

NeILF++: Inter-Reflectable Light Fields for Geometry and Material Estimation

Jingyang Zhang¹ Yao Yao^{2✉*} Shiwei Li¹ Jingbo Liu¹ Tian Fang¹
David McKinnon¹ Yanghai Tsin¹ Long Quan^{3*}

¹Apple ²Nanjing University ³The Hong Kong University of Science and Technology

¹{jingyang.zhang, shiwei, jingbo, fangtian, dmckinnon, ytsin}@apple.com

²yaoyao@nju.edu.cn ³quan@cse.ust.hk

Abstract

We present a novel differentiable rendering framework for joint geometry, material, and lighting estimation from multi-view images. In contrast to previous methods which assume a simplified environment map or co-located flash-lights, in this work, we formulate the lighting of a static scene as one neural incident light field (NeILF) and one outgoing neural radiance field (NeRF). The key insight of the proposed method is the union of the incident and outgoing light fields through physically-based rendering and inter-reflections between surfaces, making it possible to disentangle the scene geometry, material, and lighting from image observations in a physically-based manner. The proposed incident light and inter-reflection framework can be easily applied to other NeRF systems. We show that our method can not only decompose the outgoing radiance into incident lights and surface materials, but also serve as a surface refinement module that further improves the reconstruction detail of the neural surface. We demonstrate on several datasets that the proposed method is able to achieve state-of-the-art results in terms of geometry reconstruction quality, material estimation accuracy, and the fidelity of novel view rendering.

1. Introduction

Reconstructing 3D scene information from multi-view images is a persistent challenge in the field of computer vision and computer graphics. The Neural Radiance Field (NeRF) approach [13], which uses differentiable volume rendering to jointly optimize scene geometry and appearance, offers a novel solution and has demonstrated great results in novel view synthesis and neural surface reconstruction. While NeRF successfully models the outgoing radiance of a scene, it fails to disentangle the incident lighting

*Research done while at Apple.

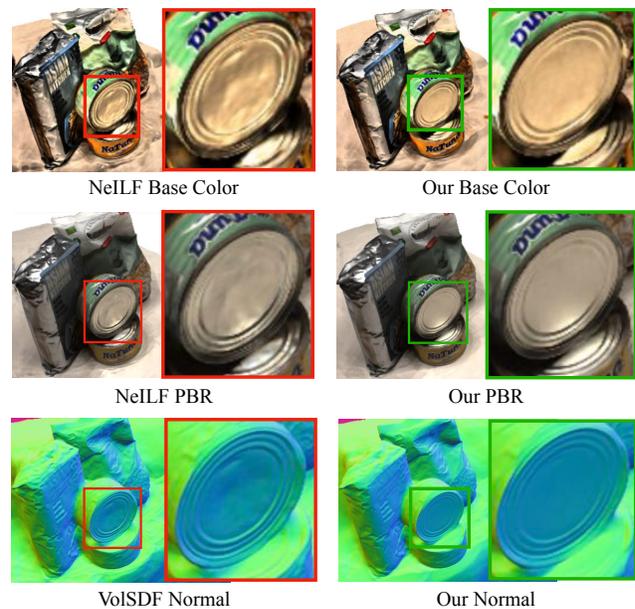


Figure 1. Compared with the baseline that simply runs NeILF [28] following VoISDF [29], the proposed approach is able to jointly optimize the scene geometry and material. The neural surface, material, and final rendering quality are improved.

and surface properties from the radiance, preventing its use in downstream applications such as object relighting and material property editing.

Recent works have acknowledged the problem and proposed to further decompose the outgoing radiance into the material and environmental lighting. Some of them adopt simplified lighting models to reduce the computational complexity, including co-located flash light [16, 3, 2, 20] and environment map [34, 36, 4, 15, 5]. These models are insufficient to model all kinds of lighting configurations, such as non-distant light sources and inter-reflections between surfaces. Recently, the Neural Incident Light Field (NeILF) [28] instead models arbitrary static lighting con-

ditions by encoding spatially-varying incident lighting in a neural network. However, it only provides the modeling capability of, but does not introduce an explicit constraint on the inter-reflection issue. Moreover, NeILF requires an object mesh as input, whose quality has a strong influence on the estimated material.

In this work, we demonstrate that suitable lighting representation is the key to joint geometry and material estimation from multi-view images. Without loss of generality, the lighting of a typical scene can be represented by the incident light to the object surface and the outgoing radiance from the surface. Based on the observation, we propose to model the light fields, geometry, and materials of the scene as four separated fields, namely 1) one outgoing radiance field [13], 2) one incident light field [28], 3) one signed distance field as the scene geometry, and 4) one field of bidirectional reflectance distribution function (BRDF) parameters as the surface material. The key insight of the proposed method is that the incident light and the outgoing radiance can be naturally unified through physically-based rendering and inter-reflections between surfaces. We demonstrate that with the union of the outgoing radiance and the incident light field, the proposed method not only can generate high-quality BRDF estimation for object relighting, but also improves the reconstruction detail of the surface geometry. The proposed approach can be easily applied to different NeRF systems for material decomposition and surface detail refinement.

To better evaluate the quality of estimated material from physically-based rendering, we construct a real-world linear high-dynamic range (HDR) dataset called NeILF-HDR. With this, we can avoid adding an HDR-LDR conversion module and directly train in HDR color space. We have extensively studied our approach on the DTU dataset [8], the NeILF synthetic dataset [28], and the NeILF-HDR dataset. We show that the proposed method is able to achieve state-of-the-art results in geometry reconstruction quality, material estimation accuracy, and novel view rendering quality. To summarize, major contributions of this paper include:

- A general light field representation by marrying one incident light field and one outgoing radiance field via PBR and inter-reflections between surfaces.
- An optimization scheme for joint geometry, material, and lighting estimation. The proposed method can be easily applied to the prevalent NeRF family for material decomposition and neural surface refinement.
- A real-world linear HDR dataset for the evaluation of material estimation and other neural rendering tasks.

