

# 000 RADIOMETRICALLY CONSISTENT GAUSSIAN SURFELS 001 002 FOR INVERSE RENDERING 003 004

005 **Anonymous authors**

006 Paper under double-blind review

## 007 008 ABSTRACT 009

010  
011 Inverse rendering with Gaussian Splatting has advanced rapidly, but accurately  
012 disentangling material properties from complex global illumination effects, par-  
013 ticularly indirect illumination, remains a major challenge. Existing methods of-  
014 ten query indirect radiance from Gaussian primitives pre-trained for novel-view  
015 synthesis. However, these pre-trained Gaussian primitives are supervised only  
016 towards limited training viewpoints, thus lack supervision for modeling indirect  
017 radiances from unobserved views. To address this issue, we introduce radiometric  
018 consistency loss, a **novel physically-based constraint** that provides supervision to-  
019 wards unobserved views by minimizing the residual between each Gaussian prim-  
020 itive's learned radiance and its physically-based rendered counterpart. Minimiz-  
021 ing the residual for unobserved views establishes a self-correcting feedback loop  
022 that provides supervision from both physically-based rendering and novel-view  
023 synthesis, enabling accurate modeling of inter-reflection. We then propose Radio-  
024 metrically Consistent Gaussian Surfels (RadioGS), an inverse rendering frame-  
025 work built upon our principle by efficiently integrating radiometric consistency  
026 by utilizing Gaussian surfels and 2D Gaussian ray tracing. We further propose a  
027 finetuning-based relighting strategy that adapts Gaussian surfel radiances to new  
028 illuminations within minutes, achieving low rendering cost (<10ms). Extensive  
029 experiments on existing inverse rendering benchmarks show that RadioGS outper-  
030 forms existing Gaussian-based methods in inverse rendering, while retaining the  
031 computational efficiency.

## 032 1 INTRODUCTION

033  
034 Inverse rendering, a long-standing task in computer vision and graphics, seeks to recover scene  
035 properties such as geometry, material, and illumination from one or more input images. Despite  
036 its significance, this problem remains non-trivial due to the complex interactions between light and  
037 materials, as well as the uncertainty of lighting conditions. Inspired by the remarkable success  
038 of neural radiance fields in novel view synthesis (NVS) (Mildenhall et al., 2021), recent inverse  
039 rendering techniques have adopted these implicit neural representations (Zhang et al., 2021b; Boss  
040 et al., 2021; Zhang et al., 2021a; Liang et al., 2023). More recently, Gaussian Splatting (Kerbl et al.,  
041 2023) has emerged as a powerful alternative to overcome the computational demands of implicit  
042 neural representations.

043 While Gaussian Splatting allows faster optimization and real-time rendering, modeling complex  
044 global illumination effects, particularly indirect illumination and inter-reflections between surfaces,  
045 remains a significant challenge. Existing Gaussian-based inverse rendering approaches often ad-  
046 dress indirect illumination as learnable residual light (Gao et al., 2024; Liu et al., 2024; Bi et al.,  
047 2024), or obtain incident radiances from NVS-trained Gaussian primitives (Liang et al., 2024; Shi  
048 et al., 2023; Sun et al., 2025; Gu et al., 2024). These approaches, however, lack supervision for  
049 indirect radiances from unobserved directions due to the limited viewpoints available during NVS  
050 training. Inaccurate indirect radiances from unobserved directions may lead to incorrect surface and  
051 illumination decomposition, such as baking indirect lighting effects into the surface.

052 In this work, we introduce a **novel physically-based constraint** for Gaussian-based inverse rendering,  
053 termed **radiometric consistency loss**, inspired by principles from self-training neural radiance caches  
for global illumination (Hadadan et al., 2021; Müller et al., 2021). Radiometric consistency aims to

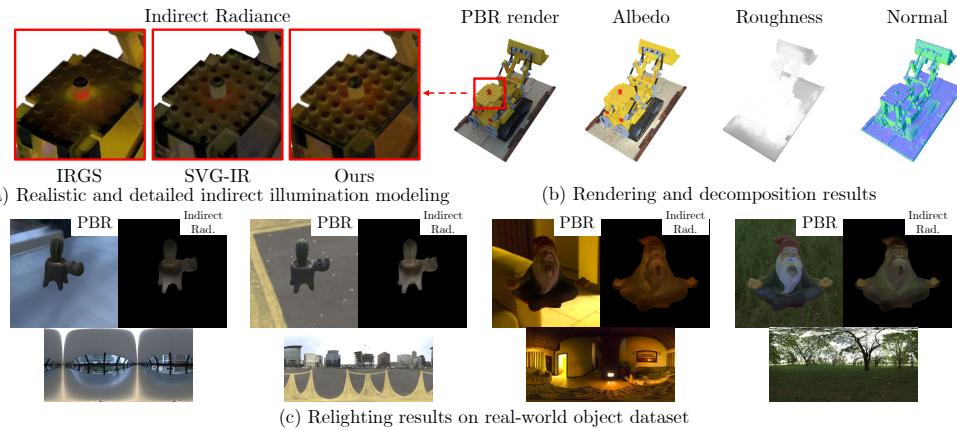


Figure 1: We introduce RadioGS, a novel inverse rendering framework that models accurate indirect illumination by providing a **novel physically-based** supervision on unobserved directions. (a) Compared to existing Gaussian-based methods (Gu et al., 2024; Sun et al., 2025), our method provides realistic inter-reflection between the red bulb and the blobs on the yellow lego surface, (b) leading to robust decomposition of scene properties. (c) Our method can also generate realistic indirect illumination on new lighting conditions for real objects from Stanford-ORB dataset (Kuang et al., 2023).

reduce the residual between the learned radiances of Gaussian primitives and the physically-based rendered (PBR) radiances for unobserved directions, generating a self-correcting guidance between view-constrained Gaussian radiances and physical principle induced from the PBR radiances. **Our physically-based constraint** allows Gaussian primitives to self-correct their radiances to match the consistency and provide accurate indirect illumination for unobserved viewpoints.

We further propose RadioGS, an inverse rendering framework with efficiently integrated radiometric consistency by employing Gaussian surfels and differentiable Gaussian ray tracing. Furthermore, we introduce an efficient relighting strategy that leverages radiometric consistency to rapidly adapt Gaussian surfel radiances under novel lighting conditions, enabling per-frame rendering time below 10ms by directly utilizing adapted surfel radiances. Thorough experiments on multiple inverse rendering benchmarks demonstrate that RadioGS shows enhanced disentanglement of inter-reflections from material and geometry reconstruction, leading to superior relighting performance compared to existing Gaussian-based inverse rendering methods both quantitatively and qualitatively. In summary, our main contributions are as follows:

- Radiometric consistency, a **novel physically-based constraint** that guides Gaussian surfels to self-correct their radiance by enforcing consistency between surfel radiance and physically rendered radiance for unobserved viewpoints.
- RadioGS, a novel inverse rendering framework that efficiently integrates radiometric consistency based on 2D Gaussian ray tracing to accurately model indirect illumination for enhanced inverse rendering performance.
- An efficient relighting method that adapts Gaussian surfel radiances under new lighting conditions within a few minutes, achieving notable reduction in rendering time (<10ms).

## 2 RELATED WORKS

**Inverse Rendering with Neural Radiance Fields.** Inverse rendering aims to recover and decompose scene properties such as geometry, material, and lighting conditions from images (Marschner, 1998). Inspired by the success of neural radiance fields (NeRFs) (Mildenhall et al., 2021) for novel view synthesis (NVS), several works leverage NeRF-like neural representations to optimize scene properties (Zhang et al., 2021a;b; Srinivasan et al., 2021) guided by the rendering equation (Kajiya, 1986). Another line of research focuses on modeling indirect illumination to achieve an

improved disentanglement of lighting conditions using NeRFs (Zhang et al., 2022; Yao et al., 2022; Zhang et al., 2023; Li et al., 2024) or directly queries indirect radiance from NVS pre-trained radiance fields (Jin et al., 2023). Subsequent works deploy path tracing (Wu et al., 2023; Dai et al., 2024) or devise enhanced sampling strategies (Attal et al., 2024) to query incident radiances for physically-based rendering with NeRFs. However, modeling indirect illumination with NeRFs is computationally intensive due to its reliance on volumetric ray marching and numerous neural network queries per ray. Our method provides an efficient representation for modeling indirect illumination with **physically-based** constraints for inverse rendering.

**Inverse Rendering with Gaussian Splatting.** Recent advances leverage Gaussian primitives (Kerbl et al., 2023) to encode geometry and material information, enabling fast optimization for inverse rendering (Liu et al., 2024; Bi et al., 2024). However, modeling indirect illumination with Gaussian primitives remains a key challenge. Existing methods model indirect radiances as per-Gaussian learnable parameters (Gao et al., 2024; Bi et al., 2024; Ye et al., 2025), but the unconstrained optimization may lead to ambiguous decomposition of illumination and material information. Another line of work queries indirect radiances from NVS-pretrained Gaussian primitives by baking irradiance volumes (Liang et al., 2024) or using point-based ray tracing (Sun et al., 2025). Yet, NVS-pretrained Gaussian primitives are supervised only towards observed directions, lacking supervision along arbitrary directions for indirect radiances. Recent work (Gu et al., 2024) leverages differentiable Gaussian ray tracing to optimize indirect radiances, but the training signal is still derived from synthesizing images of observed views. In contrast to these approaches, our work introduces **physically-based** supervision for unobserved viewpoints by enforcing all Gaussian radiances to satisfy the principle of the rendering equation.

**Self-training Radiance Caches for Global Illumination.** Efficiently evaluating the rendering equation (Kajiya, 1986) is central to both rendering and inverse rendering. Classical radiosity (Goral et al., 1984; Immel et al., 1986) solves a simplified, diffuse form of rendering equation via linear systems, while radiance caching (Krivánek et al., 2005; Krivanek & Gautron, 2022) amortizes the computational cost by storing and interpolating light samples. Recent work shows that neural caches can be self-trained to satisfy the rendering equation by iteratively minimizing the rendering-equation residual (Müller et al., 2021; Hadadan et al., 2021), and that such caches provide effective supervision on global illumination for differentiable rendering (Hadadan et al., 2023). We therefore propose an inverse rendering framework that extends these principles to Gaussian primitives, efficiently guiding Gaussian primitives to represent global illumination.

### 3 PRELIMINARIES

**Gaussian Surfels**, also termed 2D Gaussian Splatting (2DGS) (Huang et al., 2024), represent a scene with disk-like 2D Gaussian primitives, which are derived from 3D Gaussian primitives. A Gaussian surfel is expressed using a transformation matrix  $\mathbf{H} \in \mathbb{R}^{4 \times 4}$  that transforms the surfel's local UV space to world space as below:

$$\mathbf{H} = \begin{bmatrix} s_u \mathbf{t}_u & s_v \mathbf{t}_v & \mathbf{0} & \mathbf{p} \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad (1)$$

where  $\mathbf{t}_u$ ,  $\mathbf{t}_v$ ,  $s = (s_u, s_v)$ , and  $\mathbf{p}$  refer to the two principal tangential vectors, the scaling vector, and the center position, respectively.

Ray-splat intersection is employed to determine the contribution of surfels for final rendering. A Gaussian surfel contains an opacity  $\alpha$  and a view-dependent radiance attribute  $c$  parameterized by learnable spherical harmonics coefficients  $\text{SH}_j$ . Each pixel is rendered by alpha-blending of  $N$  depth-sorted Gaussian surfels:

$$\mathcal{C} = \sum_{j=1}^N T_j \alpha_j c_j, \quad T_j = \prod_{k=1}^{j-1} (1 - \alpha_k), \quad c_j = \text{SH}_j(\omega_o), \quad (2)$$

where  $\mathcal{C}$  is the final pixel color,  $T_j$  is the accumulated transmittance, and  $\text{SH}_j$  is the spherical harmonics coefficients parameterization of  $c_j$ . We utilize the Gaussian surfels as the baseline for our inverse rendering framework for robust geometry recovery, and its integration with Gaussian ray tracing (described in Sec. 4.1.2).

162 **Physically-based Rendering** (PBR) models the interaction between light and surfaces in a scene  
 163 via the rendering equation (Kajiya, 1986). The outgoing radiance  $L(x, \omega_o)$  at a surface point  $x$  in  
 164 direction  $\omega_o$  is defined as follows:

$$165 \quad 166 \quad L(x, \omega_o) = \int_{\Omega} f_r(x, \omega_o, \omega_i) L_i(x, \omega_i) (\omega_i \cdot n_x) d\omega_i \quad (3)$$

167 where  $f_r$  is the bidirectional reflectance distribution function (BRDF),  $n_x$  is the normal at point  $x$ ,  
 168 and  $L_i(x, \omega_i)$  is the incoming radiance at the point  $x$  in direction  $\omega_i$ .  
 169

170 We assume the target materials for inverse rendering are mostly dielectric, where the diffuse and  
 171 specular reflectance,  $f_d$  and  $f_s$ , of a surface point  $x$  are governed by diffuse albedo  $a(x)$  and rough-  
 172 ness  $r(x)$ , respectively. These parameters define the total reflectance  $f_r$  based on a simplified Disney  
 173 BRDF (Burley & Studios, 2012) to model the reflectance as below:

$$174 \quad 175 \quad f_r(x, \omega_o, \omega_i) = f_d(x) + f_s(x, \omega_o, \omega_i) = \frac{a(x)}{\pi} + \frac{DFG}{4(n_x \cdot \omega_i)(n_x \cdot \omega_o)}, \quad (4)$$

176 where  $D$ ,  $F$ , and  $G$  are the normal distribution function, the Fresnel term, and the geometry term,  
 177 respectively, which depend on roughness  $r(x)$ .  
 178

179 Incoming radiance  $L_i$  may result directly from the light source or through indirect bounces of other  
 180 surfaces, depending on the visibility at the surface. Thus, we model the incoming radiance as below:  
 181

$$180 \quad L_i(x, \omega_i) = V(x, \omega_i) \cdot L_{dir}(x, \omega_i) + L_{ind}(x, \omega_i), \quad (5)$$

182 where  $V$  is the visibility at the surface point  $x$  with respect to the direction  $\omega_i$ , and  $L_{dir}$  and  $L_{ind}$   
 183 are the corresponding direct and indirect incident radiance terms. We note that  $L_{dir}$  is independent  
 184 of  $x$  when light sources are distant.  
 185

## 4 METHOD

186 In this section, we first introduce our novel **physically-based** regularization termed radiometric  
 187 consistency. Building on this, we present our inverse rendering framework called Radiometrically  
 188 Consistent Gaussian Surfels (RadioGS), followed by our efficient relighting method based on our radio-  
 189 metric consistency.  
 190

### 4.1 RADIOMETRIC CONSISTENCY FOR GAUSSIAN SURFELS

191 Modeling accurate indirect illumination and inter-reflections between Gaussian surfels is crucial  
 192 for robust decomposition of lighting and material information. Recent GS-based inverse rendering  
 193 methods query indirect radiance directly from Gaussian surfels, but the surfel radiances are super-  
 194 vised only through reconstruction from the training images. As a result, surfel radiances along  
 195 directions that are unseen by camera rays can lead to arbitrary values while still fitting the training  
 196 images, degrading the stability and accuracy of indirect illumination (top-right diagram of Fig. 2).  
 197 To address this issue, we introduce radiometric consistency, a novel **physically-based** constraint that  
 198 guides Gaussian surfel radiances for unobserved directions based on the physically-based rendering  
 199 process (bottom-right diagram of Fig. 2).  
 200

#### 4.1.1 FORMULATION

201 We consider a set of Gaussian surfels  $\mathbf{G} = \{\mathcal{G}_j\}$  pretrained for novel-view synthesis (NVS) with  
 202 each surfel  $\mathcal{G}_i$ . Let us denote the surfel radiance at position  $x$  towards direction  $\omega_o$ , as  $L_{\mathbf{G}}(x, \omega_o)$ ,  
 203 which is queried from spherical harmonics coefficients of the corresponding Gaussian surfel. Each  
 204 surfel has optimizable parameters for albedo and roughness. Direct illumination  $L_{dir}(\omega_i)$  is repre-  
 205 sented by an environment cubemap for inverse rendering.  
 206

