

PBR-NeRF: Inverse Rendering with Physics-Based Neural Fields

Sean Wu¹, Shamik Basu², Tim Brödermann¹, Luc Van Gool^{1,3}, and Christos Sakaridis¹

¹ETH Zürich, ²University Of Bologna, ³INSAIT, Sofia University

Abstract

We tackle the ill-posed inverse rendering problem in 3D reconstruction with a Neural Radiance Field (NeRF) approach informed by Physics-Based Rendering (PBR) theory, named PBR-NeRF. Our method addresses a key limitation in most NeRF and 3D Gaussian Splatting approaches: they estimate view-dependent appearance without modeling scene materials and illumination. To address this limitation, we present an inverse rendering (IR) model capable of jointly estimating scene geometry, materials, and illumination. Our model builds upon recent NeRF-based IR approaches, but crucially introduces two novel physics-based priors that better constrain the IR estimation. Our priors are rigorously formulated as intuitive loss terms and achieve state-of-the-art material estimation without compromising novel view synthesis quality. Our method is easily adaptable to other inverse rendering and 3D reconstruction frameworks that require material estimation. We demonstrate the importance of extending current neural rendering approaches to fully model scene properties beyond geometry and view-dependent appearance. Code is publicly available at: <https://github.com/s3anwu/pbrnerf>.

1. Introduction

Recent advances in neural rendering and 3D reconstruction using Neural Radiance Fields (NeRF) [23] and 3D Gaussian Splatting (3DGS) [18] have driven massive research interest. However, reconstructing accurate scene properties from posed multi-view images remains a challenging, open problem due to inherent ambiguities. The material-lighting ambiguity makes the inverse rendering problem fundamentally ill-posed, with identical images being explainable by infinitely many tuples of scene materials, illumination, and geometry. Addressing this ill-posed inverse problem requires strong priors to narrow the possible solution space.

To overcome this issue, we leverage Physics-Based Rendering (PBR) theory [17, 26] from computer graphics. PBR




Figure 1. **Improving inverse rendering with physics-based priors.** The proposed PBR-NeRF significantly outperforms the NeILF++ [51] baseline by simply using our novel Conservation of Energy and NDF-weighted Specular Losses. Our physics-based losses correct “baked-in” specular highlights misrepresented in the diffuse lobe (highlighted areas) by (1) enforcing energy conservation and (2) accurately separating specular and diffuse reflections. The result is a more realistic, physically consistent material and lighting estimation.

provides a framework for generating physically accurate images given known scene materials, illumination, and geometry. By reversing this process, we derive physics-based priors to better constrain neural forward and inverse ren-

dering and enhance estimation robustness. More specifically, while NeRF, 3DGS, and their neural field derivatives achieve state-of-the-art performance in modeling scene geometry and view-dependent appearance, they do so by treating scenes as “black boxes” that ignore the underlying physics of light transport. These methods represent light by volume-rendering millions of translucent particles, conditioning each particle’s emitted radiance on its position and viewing direction. This formulation significantly improves expressiveness over traditional methods [1, 32], but it does not guarantee physically accurate results. For example, reflective surfaces may be improperly modeled, causing artifacts such as “baked-in” specular highlights, where particles emit more light in outgoing directions than physically possible. Ultimately, the performance of NeRF and 3DGS in 3D reconstruction is inherently limited by their inability to accurately and consistently model diverse light interactions.

In this paper, we solve the complete inverse rendering problem by jointly estimating the geometry, illumination, and materials to accurately model view-dependent appearance. We extend the NeILF/NeILF++ [44, 51] neural fields framework, combining the expressiveness of deep neural networks with the theoretical guarantees of PBR theory. Specifically, our PBR-NeRF approach improves NeILF/NeILF++’s reflection model by leveraging its underlying Disney Bidirectional Reflectance Distribution Function (BRDF) model [11]. Our method enforces energy conservation and promotes the disentanglement of the diffuse and specular lobes of the BRDF, mitigating issues such as highlights “baked in” the diffuse albedo that are commonly observed in NeRF and 3DGS reconstructions.

We achieve these goals through two novel physics-based losses applied directly to the Disney BRDF. Our first loss enforces energy conservation by penalizing physically invalid BRDFs that reflect more energy than received. Our second loss promotes the separation of specular and diffuse BRDF lobes, which is crucial for accurately modeling highly specular surfaces. Without constraints, the diffuse lobe (which defines the albedo) tends to overcompensate near symmetric reflection angles, where both lobes contribute to the aggregate BRDF. By penalizing the diffuse lobe’s magnitude, weighted by the normal distribution function to target specular angles, this loss indirectly encourages the *specular lobe* to capture and fully “explain” specular highlights effectively.

We thoroughly validate PBR-NeRF on two major benchmarks for inverse rendering and novel view synthesis: the NeILF++ [51] and DTU [15] datasets. Our experiments show that PRB-NeRF sets the new state of the art for material estimation in inverse rendering, while maintaining or surpassing the competitive novel view synthesis performance of its baseline. Crucially, our model delivers faithful albedo estimates that remain consistent for a given scene

across diverse illuminations, thanks to its effective diffuse-specular decomposition. This decomposition also improves metallicness and roughness estimates as an additional benefit, further highlighting the effectiveness of our approach.

2. Related Work

Neural fields for novel view synthesis have achieved impressive progress by representing 3D scenes as continuous volumetric fields [23]. This family of methods achieves photorealistic synthesis from unseen viewpoints through differentiable volume rendering, which estimates outgoing radiance by alpha-compositing color over sampled points along each camera ray. We concisely review representative works here and refer the reader to [41] for a comprehensive overview of the associated literature. Scene geometry is commonly represented using density fields [4–6, 23, 24, 28, 42, 52], 3D Gaussians [12, 18, 39], or signed distance fields (SDFs) [36, 45, 46, 49, 50] for images and other optical modalities [7, 14, 30]. While the density fields in NeRF and 3D Gaussians can model arbitrary complex scene geometries, they fundamentally lack well-defined surfaces and normals. These particle-centric representations result in ambiguous surface definitions, which hinder their applicability to PBR. In contrast, SDFs explicitly define surfaces as the zero-level set of a neural network, providing accurate surface and normal estimation. This surface-centric representation makes SDFs inherently more compatible with inverse rendering tasks which rely on PBR and require precise surface definitions.

Neural inverse rendering predates neural fields [19, 22, 31, 33, 37, 48] and learns to predict 3D geometry, materials, and illumination from 2D images. Decomposed neural fields [27, 35, 38, 47] similarly extend conventional neural fields to estimate physical scene properties such as materials and illumination, through inverse rendering or intrinsic decomposition. Compared to NeRF and 3DGS, the above additional material and lighting estimation enables complex downstream tasks such as relighting [21, 29, 43], material editing, and appearance manipulation. The decomposition of reflected radiance into materials and illumination also helps to address issues such as “baked-in” highlights in predicted colors by properly modeling view-dependent effects. However, inverse rendering introduces significant complexity due to the material-lighting-geometry ambiguity, which is harder to lift than the geometry-radiance ambiguity of standard novel view synthesis. Many approaches simplify this task by assuming known lighting, environment map lighting [8, 9, 25, 53, 54], known geometry [44], or basic Lambertian material models. State-of-the-art inverse rendering approaches jointly model (1) direct and indirect lighting from near and far sources [40, 44, 51], and (2) spatially-varying microfacet




Figure 2. **Overview of our PBR-NeRF architecture for neural inverse rendering.** Our two novel physics-based material losses, i.e. $\mathcal{L}_{\text{cons}}$ for conservation of energy and $\mathcal{L}_{\text{spec}}$ for disentangling the diffuse and specular BRDF components, are derived in Sec. 3.2. The complete PBR-NeRF model comprises multiple neural fields, which are optimized in a stage-wise fashion: a standard NeRF+SDF modeling geometry and radiance, a neural incident light field (NeILF) modeling spatially-varying illumination, and a BRDF field modeling materials via the Disney BRDF [11]. Our physics-based losses provide the BRDF field with valuable inductive biases for improved material estimation, which help resolve to a large extent the inherent material-lighting ambiguity in inverse rendering and thus benefit the incident light field as well. Consequently, novel scene views synthesized with our model enjoy state-of-the-art quality.