2. Related Works

The Rendering Equation. With the given BRDF properties of a surface point \mathbf{x} , the outgoing light modeled by physically-based rendering L_o^P from \mathbf{x} can be computed via the incident light to the surface [9]:

$$L_o^P(\omega_o, \mathbf{x}) = \int_{\Omega} f(\omega_o, \omega_i, \mathbf{x}) L_i(\omega_i, \mathbf{x})(\omega_i \cdot \mathbf{n}) d\omega_i, \quad (1)$$

where ω_o is the viewing direction of the outgoing light, \mathbf{n} is the surface normal, L_i is the incident light from direction ω_i , and f is the BRDF properties of the surface point.

The goal of joint geometry, material, and lighting estimation is to recover the surface location \mathbf{x} , the scene lighting L_i and the BRDF function f in the above equation. Below we give a brief review of the physically-based optimization approaches with the multi-view image inputs.

Lighting Modeling for Material Estimation. Due to the high-dimensional nature of lighting and material, it is rather difficult to jointly recover both from multi-view image observations. Previous methods usually apply controlled lighting to simplify the problem. For example, point/collocated flash lights [16, 3, 2, 20, 10], turn-table models [6, 26] and the environment map assumption [34, 36, 4] are applied to reduce the complexity of the problem. Also, some works apply additional sensors [7, 1, 19] to facilitate the optimization process. However, such mitigations will inevitably limit these methods to real-world applications.

Other works instead adopt more complex lighting models for material and lighting estimation. NeRV [22] proposes to use reflectance and visibility fields to model the indirect light. However, it requires the environment map to be given in advance and its pipeline is computationally expensive. ShadowNeuS [11] and NeRFactor [36] exploit light visibility for material and incoming light decomposition, but it fails to model the inter-reflections between surfaces. A joint illumination and material estimation approach is proposed by Nimier-David et al. [18], but requires the given geometry and costly multi-bounce raytracing. Recently, NeILF [28] is proposed to model arbitrary static lighting by using a spatially-varying incident light field. However, NeILF also requires the geometry to be given and its incident light setting does not explicitly take inter-reflections into account. In this work, we propose to model both incident and outgoing lights of the surface, where the two light fields are further unified through PBR and inter-reflections.

Surface Optimization by Differentiable Rendering. Recent neural rendering methods have shown promising results in scene geometry recovery. The scene geometry is usually represented by an implicit function, such as a density field [13, 35, 36, 23, 4], an occupancy field [17], or a signed distance field [30, 29, 25, 33, 31]. During the differentiable rendering optimization, both the scene geometry



Figure 2. Illustration on the proposed light field representation. The incident light and the outgoing light are constrained by physically-based rendering (as in NeILF) and inter-reflection between surfaces. The inter-reflection constraint ensures the consistency between the two lighting representations during the joint optimization of the four fields. The losses L_{phys} , L_{vol} and L_{ref} are for the physically rendered image, the volume rendered image, and the inter-reflection constraint respectively, which will be elaborated in Sec. 4.

and appearance will be recovered by minimizing the difference between rendered images and input images.

Recent neural surface reconstruction methods can be categorized into 1) differentiable surface rendering and 2) differentiable volume rendering approaches. To render a pixel, differentiable surface rendering will first find the surface intersection using root-finding algorithms and then query the color of the intersection. To make the process differentiable, the first-order approximation of the surface intersection is usually applied [17, 30]. The method is fast, however, can only refine the geometry near the surface, and often requires additional geometric priors (e.g., silhouette [30, 34], depth [33] or point cloud [31]) for robust estimation. The second type of method uses volume rendering to accumulate all the radiance along the ray weighted by the density and visibility of matter [13, 25, 29]. The method is good at recovering the topology and has gained much popularity recently due to its simplicity. In this work, we show that the surface detail from a volume rendering system can be further refined by combining the differentiable physically-based rendering system (i.e., NeILF).

3. Inter-reflectable Light Fields

3.1. Recap on Incident Light and BRDF Modeling

Neural Incident Light Field. As proposed by NeILF [28], the incoming lights in the scene can be formulated as a neural incident light field, which is recorded by a multi-layer perceptron (MLP). The MLP takes a point location \mathbf{x} and a direction ω as inputs, and returns an incident light L :

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

Compared with environment maps, the spatially-varying illumination given by incident light fields is capable of modeling direct/indirect light and occlusions of any static scenes, and thus facilitates the material estimation in scenes with such complex geometry and lighting condition.

BRDF parameterization. We use a simplified Disney principled BRDF parameterization as in [28]. For a surface point \mathbf{x} , the BRDF field $\mathbb{B} : \mathbf{x} \rightarrow \{\mathbf{b}, r, m\}$ stores a base color $\mathbf{b} \in [0, 1]^3$, a roughness $r \in [0, 1]$ and a metallic $m \in [0, 1]$. The BRDF f in Eq. 1 is the summation of a diffuse term f_d and a specular term f_s . The diffuse term can be calculated as $f_d = \frac{1-m}{\pi} \cdot \mathbf{b}$, and the specular term can be computed as:

$$f_s(\omega_o, \omega_i) = \frac{D(\mathbf{h}; r) \cdot F(\omega_o, \mathbf{h}; \mathbf{b}, m) \cdot G(\omega_i, \omega_o, \mathbf{h}; r)}{4 \cdot (\mathbf{n} \cdot \omega_i) \cdot (\mathbf{n} \cdot \omega_o)}, \quad (3)$$

where \mathbf{h} is the half vector between the incident direction ω_i and the viewing direction ω_o . D , F , and G refer to the normal distribution function, the Fresnel term, and the geometry term respectively. We adopt similar implementation of D , F , and G as in previous works [24, 34, 28] and details are provided in the supplementary material.

3.2. Unifying Incident and Outgoing Lights with Inter-reflection

NeILF represents the incoming light to the surface but does not explicitly model the outgoing light. On the other hand, as reported in previous works, the outgoing light can be efficiently recovered in differentiable rendering frameworks. For example, NeRF [12] recovers the radiance in space, while the implicit differentiable renderer (IDR [30]) optimizes the surface light field of an object. To model a complete light field, we propose to additionally add one outgoing radiance field to the NeILF framework.

A straightforward solution to recover the proposed light fields is to optimize a NeILF [28] system and an IDR [30] system independently. However, such an optimization strategy would lead to inconsistent incident and outgoing lights. In fact, the two light fields can be unified through inter-reflections between surfaces. As illustrated in Fig. 2, assuming that the space between object surfaces is empty, the

outgoing light from \mathbf{x}_2 in the helmet to \mathbf{x}_1 should be equal to the incident light to \mathbf{x}_1 in the cube.