207 The core principle of our radiometric consistency is that the learned outgoing radiance of a surfel  
 208 should match its physically-rendered radiance, as dictated by the rendering equation (Eq 3). We  
 209 formulate our principle as a residual minimization problem. Following Eq. 3, the residual  $r_{\mathbf{G}}$  can  
 210 be expressed as the difference between the surfel radiance  $L_{\mathbf{G}}$  and the physically-based rendered  
 211 radiance  $L_{\mathbf{G}}^{\text{PBR}}$  as below:  
 212

$$213 \quad 214 \quad L_{\mathbf{G}}^{\text{PBR}}(x, \omega_o) = \int_{\Omega} f_r(x, \omega_o, \omega_i; \mathbf{G}) (V(x, \omega_i; \mathbf{G}) L_{dir}(\omega_i) + L_{ind}(x, \omega_i; \mathbf{G})) (\omega_i \cdot n_x) d\omega_i, \quad (6)$$

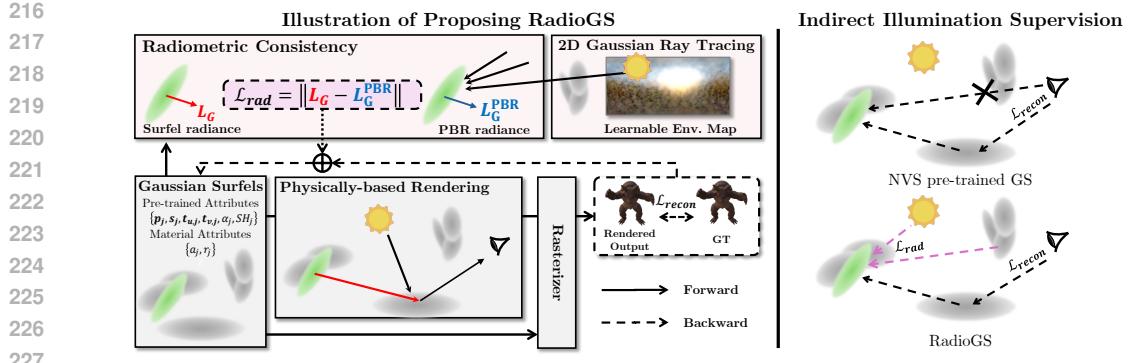


Figure 2: **Overview of our RadioGS.** **Left:** Our radiometric consistency loss  $\mathcal{L}_{rad}$  provides **physically-based** supervision on indirect radiances from views unobserved by image reconstruction loss  $\mathcal{L}_{recon}$ , by enforcing consistency between surfel radiance  $L_G$  and physically-based rendered (PBR) radiance  $L_G^{PBR}$  of Gaussian surfels. Radiometric consistency is seamlessly integrated into the inverse rendering framework, guiding Gaussian surfels to obtain physically-based radiance for delivering realistic indirect radiance to other surfels. 2D Gaussian ray tracing is deployed to jointly optimize ray-traced Gaussian surfels with our radiometric consistency loss. **Right:** Black-dotted arrows show NVS supervision, which leaves the occluded green Gaussian unconstrained, while pink-dotted arrows show our radiometric consistency providing additional supervision on its outgoing radiance along unseen directions (e.g., towards other Gaussian surfels).

$$\mathcal{R}_G(x, \omega_o) = L_G(x, \omega_o) - L_G^{PBR}(x, \omega_o), \quad (7)$$

where  $L_G^{PBR}$  is the radiance calculated by physically-based rendering,  $f_r(\cdot; G)$ ,  $V(\cdot; G)$ , and  $L_{ind}(\cdot; G)$  are the BRDF, visibility, and indirect light induced by Gaussian surfels  $G$  based on Eq. 4 and Eq. 5, respectively.

Our radiometric consistency aims to reduce the  $l_1$ -norm of the residual over all Gaussian surfels and all possible directions  $\omega_o$  denoted as  $\mathcal{L}_{rad}$ :

$$\mathcal{L}_{rad}(G) = \mathbb{E}_{j, \omega_o} [\|\mathcal{R}_G\|_1]. \quad (8)$$

Minimizing the residual norm  $\|\mathcal{R}_G\|_1$  establishes a self-correcting feedback loop based on the rendering equation. On one hand, the physically-rendered radiance  $L_G^{PBR}$  serves as a physically grounded target, guiding the surfel radiance  $L_G$  to represent global illumination for unobserved viewpoints, based on the rendering equation. On the other hand, the well-constrained surfel radiances  $L_G$  towards camera viewpoints provide a strong supervisory signal that is propagated to the surfel radiances contributing to the indirect illumination term  $L_{ind}$  of Eq. 6. This synergistic process allows the Gaussian surfels to obtain physically grounded radiances, thereby providing physically-induced illumination for other surfels.

#### 4.1.2 2D GAUSSIAN RAY TRACING AND MONTE CARLO SAMPLING

Obtaining the visibility  $V(\cdot; G)$  and indirect radiance  $L_{ind}(\cdot; G)$  from Gaussian surfels is critical for creating our self-correcting feedback loop based on the inter-reflection among surfels. Point-based ray tracing has been applied to precompute visibility (Gao et al., 2024; Guo et al., 2024) and to query indirect radiance (Sun et al., 2025) from Gaussian primitives, but lacks the differentiability and speed required for use during optimization. Inspired by recent works leveraging differentiable Gaussian ray tracing (Moenne-Loccoz et al., 2024; Xie et al., 2024), we deploy a 2D Gaussian ray tracer from IRGS (Gu et al., 2024) to leverage optimization through ray-traced surfels for radiometric consistency. 2D Gaussian ray tracing brings seamless integration with our Gaussian surfels by sharing the same ray-splat intersection that defines the contribution of Gaussian surfels.

Given a ray with the origin  $x$  and the direction  $\omega_i$ , our ray tracer  $\text{Trace}(x, \omega_i; G) = (L_{trace}, T_{trace})$  gathers Gaussian surfels intersecting the ray and returns accumulated radiance  $L_{trace}$  and the final transmittance  $T_{trace}$  following the alpha-blending process of Eq. 2. We use ray-traced radiance  $L_{trace}$  directly as indirect radiance  $L_{ind}(x, \omega_i; G)$  and the complement of transmittance  $1 - T_{trace}$

270 as visibility  $V(x, \omega_i; \mathbf{G})$ ), respectively. Using our ray tracer, we acquire the Monte Carlo estimate  
 271 of the integral in Eq. 7 as below:  
 272

$$273 L_{\mathbf{G}}^{\text{PBR}}(x, \omega_o) \approx \frac{2\pi}{N_s} \sum_{i=1}^{N_s} f_r(x, \omega_o, \omega_i; \mathbf{G}) (V(x, \omega_i; \mathbf{G}) L_{\text{dir}}(\omega_i) + L_{\text{ind}}(x, \omega_i; \mathbf{G})) (\omega_i \cdot n_x), \quad (9)$$

276 where we uniformly sample  $N_s$  incident directions  $\omega_i$  over the hemisphere defined by the surfel  
 277 normal  $n_x$ .

278 We also perform Monte Carlo sampling on Gaussian surfels  $\mathbf{G}$  and direction  $\omega_o$  for residual estimation.  
 279 We randomly sample  $N_g$  surfels for each optimization step, and also sample random directions  
 280 on the hemisphere defined by the normal of each sampled surfel to generate guidance towards un-  
 281 observed directions. In addition, we additionally sample the directions towards camera viewpoint  
 282 to propagate well-constraint supervisory signal to ray-traced Gaussian surfels. In conclusion, our  
 283 design for residual estimation allows us to efficiently deploy radiometric consistency, generating  
 284 self-correcting training signals explicitly for surfel radiance  $L_{\mathbf{G}}$  and PBR radiance  $L_{\mathbf{G}}^{\text{PBR}}$  to satisfy  
 285 the physical constraint of the rendering equation.

#### 286 4.2 INVERSE RENDERING WITH RADIOMETRICALLY CONSISTENT GAUSSIAN SURFELS

288 In this section, we introduce our inverse rendering framework RadioGS, optimizing Gaussian surfels  
 289 for inverse rendering under the **physically-based** constraints from our radiometric consistency. Our  
 290 framework operates in two stages to ensure both stable training and accuracy. We then introduce our  
 291 efficient relighting strategy based on radiometric consistency.

292 **Initialization.** Existing works initialize geometry via NVS pre-training prior to tackling inverse  
 293 rendering. To incorporate our **physically-based** constraint during initialization, we additionally intro-  
 294 duce a simplified version of our radiometric consistency loss, using an efficient split-sum approxima-  
 295 tion (Munkberg et al., 2022) instead of the Monte Carlo estimate. Our approximation avoids training  
 296 instability from oscillating geometry during the early optimization stage, resulting in a robust  
 297 geometric foundation that is efficiently regularized based on our **physically-based** constraint (see the  
 298 table of Figure 6 for ablation). Following 2DGS (Huang et al., 2024), we apply image reconstruc-  
 299 tion loss  $\mathcal{L}_{\text{recon}}$  to images rasterized by surfel radiance  $L_{\mathbf{G}}$ , depth distortion loss  $\mathcal{L}_{\text{dist}}$ , normal-depth  
 300 consistency loss  $\mathcal{L}_n$ , normal smoothing loss  $\mathcal{L}_{ns}$ , and mask-entropy loss  $\mathcal{L}_{\text{mask}}$ . We also add image  
 301 reconstruction loss  $\mathcal{L}_{\text{recon}}^{\text{PBR}}$  to images rasterized by physically-rendered radiance  $L_{\mathbf{G}}^{\text{PBR}}$ , which is  
 302 approximated using the split-sum approximation. Thus, the total loss for the initialization stage is a  
 303 weighted sum of the loss components as below:  
 304

$$\mathcal{L}_{\text{init}} = \mathcal{L}_{\text{recon}} + \mathcal{L}_{\text{recon}}^{\text{PBR}} + \lambda_{\text{rad}} \mathcal{L}_{\text{rad}} + \lambda_{\text{dist}} \mathcal{L}_{\text{dist}} + \lambda_n \mathcal{L}_n + \lambda_{ns} \mathcal{L}_{ns} + \lambda_m \mathcal{L}_m. \quad (10)$$

306 **Inverse Rendering.** With our initialized Gaussian surfels, we proceed to the main inverse rendering  
 307 stage by leveraging the full Monte Carlo-estimated radiometric consistency loss  $\mathcal{L}_{\text{rad}}$  to accurately  
 308 model complex inter-reflections. We additionally use smoothing losses for rasterized albedo and  
 309 roughness, denoted as  $\mathcal{L}_{as}$  and  $\mathcal{L}_{rs}$ , to encourage spatial coherence of material features. Finally,  
 310 a light prior loss  $\mathcal{L}_{\text{light}}$  (Liu et al., 2023) is applied to encourage the rendered incident diffuse  
 311 illumination to adopt a natural white appearance. Thus, the total optimization objective for inverse  
 312 rendering is a weighted sum of loss components as below:  
 313

$$\mathcal{L}_{\text{inv}} = \mathcal{L}_{\text{init}} + \lambda_{as} \mathcal{L}_{as} + \lambda_{rs} \mathcal{L}_{rs} + \lambda_{\text{light}} \mathcal{L}_{\text{light}}. \quad (11)$$

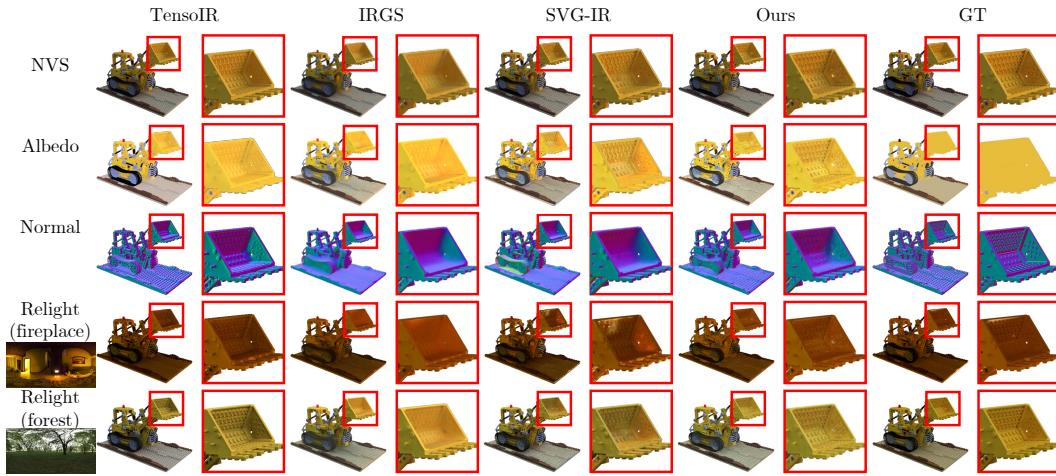
314 Please refer to the supplementary for the details of additional loss functions  $\mathcal{L}_{\text{recon}}$ ,  $\mathcal{L}_{\text{recon}}^{\text{PBR}}$ ,  $\mathcal{L}_{ns}$ ,  
 315  $\mathcal{L}_m$ ,  $\mathcal{L}_{as}$ ,  $\mathcal{L}_{rs}$ , and  $\mathcal{L}_{\text{light}}$ , and the learning rates of Gaussian surfel parameters.  
 316

317 **Relighting.** Once the lighting condition changes, surfel radiances cannot provide indirect illumina-  
 318 tion, since they are specifically optimized for the previous lighting condition. Instead, we query indi-  
 319 rect radiances following IRGS (Gu et al., 2024) by alpha-blending the normal, albedo, and roughness  
 320 towards the incident direction using Gaussian ray tracing, and applying a split-sum approximation to  
 321 efficiently estimate the incident radiance. However, storing numerous incident radiances per surfel  
 322 and re-estimating outgoing radiances based on Eq. 9 consumes additional rendering time.  
 323

To this end, we introduce a finetuning-based relighting approach that leverages radiometric consistency. Radiometric consistency allows surfel radiances to rapidly adapt to new lighting conditions.

Table 1: **Quantitative comparisons on TensoIR dataset** (Jin et al., 2023). The results are colored in rank as 1st, 2nd, and 3rd. Our method surpasses existing Gaussian-based methods and a NeRF-based method in most metrics, while maintaining the computational efficiency with the average training time of 1 hour. We report our relighting metric using Gaussian ray tracing (Ours) and finetuning-based method (Ours\*).

Method	Novel View Synthesis			Normal MAE ↓	Albedo			Relight SSIM ↑	Relight LPIPS ↓	Training hours
	PSNR ↑	SSIM ↑	LPIPS ↓		PSNR ↑	SSIM ↑	LPIPS ↓			
TensoIR	35.09	0.976	0.040	4.100	29.27	0.950	0.085	28.58	0.944	0.081
GS-IR	35.33	0.974	0.039	4.948	29.94	0.921	0.100	24.37	0.885	0.096
GI-GS	36.75	0.972	0.037	5.253	29.90	0.921	0.099	24.70	0.886	0.106
R3DG	33.35	0.964	0.041	5.927	29.27	0.951	0.078	27.37	0.909	0.083
IRGS	35.43	0.964	0.049	4.209	30.62	0.956	0.072	29.91	0.935	0.076
SVG-IR	36.71	0.976	0.033	4.358	30.48	0.950	0.074	31.10	0.946	0.056
Ours	37.86	0.980	0.027	3.689	31.05	0.952	0.072	32.09	0.953	0.048
Ours*								31.41	0.948	0.052



**Figure 3: Qualitative result on the “lego” scene of TensoIR dataset.** Our method provides enhanced decomposition and realistic relighting results compared to Gaussian-based methods. Specifically, our method shows noticeably robust performance on regions with high geometric complexity, such as the highlighted bucket. Best viewed in zoom.

Given a new lighting condition, we perform a few finetuning iterations exclusively on the surfel radiances by minimizing our radiometric consistency loss  $\mathcal{L}_{rad}$ . Once finetuning is complete, the scene can be rendered from any viewpoint using only surfel radiances.

