

MLI-NeRF: Multi-Light Intrinsic-Aware Neural Radiance Fields

Yixiong Yang¹, Shilin Hu², Haoyu Wu², Ramon Baldrich¹, Dimitris Samaras², Maria Vanrell¹

¹ Universitat Autònoma de Barcelona, ² Stony Brook University

{yixiong, ramon, maria}@cvc.uab.cat, {shilhu, haoyuwu, samaras}@cs.stonybrook.edu

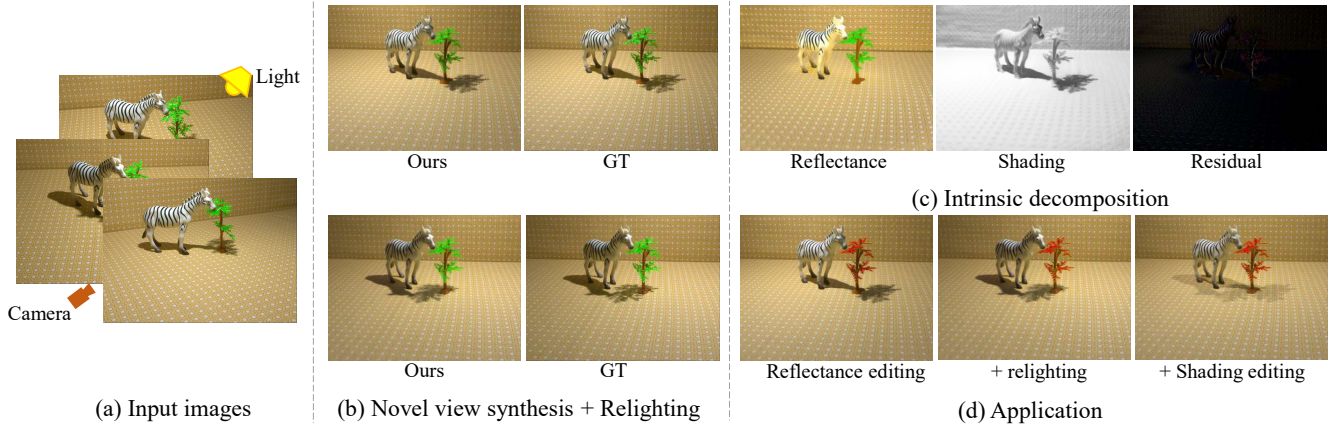


Figure 1. Given real-world images from ReNe dataset [39] (a), our method learns the neural radiance fields that enable novel view synthesis and relighting (b), and intrinsic decomposition (c) simultaneously. Image editing applications (d) can also be employed, such as reflectance editing, reflectance editing plus relighting or shading editing (simulating two lights).

Abstract

Current methods for extracting intrinsic image components, such as reflectance and shading, primarily rely on statistical priors. These methods focus mainly on simple synthetic scenes and isolated objects and struggle to perform well on challenging real-world data. To address this issue, we propose MLI-NeRF, which integrates Multiple Light information in Intrinsic-aware Neural Radiance Fields. By leveraging scene information provided by different light source positions complementing the multi-view information, we generate pseudo-label images for reflectance and shading to guide intrinsic image decomposition without the need for ground truth data. Our method introduces straightforward supervision for intrinsic component separation and ensures robustness across diverse scene types. We validate our approach on both synthetic and real-world datasets, outperforming existing state-of-the-art methods. Additionally, we demonstrate its applicability to various image editing tasks. Code and data are available at <https://github.com/liulisixin/MLI-NeRF>.

1. Introduction

Neural radiance fields (NeRF) have enabled significant strides in novel view synthesis (NVS) [24, 29] including

efforts towards scene editing [41], such as recoloring [45] and relighting [25, 46]. Scene editing becomes easier when the scene can be decomposed into editable sub-attributes. There are two related approaches to scene editing [14]: inverse rendering and intrinsic image decomposition.

The first approach [18, 44, 49, 50] integrates inverse rendering with neural rendering methods for scene decomposition. They often employ the BRDF [5] to model material properties and jointly optimize geometry, materials, and environmental lighting. However, inverse rendering presents a highly ill-posed challenge: separating material properties and illumination in images often yields ambiguous results, and tracing light within scenes is computationally intensive. These factors limit inverse rendering to object-specific scenarios. The second approach [6, 9, 45], based on intrinsic image decomposition [2], aims to provide an explainable representation of a scene in terms of components such as reflectance and shading. In general, intrinsic image decomposition is more applicable to a broader range of scenarios, including individual objects and more complex scenes with backgrounds. While [45] has pioneered the integration of intrinsic decomposition within NeRF, it has not fully leveraged the 3D information available through neural rendering.

Mineralogists illuminate their specimens from different angles to reveal their features. Similarly, varying the light

source position is essential for uncovering the intrinsic details of a scene. We aim to enhance intrinsic decomposition quality and expand scene editing capabilities by leveraging multiple light sources to build an intrinsic-aware NeRF. The connection between varying lighting conditions and intrinsic decomposition has been discussed for 2D images [20, 21], but not yet in neural rendering, even though there is interest in relighting using neural rendering [36, 46].

In this paper, we introduce **MLI-NeRF**, a two-stage method to learn an intrinsic-aware NeRF. In Stage 1, we extend NeRF to incorporate light position information and learn a relightable scene using images captured from various camera angles and light source positions. In the subsequent post-processing, we begin by obtaining normals and light visibility maps for images under multiple lighting conditions using the model from Stage 1 and sphere tracing [8]. We then generate pseudo shadings from the normals, light rays, and light visibility maps. Finally, for each camera pose, pseudo reflectance is generated by combining cues from multiple lighting conditions. In Stage 2, we make our model intrinsic-aware by introducing additional modules for reflectance and shading while restricting light position input to the shading module only, ensuring the independence of the reflectance and light. In this paper, we forego potentially oversimplifying statistical constraints on various illumination-related factors [6, 20, 23, 45] to instead use the physics-based disentanglement of reflectance and shading and achieve high-quality results.

