

NeILF: Neural Incident Light Field for Physically-based Material Estimation

Yao Yao¹, Jingyang Zhang², Jingbo Liu¹, Yihang Qu¹, Tian Fang¹, David McKinnon¹, Yanghai Tsin¹, and Long Quan¹

¹ Apple

{yaoyao, jingbo, yihang, fangtian, dmckinnon, ytsin, quan_long}@apple.com

² HKUST

jzhangbs@cse.ust.hk

Abstract. We present a differentiable rendering framework for material and lighting estimation from multi-view images and a reconstructed geometry. In the framework, we represent scene lightings as the Neural Incident Light Field (NeILF) and material properties as the surface BRDF modelled by multi-layer perceptrons. Compared with recent approaches that approximate scene lightings as the 2D environment map, NeILF is a fully 5D light field that is capable of modelling illuminations of any static scenes. In addition, occlusions and indirect lights can be handled naturally by the NeILF representation without requiring multiple bounces of ray tracing, making it possible to estimate material properties even for scenes with complex lightings and geometries. We also propose a smoothness regularization and a Lambertian assumption to reduce the material-lighting ambiguity during the optimization. Our method strictly follows the physically-based rendering equation, and jointly optimizes material and lighting through the differentiable rendering process. We have intensively evaluated the proposed method on our in-house synthetic dataset, the DTU MVS dataset, and real-world BlendedMVS scenes. Our method outperforms previous methods by a significant margin in terms of novel view rendering quality, setting a new state-of-the-art for image-based material and lighting estimation.

Keywords: differentiable rendering, physically-based rendering, BRDF estimation, incident light field

1 Introduction

Material estimation from a set of sparse images is a challenging task in both computer vision and computer graphics. The problem is usually approached by inverse rendering, where the spatially-varying bidirectional reflectance distribution functions (SV-BRDFs) and lightings of the scene are jointly optimized by minimizing the rendering loss. However, the problem is hard to solve due to the complex form of the BRDF and the high-dimensional nature of scene illuminations. To mitigate the problem, previous methods usually apply simplified material and lighting models. For example, non-spatially varying BRDF [47]

is applied for certain types of objects; approximated illuminations, such as co-located flash lights [4, 5, 26, 32] and environment maps [6, 7, 25, 47, 49], are applied to reduce the complexity of the scene lighting. In most scenarios, special capturing devices or environments are required to assist the estimation, limiting these methods to real-world applications. As the result, a practical material estimator is still missing.

On the other hand, recent progress on neural representation has shown promising results for lighting modelling. NeRF [23] jointly optimizes a neural radiance field and a density field, which has demonstrated great success for novel view synthesis. The surface light field is applied to model the outgoing light from the surface, which has been widely applied to neural surface reconstructions [42, 46]. Other methods further decompose observed lights into neural material properties and environmental lightings. However, similar to classical methods, they either use simplified lighting representations [4, 6, 7, 25], or apply approximated occlusion and indirect light handling [34, 49]. Until now, lighting modelling is still an open problem in image-based material estimation.

In this work, we address this long-standing problem by representing scene lightings as the neural incident light field. Without losing generality, the proposed NeILF is capable of modelling lighting conditions of any static scenes. Also, occlusions and indirect lights could be naturally handled in the proposed framework without tracing multiple bounces of rays. For material properties, we consider a simplified Disney BRDF model [8] consisting of base color, roughness and metallic. Implementation-wise, we use multi-layer perceptrons (MLPs) to represent both the incident light field and the BRDF. The NeILF network takes a 5D vector of location and incident direction as inputs, and returns as output a RGB value of the incident light; the material network takes a 3D location as input, and outputs a 5D vector of surface BRDF properties. Meanwhile, to reduce the ambiguity between the material and the scene lighting, we propose two regularization terms, namely the bilateral smoothness and the Lambertian assumption, to constrain the optimization of roughness and metallic. Finally, we analyze similarities between NeILF for material estimation and NeRF for novel view synthesis [23], providing readers an intuitive explanation of the complexity and solvability of the problem.

We demonstrate in several datasets that our method significantly outperforms previous SOTA methods in terms of novel view rendering accuracy. Our method is able to recover the surface BRDF even for scenes with complex lightings and geometries, which cannot be handled by previous environment map based methods. To summarize, main contributions of the paper include:

- Representing scene lightings by the neural incident light field, where occlusions and indirect lights of the scene can be naturally handled.
- A differentiable framework for joint material and lighting estimation, which significantly outperforms previous state-of-the-art methods in various datasets.
- A bilateral smoothness and a Lambertian assumption to constrain the roughness and the metallic, reducing the material-lighting ambiguity during the network optimization.

2 Related Works

2.1 The Rendering Equation

The rendering equation [14] computes the emitted radiance from a surface point \mathbf{x} along a viewing direction $\boldsymbol{\omega}_o$:

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

where \mathbf{n} is the normal of the surface, L_i is the incoming light from direction $\boldsymbol{\omega}_i$, and f is the BRDF function to describe the reflectance property, which is usually decomposed into a diffuse term and a specular term $f = f_d + f_s$. The integration is performed over all incident direction $\boldsymbol{\omega}_i$ on the hemisphere Ω where $\boldsymbol{\omega}_i \cdot \mathbf{n} > 0$.

The goal of material estimation is to recover continuous functions of the scene lighting L_i and the BRDF property f in the above equation. Due to the complex form of the scene lighting and the material property, it is crucial to select suitable representations for L_i and f . In this paper, we propose to use a neural incident light field to model L_i (Sec. 3.1), and apply a simplified Disney BRDF [8] model to approach the BRDF f (Sec. 3.2). Below we give a brief review on the physically-based material estimation from multi-view images.

2.2 Differentiable Rendering

Unlike classical approaches that recover 3D scene parameters in a forward reconstruction manner, differentiable rendering [2, 15] inverts the rendering process in graphics, and optimizes all parameters by minimizing the difference between rendered and input images.

Recently, the technique has been combined with neural representations and has shown promising results for image-based 3D problems. NeRF [23] and follow-up works [19, 22, 48] decouple a 3D scene into a density field and a radiance field. Other methods also apply implicit functions to model different geometry [20, 21, 46] and appearance representations [27, 42]. In another line of works, the received radiance is further decomposed into BRDF properties and input light sources [6, 34, 47, 49]. Our method follows this practice and applies a neural BRDF model to approach the material property of the surface.