## 5 EXPERIMENTS

## 5.1 EXPERIMENTAL SETUPS

**Dataset and Metric.** We evaluate our method’s novel view synthesis (NVS), inverse rendering, and relighting capabilities using two synthetic datasets: TensoIR (Jin et al., 2023) and Synthetic4Relight (Zhang et al., 2022). These two synthetic datasets provide diverse lighting conditions and ground truth for geometry and material evaluation. We employ PSNR, SSIM, and LPIPS for evaluating NVS, albedo, and relighting. Normal reconstruction is evaluated using Mean Angular Error (MAE), and roughness is evaluated using Mean Square Error (MSE). We also provide qualitative relighting results on a real-world object dataset Stanford-ORB (Kuang et al., 2023) in Figure 1-(c).

**Implementation Details.** For our radiometric consistency loss  $\mathcal{L}_{rad}$ , we set the weight  $\lambda_{rad} = 0.2$ . We sample  $N_g = 4096$  Gaussian surfels and  $N_s = 64$  incident rays per surfel, resulting in  $2^{18}$  rays traced through Gaussian surfels to calculate the radiometric consistency loss at every training iteration. We store the ray-traced results on sampled Gaussians at each step for use in physically-rendered image. For our relighting method, we set the weight  $\lambda_{rad} = 1.0$ , and discard all other

378 losses. Experiments were conducted on an NVIDIA RTX 4090 GPU, with total optimization taking  
 379 approximately 60 minutes (30 for initialization and 30 for inverse rendering), and the finetuning  
 380 process taking approximately 2 minutes. Please refer to the appendix for further details.  
 381

382 **Baselines.** We compare our method against prior Gaussian Splatting (GS)-based methods: GS-  
 383 IR (Liang et al., 2024), GI-GS (Chen et al., 2024), R3DG (Gao et al., 2024), IRGS (Gu et al., 2024),  
 384 and SVG-IR (Sun et al., 2025). We also include TensoIR (Jin et al., 2023), an efficient NeRF-based  
 385 approach. Quantitative and qualitative results are reproduced using the publicly available code.  
 386

## 387 5.2 INVERSE RENDERING PERFORMANCE COMPARISONS

388 **TensoIR.** On the TensoIR dataset (Table 1), our approach  
 389 demonstrates superior performance on various metrics in-  
 390 cluding novel-view synthesis (NVS), normal estimation,  
 391 and relighting compared to existing methods. Notably,  
 392 our method outperforms other ray-tracing based meth-  
 393 ods on Gaussian primitives (Gu et al., 2024; Sun et al.,  
 394 2025), reflecting the necessity of **physically-based** con-  
 395 straints on surfel radiances in inverse rendering. More-  
 396 over, our finetuning-based relighting method outperforms  
 397 existing relighting methods, indicating the effectiveness  
 398 of our self-correcting guidance from radiometric consis-  
 399 tency.

400 Qualitative results on Figure 3 illustrate our method’s per-  
 401 formance on reconstructing finer geometric details for nor-  
 402 mal reconstruction and NVS, which leads to more realistic  
 403 relighting results. Figure 1-(a) showcases the realistic and  
 404 detailed indirect illumination modeled by our method on  
 405 the same scene compared to the other competitors.

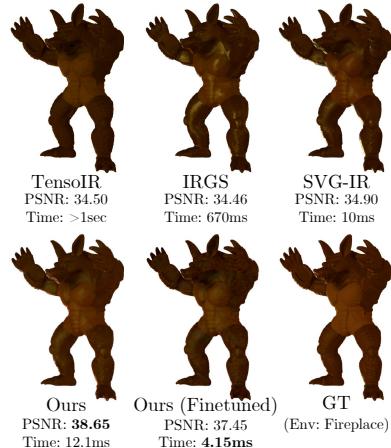
406 Additional comparisons on relighting (Figure 4) show that  
 407 both of our relighting methods achieve realistic relight-  
 408 ing results, showing real-time rendering capabilities. Es-  
 409 pecially, our finetuning-based method shows the fastest  
 410 rendering time, with a minor compromise in quality com-  
 411 pared to ray-tracing based relighting. This is because finetuning process accumulates minor errors  
 412 from estimated geometry and material properties of Gaussian surfels into surfel radiances, leading  
 413 to the trade-off in visual quality.

414 **Synthetic4Relight.** Results on the Synthetic4Relight  
 415 dataset (Table 2) further validate the capabilities of  
 416 our method, outperforming existing methods in NVS,  
 417 albedo reconstruction and relighting, while showing  
 418 comparable performance on roughness estimation. Vi-  
 419 sual comparisons on Figure 5 demonstrate how our re-  
 420 alistic modeling of indirect illumination leads to en-  
 421 hanced albedo reconstruction and NVS. The inter-  
 422 reflecting directions of the highlighted region are over-  
 423 looked during novel-view synthesis training, whereas  
 424 our radiometric consistency provides **physically-based** constraint on surfel radiances towards reflect-  
 425 ing directions, resulting in realistic indirect illumination.

## 426 5.3 ABLATION STUDIES ON RADIOMETRIC CONSISTENCY

427 We report ablation studies on components of our radiometric consistency on the TensoIR dataset.  
 428 The table of Figure 6 shows the PSNR metrics for three categories, NVS, albedo reconstruction, and  
 429 relighting, in our ablation studies.

430 **Absence of Radiometric Consistency.** We perform ablation studies on the radiometric consistency  
 431 by removing the radiometric consistency loss during the inverse rendering stage (see the left sub-  
 432 figure of Figure 6 and “ $\lambda_{rad} = 0$ ” on the right table of Figure 6). The absence of radiometric



433 **Figure 4: Relighting results on**  
 434 **the “armadillo” scene of TensoIR**  
 435 **dataset.**

436 **Table 2: Quantitative comparisons on**  
 437 **Synthetic4Relight dataset.**

Method	NVS PSNR $\uparrow$	Roughness MSE $\downarrow$	Albedo PSNR $\uparrow$	Relight PSNR $\uparrow$
R3DG	34.10	0.010	28.65	33.12
IRGS	34.44	<b>0.008</b>	30.50	34.35
SVG-IR	34.14	0.009	29.06	32.59
Ours	<b>34.98</b>	0.011	<b>30.69</b>	<b>34.87</b>

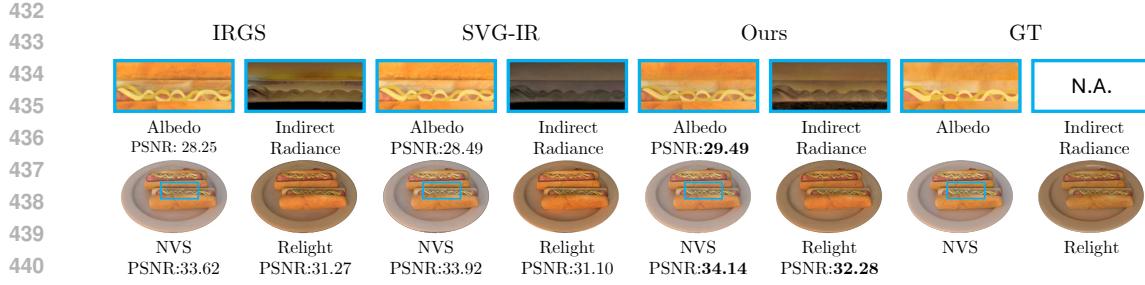


Figure 5: Qualitative results on the “hotdog” scene of Synthetic4Relight (Zhang et al., 2022) dataset. Our method models natural inter-reflection between the sausages and the buns, showing superior reconstruction performance on highlighted regions. IRGS shows relatively bright and fluctuating indirect illumination, which led to darker albedo reconstruction. SVG-IR models relatively darker indirect illumination, returning brighter albedo reconstruction. Best viewed in zoom.

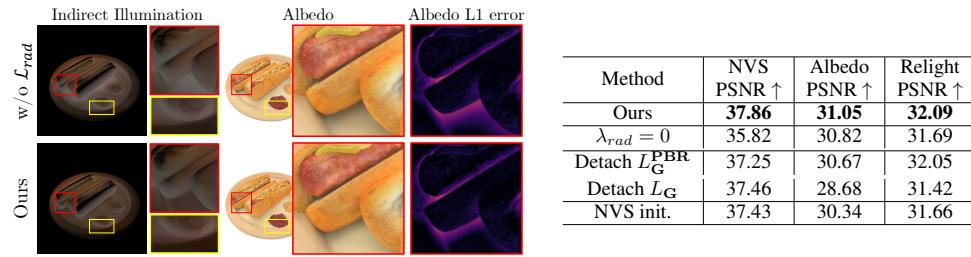


Figure 6: Ablation studies on our radiometric consistency. The left sub-figure demonstrates how our radiometric consistency loss  $\mathcal{L}_{rad}$  provides guidance on radiances towards unobserved views such as the interstices, leading to enhanced albedo reconstruction (red box). Also, our method guides the generation of inter-reflections between the ketchup and the plate (yellow box). The right table contains PSNR metrics for the ablation studies.

consistency provides incorrect indirect radiances on unobserved directions, degrading the albedo reconstruction on the corresponding regions and leading to significant performance degradation in all three categories.

**Detaching Gradient Flows from  $\mathcal{L}_{rad}$ .** We ablate on the self-correcting gradient flow by detaching the gradients towards Gaussian surfels during the calculation of the surfel radiance  $L_G$  and physically-based rendered radiance  $L_G^{PBR}$  on Eq. 7. Detaching either gradient leads to an overall performance drop. Detaching gradient from  $L_G$  cause noticeable degradation on albedo reconstruction, while detaching gradient from  $L_G^{PBR}$  degrades NVS. Such degradation reflects the importance of the view-constrained supervision signal from  $L_G$ , and the **physically-based** constraint delivered by  $L_G^{PBR}$ .

**Initialization.** We found that removing radiometric consistency during initialization degrades overall performance, highlighting the contribution of our radiometric consistency as beneficial **physically-based** guidance for initialization.

## 6 CONCLUSION AND FUTURE WORKS

We introduced a novel **physically-based** supervision called radiometric consistency, which addresses the key challenge of modeling indirect illumination in Gaussian-based representations by guiding Gaussian surfels to learn accurate indirect illumination towards unobserved directions. We then introduced Radiometrically Consistent Gaussian Surfels (RadioGS), a novel inverse rendering framework that efficiently leverages radiometric consistency by utilizing 2D Gaussian ray tracing. We also presented a new relighting method that leverages our constraint to quickly adapt surfel radiances to

486 new lighting environments, achieving a rendering time below 10ms per frame. Experiments demon-  
 487 strated that RadioGS outperforms existing Gaussian-based methods on two synthetic benchmarks,  
 488 based on accurate and realistic indirect illumination. Since the current method only supports dielec-  
 489 tric materials, extending radiometric consistency to more complex materials, such as anisotropic or  
 490 highly-reflective surfaces, would be an interesting future direction.  
 491

492 **REFERENCES**  
 493

494 Benjamin Attal, Dor Verbin, Ben Mildenhall, Peter Hedman, Jonathan T Barron, Matthew O’Toole,  
 495 and Pratul P Srinivasan. Flash cache: Reducing bias in radiance cache based inverse rendering.  
 496 In *European Conference on Computer Vision*, pp. 20–36. Springer, 2024.  
 497

498 Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf  
 499 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF conference*  
 500 *on computer vision and pattern recognition*, pp. 5470–5479, 2022.

501 Zoubin Bi, Yixin Zeng, Chong Zeng, Fan Pei, Xiang Feng, Kun Zhou, and Hongzhi Wu. Gs3:  
 502 Efficient relighting with triple gaussian splatting. In *SIGGRAPH Asia 2024 Conference Papers*,  
 503 pp. 1–12, 2024.  
 504

505 Mark Boss, Raphael Braun, Varun Jampani, Jonathan T Barron, Ce Liu, and Hendrik Lensch. Nerd:  
 506 Neural reflectance decomposition from image collections. In *Proceedings of the IEEE/CVF In-  
 507 ternational Conference on Computer Vision*, pp. 12684–12694, 2021.

508 Brent Burley and Walt Disney Animation Studios. Physically-based shading at disney. In *Acm  
 509 Siggraph*, volume 2012, pp. 1–7. vol. 2012, 2012.  
 510

511 Hongze Chen, Zehong Lin, and Jun Zhang. Gi-gs: Global illumination decomposition on gaussian  
 512 splatting for inverse rendering. *arXiv preprint arXiv:2410.02619*, 2024.

513 Yuxin Dai, Qi Wang, Jingsen Zhu, Dianbing Xi, Yuchi Huo, Chen Qian, and Ying He. Mirres: Multi-  
 514 bounce inverse rendering using reservoir sampling. *arXiv preprint arXiv:2406.16360*, 2024.  
 515

516 Jian Gao, Chun Gu, Youtian Lin, Zhihao Li, Hao Zhu, Xun Cao, Li Zhang, and Yao Yao. Re-  
 517 lightable 3d gaussians: Realistic point cloud relighting with brdf decomposition and ray tracing.  
 518 In *European Conference on Computer Vision*, pp. 73–89. Springer, 2024.

519 Cindy M Goral, Kenneth E Torrance, Donald P Greenberg, and Bennett Battaile. Modeling the  
 520 interaction of light between diffuse surfaces. *ACM SIGGRAPH computer graphics*, 18(3):213–  
 521 222, 1984.  
 522

523 Chun Gu, Xiaofei Wei, Zixuan Zeng, Yuxuan Yao, and Li Zhang. Irgs: Inter-reflective gaussian  
 524 splatting with 2d gaussian ray tracing. *arXiv preprint arXiv:2412.15867*, 2024.  
 525

526 Yijia Guo, Yuanxi Bai, Liwen Hu, Ziyi Guo, Mianzhi Liu, Yu Cai, Tiejun Huang, and Lei Ma.  
 527 Prtgs: Precomputed radiance transfer of gaussian splats for real-time high-quality relighting. In  
 528 *Proceedings of the 32nd ACM International Conference on Multimedia*, pp. 5112–5120, 2024.

529 Saeed Hadadan, Shuhong Chen, and Matthias Zwicker. Neural radiosity. *ACM Transactions on  
 530 Graphics (TOG)*, 40(6):1–11, 2021.  
 531

532 Saeed Hadadan, Geng Lin, Jan Novák, Fabrice Rousselle, and Matthias Zwicker. Inverse global  
 533 illumination using a neural radiometric prior. In *ACM SIGGRAPH 2023 Conference Proceedings*,  
 534 pp. 1–11, 2023.

535 Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2d gaussian splatting  
 536 for geometrically accurate radiance fields. In *ACM SIGGRAPH 2024 conference papers*, pp.  
 537 1–11, 2024.

538 David S Immel, Michael F Cohen, and Donald P Greenberg. A radiosity method for non-diffuse  
 539 environments. *Acm Siggraph Computer Graphics*, 20(4):133–142, 1986.

540 Haian Jin, Isabella Liu, Peijia Xu, Xiaoshuai Zhang, Songfang Han, Sai Bi, Xiaowei Zhou, Zex-  
 541iang Xu, and Hao Su. Tensoir: Tensorial inverse rendering. In *Proceedings of the IEEE/CVF*  
 542 *Conference on Computer Vision and Pattern Recognition*, pp. 165–174, 2023.

543

544 James T Kajiya. The rendering equation. In *Proceedings of the 13th annual conference on Computer*  
 545 *graphics and interactive techniques*, pp. 143–150, 1986.

546 Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splat-  
 547 ting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1, 2023.

548 Jaroslav Krivanek and Pascal Gauthron. *Practical global illumination with irradiance caching*.  
 549 Springer Nature, 2022.

550

551 Jaroslav Krivánek, Pascal Gauthron, Sumanta Pattanaik, and Kadi Bouatouch. Radiance caching  
 552 for efficient global illumination computation. *IEEE Transactions on Visualization and Computer*  
 553 *Graphics*, 11(5):550–561, 2005.

554 Zhengfei Kuang, Yunzhi Zhang, Hong-Xing Yu, Samir Agarwala, Elliott Wu, Jiajun Wu, et al.  
 555 Stanford-orb: a real-world 3d object inverse rendering benchmark. *Advances in Neural Informa-*  
 556 *tion Processing Systems*, 36:46938–46957, 2023.

557

558 Jia Li, Lu Wang, Lei Zhang, and Beibei Wang. Tensosdf: Roughness-aware tensorial representation  
 559 for robust geometry and material reconstruction. *ACM Transactions on Graphics (TOG)*, 43(4):  
 560 1–13, 2024.