As illustrated in Fig. 1, our method achieves high-quality intrinsic decomposition results (Fig. 1(c)), as well as NVS and relighting results (Fig. 1(b)). It also enables applications such as reflectance editing, relighting, and shading editing (Fig. 1(d)). Furthermore, our method is applicable across various datasets, including the object-only synthetic NeRF [29] dataset, the real object dataset [13, 46], and the ReNe [39] dataset with real-world full scenes. Our contributions are summarized as follows:

- A novel intrinsic-aware NeRF model that integrates multiple light information, enabling applications such as NVS, lighting modification, and scene editing.
- A method that separates intrinsic components by using supervision from generated pseudo intrinsic images. We introduce straightforward constraints based on 3D information to eliminate the need for statistical priors required by traditional approaches. Our method ensures robustness across various scene types.
- Experimental results across three different datasets demonstrate our method’s superior performance in intrinsic decomposition compared to existing state-of-the-art methods, showing advancements not only in synthetic object-only scenes but also in challenging real scenes with backgrounds and cast shadows.

2. Related Works

Intrinsic decomposition. Intrinsic decomposition is a classical challenge in computer vision [2], with much of the previous research focused on the 2D image [1, 6, 9, 22]. A key difficulty in this area is the scarcity of real datasets, which need complicated and extensive annotation. This limitation has spurred interest in semi-supervised and unsupervised techniques [20, 21, 27]. IntrinsicNeRF [45] has been a pioneer in applying intrinsic decomposition to 3D neural rendering. Similar to previous unsupervised methods in 2D, it utilizes statistical priors, including chromaticity and semantic constraints, for guidance. However, these constraints do not accurately reflect physical principles and often fall short in complex scenarios. Our approach leans on 3D information and physical constraints (e.g., variations in illumination) to achieve superior results.

Relighting. Relighting has recently garnered attention from various perspectives within the field [10]. Data-driven approaches have been explored, with research focusing on portrait scenes [17, 32, 33, 38, 51] and extending to more complex scenarios [11, 16, 31, 35, 40]. Kocsis *et al.* [19] have also investigated lighting control within diffusion models, enabling the generation of scenes under varying lighting conditions. Meanwhile, relighting has also received widespread attention within the field of neural rendering [13, 37, 39, 46], achieving impressive relighting outcomes within individual scenes. Among them, Toschi *et al.* [39] proposed the ReNe dataset, which consists of images captured with various cameras and light poses under controlled lab conditions. Zeng *et al.* [46] enhanced NeRF relighting with visibility and specular hints. Yang *et al.* [43] enhance outdoor scene relighting by combining NeRF with a CNN for intrinsic image prediction. Chang *et al.* [7] use location and time to estimate the direction of sunlight in their approach to outdoor relighting. However, the potential of using information from multiple lights for 3D scene understanding remains unexplored.

Inverse rendering. Inverse rendering [14] is an alternative approach to recovering the fundamental properties of a scene, which aims at extracting the geometry materials, and lighting of a 3D scene. Recently, the study of inverse rendering methods based on NeRF has become a popular topic. NeRFactor [49] introduced a method to improve the geometric quality of NeRF and incorporated a data-driven BRDF prior. More methods have been developed to address the scenes with different light conditions, including fixed illumination conditions [47], varying light sources [3, 4]. Invrender [50] proposed a method for predicting indirect light. L-Tracing [8] introduced an efficient algorithm for estimating visibility without training. SIRE-IR [44] introduced a method for high-illuminance scenes, addressing the issue where previous methods struggled under prominent cast shadows. Liu *et al.* [26] propose the OpenIllumination

dataset with multi-illumination, which focuses on inverse rendering evaluation on real objects. However, inverse rendering methods are primarily based on individual objects and are challenging to extend to large, complex scenes, such as those with backgrounds.

3. Method

We propose a two-stage method, with the overall framework illustrated in Fig. 2. In previous inverse rendering methods [50], accounting for indirect illumination and shadow has been a critical challenge. Our strategy is to leverage multiple lighting conditions to make the corresponding information more accessible. In the first stage, we train our model to represent scenes under varying camera positions and lighting conditions, enabling NVS and lighting modification. In the subsequent post-processing, we use the model from Stage 1 to generate pseudo intrinsic images of reflectance and shading for each image. Specifically, pseudo shadings are derived from normals and light visibility, while pseudo reflectance is obtained by dividing the image by the shading. Leveraging multiple lighting conditions allows us to effectively gather more detailed information about the scene, thus resulting in high-fidelity pseudo labels. In Stage 2, we retrain the network with added modules for predicting reflectance and shading, restricting light information input solely to the shading module. Pseudo images guide the separation of intrinsic components in this phase.

3.1. Preliminaries

Some traditional methods decompose images into reflectance and shading [2, 6, 12], primarily targeting diffuse components. Since the real world contains many non-diffuse effects, recent approaches[14, 45] add a residual term to account for discrepancies. We follow this setup and model intrinsic decomposition as follows:

$$I(i, j) = R(i, j) \otimes S(i, j) + Re(i, j) \quad (1)$$

where R , S and Re denote Reflectance, Shading and Residual, respectively. Additionally, following the Lambert’s cosine law, a shading can be computed as below:

$$S = \vec{N} \cdot \vec{L} \quad (2)$$

where \vec{N} is the normal ray and \vec{L} is the light ray. This physical illumination model guides our following pseudo shading generation and intrinsic-aware NeRF training.

3.2. Stage 1: Neural Multi-Light Implicit Representation

We follow the structure of Neuralangelo [24], which achieves promising results in both small and large scenes. It uses 3D hashing encoding combined with Signed Distance

Function (SDF) [30] to represent the implicit geometry, and then employs an MLP to model the color information. The original Neuralangelo does not support light position as an input, so, besides the original MLP input, we incorporate the light position encoded with spherical encoding as an additional input. We illustrate Stage 1 in Fig. 2 (top-left), with formulas as follows:

$$sdf = f(\mathbf{x}), \quad \mathbf{c} = \text{MLP}_{color}(\mathbf{x}, \mathbf{n}, \text{feat}, \mathbf{d}, \mathbf{l}) \quad (3)$$

where $f(\cdot)$ is the geometry network that predicts SDF and $\text{MLP}_{color}(\cdot)$ is the color network. \mathbf{x} is the spatial position, \mathbf{n} and feat are the normal and the features from the SDF network, \mathbf{d} is the view direction, and \mathbf{l} is the light position. Following [24], the loss for Stage 1 is:

$$\mathcal{L}_{S1} = w_{rgb}\mathcal{L}_{rgb} + w_{eik}\mathcal{L}_{eik} + w_{curv}\mathcal{L}_{curv} \quad (4)$$

where \mathcal{L}_{rgb} is the loss of the rendered image, \mathcal{L}_{eik} represents the Eikonal loss [15], and \mathcal{L}_{curv} is the curvature loss [24]. The terms w_{eik} and w_{curv} are the corresponding weights.