BRDFs [34, 44, 51]. For computational efficiency, techniques such as light caching [3, 16, 20, 44, 51, 55] and incident light sampling [3, 44, 51] are crucial to avoid costly recursive light tracing.

NeILF [44] proposed joint incident light and BRDF fields, improving on environment-map-based approaches such as NeRV [34] by handling near-field lights and mixed lighting settings in addition to indirect lighting and occlusions. Building on this, NeILF++ [51] introduced inter-reflectable light fields, combining the NeILF and BRDF fields with an SDF for geometry estimation. This approach uses the SDF’s volume-rendered outgoing radiance to supervise the incident radiance predicted by the NeILF. The reflection loss of [51] harnesses the SDF’s strong novel view synthesis performance to steer the ill-posed inverse rendering, achieving state-of-the-art material estimation among neural inverse rendering works. While largely effective, current inverse rendering methods including [44, 51] lack informative material priors for disambiguating specular from diffuse properties of scene surfaces. This typically results in incorrectly attributing specular highlights to alleged changes in diffuse albedo, even though such errors do not affect view synthesis results.

Material priors for inverse rendering are crucial for constraining material estimation and reducing material-lighting

ambiguities. Basic BRDF smoothness priors are used in [3, 16, 44, 51] to discourage abrupt spatial changes in material properties. NeILF [44] and NeILF++ [51] also use a Lambertian prior for regularizing the Disney BRDF’s metallicness and roughness parameters. However, the strong assumption of perfectly rough, non-metallic materials induced by this Lambertian prior limits BRDF expressiveness. For instance, specular effects often appear in the predicted BRDF’s diffuse lobe due to the Lambertian prior suppressing the specular lobe. We address these limitations by proposing two novel physics-based losses that enhance the physical validity and accuracy of estimated BRDFs without sacrificing expressiveness.

3. Method

We present an overview of our PBR-NeRF method with our novel Conservation of Energy and NDF-weighted Specular losses in Fig. 2.

3.1. Background

We build upon the NeILF++ [51] implicit differential renderer (IDR) as the foundation for our physics-based contributions. The IDR estimates outgoing radiance $L_o(\mathbf{x}, \omega_o)$

by solving the Rendering Equation [17]:

$$L_o(\mathbf{x}, \boldsymbol{\omega}_o) = \int_{\Omega} f_r(\mathbf{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o) L_i(\mathbf{x}, \boldsymbol{\omega}_i) (\boldsymbol{\omega}_i \cdot \mathbf{n}) d\boldsymbol{\omega}_i, \quad (1)$$

where $L_o(\mathbf{x}, \boldsymbol{\omega}_o)$ is the outgoing radiance from a surface point \mathbf{x} in direction $\boldsymbol{\omega}_o$, $f_r(\mathbf{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o)$ is the BRDF, and integration is performed over incident directions $\boldsymbol{\omega}_i$ in the positive hemisphere Ω .

To approximate the BRDF in (1), we use the simplified Disney BRDF model [11] with three spatially varying parameters: albedo (base color) $b \in [0, 1]^3$, metallicness $m \in [0, 1]$, and roughness $r \in [0, 1]$. The Disney BRDF is decomposed into a diffuse lobe $f_d(\mathbf{x})$ and a specular lobe $f_s(\mathbf{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o)$:

$$f_r(\mathbf{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o) = f_d(\mathbf{x}) + f_s(\mathbf{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o). \quad (2)$$

The diffuse lobe f_d models the view-independent appearance based on the albedo b and metallicness m as

$$f_d(\mathbf{x}) = \frac{1 - m(\mathbf{x})}{\pi} b(\mathbf{x}). \quad (3)$$

The specular lobe f_s models light L_o reflected in direction $\boldsymbol{\omega}_o$ due to incident radiance L_i from direction $\boldsymbol{\omega}_i$ being reflected across the halfway vector $\boldsymbol{\omega}_h = \frac{\boldsymbol{\omega}_o + \boldsymbol{\omega}_i}{\|\boldsymbol{\omega}_o + \boldsymbol{\omega}_i\|}$. The Disney BRDF approximates rough surface effects in the specular lobe using the microfacet BRDF model [11]. This specular lobe of the BRDF is expressed as

$$f_s(\mathbf{x}, \boldsymbol{\omega}_o, \boldsymbol{\omega}_i) = \frac{D(\boldsymbol{\omega}_h) F(\boldsymbol{\omega}_o, \boldsymbol{\omega}_h) G(\boldsymbol{\omega}_i, \boldsymbol{\omega}_o, \mathbf{n})}{4(\mathbf{n} \cdot \boldsymbol{\omega}_i)(\mathbf{n} \cdot \boldsymbol{\omega}_o)}, \quad (4)$$

where the normal distribution function (NDF) D models the distribution of microfacet orientations, the Fresnel term F models Fresnel reflections, and G models geometric occlusion. We adopt the Spherical Gaussian approximation from [44, 51, 53], in which roughness r controls the sharpness of the NDF through

$$D(\boldsymbol{\omega}_h) = \frac{1}{\pi r^4} \exp\left(\frac{2}{r^4}(\boldsymbol{\omega}_h \cdot \mathbf{n} - 1)\right). \quad (5)$$

Following NeILF++ [51], our IDR framework uses three implicit networks (cf. Fig. 2):

1. **BRDF MLP**: predicts the Disney BRDF parameters b , r , and m to compute the BRDF $f_r(\mathbf{x}, \boldsymbol{\omega}_o, \boldsymbol{\omega}_i)$.
2. **NeILF MLP**: predicts incident radiance $L_i(\mathbf{x}, \boldsymbol{\omega}_i)$.
3. **NeRF SDF**: predicts density σ (geometry) and color c , where c estimates the outgoing radiance $L_o(\mathbf{x}, \boldsymbol{\omega}_o)$.

We approximate the rendering integral in (1) according to NeILF/NeILF++ [44, 51] using a fixed set S_L of incident directions $\boldsymbol{\omega}_i$ generated by Fibonacci sampling:

$$L_o(\mathbf{x}, \boldsymbol{\omega}_o) = \frac{2\pi}{|S_L|} \sum_{\boldsymbol{\omega}_i \in S_L} f_r(\mathbf{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o) L_i(\mathbf{x}, \boldsymbol{\omega}_i) (\boldsymbol{\omega}_i \cdot \mathbf{n}). \quad (6)$$

To simplify notation in the following sections, we omit the arguments \mathbf{x} , $\boldsymbol{\omega}_i$, and $\boldsymbol{\omega}_o$ from f_r , f_d , f_s , and L_i .




Figure 3. **Illustration of our physics-based losses.** We constrain our Disney BRDF [11] material and NeILF [44, 51] incident light estimation with two novel physics-based losses: (1) the Conservation of Energy Loss $\mathcal{L}_{\text{cons}}$ to supervise the complete BRDF $f_r = f_s + f_d$ denoted by the dotted envelope, and (2) the NDF-weighted Specular Loss $\mathcal{L}_{\text{spec}}$ to adjust the relative magnitudes of the specular f_s (red) and diffuse f_d (blue) BRDF lobes.