2.3 Material and Lighting Estimation

Due to the difficulty of joint material and lighting estimation, previous methods usually apply additional sensors or controlled lightings to facilitate the optimization process. For example, additional sensors [3, 12, 30], co-located flash lights [4, 5, 26, 32], or turn-table settings [11, 39] are applied to capture images or scene lightings. Moreover, simplified material and lighting models, e.g., the non-spatially varying BRDF [47] or approximated illuminations of environment maps [6, 9, 47, 49], are applied to reduce the complexity of the problem. NeRV [34] introduces the visibility field to model indirect lights. However, it requires environmental lightings to be known in advance. Nevertheless,

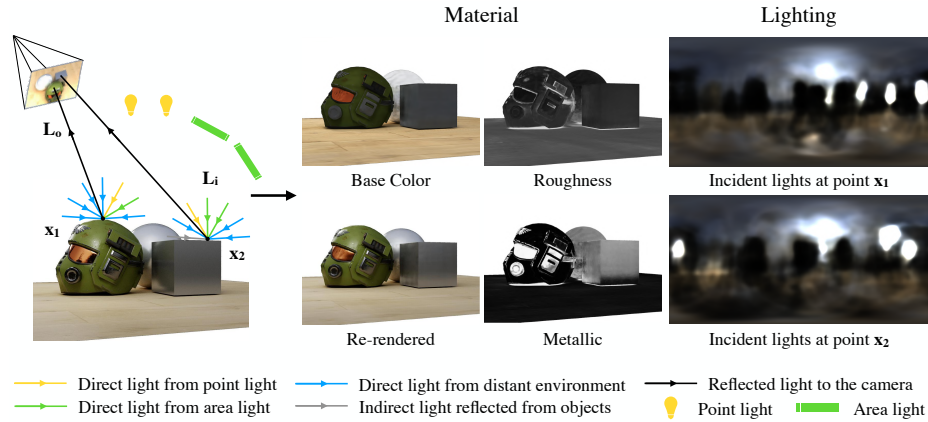


Fig. 1: Illustrations of the proposed method and our material and lighting estimation results. NeILF is capable of modelling the joint illumination of direct/indirect lights from different sources. Estimated incident lights at point x_1 and point x_2 well explain the mixed lighting of the scene, including an environment map with high-radiance sun light, two near-range point lights, and two near-range area lights.

such mitigations will inevitably limit these methods to real-world applications. Nimier-David et al. proposed a method [28] for joint illumination and material estimation. However, the method requires to trace the ray of multiple bounces. In contrast, our method applies a unified incident light field to represent different light sources in the scene, and is capable of estimating material under any lighting conditions without the need of multi-bounce raytracing.

3 Method

3.1 Neural Incident Light Field

One of the keys to invert the rendering equation is to model the incoming light L_i in a correct way. Ideally, L_i should take into account 1) *direct lights* from light sources in the scene, 2) *occlusions* that block the surface point from receiving direct lights, and 3) *indirect lights* that are reflected from other surface points. However, each of the three components is hard to model. Previous methods [6, 18, 47, 49] usually approximate direct lights as an environment map and hardly handle indirect lights as they require multi-bounce raytracing.

In contrast, we formulate incoming lights in the scene directly as a neural incident light field, where an MLP takes a point location \mathbf{x} and an incident direction $\boldsymbol{\omega}$ as inputs, and returns an incident light radiance L as output:

$$\mathbb{L} : \{\mathbf{x}, \boldsymbol{\omega}\} \rightarrow L. \quad (2)$$

Without losing generality, the proposed NeILF representation is capable of modelling the joint illumination effect of direct/indirect lights and occlusions of *any static scenes*. An illustration is shown in Fig. 1. Compared with the commonly used environment map, NeILF is able to handle the spatially-varying illumination effect, making it possible to estimate material for scenes with complex geometries and lightings.

3.2 Simplified Disney BRDF

In this section, we describe the BRDF representation used in the proposed framework. We apply a simplified Disney principled BRDF model, where the BRDF of a surface point \mathbf{x} is parameterized by a base color $\mathbf{b}(\mathbf{x}) \in [0, 1]^3$, a roughness $r(\mathbf{x}) \in [0, 1]$ and a metallic $m(\mathbf{x}) \in [0, 1]$, which is a subset of the full Disney model [8]. Similar to the neural incident lighting field, BRDF parameters are also stored using multi-layer perceptrons:

$$\mathbb{B} : \mathbf{x} \rightarrow \{\mathbf{b}, r, m\}, \quad (3)$$

where the MLP takes a 3D surface point \mathbf{x} as input, and returns the 5-channel BRDF parameters as output. Note that other representations, e.g. UV atlas or per-vertex BRDF parameters, can also be applied. Here we choose the neural representation because it has been proven to be effective for modelling continuous functions in 3D space [23, 47], and its derivative can be easily and analytically derived for our regularization computation (Sec. 3.3).

The rendering equation Given the BRDF parameterization, we now describe the concrete formulation of f in Equation 1. In the following equations, we omit notations of surface point \mathbf{x} and normal \mathbf{n} as the geometry of the scene is assumed to be given. The diffuse term can be calculated as $f_d = \frac{1-m}{\pi} \cdot \mathbf{b}$, and the specular term as:

$$f_s(\omega_o, \omega_i) = \frac{D(\mathbf{h}; r) \cdot F(\omega_o, \mathbf{h}; \mathbf{b}, m) \cdot G(\omega_i, \omega_o, \mathbf{h}; r)}{4 \cdot (\mathbf{n} \cdot \omega_i) \cdot (\mathbf{n} \cdot \omega_o)}, \quad (4)$$

where \mathbf{h} is the half vector between the incident direction ω_i and the viewing direction ω_o . The first term D is the normal distribution function of the microfacets in the surface. It is related to the roughness r and we use Spherical Gaussians to model this function as in previous methods [36, 47]. The second Fresnel term F models the portion of light that can be reflected from the surface, which is determined by the surface metallic m and the base color \mathbf{b} . The final geometry term G handles the shadow and occlusion of the microfacets, which is parameterized on the roughness r and is approximated using the GGX distribution [35]. Details of D , F and G are provided in the supplementary material.