3.3. Post-processing: Physics-based Pseudo Label Generation

Here we propose a post-processing that generate pseudo labels for reflectance and shading in three steps, as illustrated in Fig. 2 (right). It starts with generating pseudo shading from the normal and the light visibility. We then generate pseudo reflectance using multi-light shadings and images.

Step A. We derive the normal from the gradient of the SDF network. The geometry network also provides depth information, which is used to estimate the intersection points in conjunction with sphere tracing [8]. Sphere tracing is an iterative ray-marching technique that locates surface intersections using the SDF. Specifically, it marches along the ray in steps determined by the SDF values until convergence. In our case, sphere tracing is used to estimate both the intersection points and the light visibility V , which indicates whether a point is directly illuminated based on the light position and intersection points.

Step B. Since the Eq. (2) does not consider occlusion and other effects, the generation of pseudo-shading in our implementation follows the formula:

$$S^* = (\max(\vec{N} \cdot \vec{L}, 0) \otimes V)^\gamma \quad (5)$$

where $(\cdot)^\gamma$ represents gamma correction, and \otimes denotes pixel-wise product. This gamma correction is essential to adapt to the nonlinear representation of digital images. All image sensors introduce a gamma correction to accommodate the light integration to the nonlinear perception of brightness in the human eye. Since our pseudo-shading is directly created from the 3D world, we need to introduce the sensor representation.

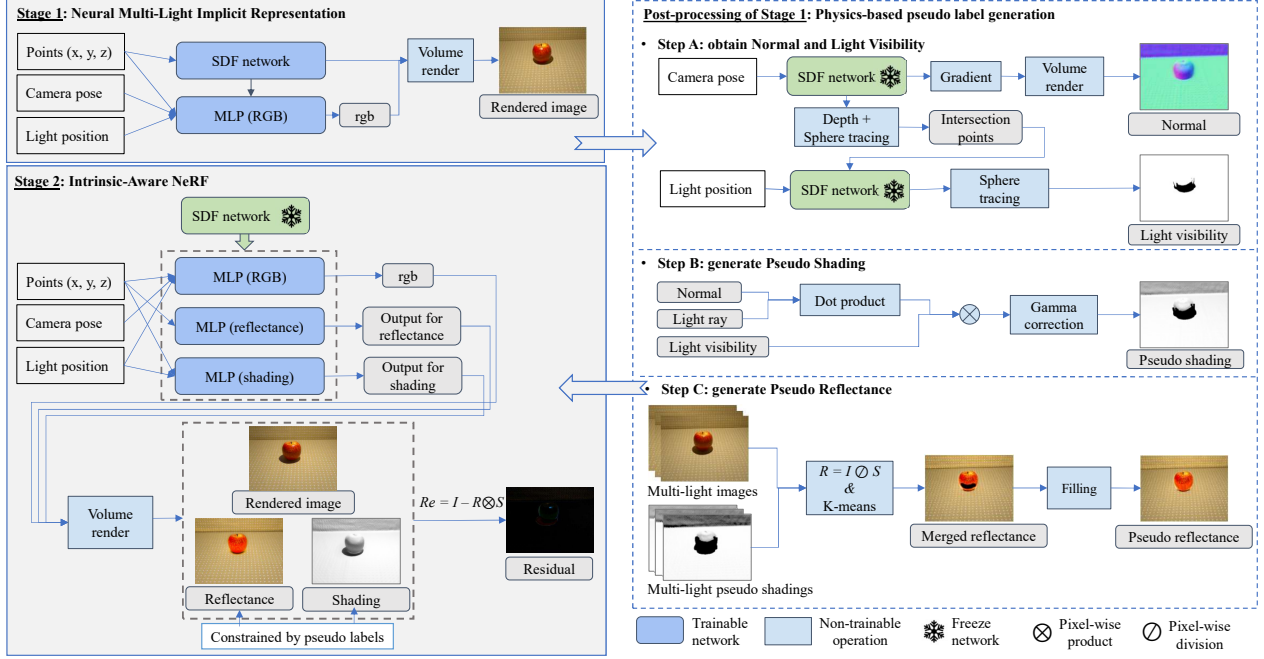


Figure 2. **Illustration of the Framework.** In Stage 1, we introduce light position as input to extend NeRF for multi-light implicit representation (top left). Following Stage 1, three post-processing steps are applied to generate pseudo labels for reflectance and shading using the proposed physics-based pipeline (right). In Stage 2, we train the intrinsic-aware NeRF based on the model from Stage 1 and the pseudo labels from post-processing (bottom left).

Step C. This step entails inferring high-fidelity pseudo reflectance from multiple pseudo shading. We use the equation $R = I \oslash S$ as a simplified version of Eq. (1) ignoring at this point the residual component that mainly entangles specular and undirected light components.

Our method in this step leverages the trained model from Stage 1 to generate multiple images under different direct light conditions, each accompanied by its corresponding pseudo shading. The flowchart is shown at the right bottom in Fig. 2, and Fig. 3 illustrates the images during the calculation. By calculating $R = I \oslash S$, we can compile the reflectance information for every pixel from multiple light conditions to diminish the influence of entangled indirect light. Additionally, under different light positions, a larger pseudo shading value indicates a smaller contribution of indirect light relative to the main light source, making the calculated reflectance more accurate. To generate a unique reflectance from multiple pseudo shading, we use the K-means algorithm [34] at the pixel level, incorporating the weights of each pseudo shading to select the most probable pseudo reflectance. This approach allows us to achieve a merged reflectance under varied lighting conditions.

However, some regions within the merged reflectance may appear as holes due to the absence of direct illumination in all lighting conditions. We address these areas with a filling strategy. We compute a weighted distance between hole pixels and non-hole pixels, considering their spatial

distance in the image, the angular difference of their normals, and the color difference in the RGB image. Then, we assign the color of the nearest non-hole pixel, based on this weighted distance, to the hole pixel. Under the multi-light setting, hole regions are typically small, and our filling strategy is a rough supplement. We reduce their weight during training, further minimizing its impact.

After our proposed three-step post-processing, we achieve the final pseudo reflectance, as shown in Fig. 3. Our pseudo reflectance offers a straightforward basis for disentanglement and has proven to be more reliable in guiding the following training stage compared to the previous statistical priors. Additionally, we compute weight maps W_R and W_S for both pseudo reflectance and pseudo shading based on the edges of pseudo shading and visibility. Areas with higher pseudo shading values, or those further from visibility edges (where visibility calculations may be prone to errors), exhibit greater credibility in their pseudo labels; conversely, areas closer to visibility edges or with lower pseudo shading values are deemed less reliable.

