

IntrinsicNeRF: Learning Intrinsic Neural Radiance Fields for Editable Novel View Synthesis

Weicai Ye^{1*} Shuo Chen^{1*} Chong Bao¹ Hujun Bao¹ Marc Pollefeys^{2,3} Zhaopeng Cui¹

Guofeng Zhang^{1†}

¹State Key Lab of CAD&CG, Zhejiang University

²ETH Zurich

³Microsoft

https://zju3dv.github.io/intrinsic_nerf

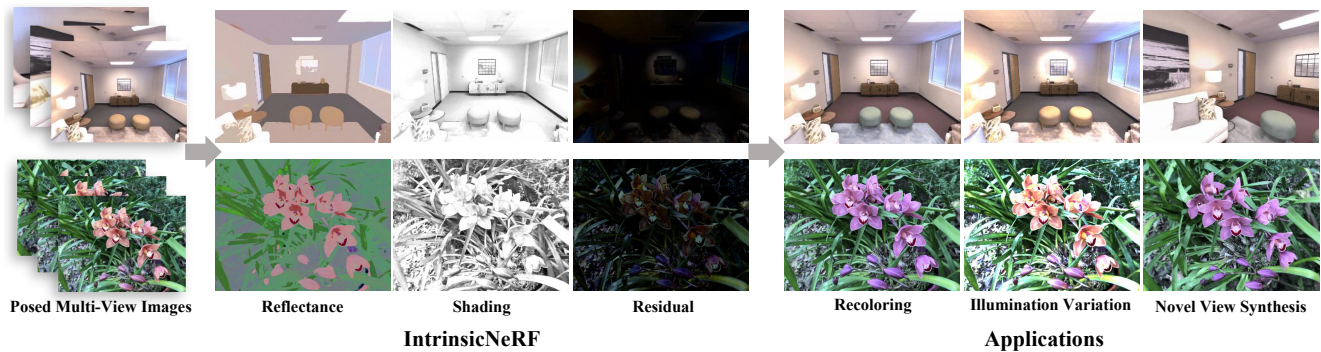


Figure 1: **Intrinsic Neural Radiance Fields (IntrinsicNeRF)**. Given multi-view posed images of static scenes, IntrinsicNeRF can factorize them into the multi-view consistent components: reflectance, shading, and residual layers. The decomposition can support online applications such as scene recoloring, illumination variation, and editable novel view synthesis.

Abstract

Existing inverse rendering combined with neural rendering methods can only perform editable novel view synthesis on object-specific scenes, while we present intrinsic neural radiance fields, dubbed *IntrinsicNeRF*, which introduce intrinsic decomposition into the NeRF-based neural rendering method and can extend its application to room-scale scenes. Since intrinsic decomposition is a fundamentally under-constrained inverse problem, we propose a novel distance-aware point sampling and adaptive reflectance iterative clustering optimization method, which enables *IntrinsicNeRF* with traditional intrinsic decomposition constraints to be trained in an unsupervised manner, resulting in multi-view consistent intrinsic decomposition results. To cope with the problem that different adjacent instances of similar reflectance in a scene are incorrectly clustered together, we further propose a hierarchical clustering method with coarse-to-fine optimization to obtain a fast hierarchical indexing representation. It supports compelling real-time augmented applications such as recoloring and illumination variation. Extensive experiments and editing samples on both object-specific/room-scale scenes

and synthetic/real-world data demonstrate that we can obtain consistent intrinsic decomposition results and high-fidelity novel view synthesis even for challenging sequences.

1. Introduction

Recently neural rendering techniques have gained increasing attention and demonstrated tremendous performance in novel view synthesis, ranging from small objects [37, 41, 46, 63] to large outdoor scenes [41, 60], but they struggle to perform further intuitive editing like realistic scene recoloring, relighting, etc, for the scenes are usually represented as neural fields implicitly and required to be decomposed into the editable properties explicitly.

Several works have proposed to fulfill this goal by introducing inverse rendering into neural rendering [74, 76, 77], where the scene is decomposed into geometry, reflectance, and illumination. However, since inverse rendering is fundamentally ambiguous and highly ill-posed, these NeRF-based inverse rendering works [74, 77] introduce many prior assumptions preventing the modeling of mutual occlusion, inter-reflection, and indirect light propagation of different objects in the scene. An accurate 3D surface recovery is also required as a prerequisite. All these factors limit their applications to object-specific scenarios.

*indicates equal contribution.

† indicates corresponding author.

To empower such editable capabilities to the scene-level neural rendering, we present intrinsic neural radiance fields, which introduce intrinsic decomposition into neural rendering, based on the fact that intrinsic decomposition can be considered as a simplification of inverse rendering designed to provide interpretable intermediate representations (i.e., reflectance and shading) that are relatively easy to solve for both in small objects and large scenes. A potential naive solution may use the trained NeRF model to generate multi-view images and then perform multi-view intrinsic decomposition, where these two tasks are separated. In contrast, extending from NeRF [46], IntrinsicNeRF (see Sec. 3.1 and Fig. 2) takes the sampled spatial coordinate point $\mathbf{x} = (x, y, z)$ and the direction $\mathbf{d} = (\theta, \phi)$ as input and regresses them into the density σ , view-independent reflectance r and shading s (Lambertian reflectance assumption) and additional view-dependent residual term re [42, 61] (Eq. 2), which naturally guarantees the multi-view consistency of decomposition after training, thanks to neural rendering.