561 Ruofan Liang, Huiting Chen, Chunlin Li, Fan Chen, Selvakumar Panneer, and Nandita Vijaykumar.  
 562 Envidr: Implicit differentiable renderer with neural environment lighting. In *Proceedings of the*  
 563 *IEEE/CVF International Conference on Computer Vision*, pp. 79–89, 2023.

564 Zhihao Liang, Qi Zhang, Ying Feng, Ying Shan, and Kui Jia. Gs-ir: 3d gaussian splatting for  
 565 inverse rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern*  
 566 *Recognition*, pp. 21644–21653, 2024.

567

568 Yuan Liu, Peng Wang, Cheng Lin, Xiaoxiao Long, Jiepeng Wang, Lingjie Liu, Taku Komura, and  
 569 Wenping Wang. Nero: Neural geometry and brdf reconstruction of reflective objects from multi-  
 570 view images. *ACM Transactions on Graphics (ToG)*, 42(4):1–22, 2023.

571 Zhenyuan Liu, Yu Guo, Xinyuan Li, Bernd Bickel, and Ran Zhang. Bigs: Bidirectional gaussian  
 572 primitives for relightable 3d gaussian splatting. *arXiv preprint arXiv:2408.13370*, 2024.

573

574 Stephen Robert Marschner. *Inverse rendering for computer graphics*. Cornell University, 1998.

575 Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and  
 576 Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications*  
 577 *of the ACM*, 65(1):99–106, 2021.

578

579 Nicolas Moenne-Loccoz, Ashkan Mirzaei, Or Perel, Riccardo de Lutio, Janick Martinez Esturo,  
 580 Gavriel State, Sanja Fidler, Nicholas Sharp, and Zan Gojcic. 3d gaussian ray tracing: Fast tracing  
 581 of particle scenes. *ACM Transactions on Graphics (TOG)*, 43(6):1–19, 2024.

582

583 Thomas Müller, Fabrice Rousselle, Jan Novák, and Alexander Keller. Real-time neural radiance  
 584 caching for path tracing. *arXiv preprint arXiv:2106.12372*, 2021.

585

586 Jacob Munkberg, Jon Hasselgren, Tianchang Shen, Jun Gao, Wenzheng Chen, Alex Evans, Thomas  
 587 Müller, and Sanja Fidler. Extracting triangular 3d models, materials, and lighting from images.  
 588 In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp.  
 589 8280–8290, 2022.

590

591 Steven G Parker, James Bigler, Andreas Dietrich, Heiko Friedrich, Jared Hoberock, David Luebke,  
 592 David McAllister, Morgan McGuire, Keith Morley, Austin Robison, et al. Optix: a general  
 593 purpose ray tracing engine. *Acm transactions on graphics (tog)*, 29(4):1–13, 2010.

594

595 Yahao Shi, Yanmin Wu, Chenming Wu, Xing Liu, Chen Zhao, Haocheng Feng, Jian Zhang, Bin  
 596 Zhou, Errui Ding, and Jingdong Wang. Gir: 3d gaussian inverse rendering for relightable scene  
 597 factorization. *arXiv preprint arXiv:2312.05133*, 2023.

594 Pratul P Srinivasan, Boyang Deng, Xiuming Zhang, Matthew Tancik, Ben Mildenhall, and  
 595 Jonathan T Barron. Nerv: Neural reflectance and visibility fields for relighting and view syn-  
 596 thesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*,  
 597 pp. 7495–7504, 2021.

598 Hanxiao Sun, Yupeng Gao, Jin Xie, Jian Yang, and Beibei Wang. Svg-ir: Spatially-varying gaussian  
 599 splatting for inverse rendering. In *Proceedings of the Computer Vision and Pattern Recognition  
 600 Conference*, pp. 16143–16152, 2025.

602 BovikAC WangZhou, Hamid R Sheikh, et al. Image qualityassessment: From errorvisibilitytostruc-  
 603 tural similarity. *IEEE Transon ImageProcessing*, 13(4):600, 2004.

604 Haoqian Wu, Zhipeng Hu, Lincheng Li, Yongqiang Zhang, Changjie Fan, and Xin Yu. Nefii: Inverse  
 605 rendering for reflectance decomposition with near-field indirect illumination. In *Proceedings of  
 606 the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4295–4304, 2023.

608 Tao Xie, Xi Chen, Zhen Xu, Yiman Xie, Yudong Jin, Yujun Shen, Sida Peng, Hujun Bao, and  
 609 Xiaowei Zhou. Envgs: Modeling view-dependent appearance with environment gaussian. *arXiv  
 610 preprint arXiv:2412.15215*, 2024.

611 Yao Yao, Jingyang Zhang, Jingbo Liu, Yihang Qu, Tian Fang, David McKinnon, Yanghai Tsin,  
 612 and Long Quan. Neilf: Neural incident light field for physically-based material estimation. In  
 613 *European conference on computer vision*, pp. 700–716. Springer, 2022.

615 Kai Ye, Chong Gao, Guanbin Li, Wenzheng Chen, and Baoquan Chen. Geosplatting: Towards  
 616 geometry guided gaussian splatting for physically-based inverse rendering. In *Proceedings of the  
 617 IEEE/CVF International Conference on Computer Vision*, pp. 28991–29000, 2025.

618 Jingyang Zhang, Yao Yao, Shiwei Li, Jingbo Liu, Tian Fang, David McKinnon, Yanghai Tsin,  
 619 and Long Quan. Neilf++: Inter-reflectable light fields for geometry and material estimation. In  
 620 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 3601–3610,  
 621 2023.

622 Kai Zhang, Fujun Luan, Qianqian Wang, Kavita Bala, and Noah Snavely. Physg: Inverse rendering  
 623 with spherical gaussians for physics-based material editing and relighting. In *Proceedings of the  
 624 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5453–5462, 2021a.

626 Xiuming Zhang, Pratul P Srinivasan, Boyang Deng, Paul Debevec, William T Freeman, and  
 627 Jonathan T Barron. Nerfactor: Neural factorization of shape and reflectance under an unknown  
 628 illumination. *ACM Transactions on Graphics (ToG)*, 40(6):1–18, 2021b.

629 Yuanqing Zhang, Jiaming Sun, Xingyi He, Huan Fu, Rongfei Jia, and Xiaowei Zhou. Modeling indi-  
 630 rect illumination for inverse rendering. In *Proceedings of the IEEE/CVF Conference on Computer  
 631 Vision and Pattern Recognition*, pp. 18643–18652, 2022.

633  
 634  
 635  
 636  
 637  
 638  
 639  
 640  
 641  
 642  
 643  
 644  
 645  
 646  
 647

648 APPENDIX  
649650 THE USE OF LLMs  
651652 The author(s) used ChatGPT for minor grammatical adjustments, and all resulting edits were care-  
653 fully reviewed and finalized by the author(s).  
654655 A IMPLEMENTATION DETAILS  
656657 In this section, we discuss additional details of the implementation of our work.  
658659 A.1 DEPTH INTERPOLATION  
660661 Depth maps of our Gaussian surfels are rendered by interpolation of the Gaussian primitives:  
662

663 
$$\mathcal{D} = \sum_{i=1}^N \frac{\alpha_i T_i}{\sum_{j=1}^N \alpha_j T_j} d_i, \quad (12)$$
  
664

665 where  $d_i$  is the depth of  $i$ -th primitive. This formulation ensures the depth map reflects the visibility-  
666 weighted contribution of all overlapping primitives.  
667668 A.2 LOSS FUNCTIONS  
669670 A.2.1 RECONSTRUCTION LOSS  
671672 The reconstruction loss is composed of a weighted sum of  $l_1$ -loss and SSIM (WangZhou et al.,  
673 2004). Following our baseline (Huang et al., 2024), we assign a weight of 0.8 to the  $l_1$ -loss and 0.2  
674 to SSIM.  
675676 A.2.2 DISTORTION LOSS  
677678 The depth distortion loss enforces geometric consistency along rays by minimizing the weighted  
679 pairwise depth differences between Gaussian primitives:  
680

681 
$$\mathcal{L}_{dist} = \sum_{i,j} \alpha_i T_i \alpha_j T_j |z_i - z_j|, \quad (13)$$
  
682

683 where  $z_i$  denotes the depth value of the  $i$ -th primitive. The loss drives Gaussian primitives to collapse  
684 into tight clusters aligned with surface geometry and enhances depth coherence.  
685686 A.2.3 NORMAL-DEPTH CONSISTENCY LOSS  
687688 This loss enforces geometric coherence by aligning Gaussian primitive normals with surface geom-  
689 etry derived from depth gradients:  
690

691 
$$\mathcal{L}_n = \sum_i \alpha_i T_i (1 - n_i^T N) \quad (14)$$
  
692

693 where  $n_i$  is the normal vector of  $i$ -th primitive and  $N$  is the surface normal at the median of inter-  
694 section  $p_s$  estimated from gradient of depth map:  
695

696 
$$N = \frac{\nabla_x p_s \times \nabla_y p_s}{\|\nabla_x p_s \times \nabla_y p_s\|}. \quad (15)$$
  
697

698 A.2.4 FIRST-ORDER EDGE AWARE SMOOTHING LOSS  
699700 We use edge-aware smoothing constraints to enhance spatial coherence while preserving structural  
701 edges for surface normal, albedo, and roughness predictions. These losses minimize the gradient of  
702 each feature and relax smoothing constraints at image edges:  
703

704 
$$\mathcal{L}_{\{n,a,r\}s} = \|\nabla \{\mathcal{N}, \mathcal{A}, \mathcal{R}\}\| \exp(-\|\nabla \mathcal{C}_{gt}\|), \quad (16)$$
  
705

706 where  $\mathcal{N}$ ,  $\mathcal{A}$ , and  $\mathcal{R}$  are rendered normal, albedo, and roughness map, respectively and  $\mathcal{C}_{gt}$  is the  
707 ground truth training image.  
708

702 A.2.5 SPARSITY LOSS  
703704 The sparsity loss drives Gaussian’s opacity towards 0 or 1:  
705

706 
$$\mathcal{L}_s = \frac{1}{|\alpha|} \sum_{\alpha_i} [\log(\alpha_i) + \log(1 - \alpha_i)] \quad (17)$$
  
707

708 It collapses the spatial distribution of Gaussian primitives into thin surface-aligned layers and accelerates ray tracing by reducing hits and sorting via early termination.  
709  
710711 A.2.6 LIGHT PRIOR LOSS  
712713 The light prior loss enforces neutral white illumination in diffuse rendering. It constrains the per-  
714 channel average intensities  $\bar{c}_i$  of estimated lighting:  
715

716 
$$\mathcal{L}_{light} = \frac{1}{3} \sum_{i=1} 3 \left| \bar{c}_i - \frac{1}{3} \sum_{j=1} 3 \bar{c}_j \right| \quad (18)$$
  
717  
718

719 A.3 2D GAUSSIAN RAY TRACER  
720721 We implemented 2D Gaussian ray tracer using Pytorch CUDA extensions and OptiX (Parker et al.,  
722 2010) following Moenne-Loccoz et al. (2024) and Gu et al. (2024). We adopt a simpler BVH  
723 construction with two triangles encapsulating the 2D Gaussian primitives from Xie et al. (2024),  
724 reducing the BVH update on each training iteration from 3ms to 2ms. The Gaussian response is  
725 achieved by analytically calculating the intersection point  $p$  between the flat 2D Gaussian primitive  
726 with the center  $\mu$  and normal  $n$  and the ray with origin  $o$  and direction  $d$  as below:  
727

728 
$$p = \left( \frac{n \cdot (\mu - d)}{n \cdot d} \right) d + o. \quad (19)$$
  
729  
730

731 Such formulation is identical to the one that of the 2DGS Huang et al. (2024) rasterizer, ensuring  
732 consistent Gaussian response between the rasterizer and the ray tracer.  
733733 To reduce the computation of depth-sorting ray-traced Gaussians, we utilize any-hit program to  
734 gather  $k$  Gaussians within the buffer. Once the buffer is full, we sort the gathered Gaussians by  
735 depth, and accumulate the radiance and transmittance based on Eq. 2. The process repeats to gather  
736 the next  $k$  Gaussians until all ray-traced Gaussians are accumulated or the transmittance reaches the  
737 threshold. We use buffer of  $K=16$  for sorting Gaussians per ray, and terminate the tracing when with  
738 the transmittance threshold of 0.03. For differentiability, we re-cast the rays to gather the same set  
739 of Gaussians, and analytically calculate the gradients.  
740741 A.4 ADDITIONAL TRAINING DETAILS  
742743 We use learning rates of 0.005, 0.005, 0.01 for albedo, roughness, and cubemap, respectively, and  
744 other hyperparameters following the configuration of 2DGS (Huang et al., 2024). We represent  
745 the optimizable environment map as cubemap with a resolution of 32. The first initialize stage is  
746 trained for 40K iterations, with loss weight hyperparameters  $\lambda_d$ ,  $\lambda_n$ ,  $\lambda_{ns}$ ,  $\lambda_s$  as 1000, 0.05, 0.02,  
747 and 0.05, respectively. The inverse rendering stage is trained for 20K iterations, with loss weight  
748 hyperparameters  $\lambda_{as}$ ,  $\lambda_{rs}$ ,  $\lambda_{light}$  as 0.2, 0.1, and 0.01, respectively. After the initialization stage,  
749 we reinitialize the albedo, roughness, and cubemap. Then, we start the inverse rendering stage with  
750 the same learning rate depicted above. For the finetuning stage, we set the same learning rate only  
751 for the spherical harmonics coefficients.  
752753 A.5 RENDERING AND RELIGHTING WITH SPLIT-SUM APPROXIMATION  
754755 Split-sum approximation is a technique for efficiently computing indirect illumination in physically  
756 based rendering. By decomposing the complex specular BRDF integral into two separable terms on  
757 Eq. 3, it avoids the computational burden of Monte Carlo sampling while preserving visual fidelity.  
758

We divide the light transport into diffuse  $L_d$  and specular  $L_s$  components each and approximate the specular light transport as below:

$$L_s(\omega_o) \approx \int_{\Omega} f_s(\omega_i, \omega_o)(\omega_i \cdot N) d\omega_i \cdot \int_{\Omega} L_i(\omega_i) D(\omega_i, \omega_o)(\omega_i \cdot N) d\omega_i. \quad (20)$$

This precomputation allows the specular contribution to be efficiently estimated at runtime by sampling the pre-filtered environment map (using the reflection vector and roughness) and the BRDF LUT.

Diffuse radiance  $L_d$  is computed more directly as the product of the surface’s diffuse reflectance (albedo) and the total incoming diffuse light. The latter is also precomputed by convolving the environment map with a cosine lobe to create an irradiance map.

We apply the split-sum approximation for the initialization stage to easily approximate the estimate of physically-rendered outgoing radiance  $L_{pbr}$  on each Gaussian primitive. For relighting, we apply the split-sum approximation to calculate the incident indirect illumination  $L_{ind}$  from the traced secondary ray using the ray-traced surface position, normal, albedo, and roughness values. The achieved  $L_{ind}$  is used for relighting integrated with the traced visibility  $V$  and the queried direct light  $L_{dir}$ .

## B COMAPARISON ON RELIGHTING PERFORMANCE AND RENDERING COST

On table 3, we report relighting performance using three configurations: (1) Gaussian ray tracing that estimates indirect radiance using the PBR split-sum approximation (PBR split-sum), (2) Gaussian ray tracing that uses indirect radiance predicted by fine-tuned surfels (PBR fine-tuned), and (3) direct rasterization with fine-tuned surfel radiances (Surfel fine-tuned). While fine-tuning introduces a slight quality drop, it enables the surfel radiances to adapt to new lighting conditions and act as physically consistent indirect illumination sources, all while achieving substantially faster rendering speeds than competing approaches. We also provide qualitative comparisons between the three configurations in Figure 20.

Table 3: Relighting Performance and Rendering cost during relighting on TensoIR dataset.