3.3 Material-Lighting Ambiguity and Regularizations

While the Disney BRDF and incident light field are capable of representing materials and lightings of different scenes, jointly optimizing both would inevitably

lead to ambiguous solutions between them. One degenerate case could be that we can force a pure reflective BRDF to all surface points, and then only optimize the incident lights to adjust the input image. Theoretically, we can still find a perfect solution that fits the given BRDF and input images: for each 3D point, whenever there is a visible camera, we set its mirror symmetric incident light equal to the viewing out-going light, and set other incident lights equal to zero. It is also reported in previous work [31] that even human observers cannot distinguish the two confounded components from only image observations.

In the proposed framework, we can still manage to recover reasonable material and lighting results as MLPs can implicitly enforce a spatial smoothness constraint [48] on the two components. However, for robust material and lighting estimation, additional regularizations are desired. In this paper, we propose two regularizations for roughness r and metallic m :

Bilateral Smoothness We encourage r and m not to change rapidly in space, and the gradient of the input image \mathbf{I} can be used as a hint to guide the smoothing process. Thus, we define the bilateral smoothness cost of r and m as:

$$l_{smooth} = \frac{1}{|S_I|} \sum_{\mathbf{p} \in S_I} (\|\nabla_{\mathbf{x}} r(\mathbf{x}_{\mathbf{p}})\| + \|\nabla_{\mathbf{x}} m(\mathbf{x}_{\mathbf{p}})\|) e^{-\|\nabla_{\mathbf{p}} \mathbf{I}(\mathbf{p})\|}, \quad (5)$$

where S_I is the set of all sampled pixels and $\mathbf{x}_{\mathbf{p}}$ is the corresponding 3D point of the sampled pixel \mathbf{p} . The image gradient $\nabla_{\mathbf{p}} \mathbf{I}(\mathbf{p})$ can be pre-calculated from the input image, and the roughness gradient $\nabla_{\mathbf{x}} r(\mathbf{x}_{\mathbf{p}})$ and metallic gradient $\nabla_{\mathbf{x}} m(\mathbf{x}_{\mathbf{p}})$ can be derived analytically by back-propagating the neural network.

Lambertian Assumption We also assume that all surfaces tend to be Lambertian if no view-dependent lighting is observed, which leads to high roughness and low metallic, and we define the Lambertian cost as:

$$l_{lambertian} = \frac{1}{|S_I|} \sum_{\mathbf{p} \in S_I} (|r(\mathbf{x}_{\mathbf{p}}) - 1| + |m(\mathbf{p})|). \quad (6)$$

The proposed two regularizations will be minimized during network training. It is noteworthy that the two losses may not necessarily improve quantitative results as they are heuristically defined for robust material and lighting estimation. We show in a later ablation study that the bilateral smoothness will lead to visually much more pleasing results for real-world reconstructions.

3.4 Loss

Similar to other differentiable rendering frameworks, we compute the L1 loss between the rendered image and the input image:

$$l_{image} = \frac{1}{|S_I|} \sum_{\mathbf{p} \in S_I} \|\mathbf{I}(\mathbf{p}) - L_o(\mathbf{x}_{\mathbf{p}}, \boldsymbol{\omega}_o)\|_1. \quad (7)$$

The final loss of the proposed system is a weighted sum of the image loss and the two regularization losses: $l = l_{image} + w_s l_{smooth} + w_l l_{lambertian}$, where the two weights are empirically set to $w_s = 10^{-4}$ and $w_l = 10^{-3}$ in all our experiments.

4 Implementations

4.1 Sphere Sampling

To compute L_o using a finite number of incident lights, we need to discretize Equation 1 as: $L_o(\boldsymbol{\omega}_o, \mathbf{x}) = \sum_{i \in S_L} f(\boldsymbol{\omega}_o, \boldsymbol{\omega}_i, \mathbf{x}) L_i(\boldsymbol{\omega}_i, \mathbf{x}) (\boldsymbol{\omega}_i \cdot \mathbf{n}) \cdot A(\boldsymbol{\omega}_i)$, where S_L is the set of incident lights sampled for point \mathbf{x} and $A(\boldsymbol{\omega}_i)$ is the solid angle that corresponds to the incident light. In computer graphics, randomized Monte-Carlo Samplings are usually applied in ray-tracing, and the solid angle $A(\boldsymbol{\omega}_i)$ is approximated by the probability distribution $P(\boldsymbol{\omega}_i)$ of ray samples.

However, in differentiable rendering, it is critical to accurately compute the solid angle $A(\boldsymbol{\omega}_i)$ for each light sample as we need to correctly pass loss gradients to network parameters. We found that using random sampling and approximating $A(\boldsymbol{\omega}_i)$ as the probability distribution $P(\boldsymbol{\omega}_i)$ will lead to erroneous BRDF results. Thus, we apply a fixed Fibonacci sampling over the half sphere to get all samples. In this case, $A(\boldsymbol{\omega}_i) = \frac{2\pi}{|S_L|}$ and the rendering equation becomes:

$$L_o(\boldsymbol{\omega}_o, \mathbf{x}) = \frac{2\pi}{|S_L|} \sum_{i \in S_L} f(\boldsymbol{\omega}_o, \boldsymbol{\omega}_i, \mathbf{x}) L_i(\boldsymbol{\omega}_i, \mathbf{x}) (\boldsymbol{\omega}_i \cdot \mathbf{n}). \quad (8)$$

4.2 Learned HDR-LDR Mapping

For real-world datasets with low dynamic range (LDR) images, we need to convert the high dynamic range (HDR) output from our renderer to LDR before computing the image loss. As such transformation is unavailable in previous MVS datasets, we apply a learned HDR-LDR mapping to mimic the conversion in our network. Note that linear transformations, including exposure and white balance, can be embedded into the incident light. Thus, we only explicitly model the gamma correction with a learnable parameter:

$$L_o^{LDR} = (L_o^{HDR})^\gamma. \quad (9)$$

4.3 Training Details

We use an 8-layer Siren [33] with feature size of 512 and a skip connection in the middle to represent the BRDF MLP. Also, the positional encoding [22] is applied to further strengthen the network. The NeILF MLP shares the same implementation as BRDF, except that 1) the feature size is downsized to 128 to reduce the VRAM usage and 2) the last layer activation function is changed from *tanh* to *exp* in order to guarantee non-negative and unbounded light intensities.