Implementation-wise, the incident and outgoing light is constrained by a consistency loss L_{ref} . As in NeILF, we sample incident rays and query their intensity during PBR. Then, we trace the reversed ray and check whether it hit any surface. If a ray $-\omega_i$ from \mathbf{x}_1 hit the surface at \mathbf{x}_2 , we query the intensity from the outgoing radiance field \mathbb{R} and encourage the two light to be the same.

$$L_{ref} = \|\mathbb{L}(\mathbf{x}_1, \omega_i) - \mathbb{R}(\mathbf{x}_2, \omega_i)\|_1. \quad (4)$$

Compared to a hard constraint that directly uses the queried outgoing light as the incident light, we do not need to trace a large amount of all incident rays. This can greatly reduce the computational cost as ray tracing is time-consuming especially when volume rendering is applied (Sec. 4).

4. Joint Geometry Optimization

In this section, we describe how the proposed inter-reflectable light fields can be used for geometry and material estimation. We first discuss our geometry representation in Sec. 4.1, and then describe details of the proposed optimization scheme in Sec. 4.2.

4.1. Recap on Geometry Representation

Signed Distance Field. Following recent neural surface reconstruction approaches [30, 25, 29, 33, 31], we represent the scene geometry as a signed distance field (SDF) in space. The surface \mathcal{S} is the zero level set of the SDF \mathbb{G} represented by a neural network with parameters θ . Let $\Omega \subset \mathbb{R}^3$ be the domain of 3D space, the network takes a point location \mathbf{x} as input and outputs its distance to the nearest surface point: $\mathbb{G} : \Omega \rightarrow \mathbb{R}$. The surface \mathcal{S} is defined as:

$$\mathcal{S} = \{\mathbf{x} \in \Omega \mid \mathbb{G}(\mathbf{x}; \theta) = 0\}. \quad (5)$$

Volume Rendering. The volume rendering technique from NeRF can globally optimize the scene geometry as the alpha composition process is dependent on points in the whole space rather than only on the surface. Let the outgoing radiance field be $\mathbb{R}(\mathbf{x}_i, \omega_o)$, we render L_o^R by alpha-blending the samples $\{\mathbf{x}_i\}_{i=1}^N$ along the ray with direction $-\omega_o$:

$$L_o^R = \sum_{i=1}^N T_i (1 - \exp(-\sigma_i \delta_i)) \mathbb{R}(\mathbf{x}_i, \omega_o), \quad (6)$$

$$\text{where } T_i = \exp\left(-\sum_{j=1}^{i-1} \sigma_j \delta_j\right),$$

In addition, recent works [25, 29] proposed to further transform an SDF to a density field. In this paper, we adopt the transformation from VolSDF [29] for its simplicity:

$$\sigma_i = \alpha \Psi_\beta(-\mathbb{G}(\mathbf{x}_i)), \quad (7)$$

Stage	Joint	Geometry					Material		
	L_{ref}	L_{vol}	L_{Eik}	L_{Hess}	L_{surf}	L_{pcd}	L_{phys}	L_{smth}	L_{Lam}
SDF Init.	✓	✓	✓	✓	✓	■	✓	✓	✓
Mat. Init.	✓	✗	✗	✗	✗	✗	✓	✓	✓
Joint	✓	✓	✓	↓	↓	✗	✓	✓	✓

Table 1. Used losses in each training stage. ■: The point cloud supervision is optional. ↓: The weight of this loss is downscaled than the previous stage.

where Ψ_β is the cumulative distribution function of the Laplace distribution with zero mean and β scale.

4.2. Joint Shape, Material and Lighting Estimation

As discussed in previous sections, a 3D scene is parameterized by a SDF field $\mathbb{G}(\mathbf{x})$ as the scene geometry, a BRDF field $\mathbb{B}(\mathbf{x})$ as the surface material, an incident light field $\mathbb{L}(\mathbf{x}, \omega_i)$ and a outgoing radiance field $\mathbb{R}(\mathbf{x}, \omega_o)$ as the scene lighting. In each iteration of joint training, the losses for joint training (shown in Fig. 2) are calculated by the following steps:

1. Perform volume rendering to get the color for a ray from the outgoing radiance field, and supervise it with ground truth (L_{vol}).
2. Find the position and normal of the intersection point between the ray and the surface by alpha blending.
3. Feed the position into the BRDF field to get the material parameters.
4. Sample incident light directions from the upper hemisphere and query the incident light from the incident light field.
5. Feed the BRDF parameters, the surface normal and the incident lights into the rendering equation to get the PBR color, and supervise it with ground truth (L_{phys}).
6. Back trace the incident directions, get the color by volume rendering from the outgoing radiance field, and minimize its difference with the incident light (L_{ref}).

In this procedure, we do not require any specific design from the geometry and the outgoing radiance field. Abstractly, the PBR requires the following outputs from these two fields for a ray: 1) volume rendered color and 2) position and normal of the intersection point with the surface. These requirements can be fulfilled by most NeRF-like systems, which guarantees the generalization of our system and the capability for future upgrades.

4.3. Training Scheme

To facilitate convergence, we propose the following 3-stage scheme to joint optimize the four fields:

SDF Initialization. First we initialize the SDF so that it recovers the basic topology of the object. The training strategy is similar to common NeRF frameworks. We only do step 1 described in Sec. 4.2. In addition, we use Eikonal loss