3.4. Stage 2: Intrinsic-Aware NeRF

As illustrated in Fig. 2 (bottom-left), we use high-fidelity pseudo intrinsic images to guide the intrinsic decomposition learning. Expanding the model from Stage 1, we add two extra MLPs dedicated to generating reflectance and shading outputs, while the geometry network is frozen. Compared




Figure 3. Illustration of the pseudo reflectance generation process in the post-processing.

to the MLP_{color} in Eq. (3), these two MLPs receive different inputs. Since reflectance is independent of lighting, the $MLP_{reflectance}$ does not take the light position as input. Additionally, both reflectance and shading are diffuse components, and a diffuse reflecting surface exhibits Lambertian reflection, indicating that it maintains equal luminance when observed from any direction. Thus, view poses are excluded from their inputs. The formulas are:

$$\begin{aligned} \mathbf{r} &= MLP_{reflectance}(\mathbf{x}, \mathbf{n}, \mathbf{feat}) \\ \mathbf{s} &= MLP_{shading}(\mathbf{x}, \mathbf{n}, \mathbf{feat}, \mathbf{l}) \end{aligned} \quad (6)$$

where \mathbf{x} is the spatial position, \mathbf{n} , \mathbf{feat} are the normal and features from the SDF network, and \mathbf{l} is the light position.

After volume rendering, we obtain RGB images, along with reflectance and shading. Subsequently, the residual is derived from Eq. (1). During training, the pseudo labels impose constraints on reflectance and shading:

$$L_{intrinsic} = W_R \cdot \|\hat{R} - R^*\|_1 + W_S \cdot \|\hat{S} - S^*\|_1 \quad (7)$$

where \hat{R} and \hat{S} represent the predicted reflectance and shading, respectively, and R^* and S^* are their corresponding pseudo labels. W_R and W_S represent weight maps for reflectance and shading, derived during pseudo label generation. As demonstrated in [45], the diffuse components dominate the scene, so it is crucial to prevent the training from converging to undesirable local minima (e.g. $R = 0, S = 0, R_e = I$). Therefore, we introduce a regularization term for R_e to ensure that the image is primarily recovered through R and S : $L_{reg} = \|\hat{R}_e\|_1$.

Finally, the Stage 2 loss is the weighted sum of:

$$\mathcal{L}_{S2} = w_{rgb} \mathcal{L}_{rgb} + w_{intrinsic} L_{intrinsic} + w_{reg} L_{reg} \quad (8)$$

where w_{rgb} , $w_{intrinsic}$ and w_{reg} are the corresponding weights.

4. Experiments

We quantitatively and qualitatively validate our method and compare it with other approaches, including traditional learning-based intrinsic decomposition methods and neural rendering methods. The comparison encompasses intrinsic decomposition and NVS with relighting. For data-driven methods, we select PIE-Net[9] and Careaga *et al.* [6] (hereafter referred to as Ordinal). For NeRF-related methods, we choose InvRender[50], TensoIR[18],

and IntrinsicNeRF[45]. Among these, InvRender and TensoIR are inverse rendering methods, while IntrinsicNeRF is an intrinsic decomposition method. Additionally, we select NRHints[46], a method focused on relighting, to validate our relighting performance.

4.1. Datasets

Our experiments are conducted on three datasets, each containing four scenes.

Synthetic Dataset. It contains synthetic data based on Blender, inspired by the synthetic dataset designed by NRHints[46] for relighting. Some scenes in this dataset are derived from NeRF[29]. Since the original dataset does not include ground truth (GT) for intrinsic decomposition, we re-render the data in Blender and export these quantities. Each scene comprises 500 images for training, 100 for validation, and 100 for testing, including intrinsic components for each image. The selection of light positions and camera poses follows the setup described by Zeng *et al.* [46], with both distributed across a hemispherical space above the scene. The light position and camera pose for each image are randomly and independently selected.

Real Object Dataset. It includes real objects from the object relighting dataset collected by [13, 46], featuring various viewpoints and lighting conditions. The training set size varies from 500 to 2000 images per scene.

ReNe Dataset. ReNe dataset[39], unlike most datasets used in previous inverse rendering research, contains complete real-world scenes with backgrounds. Notably, the camera poses and light positions in this dataset are concentrated in a specific direction rather than evenly distributed across 360 degrees, which poses challenges for both reconstruction and intrinsic decomposition. This real dataset features 2000 images across scenes, captured from 50 viewpoints under 40 lighting conditions, with lighting and camera poses grid-sampled. Following the dataset split, we use 1628 images (44 camera poses \times 37 light positions) for training. Since the test set is not publicly available, we use the validation set for inference.

4.2. Lighting and Camera Views Settings

In this section, we further clarify the lighting and camera view setups in our experiments.

Grid-sampled or non-grid-sampled. In the ReNe dataset,

Method	Light Setting	Reflectance				Shading			
		PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	MSE \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	MSE \downarrow
InvRender	Single	16.59	0.8228	0.1807	0.0271	—	—	—	—
TensoIR	Single	18.50	0.8518	0.1544	0.0213	—	—	—	—
IntrinsicNeRF	Single	18.02	0.8353	0.2142	0.0226	19.02	0.8660	0.1476	0.0168
Ours	Single	22.44	0.9012	0.0950	0.0072	24.39	0.9188	0.0843	0.0049
TensoIR	Multiple	22.70	0.8800	0.1450	0.0087	—	—	—	—
Ours	Multiple	<u>24.40</u>	<u>0.9357</u>	<u>0.0582</u>	<u>0.0051</u>	23.44	0.9225	0.0798	0.0055
PIE-Net	Random	19.59	0.8708	0.1298	0.0153	20.03	0.8868	0.1653	0.0133
Ordinal	Random	18.14	0.8716	0.1251	0.0191	16.80	0.8717	0.1679	0.0293
Ours	Random	25.48	0.9420	0.0515	0.0041	22.90	0.9024	0.1049	0.0072

Table 1. Quantitative results of the intrinsic decomposition on the **Synthetic Dataset**. We compare different methods under three light settings: single, multiple, and random. Our method outperforms others in all settings. Best results are in **bold**, second best in underlined.

camera views and light positions are grid-sampled. In contrast, the Synthetic and Real Object Datasets use independently sampled views and lights. We denote this as non-grid-sampled. When merging results to generate pseudo reflectance in the post-processing of Stage 1 (Fig. 2), for grid-sampled setups, we use all light positions in the training set. For non-grid-sampled setups, where the number of possible light positions is large, we randomly select several (typically four) light positions.