In the experiment, we use $|S_L| = 128$ incident lights to compute the output radiance during training, and use $|S_L| = 256$ incident lights to evaluate the rendered image during testing. For each training iteration, we randomly sample 16000 pixels from all images, and the network is optimized for a total of 15000 iterations. The Adam optimizer [17] with an initial learning rate of 10^{-3} is applied in our network, and the learning rate is scaled down by $\sqrt{0.1}$ at 5000 and 10000 iterations. The training process takes around 1.5 hours to finish on a Tesla V100 GPU and the VRAM consumption is around 30 GB.

5 Experiments

5.1 Baseline Methods

We compare our method with the following baselines:

PhySG* Firstly, we consider the recent PhySG [47] for material estimation. The original PhySG jointly optimizes the non-spatially varying BRDF, the environment map, and the geometry of the object. To fairly compare with the method, we fix the given geometry and optimize only the uniform BRDF and the environment map.

SG-Env This baseline is another variant of PhySG [47]. Compared with PhySG*, SG-Env applies a SV-BRDF model and a slightly different rendering formulation (the same f_d , D and G as ours). We use this baseline to directly compare NeILF with the SG environment map representation.

NeRFactor* This baseline is a variant of NeRFactor [49] with the ground truth geometry input. In the setting, we follow NeRFactor to represent the scene lighting as a 2D environment map of resolution 32×16 . The ground truth geometry are fixed during the training and we use the online raytracing to compute the ground truth visibilities of each ray.

Ne-Env Lastly, we compare our method with the neural environment map representation. This baseline shares the same implementation of the proposed NeILF, except that the positional input in the incident light field is omitted such that the incident light is only related to the incoming direction: $\{\omega\} \rightarrow \mathbf{L}$.

5.2 Benchmark on Synthetic Scenes

To quantitatively evaluate our method under different lighting conditions, we generate a set of synthetic data and compare our method with the above baselines.

Data Preparation The synthetic dataset contains three objects and their combinations: a single rough metallic sphere, a single rough metallic cube, and a helmet with spatially variant materials. The objects are placed on a plane to model the real-world object capture. We also create six lighting conditions to lit the objects, including three environment maps and three mixed lightings:

- *Env-city*: an environment map of a city;
- *Env-studio*: an environment map of a studio;
- *Env-castel*: an environment map of a castel;
- *Mix-city*: Env-city plus two point lights and two area lights;
- *Mix-studio*: Env-studio plus two point lights and two area lights;
- *Mix-castel*: Env-castel plus two point lights and two area lights.

Each scene contains 96 images, where images $\{9, 18, 30, 41, 50, 62, 73, 82, 94\}$ are used for evaluation and the remaining 87 images are used for training. The image

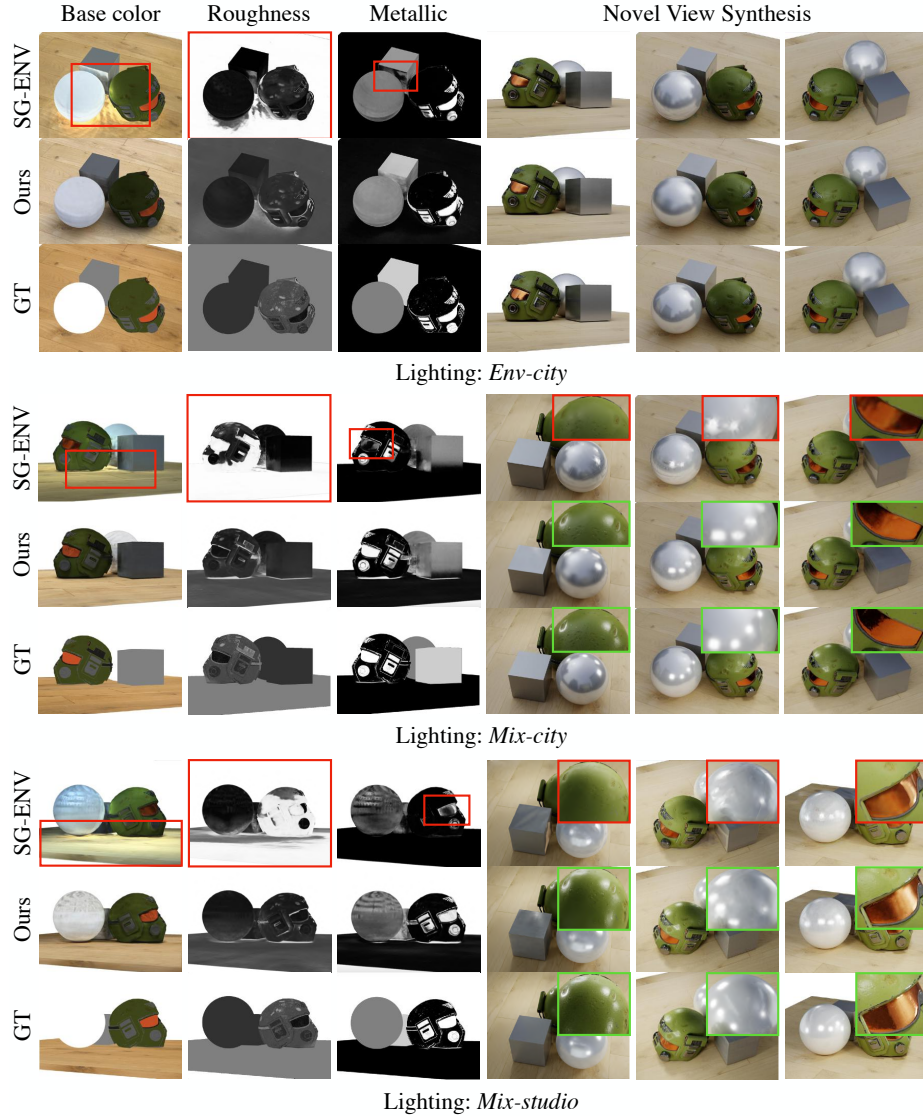


Fig. 2: Comparative results on BRDF estimation and novel view synthesis on the synthetic dataset. From left to right are images of *base color*, *roughness*, *metallic* and synthesized testing views. Our method is able to generate high-quality BRDF and novel view synthesis results under different lighting conditions. In contrast, the environment map based SG-ENV [47] produces noisy BRDF outputs especially in occluded regions. And also, highlights are wrongly recovered in novel view renderings if mixed lightings occur.