3.2. Physics-Based Losses

Since inverse rendering is ill-posed, we propose two novel losses inspired by physical priors to constrain both material and lighting estimation, which are illustrated in Fig. 3.

Conservation of Energy Loss. Our first contribution addresses the non-energy-conserving behavior inherent in the Disney BRDF [11] and other microfacet models. Despite their widespread use, these models fail to conserve energy, as demonstrated in previous works [10, 13]. This limitation poses a critical challenge for inverse rendering, as it permits a material to create energy (reflect more light than received) or destroy energy (reflect too little light). Such inaccuracies skew material estimates and disrupt lighting estimation, causing the estimated illumination to overcompensate and hence appear too bright or dark. These effects cascade to downstream tasks, such as relighting and object insertion, and pose a risk for them.

To mitigate energy creation issues, we enforce the energy conservation property for the BRDF by requiring $\int_{\Omega} f_r(\boldsymbol{\omega}_i \cdot \mathbf{n}) d\boldsymbol{\omega}_i \leq 1$. This constraint, derived from the rendering equation (1), requires that the sum of reflected radiance weights $f_r(\boldsymbol{\omega}_i \cdot \mathbf{n})$ across all incident directions in the hemisphere Ω must not exceed unity. By enforcing this intrinsic physical property, independent of incident lighting, we ensure physically correct materials.

We reformulate this constraint in our discretized setting as the Conservation of Energy Loss:

$$\mathcal{L}_{\text{cons}} = \max\left\{\left(\frac{2\pi}{|S_L|} \sum_{\boldsymbol{\omega}_i \in S_L} f_r(\boldsymbol{\omega}_i \cdot \mathbf{n})\right) - 1, 0\right\}. \quad (7)$$

This ReLU-style formulation prevents overestimation of reflected light, allowing reflected irradiance to vary but with-

out exceeding the incident irradiance. We only penalize the model when the sum of weights $f_r(\omega_i \cdot \mathbf{n})$ exceeds 1, but any value below can freely vary. As seen in Fig. 3, the Conservation of Energy Loss affects both specular and diffuse lobes, ensuring that the overall BRDF f_r conserves energy.

NDF-weighted Specular Loss. Our second physics-based loss targets the imbalance between the diffuse and specular lobes frequently observed in inverse rendering methods, such as NeILF++ [51], which often assume Lambertian reflection. Real-world materials, however, often violate idealized Lambertian behavior, resulting in “baked-in” specular highlights where the diffuse lobe f_d compensates for insufficient specular reflection. This imbalance causes incorrect diffuse behavior at non-specular angles, degrading material estimation quality.

To mitigate this issue, we propose the NDF-weighted Specular Loss to penalize excessive diffuse reflection in specular regions. When specular highlights are “baked-in”, the aggregate BRDF value f_r may be accurate for specular directions, but the oversized diffuse lobe f_d overcompensates the specular lobe f_s . We define our loss as

$$\mathcal{L}_{\text{spec}} = \frac{1}{|S_L|} \sum_{\omega_i \in S_L} \text{softmax} \left(\frac{\text{detach} \left(D \left(\frac{\omega_o + \omega_i}{\|\omega_o + \omega_i\|} \right) \right)}{T_{\text{spec}}} \right) f_d, \quad (8)$$

where D is the NDF evaluated using the halfway vector $\frac{\omega_o + \omega_i}{\|\omega_o + \omega_i\|} = \omega_h$. By using a softmax-weighted NDF with temperature T_{spec} , we selectively penalize diffuse reflections in regions where specular effects dominate. We detach the NDF term from the computation graph to prevent gradient flow through $D(\omega_h)$ during backpropagation, ensuring that only diffuse contributions to $\mathcal{L}_{\text{spec}}$ result in weight updates.

Our NDF-weighted Specular Loss complements the standard RGB rendering loss by pushing the specular lobe f_s to compensate for the penalized diffuse lobe f_d , creating a stronger separation between diffuse and specular reflections. When the specular lobe is properly recovered, the RGB rendering loss prevents the diffuse lobe from vanishing. We note that our loss is meant to fine-tune material estimation *after* initial convergence with the RGB rendering loss. A proper weighting of the opposing RGB rendering loss and other regularization terms ensures that the diffuse lobe is not arbitrarily suppressed by our specular loss. As with $\mathcal{L}_{\text{cons}}$, the specular loss also benefits lighting estimation. An oversized diffuse lobe can lead to lower estimates of light intensities, while excessive diffuse reflection forces the estimated incident light to produce complex, view-dependent specular effects. Together, these losses improve the accuracy of material and lighting estimation significantly, resulting in more realistic scene reconstructions,

as we detail in Sec. 4.

3.3. Joint Optimization

Following NeILF++ [51], we optimize the scene reconstruction in three phases. We briefly summarize the three phases here (See Appendix Sec. C.2 for full details):

1. **Geometry:** only train NeRF SDF to initialize estimated geometry using a geometry-based loss \mathcal{L}_{geo} .
2. **Material:** train NeILF and BRDF MLPs while freezing the NeRF SDF to initialize estimated illumination and materials with a material-based loss \mathcal{L}_{mat} .
3. **Joint Optimization:** train all fields (NeRF SDF, NeILF, BRDF) with all losses $\mathcal{L} = \mathcal{L}_{\text{geo}} + \mathcal{L}_{\text{mat}}$.

Each phase uses different losses that encode specific priors or constraints relevant to that phase. The full material-based loss including our complete physics-based loss $\mathcal{L}_{\text{mat,physics}}$ is

$$\mathcal{L}_{\text{mat}} = \lambda_{\text{pbr}} \mathcal{L}_{\text{pbr}} + \mathcal{L}_{\text{ref}} + \lambda_{\text{smth}} \mathcal{L}_{\text{smth}} + \mathcal{L}_{\text{mat,physics}}, \quad (9)$$

where $\mathcal{L}_{\text{mat,physics}} = \lambda_{\text{cons}} \mathcal{L}_{\text{cons}} + \lambda_{\text{spec}} \mathcal{L}_{\text{spec}}$ includes our Conservation of Energy Loss and our NDF-weighted Specular Loss, \mathcal{L}_{pbr} is a standard RGB rendering loss supervising the estimated outgoing radiance from (6), \mathcal{L}_{ref} is the NeILF++ reflection loss [51], $\mathcal{L}_{\text{smth}}$ is the NeILF/NeILF++ BRDF smoothness loss [44, 51], and $\lambda_{\text{pbr}}, \lambda_{\text{ref}}, \lambda_{\text{smth}}, \lambda_{\text{cons}}, \lambda_{\text{spec}}$ are positive weights.

4. Experiments

4.1. Experimental Setup

Implementation Details. We implement PBR-NeRF on top of NeILF++ [51] using PyTorch [2] and keep all hyperparameters from NeILF++ the same, except for $|S_L| = 256$. We train for 30K iterations with an initial learning rate of 0.002, reducing it by a factor of 5 every 10K iterations. Training runs on a single NVIDIA RTX 4090, with runtimes ranging from 1 to 3 hours. Using grid search, we identify two sets of weights for our physics-based losses: $\lambda_{\text{cons}} = 0.01$ and $\lambda_{\text{spec}} = 0.5$ for the NeILF++ dataset, and $\lambda_{\text{cons}} = 0.01$ and $\lambda_{\text{spec}} = 0.01$ for DTU. More hyperparameter details are provided in Appendix Sec. C.3.