Additional Lighting Setups for Better Comparison. Previous methods often struggle with varying lighting conditions. InvRender [50] and IntrinsicNeRF [45] operate under single lighting conditions, while TensoIR [18] supports multiple lights but recommends around four, which cannot handle hundreds of light positions. We introduce two additional settings to better compare intrinsic decomposition methods and examine the impact of different light positions: **single light** (one fixed position) and **multiple lights** (four fixed positions). TensoIR occasionally failed with four lights, so we used three instead in this case.

For the Synthetic Dataset, the original setting is referred to as **random lights**. We re-render scenes in Blender for the single and multiple light tracks, maintaining the same image counts and random seed for camera views. For the ReNe Dataset, the original setting is denoted as **all lights**. We selected 1 light condition as the single light track and 4 light conditions as the multiple lights track, thereby the training sets respectively contain 44×1 and 44×4 images. Notably, the single light setup serves as an extreme condition, diverging from our intent of utilizing multi-light information. This track was designed for fair comparison with other methods and to highlight the importance of multi-light information in intrinsic decomposition.

4.3. Implementation Details

Our model’s hyperparameters include a batch size of 2048 and each stage is trained for 500k iterations. We implement the model in PyTorch and use the AdamW [28] optimizer with a learning rate of $1e^{-3}$ for optimization. The experiments can be conducted on a single Nvidia RTX 3090. The training time of Stage 1 follows the training time of Neu-

ralangelo [24], and Stage 2 takes 19 hours. The weights of losses, w_{eik} , w_{curv} , $w_{\text{intrinsic}}$, w_{reg} are set to 0.1, $5e^{-4}$, 1.0, and 1.0, respectively. To evaluate the comparison between predicted images and GT, we employ the following metrics: Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index (SSIM) [42], Learned Perceptual Image Patch Similarity (LPIPS) [48] and Mean Squared Error (MSE).

4.4. Results on the Synthetic Dataset

Comparison with SOTA. Tab. 1 presents the quantitative comparison of intrinsic decomposition on all methods and light settings. Since the sequence of camera poses is the same across all settings when generating the dataset, the reflectance GT remains consistent, allowing for direct comparison. Our method outperforms existing methods in predicting reflectance for each light setting, with the best results under the random setting. Due to the different shading GTs, the comparisons are conducted separately. Our method also achieves the best results across all settings. A qualitative comparison is shown in Fig. 4. The single light setting demonstrates that under challenging lighting conditions, such as those with pronounced cast shadows, all methods struggle to achieve consistently good results. This indicates that using multi-light information is a practical approach. However, our method still achieves the most meaningful results in the single light setting of the FurBall scene. In the other two settings, our approach also demonstrates significantly better performance, producing high-quality reflectance and shading results, effectively handling cast shadows.

Comparison across different light settings. As shown in Tab. 1 and Fig. 4, both TensoIR and our method exhibit improved performance as the number of light positions increases. This observation aligns with the findings mentioned in the TensoIR [18]. Additionally, the results in Fig. 4 demonstrate that our method’s performance under the multiple light setting (4 lights) is very close to that of the random light setting.

NVS and relighting analysis. Tab. 2 shows the quantitative results of NVS and relighting. Since the lights are fixed and known under the Single and Multiple light settings, the

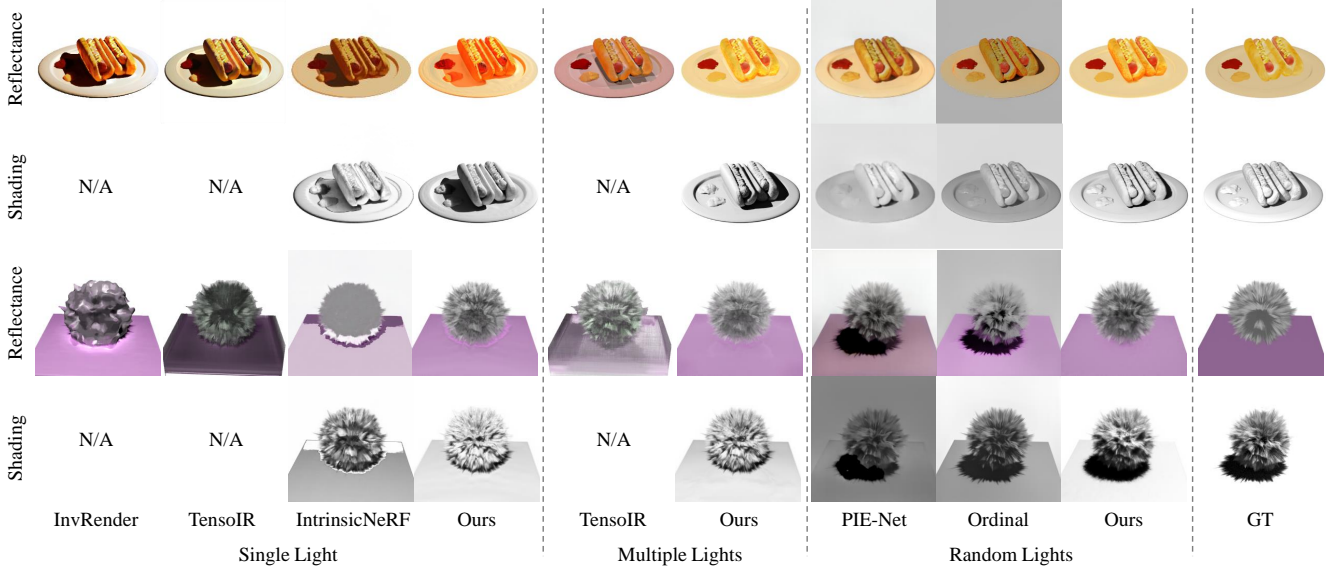


Figure 4. Qualitative Results on the **Synthetic Dataset** with all settings. The same GT reflectance applies across all settings, but GT shading differs due to varying light positions. Here for brevity, only the shading under the Random Lights setting is shown. Compared to other methods, our approach predicts the best reflectance and effectively handles cast shadows.

results pertain only to NVS, where our method achieves the best performance. Under the Random setting, the results reflect both NVS and relighting, where our method shows comparable performance to NRHints [46].