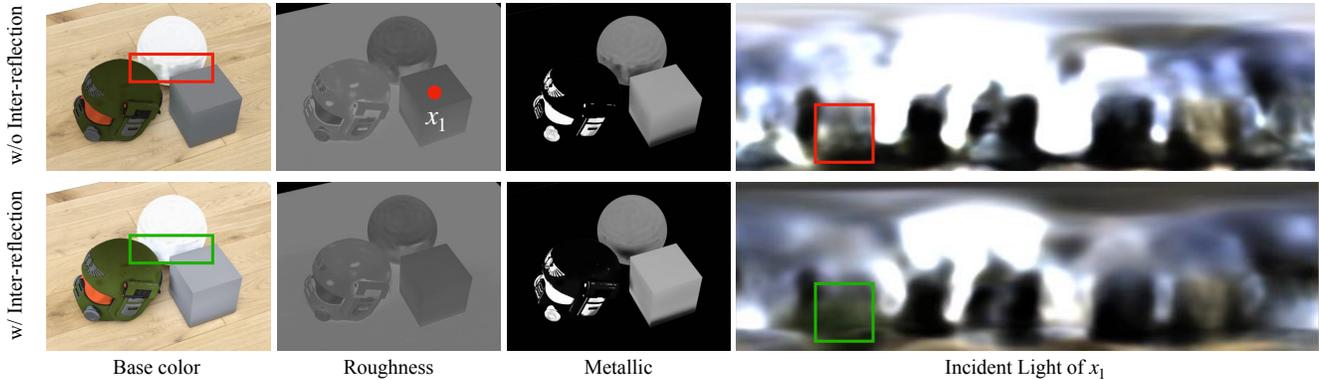


Figure 3. Qualitative ablation study on the inter-reflection constraint. The incident light map of x_1 on the cube successfully matches the outgoing radiance from the helmet. And reflections are not baked into the base color.

L_{Eik} [30], Hessian loss L_{Hess} and minimal surface loss L_{surf} [31] to regularize the surface. Optionally, an oriented point cloud can be used to directly supervise the surface (L_{pcd}) [31], which further enhances the stability of the optimization. In this stage, the material estimation module does not participate in the calculation, and thus is not optimized.

Material Initialization. In the second stage, we also initialize the material and the lighting to facilitate convergence in the joint training stage. We do steps 2-6 in Sec. 4.2. In this stage, the intersection point is treated as a constant and the geometry is fixed. Regularization including bilateral smoothness L_{smth} and Lambertian regularization L_{Lam} [28] are also added.

Joint Optimization with Inter-reflections. In this stage, we jointly optimize all four fields. Now we require the intersection points to be differentiable so that the PBR loss can optimize the geometry through the normal fed into the rendering equation and the position fed into the BRDF and incident light field. The inter-reflection loss is also added as the previous stage. Specially, we reduce some of the geometry smoothness terms to facilitate the recovery of fine details. A comprehensive list of losses is presented in Tab. 1.

4.4. HDR Rendering

In this section, we claim that linear HDR input is necessary for both the quality of the estimated material and the correctness of the inter-reflection constraint. Because the common LDR images are usually processed by unknown non-linear tone-mapping, gamma correction, and value clipping, it may result in inaccurate material and lighting estimation if we directly supervise the rendering value with the LDR ground truth. For example, if the observed intensity of a specular area is suppressed, the intensity of light sources will be underestimated. Converting the raw render to the same color space as the LDR images by either fixed or learned [28] mapping is still problematic because of potential information loss. For example, the saturated

	L_{ref}	Env			Mix			Mean
		city	studio	castel	city	studio	castel	
Base Color	✗	18.09	17.75	16.33	17.93	17.26	18.69	17.68
	✓	18.23	19.29	17.47	17.38	17.13	18.19	17.95
Roughness	✗	20.34	22.20	19.69	21.31	23.23	21.53	21.38
	✓	21.31	22.43	20.98	21.79	23.05	21.84	21.90
Metallic	✗	18.56	17.13	17.09	20.52	18.18	18.50	18.33
	✓	19.42	18.59	17.26	20.58	20.67	17.96	19.08

Table 2. Quantitative ablation study on the effect of the inter-reflection constraint on material estimation. The proposed inter-reflection improves the overall quality of the estimated material.

pixels are never supervised because the clipping operation eliminates the gradients. On the other hand, the outgoing radiance field also needs to be trained in HDR because it is used to regularize the incident light in the inter-reflection constraint. Therefore, linear HDR images are necessary for our method. In Sec. 5.4, we show that HDR supervision results in a more robust estimation than LDR.

5. Experiments

Our method has been evaluated on 1) the DTU dataset, 2) an in-house synthetic dataset (NeILF-synthetic) [28] and 3) an in-house captured linear HDR dataset (NeILF-HDR). In Sec. 5.2, we do an ablation study on the inter-reflection handling mechanism, and in Sec. 5.3 we compare the effect of joint geometry and material estimation. In Sec. 5.4, we report the geometry accuracy and fidelity of novel view rendering for real-world datasets.

5.1. Implementation

Network Architecture. For the network architecture, we test both the heavy MLP with sinusoidal positional encoding (Ours (MLP)) [13] and the lightweight MLP with hash grid encoding (Ours (NGP)) [14]. Because the hash grid based network has weaker intrinsic smoothness, we additionally use a point cloud to facilitate the convergence in the geometry initialization stage, where the point cloud can