However, it is nontrivial to design such a framework due to huge gaps in optimization between traditional intrinsic decomposition and NeRF-based methods. Traditional intrinsic decomposition methods optimize the energy equation by establishing constraints related to the image pixels, while NeRF-based methods optimize the view-dependent densities and colors of several sampled 3D points through volume rendering, which makes it hard to exploit the commonly used prior knowledge in intrinsic decomposition (see Sec. 3.2) such as chromaticity prior, reflectance sparsity, etc. To address this problem, we propose a distance-aware sampling method (see Fig. 3) that allows the sampled points not only to be random but also to establish local and global relationships between points. In this way, IntrinsicNeRF satisfies both the novel view synthesis and the better recovery of the intrinsic properties of the scene.

Moreover, to deal with the inconsistencies of similar reflectance regions [44], we present an adaptive reflectance iterative clustering method (see Sec. 3.3) with mean shift [13] to adaptively cluster color points with similar reflectance based on the scene itself, rather than K-Means used in [44], which limits the number of specific classes. A continuously updated clustering operation with the voxel grid filter is constructed to map similar reflectance colors to the same target reflectance color and then obtain the clustered category for each color point (see Fig. 5).

To settle the problem of different adjacent instances of similar reflectance in a scene being clustered together, we propose a semantic-aware reflectance sparsity constraint during training. Inspired by Semantic-NeRF [79], we add an additional semantic branch to IntrinsicNeRF, along with reflectance clustering, which yields a hierarchical reflectance iterative clustering and indexing method (see Fig. 6), optimizing the network from coarse to fine. Ex-

tensive experiments on the Blender Object and the Replica Scene dataset demonstrate our method can obtain consistent intrinsic decomposition results and high-fidelity novel view synthesis even for challenging sequences. We also develop video editing software to facilitate users to perform online scene recoloring, illumination variation, and editable novel view synthesis on both real-world and synthetic data on the CPU (see Fig. 1).

2. Related Work

Intrinsic Image Decomposition. Intrinsic decomposition [2] is a typical image layer separation problem aimed at decomposing images into reflectance, shading, etc., and has been studied for decades. To deal with this ill-posed problem, additional priors [21, 30, 56] with optimization framework have been used. Recently, deep learning methods [3, 17, 36, 40, 73, 81] have emerged to perform intrinsic decomposition, and with large datasets [34, 35, 53], they have shown further improvement. Unsupervised intrinsic image decomposition works [22, 38] have also achieved impressive results. IntrinsicNeRF considers not only the intrinsic decomposition prior but also the consistency of different perspectives in neural rendering, performing unsupervised optimization of the network.

Intrinsic Video Decomposition. Intrinsic video decomposition extends intrinsic decomposition from the image domain to the video domain and can be roughly divided into two types. One is to perform the intrinsic image decomposition first and use the motion information to establish the correlation between frames for post-processing [8, 29, 67]. The other is to directly unify the image’s local and global relations using some prior, by optimizing the energy equation [7, 44]. There are also works [16, 23, 28, 70] on intrinsic decomposition from multi-view images. These methods have some consistency in intrinsic video decomposition but are unable to perform novel view synthesis. While IntrinsicNeRF introduces traditional intrinsic decomposition prior to the neural radiance fields to achieve end-to-end optimization, which not only performs better intrinsic video decomposition than previous methods but also allows for realistic editable novel view synthesis.

Inverse Rendering. Inverse rendering [19] is another way to restore the basic properties of scene, which can be broadly divided into two categories: classic approaches [5, 50, 24], differentiable renders [32, 49, 78, 39] methods. Plenty of works combining neural rendering with inverse rendering [6, 9, 54, 65, 74, 76, 77, 80] have shown realistic view synthesis and consistent estimation of the underlying properties of the scenes. Among them, PhySG [74] and Invrender [77] rely heavily on precise geometry, limiting their applicability to specific objects. While IntrinsicNeRF introduces intrinsic decomposition into neural rendering and expands the capabilities of editable novel view synthesis from




Figure 2: **IntrinsicNeRF Framework.** IntrinsicNeRF takes the sampled spatial coordinate point and direction as input and outputs the density, reflectance, shading, and residual term. The semantic branch is optional. Unsupervised Prior and Reflectance Clustering are exploited as the loss function constraints to train the IntrinsicNeRF in an unsupervised manner.

individual objects to room-scale scenes.

3. Method

Given multi-view posed images under unknown illumination of static scenes, we aim to achieve a reliable understanding of the basic properties of the scene, such as reflectance, shading, etc, and to enable real-time editable novel view synthesis. Fig. 2 outlines the general framework.