	Light Setting	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	MSE \downarrow
InvRender	Single	23.99	0.8760	0.1109	0.0051
TensoIR	Single	35.00	0.9761	0.0343	0.0019
IntrinsicNeRF	Single	34.53	0.9794	0.0137	0.0005
Ours	Single	37.51	0.9878	0.0099	0.0003
TensoIR	Multiple	33.19	0.9645	0.0503	0.0022
Ours	Multiple	36.12	0.9836	0.0135	0.0003
NRHints	Random	32.79	0.9674	0.0369	0.0007
Ours	Random	31.20	0.9619	0.0354	0.0010

Table 2. NVS and relighting results on the **Synthetic Dataset**.

4.5. Results on the Real Object Dataset

Results on the Real Object Dataset allow us to confirm a quite accurate qualitative performance of our intrinsic decomposition. In Fig. 5(a) we show a comparison of our method versus SOTA. Our RGB-rendered images are similar to NRHints [46] and our estimations for Reflectance and Shading can be compared to PIE-Net [9] and Ordinal [6]. In particular, we want to highlight our estimation of reflectance that does not present any remaining shading effect with respect to the rest of the methods, restoring the vibrant colors of the objects and achieving significantly better results. In Fig. 5(b) we add two more examples where we zoom out two details. At the top is a hole through which our method estimates the texture of the support surface, and at the bottom, it gets a proper reflectance from a strong cast shadow. Our shading presents a better disentangling from reflectance than the rest of the methods. Although we can

show a clear overall improvement, some limitations can still be observed from some difficult areas that remain hidden from any camera point of view and light source, as we can see in the unremoved shadow we have under the object’s tail of the top zoomed window.

4.6. Results on the ReNe Dataset

The ReNe dataset presents a significant challenge for previous methods due to its full scenes with backgrounds and camera poses that are mostly concentrated in a specific area rather than being distributed 360 degrees around the scene. The NVS results in Tab. 3 show our method achieves the best performance across all settings and metrics.

	Light Setting	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	MSE \downarrow
TensoIR	Single	24.57	0.6063	0.1912	0.0040
IntrinsicNeRF	Single	24.10	0.4009	0.5392	0.0039
Ours	Single	27.63	0.7210	0.1136	0.0019
TensoIR	Multiple	26.44	0.6743	0.1572	0.0028
Ours	Multiple	27.71	0.7377	0.1078	0.0018
Ours	All	27.48	0.7300	0.1113	0.0019

Table 3. NVS and relighting on the **ReNe Dataset**.

The results from the single light setting (left of Fig. 6) show that all methods struggle to achieve good performance; however, our method produces meaningful results in the first and third scenes. The other neural rendering methods, TensoIR[18] and IntrinsicNeRF [45], fail to achieve correct decomposition, primarily attributed to the failure in distinguishing intrinsic components and also the difficulty in scene reconstruction. In the multiple light setting(middle of Fig. 6), TensoIR incorrectly leaks multiple shading components into the reflectance. In contrast, our

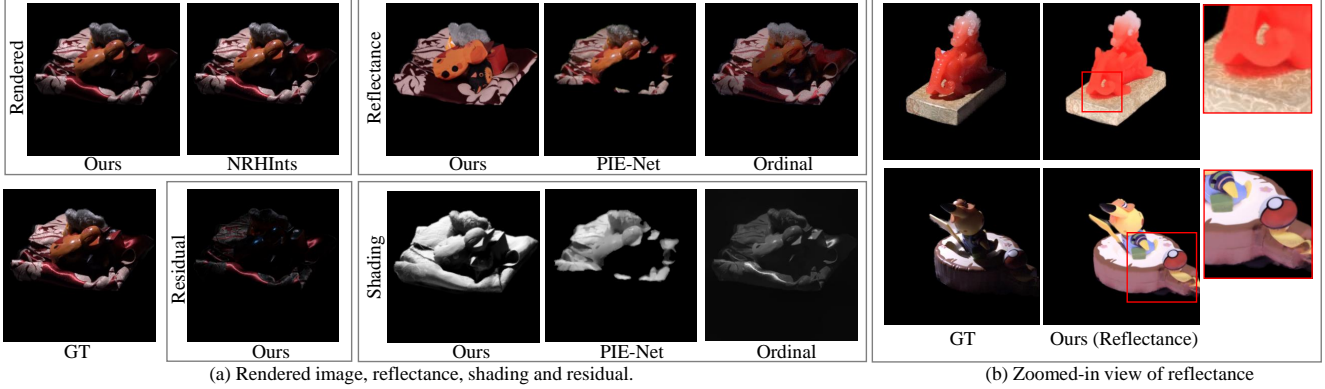


Figure 5. Qualitative Results on the **Real Object Dataset**. (a) Our method compared with NRHints for the rendered image, and PIE-Net and Ordinal for intrinsic decomposition. (b) Our reflectance estimation for two different scenes, with zoomed-in views on the object hole and cast shadow area.

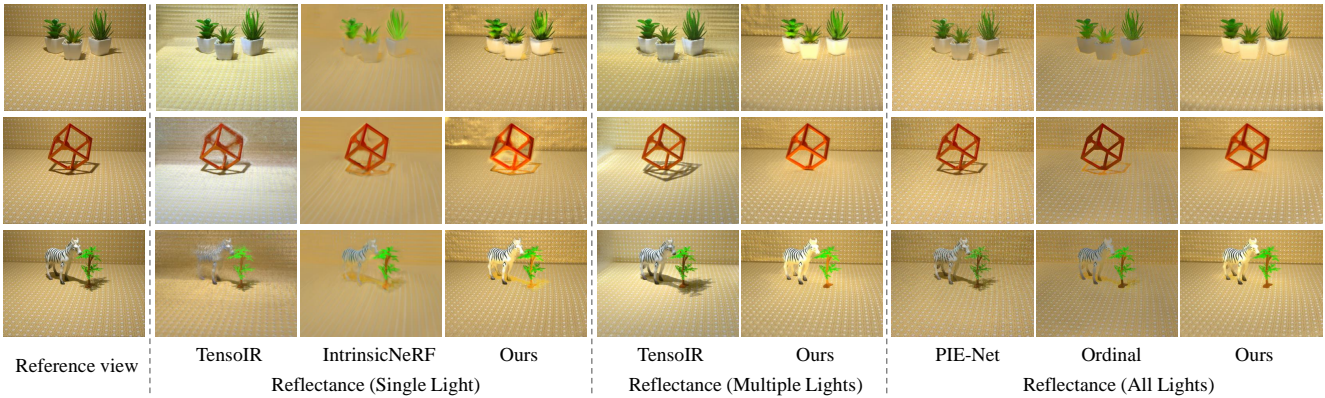


Figure 6. Qualitative Results on the **ReNe Dataset**. We show the reflectance estimation for a reference view across all settings.

method successfully combines multiple lighting conditions to produce clean reflectance. In the all lights setting, our method achieves superior results, outperforming both PIE-Net [9] and Ordinal [6], achieving sharp texture edges, vibrant colors, accurate shadow elimination, and precise background reconstruction. For more results, please refer to the supplementary material.