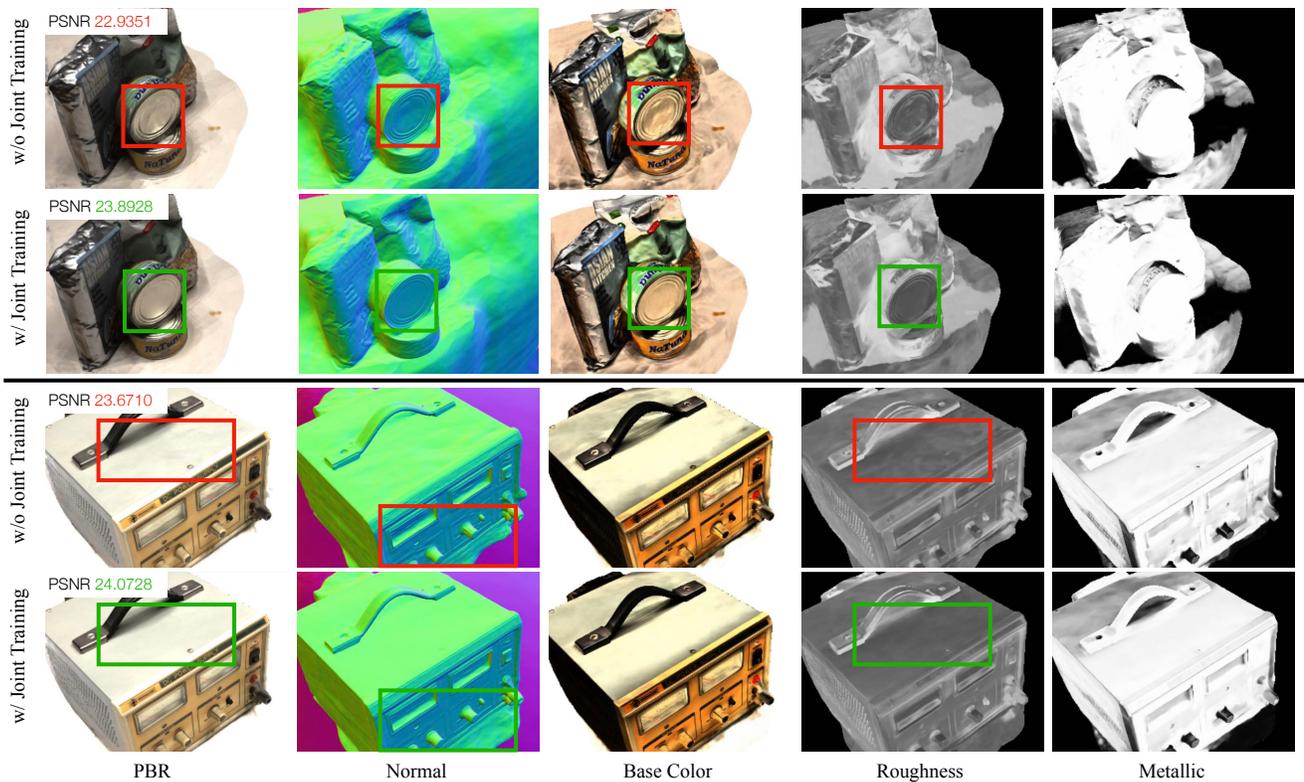


Figure 4. Qualitative ablation study on the joint optimization. For the baseline method, we use a modified VolSDF [30] to generate the geometry and use NeILF [28] to generate the surface material. With the joint refinement, the geometry and material estimation and the final rendering quality are improved.

be automatically generated from off-the-shelf multi-view stereo methods [21, 32]. Note that point cloud supervision is removed for all architecture in the joint training stage.

HDR Rendering. The dynamic range handling of the material estimation module follows NeILF [28]. The raw rendered color is directly supervised by the HDR ground truth. If the input is in LDR, we clip the rendered color to $[0, 1]$ and convert to the sRGB color space. The incident field network uses an exponential activation for the last layer for two reasons: 1) It guarantees the non-negativity of light intensity. 2) It regularizes the value distribution of pre-activated output, given that usually the light sources in the environment have extremely large intensity compared with the light from other directions. For the outgoing radiance field, we scale up the output from the Sigmoid activation to support the unbounded color. To preserve the magnitude of the gradient, we scale down the loss of the outgoing radiance field L_{vol} rendering accordingly.

Time and Memory. In each iteration, we sample 4096 rays with 64 volume rendering samples. During PBR we sample 128 incident rays for each surface point. We conduct all the experiments on a single Tesla V100 GPU. For the heavy backbone, we train the system for 80000 steps which takes 24 hours and 30GB of VRAM. For the lightweight

backbone, we train the system for 30000 steps which takes 4 hours and 10GB of VRAM.

5.2. Evaluation on Inter-reflectable Light Fields

We do an ablation study on the inter-reflection constraint. The experiments are conducted on the NeILF-synthetic [27] dataset that contains three objects placed on a plane and lit by 6 different lighting conditions. The dataset provides raw HDR images, geometry, and ground truth material for evaluation. In this experiment, we focus on material estimation so we use the given geometry and keep it fixed. For each scene, we split 9 from all 96 images as validation views and report the accuracy of BDRF parameters and rendering of these views in terms of peak signal-to-noise ratio (PSNR).

As is shown in Fig. 3, the incident light map at the point x_1 on the cube successfully captures the objects beside it. Because lights in the occlusion-severe region are correctly constrained, colors reflected from other objects are not baked into the base color anymore. Quantitative results (Tab. 2) also show the overall improvements of the estimated materials.

5.3. Evaluation on Joint Optimization

We do an ablation study on the effect of joint optimization on geometry quality. The experiments are conducted

	DTU Scan	37	97	24	40	55	63	65	69	83	105	106	110	114	118	122	Mean
Novel View Render	Nerfactor [36]	21.91	20.45	23.24	23.33	26.86	22.70	24.71	27.59	22.56	25.08	26.30	25.14	21.35	26.44	26.53	24.28
	PhysSG [34]	15.11	17.31	17.38	20.65	18.71	18.89	18.47	18.08	21.98	20.67	18.75	17.55	21.20	18.78	23.16	19.11
	Neural-PIL [5]	19.51	19.88	20.67	19.12	21.01	23.70	18.94	17.05	20.54	19.67	18.20	17.75	21.38	21.69	-	19.94
	Ours (MLP-Sep.)	23.56	23.85	26.71	28.30	29.19	28.15	26.96	29.86	25.29	27.50	31.88	28.43	27.16	31.52	33.30	28.11
PSNR	Ours (MLP)	24.17	24.60	26.40	27.24	29.85	28.16	27.39	29.82	25.50	28.19	31.84	30.20	27.71	30.87	33.62	28.37
	Ours (NGP-Sep.)	26.00	24.27	27.14	28.09	30.22	27.54	26.78	30.72	24.44	29.19	31.63	30.48	26.88	31.58	32.86	28.52
	Ours (NGP)	26.21	24.56	27.31	28.19	30.07	27.47	26.79	30.92	24.63	29.25	31.58	30.69	26.93	31.33	33.19	28.61
Geometry	NeRF [13]	1.920	1.730	1.920	0.800	3.410	1.390	1.510	5.440	2.040	1.100	1.010	2.880	0.910	1.000	0.790	1.857
	VolSDF [29]	1.140	1.260	0.810	0.490	1.250	0.700	0.720	1.290	1.180	0.700	0.660	1.080	0.420	0.610	0.550	0.857
	NeuS [25]	1.000	1.370	0.930	0.430	1.100	0.650	0.570	1.480	1.090	0.830	0.520	1.200	0.350	0.490	0.540	0.837
Chamfer Distance (mm)	Ours (MLP-Sep.)	1.593	1.998	0.919	0.498	1.116	0.899	0.813	1.509	1.274	1.109	0.723	1.852	0.428	0.654	0.653	1.069
	Ours (MLP)	1.303	1.911	0.954	0.595	1.271	0.881	0.860	1.529	1.242	1.125	0.723	2.096	0.488	0.838	0.654	1.098
	Ours (NGP-Sep.)	0.608	1.017	0.470	0.466	0.701	0.777	0.784	1.266	1.117	0.841	0.606	0.926	0.388	0.552	0.506	0.735
Normal	Ours (NGP)	0.602	1.034	0.466	0.470	0.699	0.773	0.783	1.281	1.125	0.848	0.614	0.936	0.398	0.559	0.501	0.739
	Ours (MLP-Sep.)	26.49	28.11	14.04	15.88	18.58	14.57	22.10	31.90	30.65	34.00	15.77	25.81	16.40	17.74	18.71	22.05
Angle (°)	Ours (MLP)	24.55	26.79	13.82	15.67	18.21	14.08	21.94	31.62	29.36	34.07	15.96	25.64	16.90	19.80	18.85	21.82
	Ours (NGP-Sep.)	19.98	24.27	12.10	13.31	15.50	13.82	18.98	29.87	31.23	32.55	14.47	21.66	16.01	14.66	15.88	19.62
	Ours (NGP)	19.61	23.81	11.85	13.40	15.40	13.86	18.91	29.56	29.93	32.37	14.49	21.50	16.06	14.69	15.90	19.42