Datasets. We evaluate our method on the NeILF++ [51] and DTU [15] datasets. We use the NeILF++ synthetic dataset to evaluate our material-illumination estimation, assuming known geometry. This dataset provides ground-truth geometry, RGB, albedo, metallicness, and roughness for a scene rendered under 6 static illumination patterns, including three environment maps (Env) and three mixed-lighting scenarios with additional point and area light sources (Mix). Each pattern includes 96 RGB HDR 1600×1200 images, split into 87 training and 9 validation images. Ground-truth geometry and poses are used in training, while ground-truth color, albedo, metallicness, and roughness are only used for evaluation. DTU captures 15

Table 1. **Comparison against state-of-the-art methods on the NeILF++ dataset [51].** Performance in material estimation and novel view synthesis is compared across different scenes (City, Studio, and Castel) and types of illumination (Env: global environment map, Mix: mixed lighting with environment map, point and area light sources) using the PSNR metric. †: our reproduced results.

Predicted Quantity	Method	Env			Mix			Mean
		City	Studio	Castel	City	Studio	Castel	
RGB	PhySG [53]	24.82	25.65	27.24	24.38	24.04	25.81	25.32
	SG-ENV [53]	31.01	29.46	32.34	27.20	25.88	27.70	28.93
	NeILF [44]	33.77	31.07	35.28	31.11	28.59	32.11	31.99
	NeILF++ [51]†	31.53	29.77	33.87	29.64	27.50	30.74	30.51
	PBR-NeRF (ours)	32.05	30.67	34.58	29.13	27.13	30.03	30.60
Roughness	PhySG [53]	6.62	11.29	6.22	6.27	6.83	6.14	7.23
	SG-ENV [53]	9.61	17.64	9.74	8.77	12.58	9.14	11.25
	NeILF [44]	17.48	18.30	13.40	17.05	16.27	16.44	16.49
	NeILF++ [51]†	21.22	22.01	19.91	22.19	23.43	22.20	21.83
	PBR-NeRF (ours)	22.20	23.19	22.49	22.19	22.30	22.36	22.46
Metallicness	PhySG [53]	8.72	7.97	8.35	8.67	8.95	8.76	8.57
	SG-ENV [53]	17.01	16.40	16.39	15.44	14.25	14.49	15.66
	NeILF [44]	18.22	19.11	10.29	18.42	18.43	17.34	16.98
	NeILF++ [51]†	19.49	18.17	17.41	20.80	17.98	22.20	19.34
	PBR-NeRF (ours)	21.94	21.12	21.24	21.63	17.40	20.50	20.64
Albedo	PhySG [53]	15.01	16.96	16.13	14.16	12.43	14.29	14.83
	SG-ENV [53]	22.38	20.74	22.21	16.92	13.16	16.69	18.68
	NeILF [44]	15.59	15.48	12.95	17.39	16.88	17.37	15.94
	NeILF++ [51]†	18.21	19.50	17.73	16.37	14.49	17.43	17.29
	PBR-NeRF (ours)	19.42	19.15	19.50	19.33	18.36	18.78	19.09

real-world objects in a controlled lab setting. Each scene includes 49 RGB LDR 1600×1200 images, split into 44 training and 5 validation images. We only qualitatively evaluate material estimation on DTU because it contains LDR images and does not provide ground-truth materials.

4.2. Comparison with The State of The Art

Quantitative Results. Table 1 shows that PBR-NeRF achieves state-of-the-art material estimation on the NeILF++ dataset, outperforming previous methods by 1–2 PSNR on average across albedo, metallicness, and roughness. We outperform all baselines in roughness and metallicness in most scenes across all illuminations. For albedo, we observe substantial gains over all competitors both on overall average and in Mix lighting scenarios, which are generally more complex due to their combined lighting setups. Despite performing slightly worse than NeILF [44] in novel view synthesis, we outperform the closely related NeILF++ [51] in novel view synthesis and all competitors in material estimation.

Among inverse rendering methods, PBR-NeRF achieves state-of-the-art results in novel view synthesis on DTU, as shown in Table 2. We report the average results over three training runs for both our approach and NeILF++ [51], mitigating the impact of training variability. Our method achieves an improvement of 0.37 in mean PSNR over NeILF++ and substantially outperforms all other methods. Additionally, PBR-NeRF ranks first in 10 scenes, with PSNR gains ranging from 0.12 to 2.51 over NeILF++,

demonstrating consistent improvements over our baseline.

Qualitative Results. We further validate our material estimation performance qualitatively against NeILF++ on the NeILF++ dataset in Fig. 4 (see Appendix Sec. B for additional qualitative results). Over all three challenging Mix lighting conditions, our physics-based losses improve both lighting and material estimation. Our estimated lighting exhibits lower entropy with more concentrated sources and fewer artifacts, especially in the upper halves of environment maps. We also eliminate fringes and patches generated on the sphere and cube by NeILF++ due to “baked-in” specular highlights and achieve effective separation of diffuse and specular BRDF components via our NDF-weighted Specular Loss. Most importantly, we consistently outperform NeILF++ on albedo estimation, which is a critical challenge in inverse rendering. Note that although NeILF++ surpasses our metallicness estimation on Studio Mix, it fails significantly to recover the correct albedo. In general, our predicted albedo is—correctly—less bright than for NeILF++, with the floor, cube, and sphere all free from incorrect specular highlights and artifacts.

We also compare our material estimation on DTU in Fig. 5. On scan 37, we improve metallicness and roughness, accurately predicting the scissor and clamp as metal, unlike NeILF++, which creates artifacts near the scissor handle. NeILF++ also predicts the scissor and part of the clamp as rough material instead of metal. This issue creates patchy artifacts both in the metallicness and roughness. On scan 40 with the rough bricks, our roughness estimate on shadows




Figure 4. **Qualitative comparisons on the NeILF++ dataset [51].** †: no ground-truth environment maps are provided with the dataset.

Table 2. Comparison for novel view synthesis against state-of-the-art methods on DTU [15]. Performance is compared across different scenes using the PSNR metric. †: our reproduced results. ‡: mean result over 3 runs.

Method \ DTU scan		24	37	40	55	63	65	69	83	97	105	106	110	114	118	122	Mean
RGB PSNR	NeRFactor [54]	23.24	21.91	23.33	26.86	22.70	24.71	27.59	22.56	20.45	25.08	26.30	25.14	21.35	26.44	26.53	24.28
	PhySG [53]	17.38	15.11	20.65	18.71	18.89	18.47	18.08	21.98	17.31	20.67	18.75	17.55	21.20	18.78	23.16	19.11
	Neural-PIL [9]	20.67	19.51	19.12	21.01	23.70	18.94	17.05	20.54	19.88	19.67	18.20	17.75	21.38	21.69	-	19.94
	NeILF++ [51]†‡	25.85	21.93	25.42	27.99	29.76	27.43	26.54	28.37	23.71	28.28	31.56	30.45	26.86	31.32	32.19	27.85
	PBR-NeRF (ours)‡	25.97	21.91	27.32	27.95	29.88	27.45	26.96	30.88	24.49	27.53	31.71	30.53	26.91	30.55	33.22	28.22




Figure 5. Qualitative comparisons on DTU [15]. Best viewed on a screen at full zoom.

is better but we also exhibit albedo artifacts. On scan 63, we successfully recover fine details such as the green apple stem, which NeILF++ misses in its material predictions.