5. Conclusion

We propose MLI-NeRF, a Multi-light Intrinsic-aware Neural Radiance Field. MLI-NeRF generates pseudo intrinsic images for scenes under different lighting conditions, enabling the learning of intrinsic decomposition without intrinsic GT. Our experiments demonstrate that our method achieves excellent performance across different types of scenes. Our approach relies on fundamental physical principles, highlighting its potential applicability in more complex scenes. Our method enables the simultaneous synthesis of novel views, relighting, and intrinsic decomposition, providing a versatile tool for various editing applications, such as reflectance and shading modification.

Limitations and Future Work. One limitation of our

method is the need to know the light positions. As noted in our related works, many methods have proposed ways to obtain light pose information when capturing custom data. These include recording the camera pose fixed with the light source[25, 46] and using the time of day to determine the sun’s angle in outdoor settings[7]. Another limitation is the computational efficiency: training the Stage 1 model based on Neuralangelo [24] takes around 40 hours, with Stage 2 retraining adding another 19 hours. In fact, our method is applicable to different similar baselines, and we plan to migrate to a more efficient base architecture in the future.

Acknowledgments This work was supported by Grant PID2021-128178OB-I00 funded by MCIN/AEI/10.13039/501100011033 and “ERDF/EU”. It was also supported by the Departament de Recerca i Universitats from Generalitat de Catalunya, reference 2021SGR01499. Thanks to Hassan Ahmed Sial for his assistance in generating the synthetic scenes. Yixiong Yang was supported by China Scholarship Council. Haoyu Wu was supported by the NASA Biodiversity Program (Award 80NSSC21K1027) and the National Science Foundation (NSF) Grant IIS-2212046.

References

- [1] Jonathan T Barron and Jitendra Malik. Shape, illumination, and reflectance from shading. *IEEE transactions on pattern analysis and machine intelligence*, 37(8):1670–1687, 2014. [2](#)
- [2] Harry Barrow, J Tenenbaum, A Hanson, and E Riseman. Recovering intrinsic scene characteristics. *Comput. Vis. Syst.*, 2(3-26):2, 1978. [1](#), [2](#), [3](#)
- [3] 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*, pages 12684–12694, 2021. [2](#)
- [4] Mark Boss, Andreas Engelhardt, Abhishek Kar, Yuanzhen Li, Deqing Sun, Jonathan T. Barron, Hendrik P.A. Lensch, and Varun Jampani. SAMURAI: Shape And Material from Unconstrained Real-world Arbitrary Image collections. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. [2](#)
- [5] Brent Burley and Walt Disney Animation Studios. Physically-based shading at disney. In *Acm Siggraph*, pages 1–7. vol. 2012, 2012. [1](#)
- [6] Chris Careaga and Yağız Aksoy. Intrinsic image decomposition via ordinal shading. *ACM Trans. Graph.*, 2023. [1](#), [2](#), [3](#), [5](#), [7](#), [8](#)
- [7] Yeonjin Chang, Yearim Kim, Seunghyeon Seo, Jung Yi, and Nojun Kwak. Fast sun-aligned outdoor scene relighting based on tensorf. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3626–3636, 2024. [2](#), [8](#)
- [8] Ziyu Chen, Chenjing Ding, Jianfei Guo, Dongliang Wang, Yikang Li, Xuan Xiao, Wei Wu, and Li Song. L-tracing: Fast light visibility estimation on neural surfaces by sphere tracing. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. [2](#), [3](#)
- [9] Partha Das, Sezer Karaoglu, and Theo Gevers. Pie-net: Photometric invariant edge guided network for intrinsic image decomposition. In *IEEE Conference on Computer Vision and Pattern Recognition, (CVPR)*, 2022. [1](#), [2](#), [5](#), [7](#), [8](#)
- [10] Farshad Einabadi, Jean-Yves Guillemaut, and Adrian Hilton. Deep neural models for illumination estimation and relighting: A survey. In *Computer Graphics Forum*, pages 315–331. Wiley Online Library, 2021. [2](#)
- [11] Majed El Helou, Ruofan Zhou, Sabine Susstrunk, and Radu Timofte. Ntire 2021 depth guided image relighting challenge. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 566–577, 2021. [2](#)
- [12] Qingnan Fan, Jiaolong Yang, Gang Hua, Baoquan Chen, and David Wipf. Revisiting deep intrinsic image decompositions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8944–8952, 2018. [3](#)
- [13] Duan Gao, Guojun Chen, Yue Dong, Pieter Peers, Kun Xu, and Xin Tong. Deferred neural lighting: free-viewpoint relighting from unstructured photographs. *ACM Transactions on Graphics (TOG)*, 39(6):258, 2020. [2](#), [5](#)
- [14] Elena Garces, Carlos Rodriguez-Pardo, Dan Casas, and Jorge Lopez-Moreno. A survey on intrinsic images: Delving deep into lambert and beyond. *International Journal of Computer Vision*, 130(3):836–868, 2022. [1](#), [2](#), [3](#)
- [15] Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization for learning shapes. In *Proceedings of the 37th International Conference on Machine Learning*, pages 3789–3799, 2020. [3](#)
- [16] Majed El Helou, Ruofan Zhou, Sabine Susstrunk, Radu Timofte, Mahmoud Affi, Michael S Brown, Kele Xu, Hengxing Cai, Yuzhong Liu, Li-Wen Wang, et al. Aim 2020: Scene relighting and illumination estimation challenge. *arXiv preprint arXiv:2009.12798*, 2020. [2](#)
- [17] Andrew Hou, Michel Sarkis, Ning Bi, Yiyi Tong, and Xiaoming Liu. Face relighting with geometrically consistent shadows. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4217–4226, 2022. [2](#)
- [18] Haian Jin, Isabella Liu, Peijia Xu, Xiaoshuai Zhang, Songfang Han, Sai Bi, Xiaowei Zhou, Zexiang Xu, and Hao Su. Tensor: Tensorial inverse rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. [1](#), [5](#), [6](#), [7](#)
- [19] Peter Kocsis, Julien Philip, Kalyan Sunkavalli, Matthias Nießner, and Yannick Hold-Geoffroy. Lightit: Illumination modeling and control for diffusion models. In *CVPR*, 2024. [2](#)
- [20] Louis Lettry, Kenneth Vanhoey, and Luc Van Gool. Unsupervised deep single-image intrinsic decomposition using illumination-varying image sequences. In *Computer Graphics Forum*, pages 409–419. Wiley Online Library, 2018. [2](#)
- [21] Zhengqi Li and Noah Snavely. Learning intrinsic image decomposition from watching the world. In *Computer Vision and Pattern Recognition (CVPR)*, 2018. [2](#)
- [22] Zhengqi Li and Noah Snavely. Cgintrinsics: Better intrinsic image decomposition through physically-based rendering. In *European Conference on Computer Vision (ECCV)*, 2018. [2](#)
- [23] Zhengqi Li and Noah Snavely. Learning intrinsic image decomposition from watching the world. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9039–9048, 2018. [2](#)
- [24] Zhaoshuo Li, Thomas Müller, Alex Evans, Russell H Taylor, Mathias Unberath, Ming-Yu Liu, and Chen-Hsuan Lin. Neuralangelo: High-fidelity neural surface reconstruction. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. [1](#), [3](#), [6](#), [8](#)
- [25] Jingwang Ling, Zhibo Wang, and Feng Xu. Shadowneus: Neural sdf reconstruction by shadow ray supervision, 2022. [1](#), [8](#)
- [26] Isabella Liu, Linghao Chen, Ziyang Fu, Liwen Wu, Haian Jin, Zhong Li, Chin Ming Ryan Wong, Yi Xu, Ravi Ramamoorthi, Zexiang Xu, and Hao Su. Openillumination: A multi-illumination dataset for inverse rendering evaluation on real objects. In *Advances in Neural Information Processing Systems*, pages 36951–36962. Curran Associates, Inc., 2023. [2](#)