Table 1: Quantitative results on Synthetic scenes. We compare the proposed NeILF with four baseline methods described in Sec. 5.1 using PSNR scores. Our method generates consistently the best novel view rendering for all scenes. Also, our method produces significantly better BRDF results than other methods if multiple objects and mixed lightings are given.

Scene Geometry		Single-helmet			Combined-objects					
Scene Lighting		Env-city	Env-studio	Env-castel	Env-city	Env-studio	Env-castel	Mix-city	Mix-studio	Mix-castel
Base Color	PhySG*[47]	13.43	14.87	13.95	15.01	16.96	16.13	14.16	12.43	14.29
	SG-ENV[47]	20.61	18.45	15.99	22.38	20.74	22.21	16.92	13.16	16.69
	NeRFactor*[49]	13.02	11.73	11.98	8.09	13.60	7.37	8.64	11.63	8.03
	Ne-ENV	13.43	12.65	12.45	11.68	11.98	7.66	11.87	10.90	9.54
	Ours	16.36	16.36	18.28	15.59	15.48	12.95	17.39	16.88	17.37
Metallic	PhySG*[47]	7.57	7.48	7.83	8.72	7.97	8.35	8.67	8.95	8.76
	SG-ENV[47]	21.19	21.31	21.79	17.01	16.40	16.39	15.44	14.25	14.49
	NeRFactor*[49]	12.88	17.39	16.78	12.38	13.43	11.98	11.77	11.88	11.95
	Ne-ENV	9.31	17.40	6.15	15.86	16.10	11.42	15.43	15.35	15.49
	Ours	17.79	18.52	16.82	18.22	19.11	10.29	18.42	18.43	17.34
Roughness	PhySG*[47]	6.91	11.88	6.75	6.62	11.29	6.22	6.27	6.83	6.14
	SG-ENV[47]	14.77	15.87	9.64	9.61	17.64	9.74	8.77	12.58	9.14
	NeRFactor*[49]	12.84	11.59	12.89	9.21	17.95	9.02	9.12	12.62	8.54
	Ne-ENV	11.95	14.84	9.26	15.56	14.48	12.94	16.20	14.43	14.14
	Ours	16.13	16.19	17.16	17.48	18.30	13.40	17.05	16.27	16.44
Rendering	PhySG*[47]	24.59	24.77	26.52	24.82	25.65	27.24	24.38	24.04	25.81
	SG-ENV[47]	29.73	29.86	32.13	31.01	29.46	32.34	27.20	25.88	27.70
	NeRFactor*[49]	29.30	28.78	30.58	29.31	30.10	30.66	27.50	27.00	28.11
	Ne-ENV	28.60	29.56	29.76	30.75	29.07	32.05	28.07	26.01	28.33
	Ours	31.57	30.84	34.43	33.77	31.07	35.28	31.11	28.59	32.11

resolution is set to 1600×1200 . We use Blender [10] to render the HDR images by ray tracing. Position maps and normal maps at all viewpoints are rendered to serve as the geometry input for the system. Meanwhile, per-view ground truth base color, metallic and roughness maps are provided for quantitative evaluation.

Results We use the PSNR score as our evaluation metric. Quantitative comparisons on 1) base color, 2) metallic, 3) roughness and 4) novel view synthesis are shown in Table 1. Our method consistently outperforms other methods with a large margin in terms of the novel view rendering quality. For material estimation, we found that if single objects and environment map light sources are given, SG-Env [47] is able to generate comparable results with ours. However, in the case of multiple objects or mixed light sources, the estimation quality of SG-Env will drop significantly. This is because the environment map representation cannot model mixed light sources of point and area lights. Also, indirect lights and occlusions within multiple objects are not handled by SG-Env. In contrast, our NeILF representation can robustly deal with mixed lightings and complex scene geometries. Qualitative results are shown in Fig. 2.

5.3 Test on Real-world Scenes

We then test our method on two real-world datasets, namely DTU [13] and BlendedMVS [40] datasets. DTU dataset is captured in a lab setting with a fixed lighting and camera trajectory, while BlendedMVS contains a variety of indoor and outdoor scenes captured by different users. As the two datasets provide only

Table 2: Quantitative results on DTU [13] and BlendedMVS [40] Datasets. The table shows PSNR scores of novel view renderings of test images. Our method consistently outperforms the other methods in terms of the rendering quality.

	DTU [13]					BlendedMVS [40]					
	scan-1	scan-11	scan-37	scan-75	scan-97	bull	cam	dog	gold	statue	stone
PhySG* [47]	20.40	20.78	20.30	16.03	19.86	21.64	18.11	20.70	19.06	19.74	21.22
SG-ENV [47]	22.18	21.56	21.71	18.06	21.09	22.51	20.14	22.06	19.44	20.79	22.31
NeRFactor* [49]	23.68	23.43	23.86	19.89	22.10	21.87	20.03	21.54	19.47	20.51	22.98
Ne-ENV	23.77	23.79	22.87	19.52	21.51	22.17	20.17	21.73	19.66	20.55	23.08
NeILF (Ours)	24.79	24.33	24.44	23.46	23.96	24.93	22.10	22.36	20.80	21.51	24.22