Table 3. Quantitative results of geometry accuracy and novel view rendering quality. The proposed method outperforms previous methods in terms of novel view rendering quality and geometry accuracy. For the ablation study on the joint optimization, the full system achieves better novel view rendering quality, and does not degrade geometry accuracy overall.

on the DTU dataset. Although the images are in LDR and thus the estimated material is not reliable, we consider it a suitable benchmarking dataset for geometry because it has been extensively tested by previous methods. We report the Chamfer distance between the exported mesh and the ground truth. Additionally, we evaluate the normal difference of the nearest points during the calculation of Chamfer distance like [31]. The reason is that one of the main contributions of surface refinement is the normal optimization from the physically-based rendering, which is similar to Shape-from-Shading. However, the vanilla Chamfer distance metric is not sufficient to detect the improvement of surface normal because the point cloud sampled from a bumpy surface can also have an overall small nearest neighbor distance. The additional normal metric can indicate the accuracy of geometry more comprehensively.

As shown in Fig. 4, the joint refinement improves both the geometry and material estimation quality. Dumps and dents resulted from the shape-radiance ambiguity can be smoothed by joint refinement. And it can also mitigate the case that the discontinuity of base color is wrongly imprinted on geometry. Quantitative results are shown in Tab. 3 (Ours (MLP-Sep.) v.s. Ours (MLP) and Ours (NGP-Sep.) v.s. Ours (NGP)). Although the Chamfer distance of the jointly trained system is slightly worse than the separately trained one, the normal consistency is better, which is as expected. As a result, the quality of novel view rendering is also better. In summary, the refinement effect can be detected qualitatively, but is not significant enough in quantitative evaluations.

5.4. Geometry and Material Estimation

We compare our method with previous neural implicit reconstruction methods and material estimation methods. Quantitative results on DTU are shown in Tab. 3. Our

method (Ours (NGP)) outperforms previous methods on both geometry accuracy and rendering fidelity. And the qualitative geometry and material estimation results are illustrated in Fig. 5. Our method can recover fine details of the surfaces and generate reasonable material.

NeILF-HDR. We present the NeILF-HDR dataset containing 10 scenes with multi-view linear HDR images. We capture around 100 views for each scene with a handheld camera, and the trajectory forms 3 loops with different altitudes. The scene is lit by a circular lamp above the object as well as dim ambient light. This setup is mainly for preventing the cameraman from blocking the light sources. For the color space processing, We first capture raw images, do basic denoising steps and exposure adjustment so that all images have RGB channels, and store actual radiance intensity up to a global scale.

The qualitative geometry and material results on NeILF-HDR dataset are also shown in Fig. 5. Additionally, we compare the material estimated from LDR and HDR images respectively and qualitative results are shown in Fig. 6. We find that the material in the region with extreme base color is often erroneous if supervised with LDR inputs, while the system trained by HDR images shows better robustness, which shows the necessity of HDR inputs.

6. Limitations and Future Works

Under-constrained inter-reflections. In Sec.3.2, we introduce the inter-reflection (L_{ref}) to further constrain the incident light to \mathbf{x}_1 using the reflected light from the traced point \mathbf{x}_2 . However, the outgoing light from \mathbf{x}_2 along the reflection direction cannot be directly supervised, which may still leave \mathbf{x}_1 into light-albedo ambiguity. Two cases exist: **1)** If \mathbf{x}_2 is visible from other input views, the mentioned outgoing light can be obtained because of the generalization of NeRF, and thus can still regularize the incident light