4.3. Ablation Study

We perform an ablation study of the proposed physically-based losses in Tab. 3 to individually evaluate each loss’s contribution. A qualitative analysis is provided in Appendix Sec. A. The baseline (ID 1) is NeILF++ [51] without the Lambertian loss. Adding the Conservation of Energy Loss (ID 2) improves material estimation, with albedo PSNR increasing by 0.84 and metallicness PSNR by 0.8, respectively. In our final method (ID 3), adding the specular loss $\mathcal{L}_{\text{spec}}$ further improves the albedo and metallicness. Specifically, the albedo PSNR increases by 2.02 PSNR while the albedo SSIM improves by 6.99%. Although the metallicness PSNR drops by 0.10 PSNR from ID 2, the metallicness SSIM increases by 9.45%. Overall, our complete method (ID 3) with both physics-based losses improves metallicness by 0.70 PSNR and albedo by 2.86 PSNR over the baseline (ID 1), while maintaining comparable performance in RGB and roughness.

Table 3. Ablation of our physics-based losses. Performance of ablated versions of PBR-NeRF and of its full version is compared on the *City* illumination using PSNR and SSIM for Disney BRDF parameters and novel views.

ID	$\mathcal{L}_{\text{cons}}$	$\mathcal{L}_{\text{spec}}$	RGB		Roughness		Metallicness		Albedo	
			PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
1	×	×	31.51	0.9153	21.22	0.9121	21.06	0.7596	17.62	0.8110
2	✓	×	31.59	0.9152	21.41	0.9132	21.86	0.7192	18.46	0.8462
3	✓	✓	31.50	0.9148	21.18	0.9123	21.76	0.8137	20.48	0.9161

5. Conclusion

We present PBR-NeRF, a novel inverse rendering method leveraging neural fields optimized with PBR-inspired losses. These losses provide physically valid inductive biases, enabling the neural fields to better disambiguate materials from lighting—addressing a major challenge in inverse rendering. Thorough experiments across various benchmarks demonstrate the effectiveness of PBR-NeRF in not only material estimation and inverse rendering, but also in novel view synthesis. In particular, we qualitatively show a compelling disentanglement of the diffuse and specular components in our estimated BRDFs. We believe that our physics-based approach will inspire further work on PBR-driven, decomposed neural fields.

References

- [1] Sameer Agarwal, Yasutaka Furukawa, Noah Snavely, Ian Simon, Brian Curless, Steven M Seitz, and Richard Szeliski. Building rome in a day. *Communications of the ACM*, 54(10):105–112, 2011. [2](#)
- [2] Jason Ansel, Edward Yang, Horace He, Natalia Gimelshein, Animesh Jain, Michael Voznesensky, Bin Bao, Peter Bell, David Berard, Evgeni Burovski, Geeta Chauhan, Anjali Chourdia, Will Constable, Alban Desmaison, Zachary DeVito, Elias Ellison, Will Feng, Jiong Gong, Michael Gschwind, Brian Hirsh, Sherlock Huang, Kshiteej Kalambarakar, Laurent Kirsch, Michael Lazos, Mario Lezcano, Yanbo Liang, Jason Liang, Yinghai Lu, CK Luk, Bert Maher, Yunjie Pan, Christian Puhersch, Matthias Reso, Mark Saroufim, Marcos Yukio Siraichi, Helen Suk, Michael Suo, Phil Tillet, Eikan Wang, Xiaodong Wang, William Wen, Shunting Zhang, Xu Zhao, Keren Zhou, Richard Zou, Ajit Mathews, Gregory Chanan, Peng Wu, and Soumith Chintala. PyTorch 2: Faster Machine Learning Through Dynamic Python Bytecode Transformation and Graph Compilation. In *29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2 (ASPLOS '24)*. ACM, 2024. [5](#)
- [3] Benjamin Attal, Dor Verbin, Ben Mildenhall, Peter Hedman, Jonathan T Barron, Matthew O’Toole, and Pratul P. Srinivasan. Flash Cache: Reducing bias in radiance cache based inverse rendering. In *European Conference on Computer Vision (ECCV)*, 2024. [3](#)
- [4] Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P Srinivasan. Mip-NeRF: A multiscale representation for anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. [2](#)
- [5] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-NeRF 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5470–5479, 2022.
- [6] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Zip-NeRF: Anti-aliased grid-based neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 19697–19705, 2023. [2](#)
- [7] David Borts, Erich Liang, Tim Broedermann, Andrea Ramazzina, Stefanie Walz, Edoardo Palladin, Jipeng Sun, David Brueggemann, Christos Sakaridis, Luc Van Gool, Mario Bijelic, and Felix Heide. Radar Fields: Frequency-space neural scene representations for FMCW radar. In *ACM SIGGRAPH*, 2024. [2](#)
- [8] Mark Boss, Raphael Braun, Varun Jampani, Jonathan T Barron, Ce Liu, and Hendrik Lensch. NeRD: Neural reflectance decomposition from image collections. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 12684–12694, 2021. [2](#)
- [9] Mark Boss, Varun Jampani, Raphael Braun, Ce Liu, Jonathan Barron, and Hendrik Lensch. Neural-PIL: Neural pre-integrated lighting for reflectance decomposition. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. [2](#), [8](#)
- [10] Brent Burley. Extending the Disney BRDF to a BSDF with integrated subsurface scattering. *SIGGRAPH Course: Physically Based Shading in Theory and Practice*. ACM, New York, NY, 19(7):9, 2015. [4](#)
- [11] Brent Burley and Walt Disney Animation Studios. Physically-based shading at Disney. In *ACM SIGGRAPH*, 2012. [2](#), [3](#), [4](#), [13](#)
- [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. In *The European Conference on Computer Vision (ECCV)*, 2024. [2](#)
- [13] Eric Heitz. Understanding the masking-shadowing function in microfacet-based BRDFs. *Journal of Computer Graphics Techniques*, 3(2):32–91, 2014. [4](#)
- [14] Shengyu Huang, Zan Gojicic, Zian Wang, Francis Williams, Yoni Kasten, Sanja Fidler, Konrad Schindler, and Or Litany. Neural LiDAR fields for novel view synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. [2](#)
- [15] Rasmus Jensen, Anders Dahl, George Vogiatzis, Engin Tola, and Henrik Aanæs. Large scale multi-view stereopsis evaluation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 406–413, 2014. [2](#), [5](#), [8](#), [15](#)
- [16] Haian Jin, Isabella Liu, Peijia Xu, Xiaoshuai Zhang, Songfang Han, Sai Bi, Xiaowei Zhou, Zexiang Xu, and Hao Su. TensorIR: Tensorial inverse rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 165–174, 2023. [3](#)
- [17] James T Kajiya. The rendering equation. In *ACM SIGGRAPH*, 1986. [1](#), [4](#)
- [18] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3D Gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4), 2023. [1](#), [2](#)
- [19] Zhengqin Li, Mohammad Shafiei, Ravi Ramamoorthi, Kalyan Sunkavalli, and Manmohan Chandraker. Inverse rendering for complex indoor scenes: Shape, spatially-varying lighting and SVBRDF from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. [2](#)
- [20] 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. *ACM Transactions on Graphics (TOG)*, 42(4):1–22, 2023. [3](#)
- [21] Linjie Lyu, Ayush Tewari, Thomas Leimkühler, Marc Habermann, and Christian Theobalt. Neural radiance transfer fields for relightable novel-view synthesis with global illumination. In *The European Conference on Computer Vision (ECCV)*, 2022. [2](#)
- [22] Linjie Lyu, Ayush Tewari, Marc Habermann, Shunsuke Saito, Michael Zollhöfer, Thomas Leimkühler, and Christian