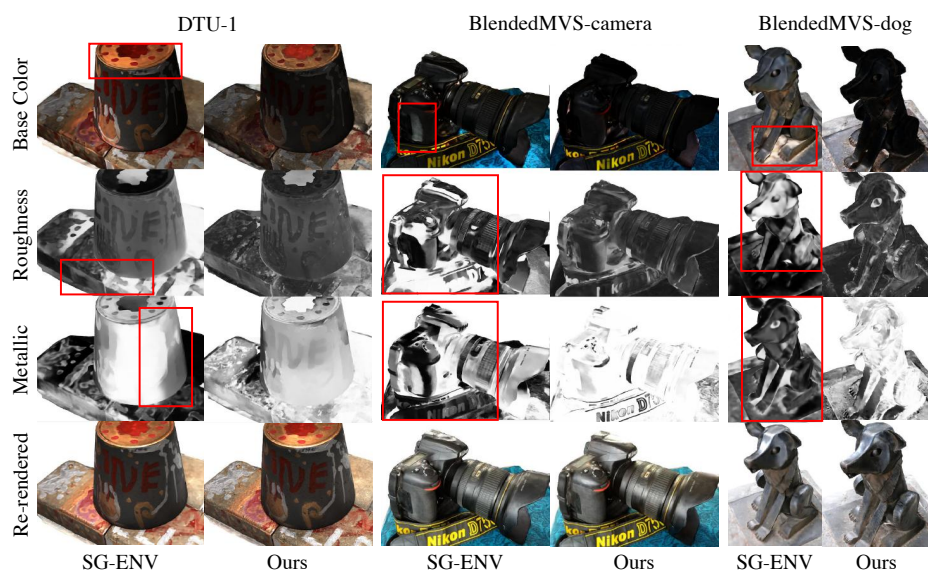


Fig. 3: Qualitative results on the real-world DTU and BlendedMVS datasets. Our method successfully removes high lights in the base color and produces visually plausible results of roughness and metallic.

LDR images, the learned HDR-LDR mapping described in Sec. 4.2 is applied for the loss computation. For each scene, we select 5 images {2, 12, 17, 30, 34} for testing and the rest of the images for training.

It is noteworthy that unlike in the synthetic dataset, here we use multi-view stereo methods to generate the geometry input rather than directly using the ground truth. For DTU dataset, we use Vis-MVSNet [45] to generate the dense 3D point cloud and SPSR [16] to recover the mesh surface. For BlendedMVS dataset, we use original images and the provided reference mesh geometry as our inputs. By doing so, our method can be viewed as an extension to nowadays 3D reconstruction pipelines.

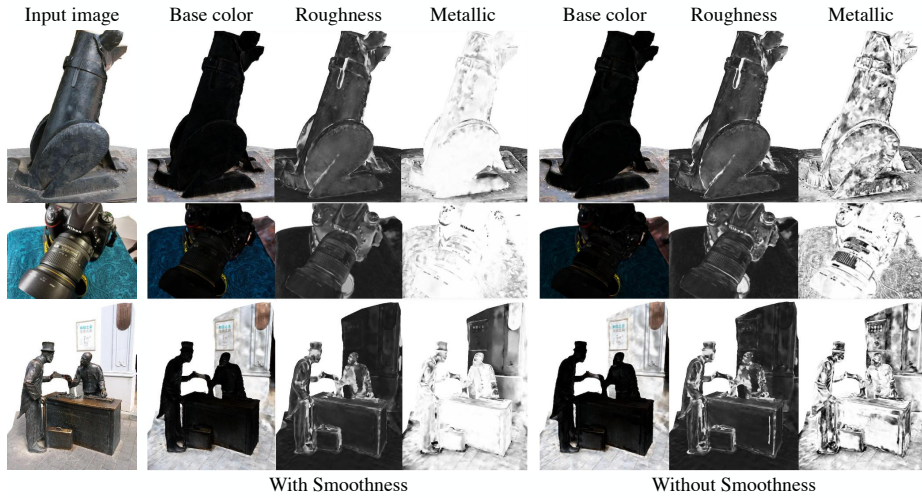


Fig. 4: Qualitative comparisons on with and without the proposed bilateral smoothness regularization.

Quantitative results are shown in Table 2 and qualitative results compared with SG-ENV are shown in Fig. 3. The proposed method produces both the best rendering PSNR and the most visually pleasing BRDF in all selected scenes. We believe the proposed method can be integrated into traditional 3D reconstruction pipelines for relightable mesh model reconstruction.

5.4 Ablation Study

In this section, we analyze several design choices of the proposed framework. The ablation studies are conducted on the synthetic dataset, and we report the average scores over all scenes to compare different settings.

Ray Sample Number We first study the influence of the ray sample number for material estimation quality. The ray sample number is decreased from $S_L = 128$ to $S_L = 64$ and $S_L = 32$. As shown in Table 3, higher sampling number will lead to better reconstruction results. In our default setting, we choose $S_L = 128$ to better balance the quality and the VRAM/runtime consumption.

Random Sample Next, we compare the fixed Fibonacci sample described in Sec. 4.1 with the random uniform sample commonly used in computer graphics. It is shown in Table 3 that the random sample would lead to worse results, showing that it is crucial to precisely discretize the rendering equation in the differentiable rendering.

Regularizations We also study the effectiveness of the two regularizations proposed in Sec. 3.3. We find that the bilateral smoothness is essential for ma-

terial estimation of real-world scenes, where the roughness and metallic will be significantly improved if the smoothness is applied (Fig. 4).

Learned Gamma Lastly, we compare the proposed learned gamma in Sec. 4.2 with the fixed $\gamma = \frac{1}{2.2}$. In BlendedMVS dataset, the learned gamma produces a mean PSNR score of 22.65, which slightly outperforms the score of the fixed gamma (22.51), showing that the learned adjustment might be a better choice.

On the other hand, we also find that the two heuristics have limited influence to quantitative results of the synthetic dataset. We believe this is because the vanilla NeILF already produces high-quality estimations for synthetic scenes. In our default setting, we keep the two regularizations for all scenes but we encourage users to selectively apply the two terms depending on different characteristics of input scenes.

	Base.	Meta.	Roug.	Rend.
S = 128	16.30	17.22	16.49	32.09
S = 64	15.20	16.47	18.88	31.40
S = 32	13.40	16.27	19.49	30.57
Rand. Samp.	12.45	15.11	18.10	29.73

Table 3: Ablation studies. Average scores among all synthetic scenes are reported.