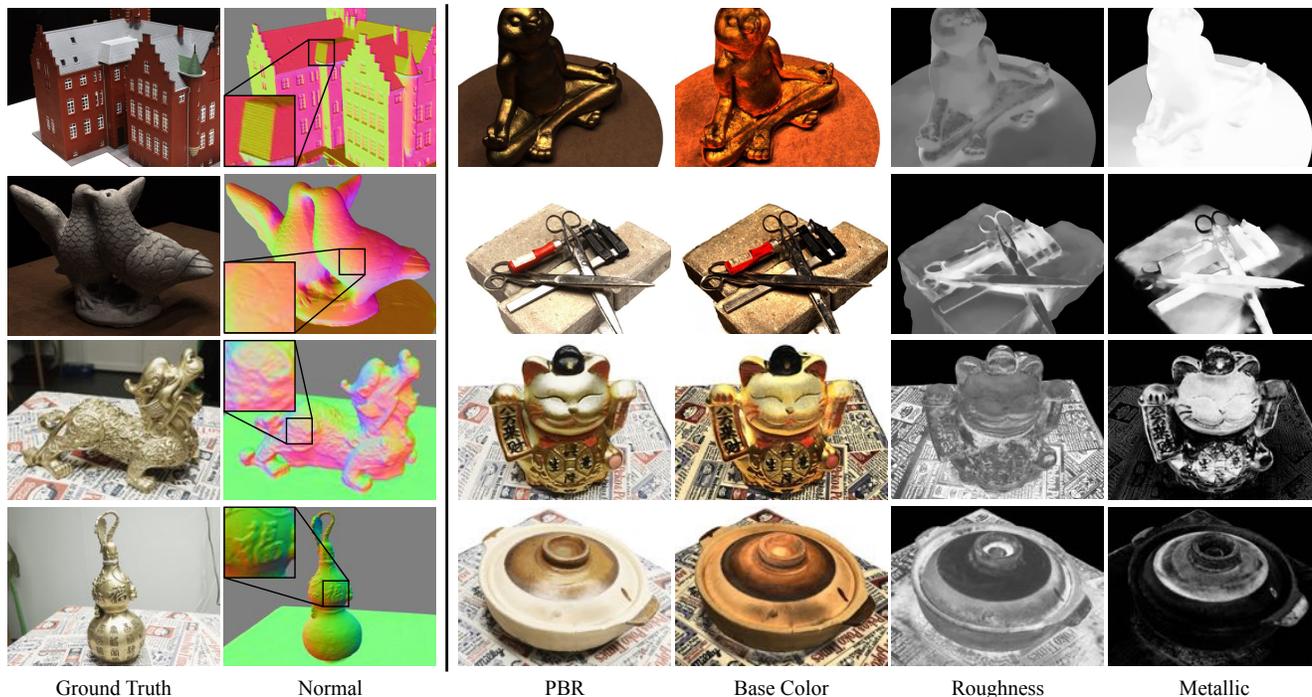


Figure 5. Qualitative geometry and material estimation results on DTU [8] and NeILF-HDR. The proposed method can generate surface with fine details and reasonable material.

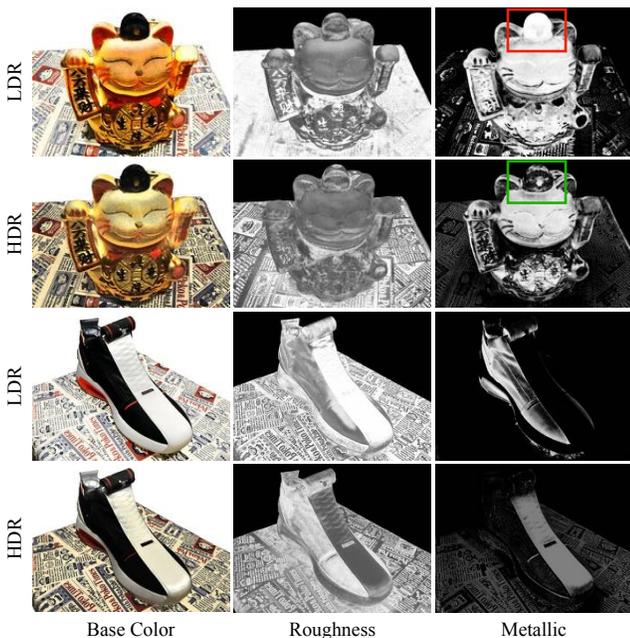


Figure 6. Comparisons of estimated material from LDR and HDR images. The results from HDR inputs are more reasonable.

to x_1 . In this case, the ambiguity will be reduced, and the baked shadow can be largely alleviated (e.g., the base color of the white ball in Fig. 3). 2) If x_2 is not visible from any input views, the radiance from x_2 cannot be recovered, which will leave the incident light to x_1 under-constrained.

In this case, the baked shadow is hard to be removed due to the light-albedo ambiguity (e.g., the scissor case in Fig. 5). In summary, although the system theoretically address the inter-reflection problem, actual results are still subject to the available ground truth. In practice, explicit shadow handling mechanism may be helpful for this issue.

Input requirements. It still requires careful photo capturing for our system to produce a decent estimation. For example, we expect the views to be dense enough to properly reconstruct the surface. Also in Sec. 4.4 we discuss the necessity of HDR inputs. Moreover, we need to guarantee that the lighting condition remains fixed during the capturing process. Although we already relax the requirement of dedicated capturing devices as in some previous works, it is still tedious for non-professional users.

Robustness. The solution space for the joint optimization of 4 fields is very large. In this work, we facilitate the convergence by pretraining each component and introducing regularization. The former costs extra time and the latter may drive the solution to a suboptimal state. In the future, we may use learned geometry and material prior to bootstrap the optimization.

Model translucency. In this work we focus on opaque objects: we use SDF as geometry representation and BRDF as surface material. However, light field estimation for translucent objects is also worth investigating. In future work, the proposed work can be extended to support non-opaque geometry and model transmittance distribution function.

7. Conclusion

We have presented a joint geometry, material, and lighting estimation system from multi-view images, among which the light field consists of both outgoing and incident light fields. The two light fields can be unified by an inter-reflection constraint that the outgoing radiance and the incident light of the two ends of the line segments between surfaces should be consistent. On the one hand, the introduction of the outgoing radiance field and the inter-reflection constraint regularizes the incident light and thus improves the estimated material. On the other hand, the geometry obtained from outgoing light optimization is also refined during the joint training with the incident light and material. The proposed method has been extensively evaluated on our in-house synthetic dataset, the real-world DTU dataset, and an in-house captured linear HDR dataset NeILF-HDR. Our method outperforms previous methods in terms of geometry accuracy and novel view rendering quality.