- Theobalt. Diffusion posterior illumination for ambiguity-aware inverse rendering. *ACM Transactions on Graphics (TOG)*, 42(6):1–14, 2023. 2
- [23] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. NeRF: Representing scenes as neural radiance fields for view synthesis. In *European Conference on Computer Vision (ECCV)*, 2020. 1, 2
- [24] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multi-resolution hash encoding. *ACM Transactions on Graphics (TOG)*, 41(4):1–15, 2022. 2
- [25] Jacob Munkberg, Jon Hasselgren, Tianchang Shen, Jun Gao, Wenzheng Chen, Alex Evans, Thomas Müller, and Sanja Fidler. Extracting triangular 3D models, materials, and lighting from images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [26] Matt Pharr, Wenzel Jakob, and Greg Humphreys. *Physically Based Rendering: From Theory to Implementation*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 3rd edition, 2016. 1
- [27] Andrea Ramazzina, Mario Bijelic, Stefanie Walz, Alessandro Sanvito, Dominik Scheuble, and Felix Heide. ScatterNeRF: Seeing through fog with physically-based inverse neural rendering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. 2
- [28] Daniel Rebain, Wei Jiang, Soroosh Yazdani, Ke Li, Kwang Moo Yi, and Andrea Tagliasacchi. DeRF: Decomposed radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [29] Viktor Rudnev, Mohamed Elgharib, William Smith, Lingjie Liu, Vladislav Golyanik, and Christian Theobalt. NeRF for outdoor scene relighting. In *The European Conference on Computer Vision (ECCV)*, 2022. 2
- [30] Viktor Rudnev, Mohamed Elgharib, Christian Theobalt, and Vladislav Golyanik. EventNeRF: Neural radiance fields from a single colour event camera. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2
- [31] Shen Sang and Manmohan Chandraker. Single-shot neural relighting and SVBRDF estimation. In *European Conference on Computer Vision (ECCV)*, 2020. 2
- [32] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2
- [33] 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019. 2
- [34] Pratul P Srinivasan, Boyang Deng, Xiuming Zhang, Matthew Tancik, Ben Mildenhall, and Jonathan T Barron. NeRV: Neural reflectance and visibility fields for relighting and view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 3
- [35] Dor Verbin, Peter Hedman, Ben Mildenhall, Todd Zickler, Jonathan T. Barron, and Pratul P. Srinivasan. Ref-NeRF: Structured view-dependent appearance for neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [36] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *arXiv preprint arXiv:2106.10689*, 2021. 2
- [37] Zian Wang, Jonah Philion, Sanja Fidler, and Jan Kautz. Learning indoor inverse rendering with 3D spatially-varying lighting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 2
- [38] Zian Wang, Tianchang Shen, Jun Gao, Shengyu Huang, Jacob Munkberg, Jon Hasselgren, Zan Gojic, Wenzheng Chen, and Sanja Fidler. Neural fields meet explicit geometric representations for inverse rendering of urban scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2
- [39] Zhiru Wang, Shiyun Xie, Chengwei Pan, and Guoping Wang. SpecGaussian with latent features: A high-quality modeling of the view-dependent appearance for 3D Gaussian splatting. In *Proceedings of the ACM International Conference on Multimedia*, 2024. 2
- [40] 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 (CVPR)*, 2023. 2
- [41] Yiheng Xie, Towaki Takikawa, Shunsuke Saito, Or Litany, Shiqin Yan, Numair Khan, Federico Tombari, James Tompkin, Vincent Sitzmann, and Srinath Sridhar. Neural fields in visual computing and beyond. In *Computer Graphics Forum*, pages 641–676. Wiley Online Library, 2022. 2
- [42] Haofei Xu, Anpei Chen, Yuedong Chen, Christos Sakaridis, Yulun Zhang, Marc Pollefeys, Andreas Geiger, and Fisher Yu. MuRF: Multi-baseline radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2
- [43] Yingyan Xu, Gaspard Zoss, Prashanth Chandran, Markus Gross, Derek Bradley, and Paulo Gotardo. ReNeRF: Relightable neural radiance fields with nearfield lighting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. 2
- [44] Yao Yao, Jingyang Zhang, Jingbo Liu, Yihang Qu, Tian Fang, David McKinnon, Yanghai Tsin, and Long Quan. NeILF: Neural incident light field for physically-based material estimation. In *European Conference on Computer Vision (ECCV)*, 2022. 2, 3, 4, 5, 6, 16
- [45] Lior Yariv, Yoni Kasten, Dror Moran, Meirav Galun, Matan Atzmon, Basri Ronen, and Yaron Lipman. Multiview neural surface reconstruction by disentangling geometry and appearance. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 2

- [46] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. [2](#)
- [47] Weicai Ye, Shuo Chen, Chong Bao, Hujun Bao, Marc Pollefeys, Zhaopeng Cui, and Guofeng Zhang. IntrinsicNeRF: Learning intrinsic neural radiance fields for editable novel view synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. [2](#)
- [48] Ye Yu and William A. P. Smith. InverseRenderNet: Learning single image inverse rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. [2](#)
- [49] Jingyang Zhang, Yao Yao, and Long Quan. Learning signed distance field for multi-view surface reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. [2](#)
- [50] Jingyang Zhang, Yao Yao, Shiwei Li, Tian Fang, David McKinnon, Yanghai Tsin, and Long Quan. Critical regularizations for neural surface reconstruction in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6270–6279, 2022. [2](#), [15](#)
- [51] Jingyang Zhang, Yao Yao, Shiwei Li, Jingbo Liu, Tian Fang, David McKinnon, Yanghai Tsin, and Long Quan. NeILF++: Inter-reflectable light fields for geometry and material estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#), [12](#), [13](#), [14](#), [15](#), [16](#)
- [52] Kai Zhang, Gernot Riegler, Noah Snaveley, and Vladlen Koltun. NeRF++: Analyzing and improving neural radiance fields. *arXiv preprint arXiv:2010.07492*, 2020. [2](#)
- [53] Kai Zhang, Fujun Luan, Qianqian Wang, Kavita Bala, and Noah Snaveley. PhysSG: Inverse rendering with spherical Gaussians for physics-based material editing and relighting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. [2](#), [4](#), [6](#), [8](#)
- [54] Xiuming Zhang, Pratul P. Srinivasan, Boyang Deng, Paul Debevec, William T. Freeman, and Jonathan T. Barron. NeRFactor: Neural factorization of shape and reflectance under an unknown illumination. *ACM Transactions on Graphics (ToG)*, 40(6):1–18, 2021. [2](#), [8](#)
- [55] Yuanqing Zhang, Jiaming Sun, Xingyi He, Huan Fu, Rongfei Jia, and Xiaowei Zhou. Modeling indirect illumination for inverse rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18643–18652, 2022. [3](#)

Appendix

A. Ablation Study Qualitative Results

We present qualitative results in Fig. 6 for the ablation study of the proposed physics-based losses on the NeILF++ dataset City scene in Tab. 3 of Sec. 4.3.

To ensure a fair comparison, all ablated versions share the same hyperparameters, varying only the inclusion of the proposed loss terms. For consistency with NeILF++ [51], we use a lower number of sampled incident light directions $|S_L| = 128$ during training instead of our final method’s $|S_L| = 256$ to allow direct comparison with our reproduced NeILF++ [51] results on the City scene in Tab. 1.

The baseline (ID 1) of NeILF++ [51] without the Lambertian loss, exhibits many “baked-in” specular highlights. As shown in Fig. 6, the baseline’s predicted albedo contains many specular artifacts being added to the diffuse lobe. Specifically, there is fringing on the sphere, shadowing artifacts on the cube (particularly on its lower regions shadowed by the other two objects), and an overly bright floor due to the entanglement of the diffuse and specular lobes. These specular effects in the albedo cause incorrect patches on the shadowed cube in the metallicness prediction and overly dark patches in the bottom left of the incident light estimation for the associated cube location.