- [27] Yunfei Liu, Yu Li, Shaodi You, and Feng Lu. Unsupervised learning for intrinsic image decomposition from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3248–3257, 2020. [2](#)
- [28] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2018. [6](#)
- [29] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. [1](#), [2](#), [5](#)
- [30] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.*, 41(4):102:1–102:15, 2022. [3](#)
- [31] Lukas Murmann, Michael Gharbi, Miika Aittala, and Fredo Durand. A dataset of multi-illumination images in the wild. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4080–4089, 2019. [2](#)
- [32] Thomas Nestmeyer, Jean-François Lalonde, Iain Matthews, and Andreas Lehrmann. Learning physics-guided face relighting under directional light. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5124–5133, 2020. [2](#)
- [33] Rohit Pandey, Sergio Orts-Escolano, Chloe Legendre, Christian Haene, Sofien Bouaziz, Christoph Rhemann, Paul E Debevec, and Sean Ryan Fanello. Total relighting: learning to relight portraits for background replacement. *ACM Trans. Graph.*, 40:43–1, 2021. [2](#)
- [34] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011. [4](#)
- [35] Densen Puthussery, Melvin Kuriakose, Jiji C V, et al. Wdrn: A wavelet decomposed relightnet for image relighting. *arXiv preprint arXiv:2009.06678*, 2020. [2](#)
- [36] Viktor Rudnev, Mohamed Elgharib, William Smith, Lingjie Liu, Vladislav Golyanik, and Christian Theobalt. Nerf for outdoor scene relighting. In *European Conference on Computer Vision (ECCV)*, 2022. [2](#)
- [37] 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*, pages 7495–7504, 2021. [2](#)
- [38] Tiancheng Sun, Jonathan T Barron, Yun-Ta Tsai, Zexiang Xu, Xueming Yu, Graham Fyffe, Christoph Rhemann, Jay Busch, Paul E Debevec, and Ravi Ramamoorthi. Single image portrait relighting. *ACM Trans. Graph.*, 38(4):79–1, 2019. [2](#)
- [39] Marco Toschi, Riccardo De Matteo, Riccardo Spezialetti, Daniele De Gregorio, Luigi Di Stefano, and Samuele Salti. Relight my nerf: A dataset for novel view synthesis and relighting of real world objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 20762–20772, 2023. [1](#), [2](#), [5](#)
- [40] Li-Wen Wang, Wan-Chi Siu, Zhi-Song Liu, Chu-Tak Li, and Daniel PK Lun. Deep relighting networks for image light source manipulation. *arXiv preprint arXiv:2008.08298*, 2020. [2](#)
- [41] Yuxin Wang, Wayne Wu, and Dan Xu. Learning unified decomposition and compositional nerf for editable novel view synthesis. In *ICCV*, 2023. [1](#)
- [42] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. [6](#)
- [43] Siqi Yang, Xuanning Cui, Yongjie Zhu, Jiajun Tang, Si Li, Zhaoifei Yu, and Boxin Shi. Complementary intrinsics from neural radiance fields and cnns for outdoor scene relighting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16600–16609, 2023. [2](#)
- [44] Ziyi Yang, Yanzhen Chen, Xinyu Gao, Yazhen Yuan, Yu Wu, Xiaowei Zhou, and Xiaogang Jin. Sire-ir: Inverse rendering for brdf reconstruction with shadow and illumination removal in high-illuminance scenes. *arXiv preprint arXiv:2310.13030*, 2023. [1](#), [2](#)
- [45] 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*, 2023. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#)
- [46] Chong Zeng, Guojun Chen, Yue Dong, Pieter Peers, Hongzhi Wu, and Xin Tong. Relighting neural radiance fields with shadow and highlight hints. In *ACM SIGGRAPH 2023 Conference Proceedings*, 2023. [1](#), [2](#), [5](#), [7](#), [8](#)
- [47] Kai Zhang, Fujun Luan, Qianqian Wang, Kavita Bala, and Noah Snavely. PhysSG: Inverse rendering with spherical gaussians for physics-based material editing and relighting. In *The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. [2](#)
- [48] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. [6](#)
- [49] 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. [1](#), [2](#)
- [50] Yuanqing Zhang, Jiaming Sun, Xingyi He, Huan Fu, Rongfei Jia, and Xiaowei Zhou. Modeling indirect illumination for inverse rendering. In *CVPR*, 2022. [1](#), [2](#), [3](#), [5](#), [6](#)
- [51] Hao Zhou, Sunil Hadap, Kalyan Sunkavalli, and David W Jacobs. Deep single-image portrait relighting. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7194–7202, 2019. [2](#)