jacobs, and Jan Kautz. Neural inverse rendering of an indoor scene from a single image. In *Int. Conf. Comput. Vis.*, pages 8598–8607, 2019. 2
- [35] Yahao Shi, Yanmin Wu, Chenming Wu, Xing Liu, Chen Zhao, Haocheng Feng, Jingtuo Liu, Liangjun Zhang, Jian Zhang, Bin Zhou, et al. Gir: 3d gaussian inverse rendering for relightable scene factorization. *arXiv preprint arXiv:2312.05133*, 2023. 2
- [36] 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 *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 7495–7504, 2021. 2
- [37] Haoyuan Wang, Wenbo Hu, Lei Zhu, and Rynson WH Lau. Inverse rendering of glossy objects via the neural plenoptic function and radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19999–20008, 2024. 2
- [38] Xin Wei, Guojun Chen, Yue Dong, Stephen Lin, and Xin Tong. Object-based illumination estimation with rendering-aware neural networks. In *Eur. Conf. Comput. Vis.*, pages 380–396. Springer, 2020. 2
- [39] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4295–4304, 2023. 2
- [40] Chen Xi, Peng Sida, Yang Dongchen, Liu Yuan, Pan Bowen, Lv Chengfei, and Zhou. Xiaowei. Intrinsically anything: Learning diffusion priors for inverse rendering under unknown illumination. 2024. 2
- [41] Wenqi Yang, Guanying Chen, Chaofeng Chen, Zhenfang Chen, and Kwan-Yee K Wong. Ps-nerf: Neural inverse rendering for multi-view photometric stereo. In *Eur. Conf. Comput. Vis.*, pages 266–284. Springer, 2022. 2
- [42] Yao Yao, Jingyang Zhang, Jingbo Liu, Yihang Qu, Tian Fang, David McKinnon, Yanghai Tsin, and Long Quan. Neirlf: Neural incident light field for physically-based material estimation. In *European Conference on Computer Vision*, pages 700–716. Springer, 2022. 3
- [43] Ye Yu and William AP Smith. Inverserendernet: Learning single image inverse rendering. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3155–3164, 2019. 2
- [44] Chong Zeng, Guojun Chen, Yue Dong, Pieter Peers, Hongzhi Wu, and Xin Tong. Relighting neural radiance fields with shadow and highlight hints. In *ACM SIGGRAPH 2023 Conference Proceedings*, pages 1–11, 2023. 2
- [45] Kai Zhang, Fujun Luan, Zhengqi Li, and Noah Snavely. Iron: Inverse rendering by optimizing neural sdf and materials from photometric images. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5565–5574, 2022. 1, 6
- [46] Kai Zhang, Fujun Luan, Qianqian Wang, Kavita Bala, and Noah Snavely. Physg: Inverse rendering with spherical gaussians for physics-based material editing and relighting. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5453–5462, 2021.
- [47] Xiuming Zhang, Pratul P Srinivasan, Boyang Deng, Paul De-bevec, William T Freeman, and Jonathan T Barron. Ner-factor: Neural factorization of shape and reflectance under an unknown illumination. *ACM Trans. Graph.*, 40(6):1–18, 2021. 1
- [48] Yuanqing Zhang, Jiaming Sun, Xingyi He, Huan Fu, Rongfei Jia, and Xiaowei Zhou. Modeling indirect illumination for inverse rendering. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 18643–18652, 2022. 2
- [49] Yiqun Zhao, Chenming Wu, Binbin Huang, Yihao Zhi, Chen Zhao, Jingdong Wang, and Shenghua Gao. Surfel-based gaussian inverse rendering for fast and relightable dynamic human reconstruction from monocular video. *arXiv preprint arXiv:2407.15212*, 2024. 2
- [50] Zhiming Zhou, Guojun Chen, Yue Dong, David Wipf, Yong Yu, John Snyder, and Xin Tong. Sparse-as-possible svbrdf acquisition. *ACM Transactions on Graphics (TOG)*, 35(6):1–12, 2016. 2
- [51] 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. 2
- [52] Zuo-Liang Zhu, Beibei Wang, and Jian Yang. Gs-ror: 3d gaussian splatting for reflective object relighting via sdf priors. *arXiv preprint arXiv:2406.18544*, 2024. 2