Adding the Conservation of Energy Loss (ID 2) mitigates these artifacts by enforcing energy conservation for the diffuse and specular lobes. As a result, the lighting estimation is also improved since the estimated light does not need to compensate for the previously non-energy-conserving (overly reflective) predicted materials. The predicted lighting is more consistent and improves the originally overly dim patches in the bottom left corner of the baseline (corresponding to the cube). Reducing specular artifacts in the lighting also improves the material estimation, as shown by the increased albedo and metallicness PSNR in Tab. 3. Qualitatively, we also see that the albedo of the helmet and cube are effectively recovered in all regions except the shadowed area. Despite these improvements, there are still fringing artifacts on the sphere as well as shadowing effects on the ground and cube.

Incorporating the specular loss $\mathcal{L}_{\text{spec}}$ in our final method (ID 3) improves the albedo and metallicness even further. We once again highlight that using both our novel physics-based losses improves metallicness by 0.7 PSNR and albedo by 2.86 PSNR over the baseline (ID 1). Qualitatively, the albedo estimation, which is most important, is significantly improved in our final method (ID 3), especially with the green albedo of the helmet and the gray albedo of the cube. However, there are still limitations with our proposed method, which highlights directions for future improvement. In particular, we fail to fully reproduce the

specular effects in the helmet’s reflective visor and metallic pieces in the predicted RGB. This limitation is likely due to our overestimation of roughness and the missing fine details in our metallicness prediction. Furthermore, we still exhibit limited specular effects in our predicted albedo with fringing on the sphere and shadowy patches on the cube and floor.

Overall, we convincingly improve the state-of-the-art material estimation over our baseline NeILF++ [51]. While challenges remain in handling high-frequency specular effects and reflective surfaces, our method convincingly advances the state of the art in joint material and lighting estimation.

B. Additional Qualitative Results

We present additional qualitative results for the *Env* lighting subset of the NeILF++ dataset [51] in Fig. 7 for the corresponding quantitative results in Tab. 1.

Using our physics-based losses, our method improves both lighting and material estimation compared to the NeILF++ baseline. These improvements are consistent with the trends observed in the more challenging *Mix* subset shown in Fig. 4 of the main paper. For the *Env* subset, we significantly reduce artifacts in the upper halves of the estimated incident light fields and enhance albedo estimation in most scenes.

Notably, in the Studio scene, while some specular patchy effects persist on the helmet, cube, sphere, and floor, our method better captures finer details in the helmet’s metallicness and roughness, which are absent in the NeILF++ baseline. Across the *Env* subset, our roughness and metallicness predictions show considerable improvements, including the removal of shadowy patches on the cube and fringe artifacts on the sphere.

Although some limitations remain, such as missing fine-grained details on the helmet, our physics-based losses deliver consistent gains in material and lighting estimation, highlighting their effectiveness across diverse lighting conditions.

C. Additional Background Details

For completeness, we provide further implementation details on the NeILF++ implicit differential renderer (Sec. C.1), joint material-illumination-geometry optimization (Sec. C.2), and hyperparameter sweeps (Sec. C.3).

C.1. Implicit Differential Renderer (IDR)

The second term of the specular BRDF f_s in Equation (4) is the Fresnel term $F(\omega_o, \omega_h) \in \mathbb{R}^3$. It models glossy reflection at glancing angles due to Fresnel reflection. Following




Figure 6. **Ablation of our physics-based losses.** Qualitative performance of ablated versions of PBR-NeRF and of its full version is compared on the *City* illumination with predicted illumination and Disney BRDF parameter estimation on novel views. †: no ground-truth environment maps are provided with the dataset.

[11], we use the Schlick Fresnel Approximation

$$F(\omega_o, \omega_h; b, m) = F_0 + (1 - F_0)(1 - \omega_o \cdot \omega_h)^5 \quad (10)$$

where

$$F_0 = 0.04(1 - m) + bm \quad (11)$$

The third term in Equation (4) is the geometry term $G(\omega_i, \omega_o, \mathbf{n}; r) \in \mathbb{R}$. It models the masking and shadowing of microfacets depending on the incident and viewing direction.

$$G(\omega_i, \omega_o, \mathbf{n}; r) = G_1(\omega_i \cdot \mathbf{n})G_1(\omega_o \cdot \mathbf{n}) \quad (12)$$

where

$$G_{GGX}(z; r) = \frac{2}{(2 - r^2)z + r^2} \quad (13)$$

as implemented in NeILF++ [51].

C.2. Joint Optimization

We provide a more detailed description of the joint material-illumination-geometry optimization discussed in Sec. 3.3. We summarize all 3 PBR-NeRF training phases and their losses in Tab. 4. The final weightings used for each loss function term are listed in Tab. 5.

Geometry phase. During the geometry phase, the NeRF SDF network is learned using the following loss.

$$\mathcal{L}_{\text{geo}} = \lambda_{\text{NeRF}} \mathcal{L}_{\text{NeRF}} + \lambda_{\text{pcd}} \mathcal{L}_{\text{pcd}} + \lambda_{\text{Eik}} \mathcal{L}_{\text{Eik}} + \lambda_{\text{Hess}} \mathcal{L}_{\text{Hess}} + \lambda_{\text{surf}} \mathcal{L}_{\text{surf}} \quad (14)$$

The NeRF rendering loss $\mathcal{L}_{\text{NeRF}}$ is identical to the original NeRF formulation. The estimated RGB color $L_{o, \text{NeRF}}$




Figure 7. Additional qualitative comparisons on the NeILF++ dataset [51]. †: no ground-truth environment maps are provided with the dataset.

Table 4. **Summary of PBR-NeRF geometry, material, and joint training phases** with their respective losses. \downarrow : loss is downscaled from the previous stage. \blacksquare : the point cloud loss is optionally used to improve geometry SDF initialization. \dagger : loss used when no ground-truth geometry is provided (e.g. DTU dataset). $*$: our novel PBR losses, namely (1) the Conservation of Energy Loss $\mathcal{L}_{\text{cons}}$ and (2) the NDF-weighted Specular Loss $\mathcal{L}_{\text{spec}}$.

Optimization Phase	Geometry-based Losses					Material-based Losses				
	$\mathcal{L}_{\text{NeRF}}$	$\mathcal{L}_{\text{pcd}} \dagger$	$\mathcal{L}_{\text{Eik}} \dagger$	$\mathcal{L}_{\text{Hess}} \dagger$	$\mathcal{L}_{\text{surf}} \dagger$	\mathcal{L}_{pbr}	\mathcal{L}_{ref}	$\mathcal{L}_{\text{smth}}$	$\mathcal{L}_{\text{cons}}^*$	$\mathcal{L}_{\text{spec}}^*$
Geometry	✓	■	✓	✓	✓	✗	✗	✗	✗	✗
Material	✗	✗	✗	✗	✗	✓	✓	✓	✓	✓
Joint	✓	✗	✓	\downarrow	\downarrow	✓	✓	✓	✓	✓

Table 5. **PBR-NeRF loss function term weighting** for the NeILF++ and DTU datasets. \dagger : the point cloud loss \mathcal{L}_{pcd} is only used in the geometry phase for NeRF-SDF initialization. $*$: our novel PBR losses, namely (1) the Conservation of Energy Loss $\mathcal{L}_{\text{cons}}$ and (2) the NDF-weighted Specular Loss $\mathcal{L}_{\text{spec}}$.