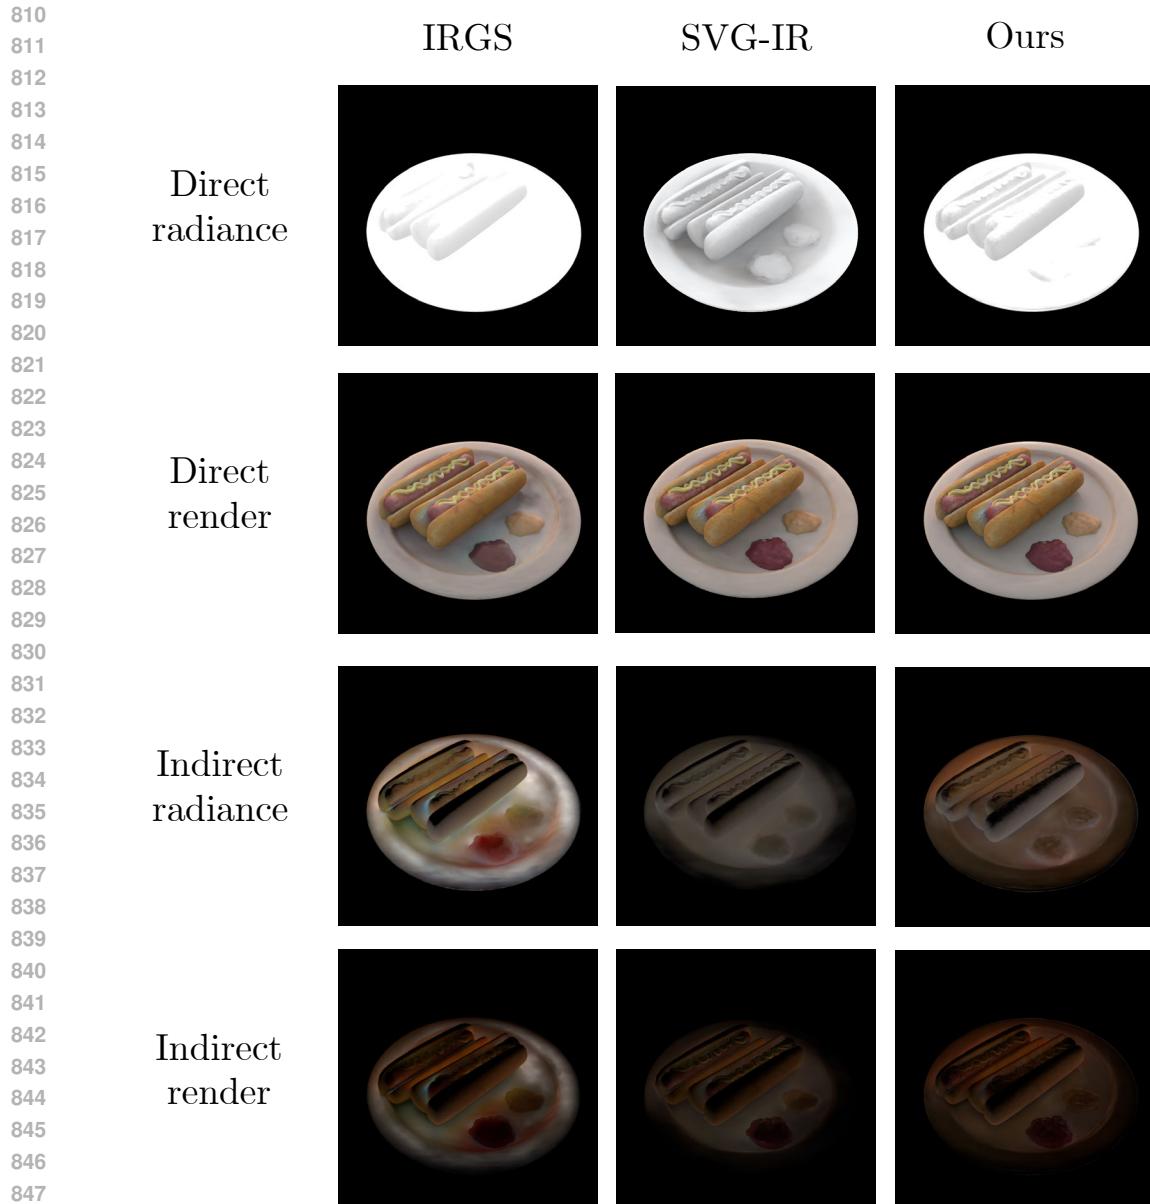
Method	PSNR $\uparrow$	Relight SSIM $\uparrow$	LPIPS $\downarrow$	Rendering ms
IRGS	29.91	0.935	0.076	1090
SVG-IR	31.10	0.946	0.056	82.48
PBR split-sum	32.09	0.953	0.048	38.64
PBR finetuned	<b>31.59</b>	<b>0.952</b>	<b>0.049</b>	<b>38.29</b>
Surfel finetuned	31.41	0.948	0.052	5.902

## C VISUAL COMPARISON ON ILLUMINATION COMPONENTS

We provide additional visual comparisons on “hotdog” and the “lego” scene from the TensoIR dataset (Jin et al., 2023) in Figure 7 and Figure 8. We visualize illumination components including incident direct and indirect radiances, and their rendered results on the datasets along with Gaussian-based methods IRGS (Gu et al., 2024) and SVG-IR (Sun et al., 2025) to compare our realistic indirect illumination. The components are the mean value of the samples during the Monte Carlo rendering. Our method provides realistic indirect illumination that maintains the fine details of inter-reflecting surfaces, while the IRGS (Gu et al., 2024) often overestimates the intensity of the indirect radiance and SVG-IR (Sun et al., 2025) often underestimates the intensity of indirect radiance due to the lack of physical guidance for indirect radiances on unobserved views.

## D ADDITIONAL VISUAL COMPARISON ON BENCHMARK DATASETS

We provide additional visual comparisons on “armadillo” scene from the TensoIR (Jin et al., 2023) dataset, and all the scenes from the Synthetic4Relight dataset from Figure 9 to 12. For the TensoIR



849 Figure 7: Qualitative comparison on illumination components on the “hotdog” scene of TensoIR  
850 dataset. Best viewed in zoom.

851  
852 dataset, we deliver comparison on novel-view synthesis (NVS), normal reconstruction, albedo re-  
853 construction and relighting with Gaussian-based methods IRGS (Gu et al., 2024) and SVG-IR (Sun  
854 et al., 2025), and NeRF-based method TensoIR (Jin et al., 2023). For the Synthetic4Relight dataset,  
855 we deliver comparison on novel-view synthesis (NVS), albedo reconstruction, roughness reconstruc-  
856 tion, and relighting with Gaussian-based methods R3DG (Gao et al., 2024), IRGS (Gu et al., 2024)  
857 and SVG-IR (Sun et al., 2025).

858  
859  
860  
861  
862  
863

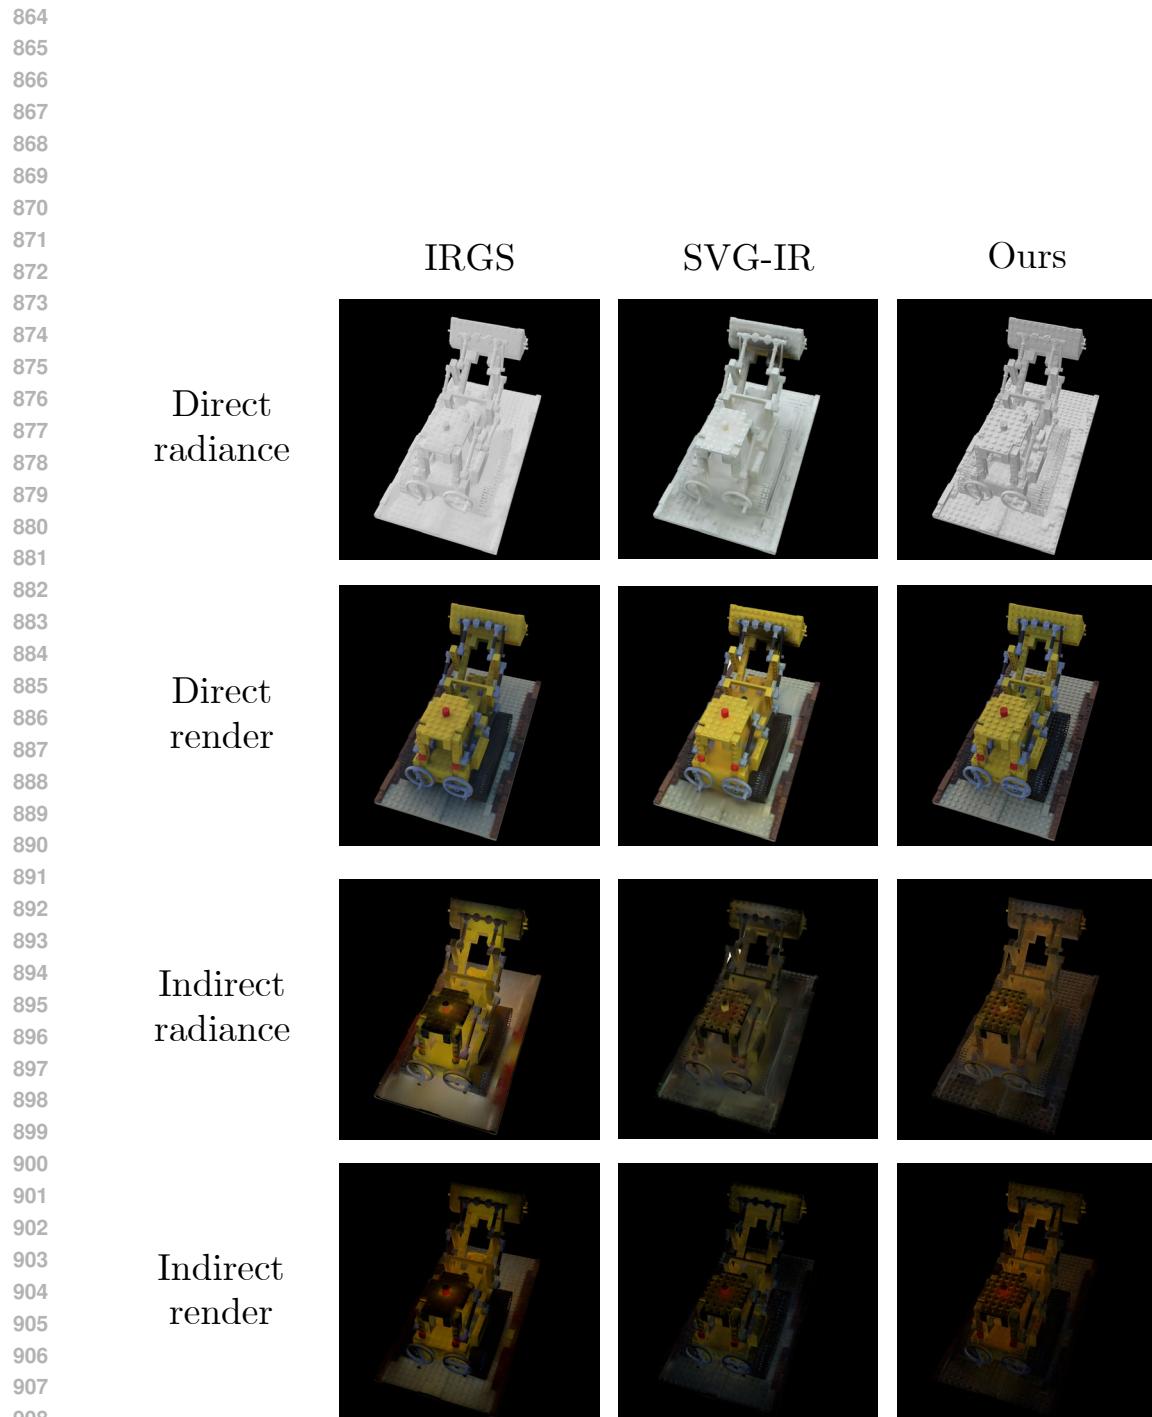
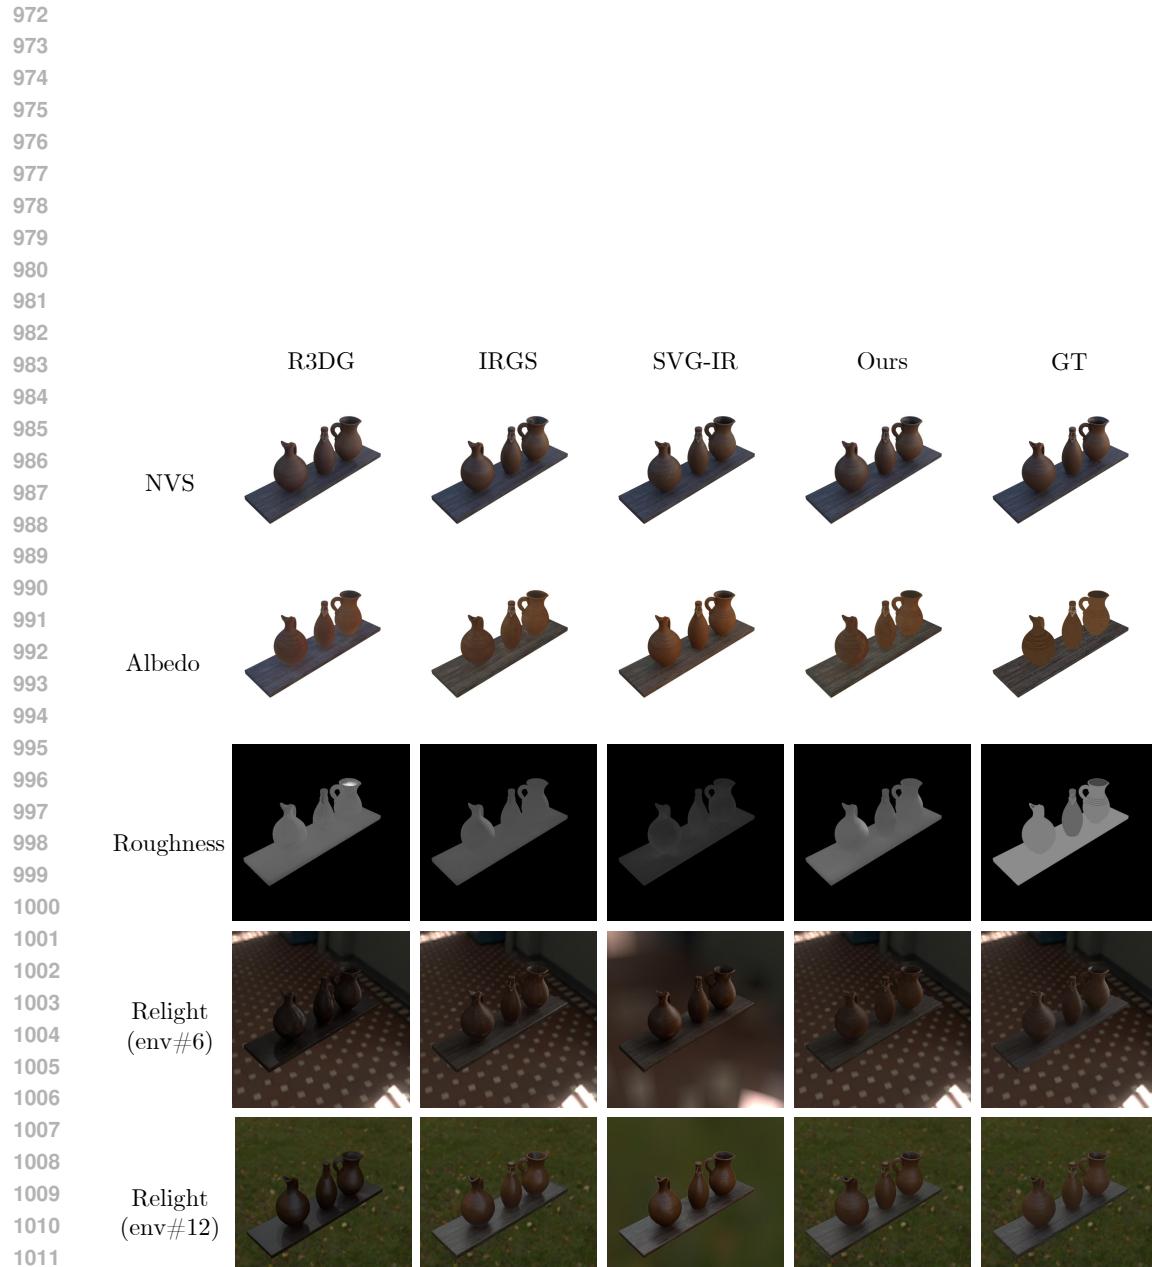


Figure 8: Qualitative comparison on illumination components on the “lego” scene of TensoIR dataset. Best viewed in zoom.