References

- [1] Dejan Azinovic, Tzu-Mao Li, Anton Kaplanyan, and Matthias Nießner. Inverse path tracing for joint material and lighting estimation. In *CVPR*, 2019. 2
- [2] Sai Bi, Zexiang Xu, Kalyan Sunkavalli, Miloš Hašan, Yannick Hold-Geoffroy, David Kriegman, and Ravi Ramamoorthi. Deep reflectance volumes: Relightable reconstructions from multi-view photometric images. In *ECCV*, 2020. 1, 2
- [3] Sai Bi, Zexiang Xu, Kalyan Sunkavalli, David Kriegman, and Ravi Ramamoorthi. Deep 3d capture: Geometry and reflectance from sparse multi-view images. In *CVPR*, 2020. 1, 2
- [4] Mark Boss, Raphael Braun, Varun Jampani, Jonathan T Barron, Ce Liu, and Hendrik Lensch. NerD: Neural reflectance decomposition from image collections. In *ICCV*, 2021. 1, 2
- [5] Mark Boss, Varun Jampani, Raphael Braun, Ce Liu, Jonathan Barron, and Hendrik PA Lensch. Neural-pil: Neural pre-integrated lighting for reflectance decomposition. In *NeurIPS*, 2021. 1, 7
- [6] Yue Dong, Guojun Chen, Pieter Peers, Jiawan Zhang, and Xin Tong. Appearance-from-motion: Recovering spatially varying surface reflectance under unknown lighting. *TOG*, 2014. 2
- [7] Kaiwen Guo, Peter Lincoln, Philip Davidson, Jay Busch, Xueming Yu, Matt Whalen, Geoff Harvey, Sergio Orts-Escolano, Rohit Pandey, Jason Dourgarian, et al. The relightables: Volumetric performance capture of humans with realistic relighting. *TOG*, 2019. 2
- [8] Rasmus Jensen, Anders Dahl, George Vogiatzis, Engil Tola, and Henrik Aanæs. Large scale multi-view stereopsis evaluation. In *CVPR*, 2014. 2, 8
- [9] James T Kajiya. The rendering equation. In *Proceedings of the 13th annual conference on Computer graphics and interactive techniques*, 1986. 2
- [10] Junxuan Li and Hongdong Li. Neural reflectance for shape recovery with shadow handling. In *CVPR*, 2022. 2
- [11] Jingwang Ling, Zhibo Wang, and Feng Xu. Shadowneus: Neural sdf reconstruction by shadow ray supervision. *arXiv preprint arXiv:2211.14086*, 2022. 2
- [12] Ricardo Martin-Brualla, Noha Radwan, Mehdi SM Sajjadi, Jonathan T Barron, Alexey Dosovitskiy, and Daniel Duckworth. Nerf in the wild: Neural radiance fields for unconstrained photo collections. In *CVPR*, 2021. 3
- [13] 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 *ECCV*, 2020. 1, 2, 3, 5, 7
- [14] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *arXiv:2201.05989*, 2022. 5
- [15] Jacob Munkberg, Jon Hasselgren, Tianchang Shen, Jun Gao, Wenzheng Chen, Alex Evans, Thomas Mueller, and Sanja Fidler. Extracting Triangular 3D Models, Materials, and Lighting From Images. *arXiv:2111.12503*, 2021. 1
- [16] Giljoo Nam, Joo Ho Lee, Diego Gutierrez, and Min H Kim. Practical svbrdf acquisition of 3d objects with unstructured flash photography. *TOG*, 2018. 1, 2
- [17] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In *CVPR*, 2020. 2, 3
- [18] Merlin Nimier-David, Zhao Dong, Wenzel Jakob, and Anton Kaplanyan. Material and lighting reconstruction for complex indoor scenes with texture-space differentiable rendering. 2021. 2
- [19] Jeong Joon Park, Aleksander Holynski, and Steven M Seitz. Seeing the world in a bag of chips. In *CVPR*, 2020. 2
- [20] Carolin Schmitt, Simon Donne, Gernot Riegler, Vladlen Koltun, and Andreas Geiger. On joint estimation of pose, geometry and svbrdf from a handheld scanner. In *CVPR*, 2020. 1, 2
- [21] Johannes L Schönberger, Enliang Zheng, Jan-Michael Frahm, and Marc Pollefeys. Pixelwise view selection for unstructured multi-view stereo. In *ECCV*, 2016. 6
- [22] 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 *CVPR*, 2021. 2
- [23] Pratul P Srinivasan, Ben Mildenhall, Matthew Tancik, Jonathan T Barron, Richard Tucker, and Noah Snavely. Lighthouse: Predicting lighting volumes for spatially-coherent illumination. In *CVPR*, 2020. 2
- [24] Jiaping Wang, Peiran Ren, Minmin Gong, John Snyder, and Baining Guo. All-frequency rendering of dynamic, spatially-varying reflectance. In *ACM SIGGRAPH Asia*, 2009. 3
- [25] 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. 2021. 2, 3, 4, 7
- [26] Rui Xia, Yue Dong, Pieter Peers, and Xin Tong. Recovering shape and spatially-varying surface reflectance under unknown illumination. *TOG*, 2016. 2

- [27] Yao Yao, Zixin Luo, Shiwei Li, Tian Fang, and Long Quan. Mvsnet: Depth inference for unstructured multi-view stereo. In *ECCV*, 2018. 6
- [28] Yao Yao, Jingyang Zhang, Jingbo Liu, Yihang Qu, Tian Fang, David McKinnon, Yanghai Tsin, and Long Quan. Neif: Neural incident light field for physically-based material estimation. In *ECCV*, 2022. 1, 2, 3, 5, 6
- [29] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. *arXiv preprint arXiv:2106.12052*, 2021. 1, 2, 3, 4, 7
- [30] Lior Yariv, Yoni Kasten, Dror Moran, Meirav Galun, Matan Atzmon, Basri Ronen, and Yaron Lipman. Multiview neural surface reconstruction by disentangling geometry and appearance. *NeurIPS*, 2020. 2, 3, 4, 5, 6
- [31] 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 *CVPR*, 2022. 2, 3, 4, 5, 7
- [32] Jingyang Zhang, Yao Yao, Shiwei Li, Zixin Luo, and Tian Fang. Visibility-aware multi-view stereo network. *BMVC*, 2020. 6
- [33] Jingyang Zhang, Yao Yao, and Long Quan. Learning signed distance field for multi-view surface reconstruction. In *ICCV*, 2021. 2, 3, 4
- [34] Kai Zhang, Fujun Luan, Qianqian Wang, Kavita Bala, and Noah Snavely. Physg: Inverse rendering with spherical gaussians for physics-based material editing and relighting. In *CVPR*, 2021. 1, 2, 3, 7
- [35] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. Nerf++: Analyzing and improving neural radiance fields. *arXiv preprint arXiv:2010.07492*, 2020. 2
- [36] 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. *TOG*, 2021. 1, 2, 7