	Geometry-based Losses					Material-based Losses				
	λ_{NeRF}	$\lambda_{\text{pcd}} \dagger$	λ_{Eik}	λ_{Hess}	λ_{surf}	λ_{pbr}	λ_{ref}	λ_{smth}	λ_{cons}^*	λ_{spec}^*
NeILF++ [51]	1.0	N/A	N/A	N/A	N/A	1.0	0.1	0.0005	0.01	0.5
DTU [15]	1.0	0.1	0.1	0.001	0.01	1.0	0.1	0.0005	0.01	0.01

is rendered using the NeRF network and compared to the ground truth color c using $\mathcal{L}_{\text{NeRF}}$ with a mean squared error.

$$\mathcal{L}_{\text{NeRF}} = \|L_{o,\text{NeRF}} - c\|_2^2 \quad (15)$$

When ground truth geometry is provided, as in the NeILF++ dataset [51], the estimated RGB color $L_{o,\text{NeRF}}$ is simply evaluated at the ground truth surface point using the viewing direction instead of using volume rendering. Since we have the ground truth geometry, the other geometry priors are not needed and their corresponding losses are assigned a zero weight. Therefore, only the NeRF rendering loss $\mathcal{L}_{\text{NeRF}}$ is used, with the only nonzero loss weight being λ_{NeRF} .

When ground truth geometry is not provided, we use additional geometry priors encoded as loss function terms.

Following [50, 51], we optimize the SDF with a point cloud loss only during the geometry phase. The point cloud loss supervises the predicted SDF distances and normals:

$$\mathcal{L}_{\text{pcd}} = |\mathbb{G}(\mathbf{x}_{\text{pcd}})| + \left(1 - \mathbf{n}_{\text{pcd}} \cdot \frac{\nabla_{\mathbf{x}} \mathbb{G}(\mathbf{x}_{\text{pcd}})}{\|\nabla_{\mathbf{x}} \mathbb{G}(\mathbf{x}_{\text{pcd}})\|} \right). \quad (16)$$

where \mathbf{x}_{pcd} is a point in the point cloud, $\mathbb{G}(\mathbf{x}_{\text{pcd}})$ represents the SDF’s predicted signed distance at \mathbf{x}_{pcd} , and \mathbf{n}_{pcd} is the point cloud normal at \mathbf{x}_{pcd} . This point cloud loss is only used to help initialize the SDF to a satisfactory quality and we perform full material-illumination-geometry optimization without point cloud input during the later joint phase. This use of point clouds is similar to the use of point clouds in 3D Gaussian Splatting for initializing 3D Gaussians.

Note that the point cloud loss \mathcal{L}_{pcd} is strictly optional, as indicated in Tab. 4. Following NeILF++ [51], we use the point cloud loss \mathcal{L}_{pcd} on the DTU dataset to improve initial SDF quality and for a fair performance comparison.

We additionally use the Eikonal loss to penalize the SDF when the gradient at the surface point $\mathbb{G}(\mathbf{x})$ does not have a magnitude of 1.

$$\mathcal{L}_{\text{Eik}} = \left| \|\nabla_{\mathbf{x}} \mathbb{G}(\mathbf{x})\| - 1 \right| \quad (17)$$

Furthermore, the Hessian loss penalizes rapidly changing gradient directions by minimizing the Hessian matrix norm.

$$\mathcal{L}_{\text{Hess}} = \|\mathbf{H}\mathbb{G}(\mathbf{x})\|_1 \quad (18)$$

Finally, the minimal surface loss encourages compact interpolation and extrapolation of unobserved surfaces by minimizing the surface elastic energy.

$$\mathcal{L}_{\text{surf}} = \delta_{\epsilon}(\mathbb{G}(\mathbf{x})) \quad (19)$$

where δ_{ϵ} is the regularized Dirac delta function parametrized by a sharpness ϵ ,

$$\delta_{\epsilon}(z) = \frac{\epsilon \pi^{-1}}{\epsilon^2 + z^2} \quad (20)$$

Material Phase. During the material phase, we train the NeILF and BRDF MLPs with the frozen NeRF SDF weights using the material phase loss:

$$\mathcal{L}_{\text{mat}} = \lambda_{\text{pbr}} \mathcal{L}_{\text{pbr}} + \lambda_{\text{smth}} \mathcal{L}_{\text{smth}} + \lambda_{\text{ref}} \mathcal{L}_{\text{ref}} + \lambda_{\text{cons}} \mathcal{L}_{\text{cons}} + \lambda_{\text{spec}} \mathcal{L}_{\text{spec}} \quad (21)$$

The physically based rendering loss \mathcal{L}_{pbr} supervises the estimated outgoing radiance from (6):

$$\mathcal{L}_{\text{pbr}} = \|L_{o,\text{NeILF++}} - c\|_2^2 \quad (22)$$

Following [44, 51], we use a bilateral smoothness loss $\mathcal{L}_{\text{smth}}$ to encode the assumption that roughness r and metallicness m at surface point x_p are smooth if the color of corresponding pixel p has no sharp gradients:

$$\mathcal{L}_{\text{smth}} = (\|\nabla_{\mathbf{x}} r(\mathbf{x}_p)\| + \|\nabla_{\mathbf{x}} m(\mathbf{x}_p)\|) \exp(-\|\nabla_p I_p\|) \quad (23)$$

where $\nabla_p I_p$ is the image gradient at pixel p .

We also use the NeILF++ inter-reflection loss to use the NeRF SDF predicted outgoing radiance to supervise the incident light predicted by the NeILF MLP along the same ray $\omega_i = -\omega_o$ between two surface points \mathbf{x}_1 and \mathbf{x}_2 .

$$\mathcal{L}_{\text{ref}} = \|L_{i,\text{NeILF}}(\mathbf{x}_2, \omega_i) - L_{o,\text{NeRF}}(\mathbf{x}_1, -\omega_i)\|_1 \quad (24)$$

Joint Phase. During the joint optimization phase, the NeRF SDF, BRDF, and NeILF MLPs have been pre-trained by the previous geometry and material phases, allowing us to jointly optimize all networks simultaneously. We reuse the geometry phase loss \mathcal{L}_{geo} and material phase loss \mathcal{L}_{mat} from (14) and (21), respectively, to obtain the overall joint phase loss

$$\mathcal{L}_{\text{joint}} = \mathcal{L}_{\text{geo}} + \mathcal{L}_{\text{mat}} \quad (25)$$

C.3. Hyperparameter Sweeps

We now specify how we picked various hyperparameters introduced with our novel Conservation of Energy Loss (7) and NDF-weighted Specular Loss (8).

The Conservation of Energy Loss $\mathcal{L}_{\text{cons}}$ does not introduce any new hyperparameters. We increase the number of incident light directions from the default $|S_L| = 128$ in NeILF++ to $|S_L| = 256$ to better enforce energy conservation and separate the diffuse and specular lobes using more samples.

For the NDF-weighted Specular Loss, we use a temperature $T_{\text{spec}} = 1$. We include the softmax temperature T_{spec} in our NDF-weighted Specular Loss definition (7) for full generality and to provide an option to control softmax sharpness.

To determine the loss term weightings for our novel physics-based losses, we perform a grid search with $\lambda_{\text{cons}} \in \{1.5, 1.0, 0.5, 0.1, 0.05, 0.01, 0.005\}$ and $\lambda_{\text{spec}} \in \{1.5, 1.0, 0.5, 0.1, 0.05, 0.01, 0.005\}$.

We also remove the Lambertian loss \mathcal{L}_{Lam} used in NeILF++ by choosing $\lambda_{\text{Lam}} = 0$ to remove the strong limitations on material estimation that the Lambertian assumption causes.

The remaining hyperparameters are identical to NeILF++ [51].