918  
919  
920  
921  
922  
923  
924  
925  
926  
927      TensoIR      IRGS      SVG-IR      Ours      GT  
928  
929  
930      NVS            
931  
932  
933  
934  
935  
936      Albedo            
937  
938  
939  
940  
941  
942  
943      Normal            
944  
945  
946  
947  
948  
949  
950      Relight  
951      (fireplace)            
952  
953  
954  
955  
956  
957      Relight  
958      (forest)            
959  
960  
961  
962  
963  
964  
965  
966  
967  
968  
969  
970  
971

Figure 9: Qualitative comparison on NVS, albedo reconstruction, normal reconstruction, and re-lighting on the “armadillo” scene of the TensoIR dataset. Best viewed in zoom.



1014 Figure 10: Qualitative comparison on NVS, albedo reconstruction, roughness reconstruction, and  
1015 relighting on the “jugs” scene of the Synthetic4Relight dataset. Best viewed in zoom.  
1016  
1017  
1018  
1019  
1020  
1021  
1022  
1023  
1024  
1025



1068  
1069 Figure 11: Qualitative comparison on NVS, albedo reconstruction, roughness reconstruction, and  
1070  
1071  
1072  
1073  
1074  
1075  
1076  
1077  
1078  
1079  
relighting on the “air balloons” scene of the Synthetic4Relight dataset. Best viewed in zoom.

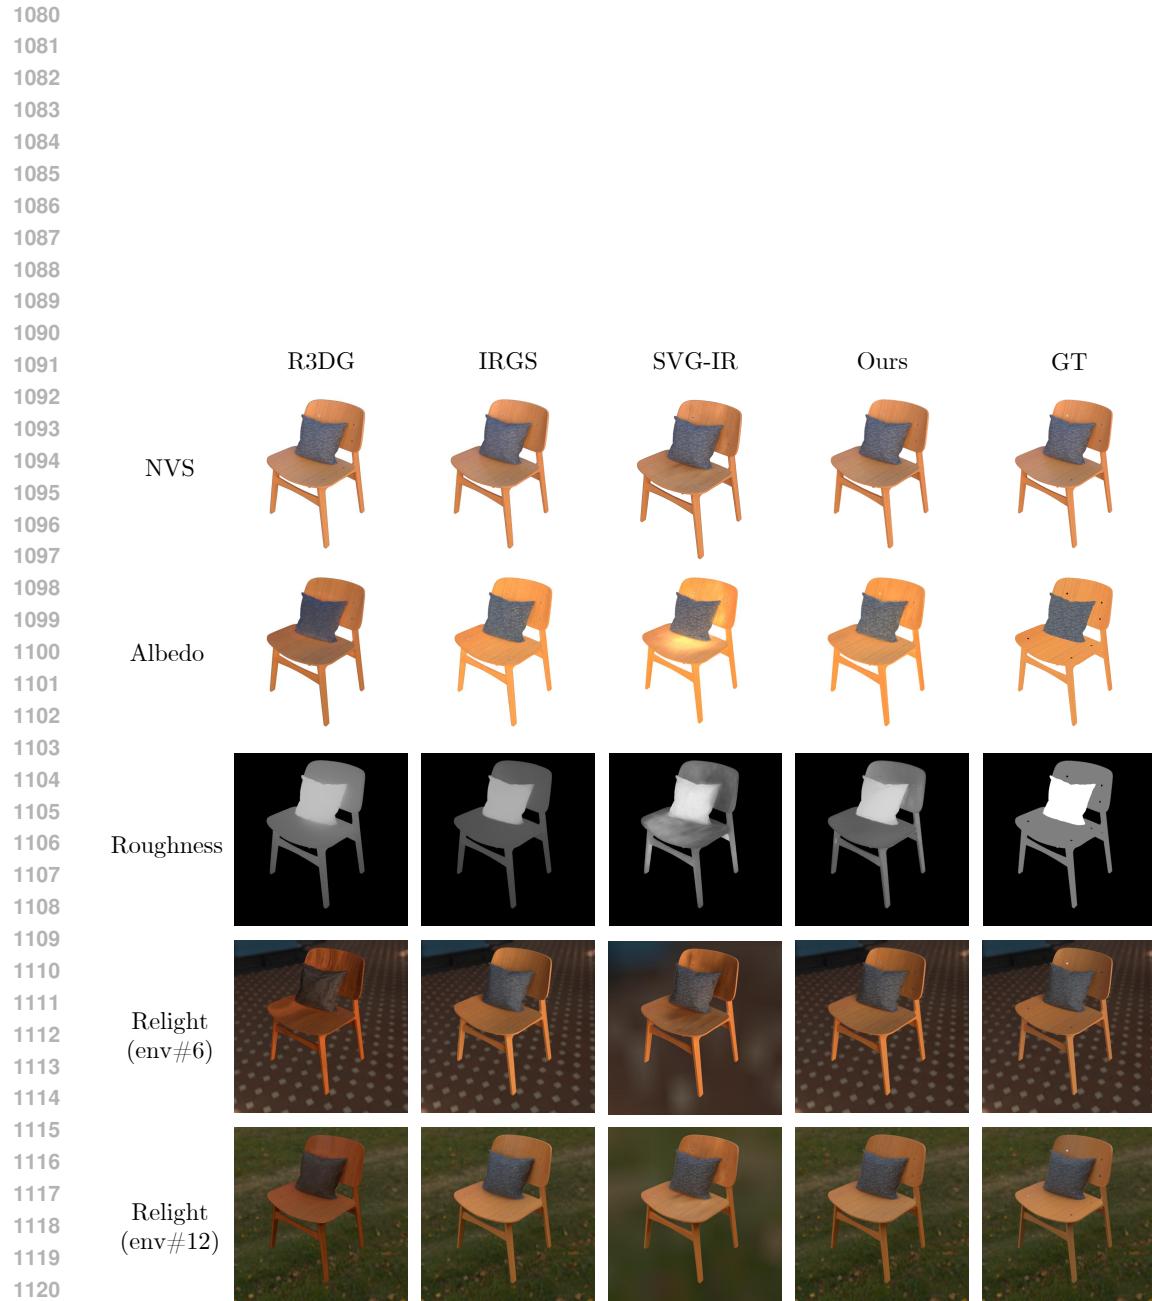


Figure 12: Qualitative comparison on NVS, albedo reconstruction, roughness reconstruction, and relighting on the “chair” scene of the Synthetic4Relight dataset. Best viewed in zoom.

1124  
1125  
1126  
1127  
1128  
1129  
1130  
1131  
1132  
1133

1134  
1135  

## E EVALUATION OF INDIRECT ILLUMINATION

1136 Current inverse rendering benchmarks, including TensoIR Jin et al. (2023), Synthetic4Relight Zhang  
 1137 et al. (2022), and Stanford-ORB Kuang et al. (2023), do not provide ground truth (GT) for indirect  
 1138 illumination for quantitative and qualitative evaluation. To address this limitation, we generated  
 1139 a new evaluation dataset with explicit GT Indirect Illumination. We utilized the original Blender  
 1140 files from the TensoIR dataset to generate high-fidelity ground truth. We used Blender Cycles path  
 1141 tracing engine to render the indirect illumination pass along with the original render pass. To ensure  
 1142 noise-free references, especially for indirect illumination, we set the sampling rate to 256 spp and  
 1143 applied the OIDN denoiser [3]. This allows direct quantitative evaluation of the indirect illumination  
 1144 components. We trained our model and the ablation model discarding the radiometric consistency  
 1145 loss (Ours w/o  $\mathcal{L}_{rad}$ ) using the same hyperparameters described in the paper. We also trained two  
 1146 baselines, IRGS and SVG-IR, for comparison.

1147 Table 4 presents the quantitative comparison against baselines (IRGS, SVG-IR) and our ablation  
 1148 model on our new dataset. Our method significantly outperforms all baselines in indirect illumina-  
 1149 tion reconstruction, confirming that our method accurately models the physical transport of indirect  
 1150 light. Our accurate indirect illumination leads to superior performance in most other metrics. When  
 1151 radiometric consistency is removed (Ours w/o  $\mathcal{L}_{rad}$ ), the indirect PSNR drops significantly, show-  
 1152 ing that the performance gain comes from our proposed framework utilizing radiometric consistency,  
 1153 which effectively supervises indirect radiance from unobserved views.

1154 We also provide qualitative comparisons in Figure 13. As shown, our method faithfully reconstructs  
 1155 indirect illumination compared to baseline models and our ablation model. Overall, our method pro-  
 1156 duces indirect illumination closest to the ground truth, while IRGS produces overestimated intensity,  
 1157 and SVG-IR tends to underestimate the intensity of the indirect radiances. Similar phenomena can  
 1158 also be observed in the qualitative results in Figures 7 and 8 of our paper. Our ablation model (Ours  
 1159 w/o  $\mathcal{L}_{rad}$ ) tends to produce white blurs on inter-reflecting regions compared to our method due to  
 1160 the lack of supervision on unseen views. Finally, we provide additional qualitative comparisons  
 1161 of indirect illumination during relighting in Figure 14, where our method produces more realistic,  
 1162 accurate indirect illumination than the baseline models.

1163 Table 4: Quantitative comparison against baselines (IRGS, SVG-IR) and our ablation model on our  
 1164 new dataset. Our method significantly outperforms all baselines in indirect illumination reconstruc-  
 1165 tion.

Method	NVS			Indirect Illumination			Geometry		Albedo		
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	Normal MAE $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	
IRGS	35.0982	0.9660	0.0436	24.2219	0.8792	0.1092	4.1835	30.0931	0.9521	0.0752	
SVG-IR	36.9634	0.9786	0.0266	30.9747	0.9134	0.0843	4.2624	29.7354	0.9309	0.0822	
Ours w/o $\mathcal{L}_{rad}$	36.3471	0.9764	0.0276	30.0954	0.9161	0.0752	3.8332	30.3425	0.9470	0.0760	
<b>Ours</b>	<b>37.8519</b>	<b>0.9822</b>	<b>0.0212</b>	<b>32.8832</b>	<b>0.9266</b>	<b>0.0726</b>	<b>3.6048</b>	<b>30.6224</b>	<b>0.9502</b>	<b>0.0744</b>	

1172  
 1173  
 1174  
 1175  
 1176  
 1177  
 1178  
 1179  
 1180  
 1181  
 1182  
 1183  
 1184  
 1185  
 1186  
 1187

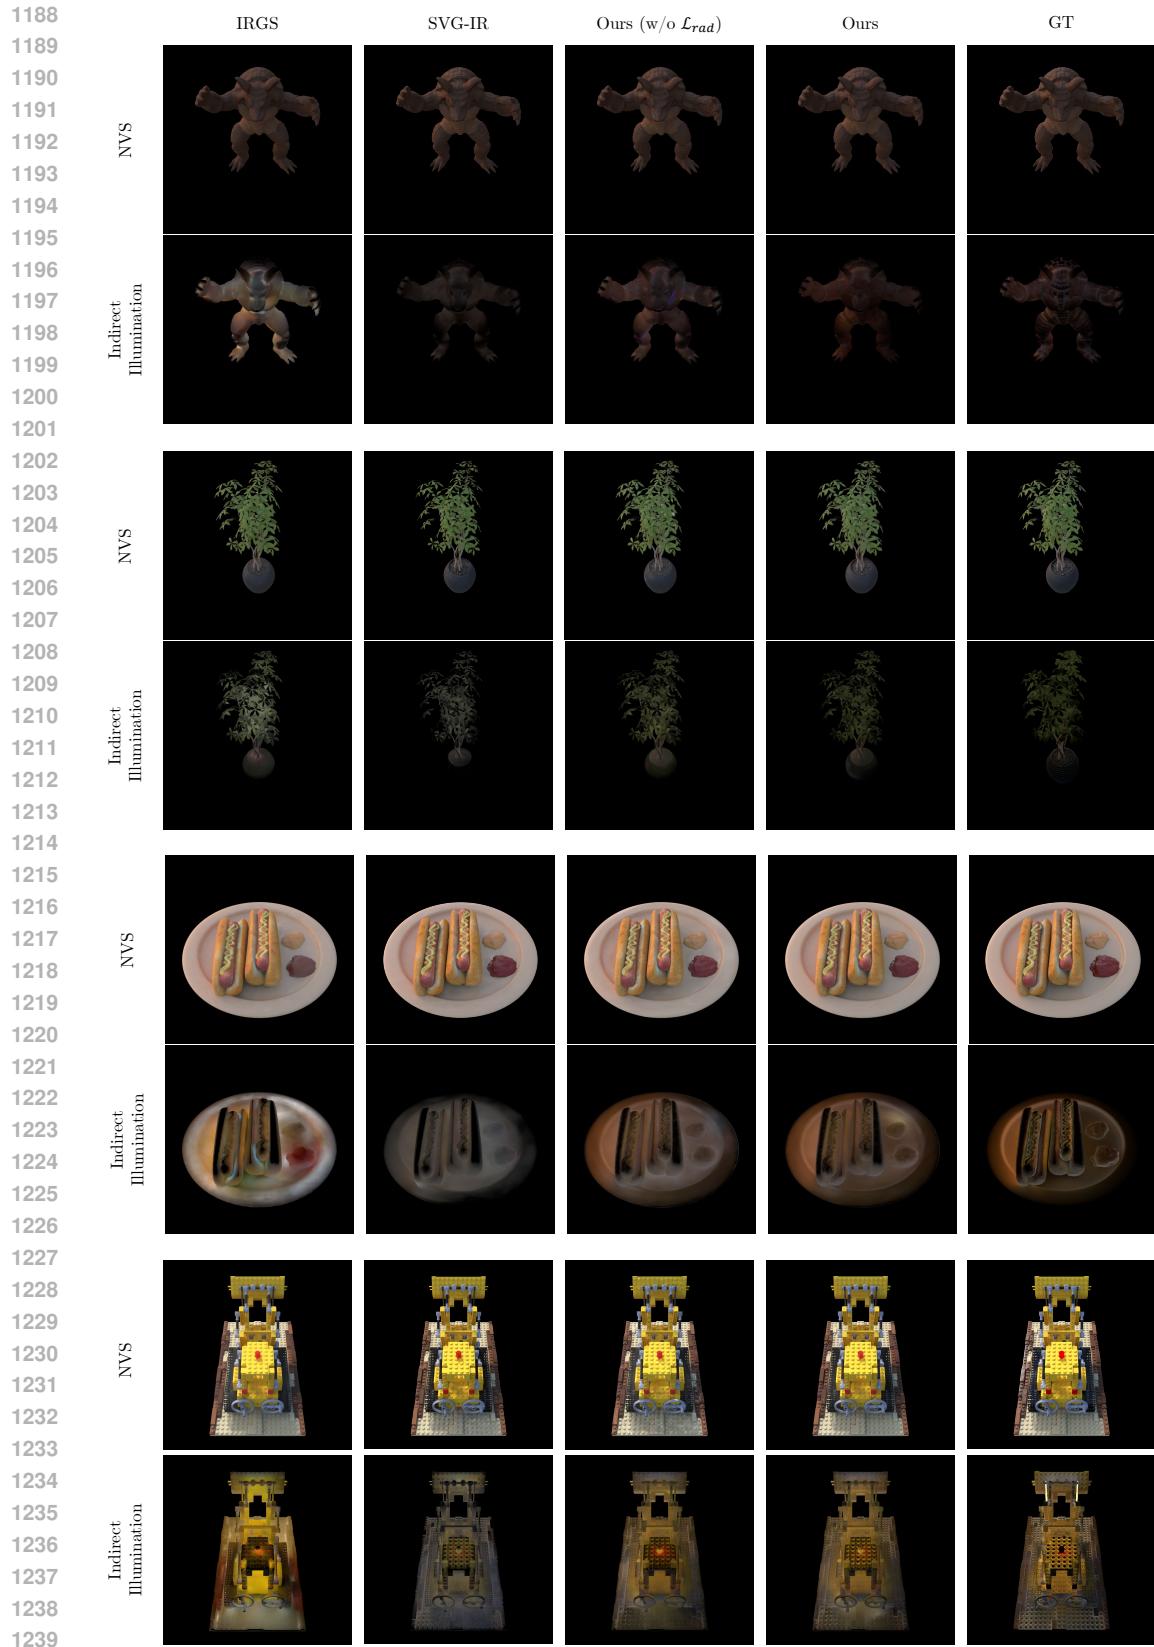
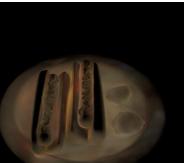
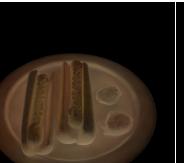
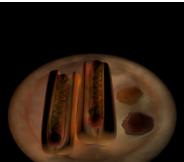
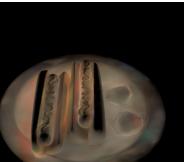
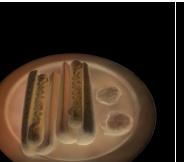
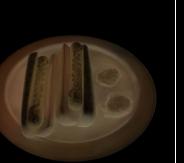


Figure 13: Qualitative comparison on novel-view synthesis and indirect illumination on our dataset. Best viewed in zoom.