6 Discussions

6.1 Comparison with NeRF Optimization

In this section, we compare the proposed NeILF with the neural radiance field [22]. We show that the two frameworks share similarities in many aspects, and thus provide readers an intuitive explanation of why NeILF can successfully disentangle the complex material and lighting in the joint optimization.

Lighting Representations NeRF [22] represents the scene appearance as the neural radiance field. While the radiance is physically different with the incident light, their complexities are completely the same: both NeILF and NeRF take a 3D position \mathbf{x} and a direction ω as inputs, and returns a RGB value as output.

Spatially-varying Properties NeILF aims to recover surface materials as BRDF properties, while NeRF jointly optimizes the scene geometry as a density field. Both our BRDF and NeRF’s density MLPs take only a 3D position \mathbf{x} as input, and return different spatial properties as outputs. Implementation-wise, the only difference is that our BRDF is consist of a 5D parameter vector, while the density value is a 1D scalar. Nevertheless, the two spatially-varying properties are very similar and their complexities are comparable.

Rendering Formulations Our method applies the surface rendering to compute the reflected light from a surface point, while NeRF adopts the volume rendering to get the accumulated color along a viewing ray. On the one hand, NeILF requires the incident light integration over the hemisphere; on the other

hand, NeRF requires alpha composition along the ray. To render a pixel, NeILF needs to sample the BRDF MLP once and the incident light MLP for multiple times, while NeRF does the same on the density and radiance MLPs.

Reconstruction Ambiguities The geometry-appearance ambiguity is addressed in NeRF++ [48]. Similarly, we analyze the material-lighting ambiguity in the NeILF optimization. It has been reported that with proper geometry regularizations [29, 37, 41], NeRF is able to produce high-quality geometry outputs. In contrast, we also show that the proposed bilateral smoothness can significantly improve the roughness and metallic quality for real-world scenes (Fig. 4).

6.2 Limitations and Future Works

While the proposed method has already shown promising results for material estimation, the current pipeline still contains several limitations that could be further addressed in future works.

Geometry Optimization The major limitation of our method is that the geometry is required to be given in advance. Although we have shown that meshes from 3D reconstruction pipelines are qualified for real-world DTU and BlendedMVS scenes, it is still desired that we can jointly optimize the scene geometry during training. Possible directions include displacement/normal map estimation and the differentiable surface refinement [42, 44, 46, 47].

High-intensity Singularities in Light Sources Rippled noises would sometimes occur in our estimation if direct lights contain high-intensity singularities. One example is the metallic result of *Mix-studio* in Fig. 2. It is worth investigating better training losses/strategies to mitigate this problem in the future.

Running Speed Similar to NeRF, our method requires frequent MLP samplings and the training process is time consuming. Our current implementation takes around 1.5 hours to estimate the BRDF of a given scene (details in Sec. 4.3). We hope that in the future, explicit Octree [43], spherical harmonics [1, 38] or neural hashing [24] can be applied to speed up the NeILF optimization.

7 Conclusions

We have presented a differentiable rendering framework for material and lighting estimation. Compared with the environment map approximation, the proposed neural incident light field is capable of modelling the lighting condition of any static scenes, making it possible to estimate qualified material properties even for scenes with complex lightings and geometries. The proposed method strictly follows the physically-based rendering equation, and jointly optimizes material and lighting through the differentiable rendering process. We have intensively evaluated our method on our in-house synthetic dataset, the DTU MVS dataset, and the real-world BlendedMVS scenes. Our method is able to outperform previous methods by a significant margin in terms of the novel view rendering quality, setting a new state-of-the-art for image-based material and lighting estimation.

Bibliography

- [1] Alex Yu and Sara Fridovich-Keil, Tancik, M., Chen, Q., Recht, B., Kanazawa, A.: Plenoxels: Radiance fields without neural networks. arXiv:2112.05131 (2021)
- [2] Azinovic, D., Li, T.M., Kaplanyan, A., Niessner, M.: Inverse path tracing for joint material and lighting estimation. In: CVPR (June 2019)
- [3] Azinovic, D., Li, T.M., Kaplanyan, A., Nießner, M.: Inverse path tracing for joint material and lighting estimation. In: CVPR (2019)
- [4] Bi, S., Xu, Z., Sunkavalli, K., Hašan, M., Hold-Geoffroy, Y., Kriegman, D., Ramamoorthi, R.: Deep reflectance volumes: Relightable reconstructions from multi-view photometric images. In: ECCV (2020)
- [5] Bi, S., Xu, Z., Sunkavalli, K., Kriegman, D., Ramamoorthi, R.: Deep 3d capture: Geometry and reflectance from sparse multi-view images. In: CVPR (2020)
- [6] Boss, M., Braun, R., Jampani, V., Barron, J.T., Liu, C., Lensch, H.: Nerd: Neural reflectance decomposition from image collections. In: ICCV (2021)
- [7] Boss, M., Jampani, V., Braun, R., Liu, C., Barron, J., Lensch, H.P.: Neural-pil: Neural pre-integrated lighting for reflectance decomposition. In: NeurIPS (2021)
- [8] Burley, B., Studios, W.D.A.: Physically-based shading at disney. In: ACM SIGGRAPH (2012)
- [9] Chen, W., Litalien, J., Gao, J., Wang, Z., Tsang, C.F., Khalis, S., Litany, O., Fidler, S.: DIB-R++: Learning to predict lighting and material with a hybrid differentiable renderer. In: NeurIPS (2021)
- [10] Community, B.O.: Blender - a 3D modelling and rendering package. Blender Foundation, Stichting Blender Foundation, Amsterdam (2018), <http://www.blender.org>
- [11] Dong, Y., Chen, G., Peers, P., Zhang, J., Tong, X.: Appearance-from-motion: Recovering spatially varying surface reflectance under unknown lighting. TOG (2014)
- [12] Guo, K., Lincoln, P., Davidson, P., Busch, J., Yu, X., Whalen, M., Harvey, G., Orts-Escolano, S., Pandey, R., Dourgarian, J., et al.: The relightables: Volumetric performance capture of humans with realistic relighting. TOG (2019)
- [13] Jensen, R., Dahl, A., Vogiatzis, G., Tola, E., Aanæs, H.: Large scale multi-view stereopsis evaluation. In: CVPR (2014)
- [14] Kajiya, J.T.: The rendering equation. In: Proceedings of the 13th annual conference on Computer graphics and interactive techniques (1986)
- [15] Kato, H., Beker, D., Morariu, M., Ando, T., Matsuoka, T., Kehl, W., Gaidon, A.: Differentiable rendering: A survey. arXiv preprint arXiv:2006.12057 (2020)
- [16] Kazhdan, M., Hoppe, H.: Screened poisson surface reconstruction. TOG (2013)