1242  
 1243  
 1244  
 1245  
 1246  
 1247  
 1248  
 1249  
 1250  
 1251  
 1252  
 1253  
 1254  
 1255  
 1256  
 1257  
 1258  
 1259      Bridge  
 1260        
 1261        
 1262        
 1263        
 1264        
 1265  
 1266      City  
 1267        
 1268        
 1269        
 1270        
 1271        
 1272  
 1273      Forest  
 1274        
 1275        
 1276        
 1277        
 1278      

1279      Figure 14: Qualitative comparison on indirect illumination for four different lighting conditions on  
 1280      the “hotdog” scene of our dataset. Best viewed in zoom.  
 1281  
 1282  
 1283  
 1284  
 1285  
 1286  
 1287  
 1288  
 1289  
 1290  
 1291  
 1292  
 1293  
 1294  
 1295

## 1296 F ABLATION STUDIES ON RADIOMETRIC CONSISTENCY

1297  
 1298 In this section, we discuss the contributions of the two main components of our framework, differ-  
 1299 entiable Gaussian ray tracing and supervision on unobserved (*i.e.* unseen) direction.  
 1300

### 1301 F.1 RADIOMETRIC CONSISTENCY AND DIFFERENTIABLE GAUSSIAN RAY TRACING

1302  
 1303 The primary challenge in applying radiometric consistency to existing GS-based inverse rendering  
 1304 is that existing pipelines cannot dynamically query surfel radiances as indirect radiances during op-  
 1305 timization. Existing pipelines employ point-based ray tracing or baked volumes, which precompute  
 1306 indirect radiance from NVS-pretrained Gaussian primitives in a non-differentiable manner. These  
 1307 precomputed values remain fixed during the optimization. When these values do not reflect updated  
 1308 Gaussian attributes, the supervision signal from radiometric consistency may become inconsistent.  
 1309 Our framework addresses this issue by employing differentiable Gaussian ray tracing to query indi-  
 1310 rect radiance for sampled surfels at every iteration.

1311 We further conducted ablation studies to show how our framework enhances the contribution of  
 1312 radiometric consistency in Gaussian-based inverse rendering. Starting from the model initialized  
 1313 with our method, we trained three ablation models using different methods for querying indirect  
 1314 radiances and rendering, while utilizing our radiometric consistency loss during inverse rendering  
 1315 optimization.

- 1316 • “Split-sum”: Applies split-sum approximation to calculate physically-based radiance dur-  
 1317 ing the inverse rendering optimization, which does not involve indirect illumination.
- 1318 • “RT precompute”: Applies Monte Carlo estimate (Eq.(8)) to calculate and precomputes  
 1319 indirect radiances via ray tracing, freezing the indirect illumination estimate during the  
 1320 optimization.
- 1321 • “RT w/o diff.”: Dynamically updates indirect radiances via ray tracing during optimization,  
 1322 but without updating through ray-traced surfels by detaching the gradient on the ray-traced  
 1323 results.

1324  
 1325 The quantitative results of the three ablation models and our method are depicted in Table 5. “Split-  
 1326 sum” shows severe degradation in normal, albedo, and relighting accuracy due to the lack of illu-  
 1327 mination effects from surfels. “RT Precompute” shows enhanced normal, albedo, and relighting  
 1328 accuracy by utilizing precomputed indirect illumination, but yields the lowest novel-view syn-  
 1329 thesis performance among all methods. While “RT w/o Diff” improves overall performance through  
 1330 dynamic updates to these values, it still falls short of our method across all metrics. Ours method  
 1331 achieves the best performance across all metrics, demonstrating that the fully differentiable self-  
 1332 correcting feedback loop is essential for robust disentanglement.

1333  
 1334 **Table 5: Ablation study on radiometric consistency strategies.** We compare our method with  
 1335 baselines using different indirect illumination handling. Best results are highlighted in **bold**.

1336 1337 Method	NVS			Geometry Normal MAE ↓	Albedo			Relighting		
	PSNR ↑	SSIM ↑	LPIPS ↓		PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓
1338 Split-sum	36.659	0.9800	0.0254	4.0066	27.471	0.9339	0.0828	28.816	0.9371	0.0575
1339 RT Precompute	35.547	0.9728	0.0332	3.7138	30.523	0.9481	0.0750	31.965	0.9516	0.0496
1340 RT w/o Diff.	37.252	0.9789	0.0275	3.6946	30.673	0.9507	0.0732	32.050	0.9524	0.0484
1341 <b>Ours</b>	<b>37.858</b>	<b>0.9801</b>	<b>0.0266</b>	<b>3.6889</b>	<b>31.048</b>	<b>0.9523</b>	<b>0.0719</b>	<b>32.092</b>	<b>0.9533</b>	<b>0.0478</b>

### 1342 F.2 SUPERVISION ON UNOBSERVED DIRECTION

1343  
 1344 To further support our interpretation that the radiometric consistency supervises unseen directions,  
 1345 we conducted an additional ablation where we train both “Ours” and “Ours w/o ” on only 50% and  
 1346 25% of the randomly subsampled training views of our new dataset on Appendix E. The numeri-  
 1347 cal results on NVS and indirect illumination reconstruction performance are in Table 6. We have  
 1348 denoted the performance drop relative to the full training view (100%) on the right side of the met-  
 1349 rics. NVS metrics degrade for both methods when fewer views are used. However, when using

1350 25% of the training views, indirect illumination reconstruction with our model remains nearly un-  
 1351 changed (-0.17dB), whereas the ablation shows a significant drop in indirect PSNR (-2.21dB). This  
 1352 indicates that when the camera viewpoint is limited, radiometric consistency still provides effective  
 1353 supervision that cannot be provided by NVS alone.

1354

1355 **Table 6: Ablation study on training view scarcity.** We report NVS and Indirect PSNR metrics  
 1356 across different subsets of training views. Values in parentheses denote the performance drop relative  
 1357 to the 100% setting.

1358

1359

1360

1361

1362

1363

1364

1365

1366

1367

1368

1369

1370

1371

1372

1373

1374

1375

1376

1377

1378

1379

1380

1381

1382

1383

1384

1385

1386

1387

1388

1389

1390

1391

1392

1393

1394

1395

1396

1397

1398

1399

1400

1401

1402

1403

Train views	NVS PSNR		Indirect PSNR	
	Ours	Ours w/o	Ours	Ours w/o
100%	37.85	36.35	32.88	30.10
50%	37.54 (-0.31)	35.81 (-0.54)	32.79 (-0.09)	29.15 (-0.95)
25%	36.79 (-1.06)	34.65 (-1.70)	32.71 (-0.17)	27.89 (-2.21)

1404 **G ABLATION STUDY ON HYPERPARAMETERS  $N_g$  AND  $N_s$** 
1405

1406 **G.1 ANALYSIS ON TENSOIR DATASET**
1407

1408 We evaluated the performance impact of two key hyperparameters on the TensoIR dataset: the
1409 number of surfels sampled for radiometric consistency ( $N_g$ ) and the number of incident ray samples
1410 ( $N_s$ ) per surfel.

1411 First, we varied  $N_g$  from 1024 to 8192 while keeping  $N_s$  fixed at 64. As shown in Table 7, increasing
1412  $N_g$  generally improves reconstruction quality across all tasks and metrics. This improvement is
1413 attributed to the radiometric consistency loss supervising a larger number of surfels per iteration.
1414 Notably,  $N_g$  does not impact the rendering cost during inference, as it strictly controls the number
1415 of surfels supervised during the optimization step.

1416
1417 **Table 7: Ablation study on the number of surfels ( $N_g$ ) for radiometric consistency. We vary  $N_g$** 
1418 **while fixing  $N_s = 64$ . Increasing  $N_g$  improves quality without affecting rendering cost.**

1419
1420 

$N_g$	NVS			Geometry	Albedo			Relight			Render (ms)
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	Normal MAE $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	
1024	37.8206	0.9799	0.0268	3.6900	30.9137	0.9518	0.0730	32.0569	0.9529	0.0481	38.5
2048 (Ours)	37.8580	0.9801	0.0266	3.6889	31.0479	0.9521	0.0721	32.0920	0.9533	0.0478	38.6
4096	37.8707	0.9802	0.0264	3.6852	31.0495	0.9522	0.0721	32.1112	0.9532	0.0478	38.7
8192	37.8799	0.9802	0.0263	3.6834	31.0507	0.9523	0.0720	32.1294	0.9533	0.0477	38.5

1425 Next, we analyzed the effect of  $N_s$  ranging from 16 to 128 with  $N_g$  fixed at 2048. Table 8 demonstrates that
1426 reconstruction quality improves as  $N_s$  increases up to 64. However, increasing  $N_s$  further to 128 yields
1427 diminishing returns, and in some tasks (e.g., NVS and Relighting), performance slightly drops. This
1428 suggests that the additional ray samples beyond this point do not significantly resolve the variance in
1429 Monte Carlo integration for the given capacity. regarding efficiency, the rendering cost scales with
1430  $N_s$  due to the additional computation required for the Monte Carlo estimate of  $L_G^{PBR}$ . However, the
1431 rendering cost remains manageable even at  $N_s = 128$ , achieving 58.3 ms per frame ( $\sim 17.2$  fps).

1432
1433 **Table 8: Ablation study on the number of secondary ray samples ( $N_s$ ). We vary  $N_s$  while fixing**
1434  **$N_g = 2048$ . Performance saturates around  $N_s = 64$ .**

1435
1436 

$N_s$	NVS			Geometry	Albedo			Relight			Render (ms)
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	Normal MAE $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	
16	37.4199	0.9792	0.0275	3.7174	30.8656	0.9505	0.0735	32.1684	0.9518	0.0494	16.2
32	37.8125	0.9799	0.0268	3.6979	30.9014	0.9510	0.0730	32.1537	0.9528	0.0483	22.6
64 (Ours)	37.8707	0.9802	0.0264	3.6852	31.0479	0.9521	0.0721	32.1112	0.9532	0.0478	38.6
128	37.6548	0.9798	0.0268	3.6821	30.9563	0.9512	0.0728	31.9768	0.9526	0.0480	58.3

1441
1442 **G.2 RENDERING EFFICIENCY ON LARGE-SCALE SCENES (MIPNeRF360)**
1443

1444 We further provide analysis on MipNeRF360 (Barron et al., 2022) dataset to evaluate the rendering
1445 cost in complex, scene-level environments. Qualitative results corresponding to this analysis are
1446 provided in Figure 15.

1447 We further investigated the computational cost by varying  $N_g$  from 1024 to 16384 while keeping
1448  $N_s$  fixed at 64 (Table 9). Consistent with the analysis on the TensoIR dataset, varying  $N_g$  has a
1449 negligible impact on the rendering cost.

1450 Finally, we analyzed the effect of the number of incident ray samples ( $N_s$ ) on rendering time, varying
1451  $N_s$  from 16 to 64 with  $N_g$  fixed at 2048 (Table 10). As shown, the rendering cost scales with
1452  $N_s$ . This indicates that while larger scenes increase the baseline computational load, the rendering
1453 speed can be effectively controlled by adjusting  $N_s$ . Although reducing  $N_s$  improves speed, it may
1454 trade off rendering quality. However, given that increasing  $N_g$  improves quality without computa-
1455 tional overhead, our method can scale to larger scenes while maintaining efficiency by strategically
1456 balancing  $N_g$  and  $N_s$ .

1458

1459 Table 9: Rendering cost analysis on MipNeRF360 with varying  $N_g$ . We fix  $N_s = 64$ . Consistent  
1460 with TensoIR,  $N_g$  does not significantly affect rendering speed.

1461

1462

$N_g$	Average (ms)	Bonsai (ms)	Counter (ms)	Kitchen (ms)	Room (ms)
1024	81.23	77.37	72.38	86.81	88.43
2048 (Ours)	81.40	77.78	72.95	84.29	90.56
4096	81.69	78.87	72.19	85.28	90.41
8192	81.53	77.75	73.34	85.22	89.82
16384	81.86	77.75	73.24	86.08	90.38

1466

1467

1468

1469 Table 10: Rendering cost analysis on MipNeRF360 with varying  $N_s$ . We fix  $N_g = 2048$ . Reducing  
1470  $N_s$  significantly decreases rendering time, allowing for trade-offs between speed and quality.

1471

1472

1473

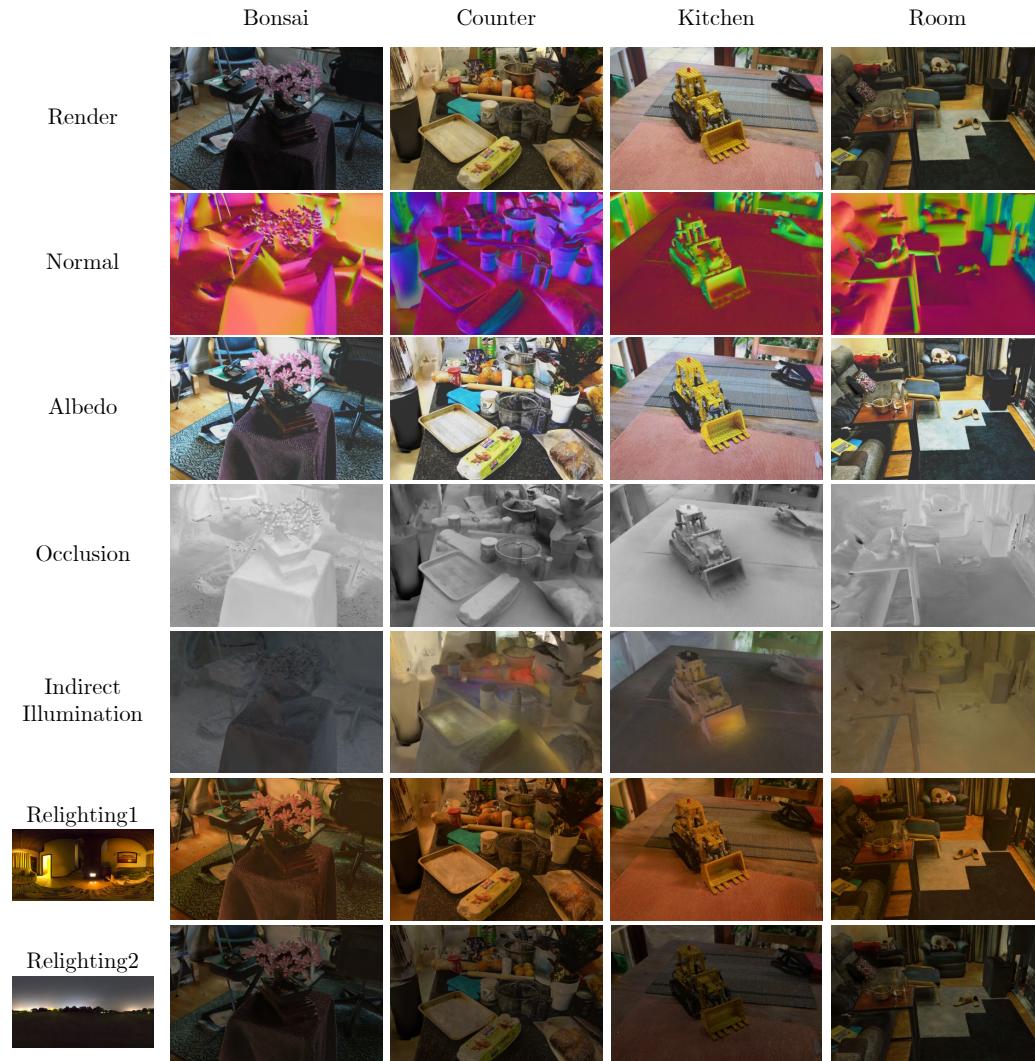
1474

$N_s$	Average (ms)	Bonsai (ms)	Counter (ms)	Kitchen (ms)	Room (ms)
16	24.57	23.49	23.86	26.07	24.85
32	43.43	42.00	39.32	47.31	45.10
64 (Ours)	81.40	77.78	72.95	84.29	90.56

1475

1476

1477



1510

1511

Figure 15: Qualitative results on the MipNeRF360 dataset. Best viewed in zoom.

## 1512 H ADDITIONAL ANALYSIS ON THE INITIALIZATION STRATEGY

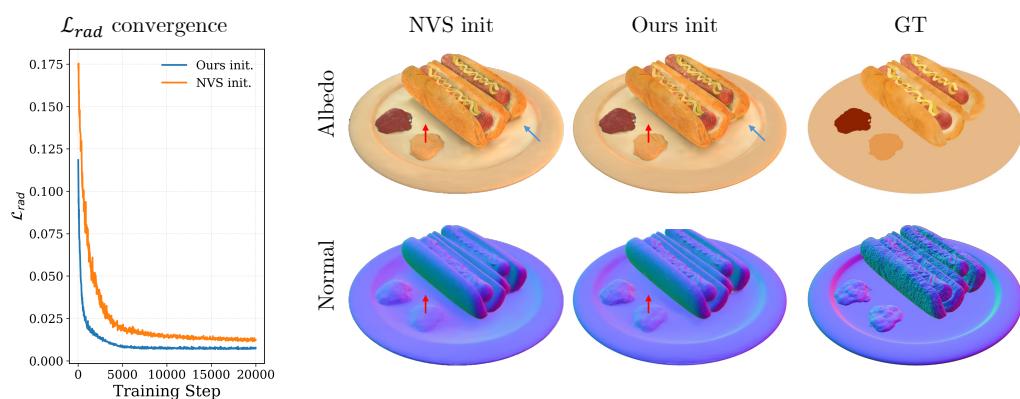
1514  
 1515 We compared our initialization strategy against NVS-based initialization methods. We define "NVS  
 1516 Init." as a model utilizing NVS pre-training for initialization followed by our full inverse rendering  
 1517 optimization. Additionally, we evaluated "NVS Init. w/o  $\mathcal{L}_{rad}$ ", where the radiometric consistency  
 1518 loss is removed throughout both the initialization and the subsequent optimization stages.

1519 Table 11 summarizes the quantitative results. The "NVS Init. w/o  $\mathcal{L}_{rad}$ " baseline exhibits the lowest  
 1520 performance across all tasks, highlighting the necessity of radiometric supervision. While introducing  
 1521 radiometric consistency after NVS pre-training ("NVS Init.") improves albedo and relighting  
 1522 quality, it still falls short of our method. Notably, even when our method is trained without radio-  
 1523 metric consistency during the main optimization phase ("Ours w/o  $\mathcal{L}_{rad}$ "), it outperforms the full  
 1524 NVS-initialized model in albedo reconstruction. This suggests that enforcing physical constraints  
 1525 from the start is crucial for robust disentanglement.

1526 Table 11: Ablation study on initialization strategies. We compare our initialization (Ours) against  
 1527 NVS-based initialization baselines. "NVS init." denotes NVS pre-training followed by our standard  
 1528 optimization. "during IR" refers to the inverse rendering optimization stage.

Method	NVS PSNR $\uparrow$	Albedo PSNR $\uparrow$	Relight PSNR $\uparrow$
NVS init. + $\lambda_{rad} = 0$ during IR	35.70	30.31	31.51
NVS init.	37.43	30.34	31.66
Ours w/o $\mathcal{L}_{rad}$ ( $\lambda_{rad} = 0$ during IR)	35.82	30.82	31.69
<b>Ours</b>	<b>37.86</b>	<b>31.05</b>	<b>32.09</b>

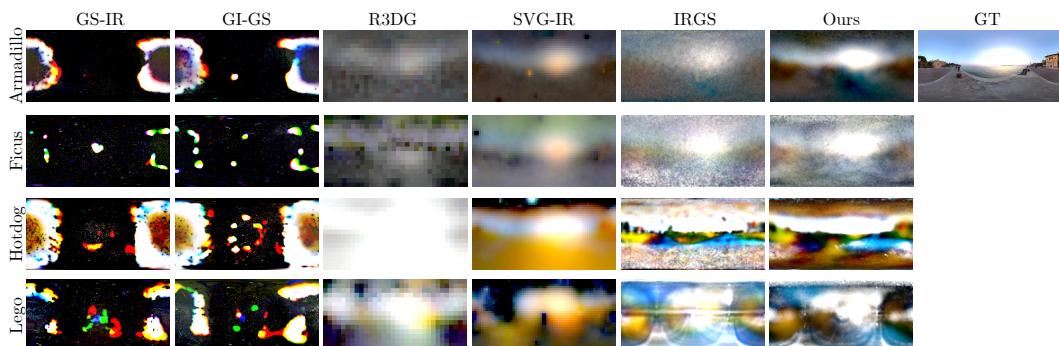
1534 We further analyze the convergence behavior and qualitative results in Figure 16. As shown in the  
 1535 convergence plot (left of Figure 16), our initialization leads to substantially faster and more stable  
 1536 convergence of the radiometric consistency loss  $\mathcal{L}_{rad}$ , whereas standard NVS initialization results  
 1537 in higher and noisier residuals. This demonstrates that conventional NVS initialization places the  
 1538 surfels into a suboptimal orientation that is misaligned with the physical decomposition required for  
 1539 inverse rendering with radiometric consistency. Qualitatively (right of Figure 16), NVS pre-training  
 1540 causes the model to misinterpret geometric cues, leading to shadows baked into the geometry (red  
 1541 arrows) and inter-reflections merged into the albedo (blue arrows). In contrast, our initialization  
 1542 preserves physical priors from the beginning, enabling enhanced geometry reconstruction and a  
 1543 cleaner separation of inter-reflection effects.



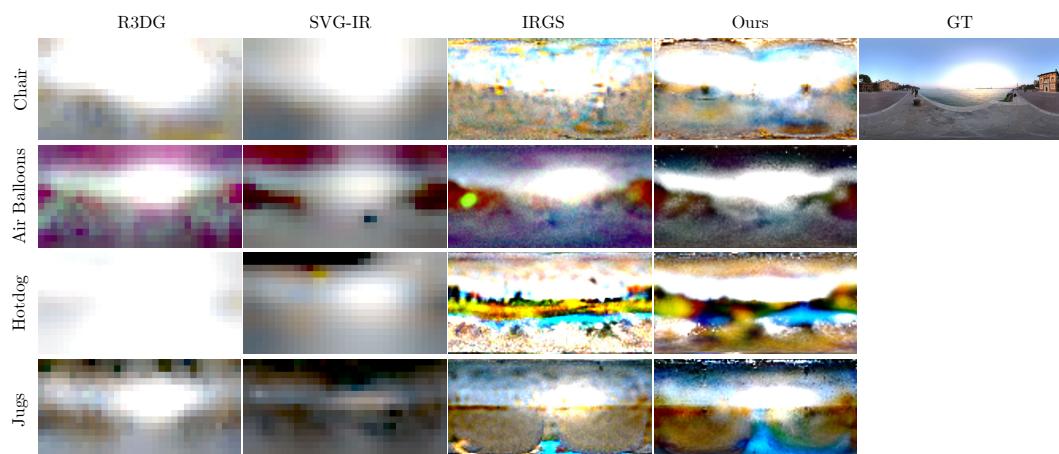
1559 Figure 16: Ablation study of our initialization method on the "hotdog" scene of TensoIR dataset.

1566 I ENVIRONMENT MAP RECONSTRUCTION ON BENCHMARK DATASETS  
1567

1568 We visualized the optimized environment maps from the TensoIR (Jin et al., 2023) and Syn-  
1569 thetic4Relight (Zhang et al., 2022) datasets and compared them with baseline methods and the  
1570 Ground Truth (GT) in Figures 17 and 18. Our method consistently recovers environment maps  
1571 that are visually closest to the ground truth environment map, reflecting the benefit of our radiomet-  
1572 ric consistency, leading to robust disentanglement of illumination for inverse rendering. In contrast,  
1573 GI-GS and GS-IR show unnatural environment maps compared to other baseline methods. Baselines  
1574 relying on point-based ray tracing (R3DG and SVG-IR) often fail to disentangle base color from the  
1575 environment map, such as for the scenes “hotdog” and “lego” from TensoIR, and “air balloons”  
1576 from the Synthetic4Relight dataset. While IRGS uses differentiable Gaussian ray tracing, it tends to  
1577 overestimate the intensity in regions that should be dark, such as the ground of the GT environment  
1578 map.



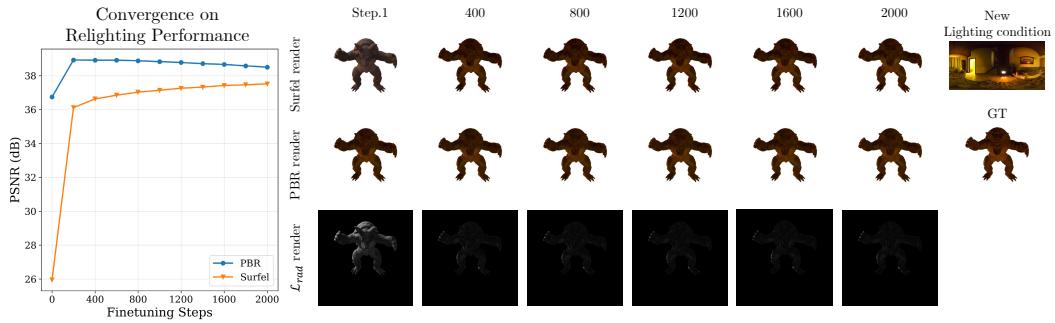
1593 Figure 17: Qualitative comparison on reconstructed environment map on TensoIR dataset. Best  
1594 viewed in zoom.



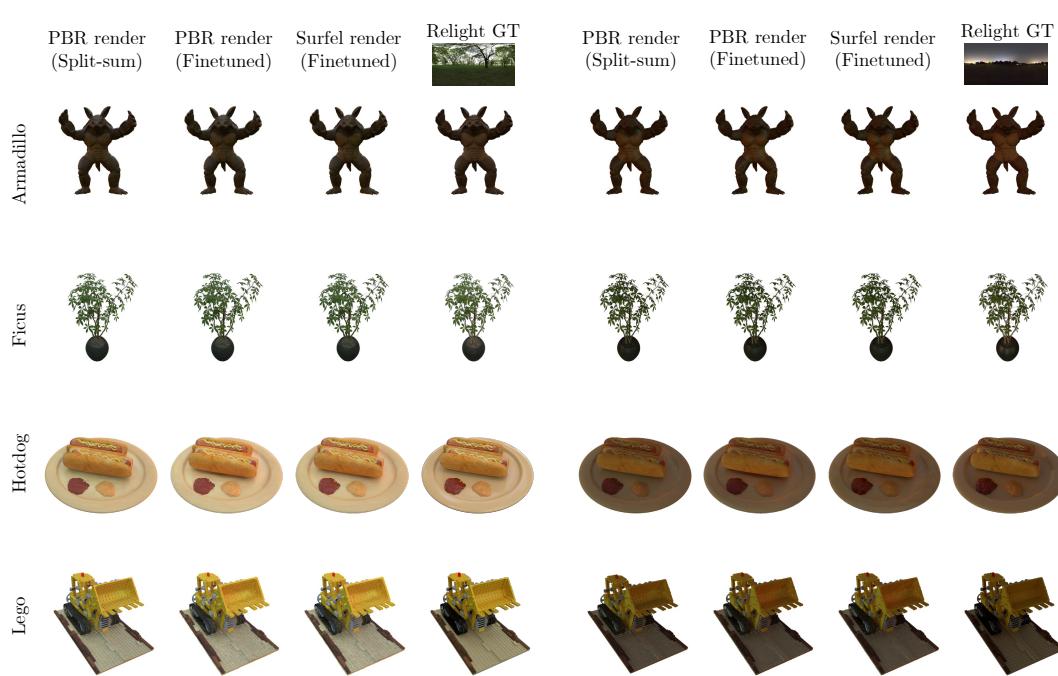
1613 Figure 18: Qualitative comparison on reconstructed environment map on Synthetic4Relight dataset.  
1614 Best viewed in zoom.

## 1620 J ADDITIONAL ANALYSIS ON FINETUNING-BASED RELIGHTING STRATEGY

1621  
 1622 We visualize the convergence and visual progression in Figure 19. As finetuning progresses, the  
 1623 image rendered by surfel radiances quickly converges towards the PBR reference (orange line, left  
 1624 plot). However, the quality of the PBR reference itself slightly degrades over time (blue line, left  
 1625 plot), due to the errors accumulated by the finetune surfel radiances that contribute as indirect il-  
 1626 lumination for physically-based rendering. Nevertheless, our method quickly adapts surfel radiances  
 1627 to new lighting conditions, allowing us to directly use learned surfel radiances to rasterize  
 1628 relighted frames. We additionally provide the corresponding visual comparisons in Figure 20 of  
 1629 the revised paper, which demonstrates that our finetuning-based relighting method shows similar  
 1630 relighting quality compared to ray-traced results.



1641  
 1642 Figure 19: Illustrative figure on how our finetuning-based relighting adapts surfel radiances for new  
 1643 lighting conditions.  
 1644



1668  
 1669 Figure 20: Qualitative ablation study on three relighting method: 1) Gaussian ray tracing that queries  
 1670 indirect radiance via split-sum approximation “PBR render (Split-sum)”, 2) Gaussian ray tracing that  
 1671 queries indirect radiance from finetuned surfels “PBR render (Finetuned)”, and 3) direct rasterization  
 1672 with finetuned surfel radiances “Surfel render (Finetuned)”.  
 1673

1672 While our proposed finetuning-based relighting introduces a pre-computation step compared to ex-  
 1673 isting relighting pipelines, it offers advantages in terms of rendering cost and memory efficiency.

1674  
 1675 We analyze the trade-off between additional finetuning cost and real-time capabilities by comparing  
 1676 average rendering speed and VRAM consumption on the TensoIR dataset. We compare our fine-  
 1677 tuned model against a baseline relighting method that utilizes Gaussian ray tracing for shading with  
 1678 varying sample counts  $N_s \in \{64, 128, 256\}$ .

1679 As reported in Table 12, although finetuning incurs an upfront computational cost, it eliminates the  
 1680 necessity of storing incident radiance for every surfel during inference. This results in significantly  
 1681 reduced memory usage and faster inference speeds ( $\sim 5.90$  ms). In contrast, baseline relighting  
 1682 methods demonstrate increasing rendering time and memory consumption proportional to the sam-  
 1683 ple count  $N$ , as they require storing dense incident ray information for shading each surfel. These  
 1684 results suggest that our method is particularly suitable for memory-constrained environments, such  
 1685 as consumer-grade GPUs or edge devices, where storing dense lighting data becomes prohibitive.

1686 Table 12: Comparison of rendering speed and memory usage. We compare our finetuned model  
 1687 against the baseline relighting method with varying ray tracing sample counts ( $N$ ). Our method  
 1688 achieves significantly lower rendering time and memory footprint.

Method	Sample Count ( $N$ )	Rendering (ms) $\downarrow$	VRAM (MB) $\downarrow$
<b>Ours (Finetuned)</b>	-	<b>5.90</b>	<b>308.2</b>
Baseline (Ray Tracing)	64	38.64	1512.6
	128	58.31	2523.3
	256	86.99	4589.5

1695  
 1696  
 1697  
 1698  
 1699  
 1700  
 1701  
 1702  
 1703  
 1704  
 1705  
 1706  
 1707  
 1708  
 1709  
 1710  
 1711  
 1712  
 1713  
 1714  
 1715  
 1716  
 1717  
 1718  
 1719  
 1720  
 1721  
 1722  
 1723  
 1724  
 1725  
 1726  
 1727