- [17] Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: ICLR (2015)
- [18] Li, Z., Xu, Z., Ramamoorthi, R., Sunkavalli, K., Chandraker, M.: Learning to reconstruct shape and spatially-varying reflectance from a single image. TOG (2018)
- [19] Liu, L., Gu, J., Lin, K.Z., Chua, T.S., Theobalt, C.: Neural sparse voxel fields. NeurIPS (2020)
- [20] Liu, S., Zhang, Y., Peng, S., Shi, B., Pollefeys, M., Cui, Z.: Dist: Rendering deep implicit signed distance function with differentiable sphere tracing. In: CVPR (2020)
- [21] Liu, S., Saito, S., Chen, W., Li, H.: Learning to infer implicit surfaces without 3d supervision. arXiv preprint arXiv:1911.00767 (2019)
- [22] Martin-Brualla, R., Radwan, N., Sajjadi, M.S., Barron, J.T., Dosovitskiy, A., Duckworth, D.: Nerf in the wild: Neural radiance fields for unconstrained photo collections. In: CVPR (2021)
- [23] Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: ECCV (2020)
- [24] Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. arXiv:2201.05989 (2022)
- [25] Munkberg, J., Hasselgren, J., Shen, T., Gao, J., Chen, W., Evans, A., Mueller, T., Fidler, S.: Extracting Triangular 3D Models, Materials, and Lighting From Images. arXiv:2111.12503 (2021)
- [26] Nam, G., Lee, J.H., Gutierrez, D., Kim, M.H.: Practical svbrdf acquisition of 3d objects with unstructured flash photography. TOG (2018)
- [27] Niemeyer, M., Mescheder, L., Oechsle, M., Geiger, A.: Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In: CVPR (2020)
- [28] Nimier-David, M., Dong, Z., Jakob, W., Kaplanyan, A.: Material and lighting reconstruction for complex indoor scenes with texture-space differentiable rendering (2021)
- [29] Oechsle, M., Peng, S., Geiger, A.: Unisurf: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. In: ICCV (2021)
- [30] Park, J.J., Holynski, A., Seitz, S.M.: Seeing the world in a bag of chips. In: CVPR (2020)
- [31] Pont, S.C., Te Pas, S.F.: Material—illumination ambiguities and the perception of solid objects. Perception (2006)
- [32] Schmitt, C., Donne, S., Riegler, G., Koltun, V., Geiger, A.: On joint estimation of pose, geometry and svbrdf from a handheld scanner. In: CVPR (2020)
- [33] Sitzmann, V., Martel, J., Bergman, A., Lindell, D., Wetzstein, G.: Implicit neural representations with periodic activation functions. NeurIPS (2020)
- [34] Srinivasan, P.P., Deng, B., Zhang, X., Tancik, M., Mildenhall, B., Barron, J.T.: Nerv: Neural reflectance and visibility fields for relighting and view synthesis. In: CVPR (2021)

- [35] Walter, B., Marschner, S., Li, H., Torrance, K.: Microfacet models for refraction through rough surfaces. In: EGSR (2007)
- [36] Wang, J., Ren, P., Gong, M., Snyder, J., Guo, B.: All-frequency rendering of dynamic, spatially-varying reflectance. In: ACM SIGGRAPH Asia (2009)
- [37] Wang, P., Liu, L., Liu, Y., Theobalt, C., Komura, T., Wang, W.: Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction (2021)
- [38] Wizadwongsa, S., Phongthawee, P., Yenphraphai, J., Suwajanakorn, S.: Nex: Real-time view synthesis with neural basis expansion. In: CVPR (2021)
- [39] Xia, R., Dong, Y., Peers, P., Tong, X.: Recovering shape and spatially-varying surface reflectance under unknown illumination. TOG (2016)
- [40] Yao, Y., Luo, Z., Li, S., Zhang, J., Ren, Y., Zhou, L., Fang, T., Quan, L.: Blendedmvs: A large-scale dataset for generalized multi-view stereo networks. In: CVPR (2020)
- [41] Yariv, L., Gu, J., Kasten, Y., Lipman, Y.: Volume rendering of neural implicit surfaces. arXiv preprint arXiv:2106.12052 (2021)
- [42] Yariv, L., Kasten, Y., Moran, D., Galun, M., Atzmon, M., Ronen, B., Lipman, Y.: Multiview neural surface reconstruction by disentangling geometry and appearance. NeurIPS (2020)
- [43] Yu, A., Li, R., Tancik, M., Li, H., Ng, R., Kanazawa, A.: PlenOctrees for real-time rendering of neural radiance fields. In: ICCV (2021)
- [44] Zhang, J., Yao, Y., Li, S., Fang, T., McKinnon, D., Tsin, Y., Quan, L.: Critical regularizations for neural surface reconstruction in the wild. In: CVPR (2022)
- [45] Zhang, J., Yao, Y., Li, S., Luo, Z., Fang, T.: Visibility-aware multi-view stereo network. BMVC (2020)
- [46] Zhang, J., Yao, Y., Quan, L.: Learning signed distance field for multi-view surface reconstruction. In: ICCV (2021)
- [47] Zhang, K., Luan, F., Wang, Q., Bala, K., Snavely, N.: Physg: Inverse rendering with spherical gaussians for physics-based material editing and re-lighting. In: CVPR (2021)
- [48] Zhang, K., Riegler, G., Snavely, N., Koltun, V.: Nerf++: Analyzing and improving neural radiance fields. arXiv preprint arXiv:2010.07492 (2020)
- [49] Zhang, X., Srinivasan, P.P., Deng, B., Debevec, P., Freeman, W.T., Barron, J.T.: NeRFactor: Neural Factorization of Shape and Reflectance Under an Unknown Illumination. TOG (2021)