3.1. Intrinsic Neural Radiance Fields

Preliminaries: Intrinsic Decomposition. Lambertian and grayscale shading assumptions [17] are commonly used and introduced to simplify this inverse problem, achieving good approximations of most scenarios. Based on Lambertian assumption, Intrinsic decomposition [17] presents an input image I as the pixel-wise product of the illumination-invariant reflectance $R(I)$, and the illumination-varying shading $S(I)$:

$$C(I) = R(I) \odot S(I), \quad (1)$$

where \odot is channel-wise multiplication. However, the Lambertian assumption is difficult to be satisfied in realistic scenes, and the intrinsic residual model [61, 42] introduces view-independent reflectance and shading with an additional view-dependent residual term $Re(I)$ to model scenes that do not satisfy the Lambertian assumption, such as glossy reflections, metallic materials:

$$C(I) = R(I) \odot S(I) + Re(I). \quad (2)$$

Our representation. IntrinsicNeRF takes the sampled coordinate point $\mathbf{x} = (x, y, z)$ and direction $\mathbf{d} = (\theta, \phi)$ as input, and outputs the view-independent reflectance r and shading s , the view-dependent intrinsic residual term re and the volume density σ through an MLP network F_{Θ} :

$$(r, s, re, \sigma) = F_{\Theta}(\mathbf{x}, \mathbf{d}). \quad (3)$$

The predicted color c of each spatial point can be obtained by Eq. 2 and the target color $C(\mathbf{r})$ of camera ray \mathbf{r} is:

$$\hat{C}(\mathbf{r}) = \sum_{k=1}^K \hat{T}_k \alpha_k c_k \text{ and } \hat{T}_k = \exp\left(-\sum_{k'=1}^{k-1} \sigma_k \delta_k\right), \quad (4)$$

where $\alpha_k = 1 - \exp(-\sigma_k \delta_k)$, and δ_k is the distance between two adjacent sampled points. We follow NeRF's coarse-to-fine training policy and train IntrinsicNeRF from scratch with the photometric loss L_{pho} in NeRF [46].




Figure 3: **Distance-Aware Point Sampling.** We first randomly sample 512 points, and then randomly sample the remaining 512 points in the eight neighborhoods of each sampled point to construct the unsupervised constraint term for the intrinsic decomposition.

Distance-Aware Point Sampling. NeRF [46] randomly samples batches of camera rays from the image pixel set (roughly 1024 points) in each optimization, where these points are random, and no relationship is established between them. It is not applicable in IntrinsicNeRF, for the introduction of ill-posed intrinsic decomposition into NeRF makes the whole optimization process stochastic, as shown in Fig. 8 (Baseline column). To this end, we make a sophisticated design of the sampling policy (see Fig. 3) which helps to construct intrinsic prior constraints (see Sec. 3.2), and the network can be trained unsupervised.

3.2. Unsupervised Prior Guidance

Following intrinsic decomposition works [44], we adopt the grayscale shading assumption to simplify this inverse problem, so that the shading layer is single-channel and the

reflectance chromaticity of the image I is approximated to $c(\mathbf{x}) = I(\mathbf{x})/|I(\mathbf{x})|$. We define the chromaticity similarity weight $\omega_{cs}(\mathbf{x}, \mathbf{y})$ [44] that is associated with many priors:

$$\omega_{cs}(\mathbf{x}, \mathbf{y}) = \exp(-\alpha_{cs} \|c(\mathbf{x}) - c(\mathbf{y})\|_2^2), \quad (5)$$

where \mathbf{x} and \mathbf{y} are the image pixel coordinates. Coefficient $\alpha_{cs} = 60$ produces the best decomposition results.

Chromaticity Prior. Due to the residual term, the chromaticity of the unknown reflectance R and the input image I are not the same. We want them to be as close as possible:

$$L_{chrom}(\mathbf{x}) = \|c_r(\mathbf{x}) - c(\mathbf{x})\|_2^2, \quad (6)$$

where c and c_r are the chromaticity of the input sample points and the sampled points' reflectance, respectively.

Reflectance Sparsity. Two pixels that are similar in spatial location and chromaticity, have converging reflectance r , which leads to reflectance sparsity. Following [44], we minimize the reflectance gradients magnitude independently:

$$L_{reflect}(\mathbf{x}) = \sum_{\mathbf{y} \in \mathcal{N}(\mathbf{x})} \omega_{cs}(\mathbf{x}, \mathbf{y}) \|r(\mathbf{x}) - r(\mathbf{y})\|_2^2, \quad (7)$$

where $\mathcal{N}(\mathbf{x})$ is the neighbourhood of pixel \mathbf{x} . Specifically, in IntrinsicNeRF, the sampled points in the first half will be adjacent to the second half, shown in Fig. 3.

Non-Local Reflectance Sparsity. In natural and man-made scenes, two distant spatial points may also have the same reflectance, such as a wall and floor that occupy a larger image area, which requires non-local reflectance sparsity. In the sampling of IntrinsicNeRF, the first half of the points are randomly sampled, so the distance between any two points can be very far. We simply bisect the first half of the points and construct a non-local reflectance sparsity constraint (following [44]) on the points in the first 1/4 segment and the corresponding points in the next 1/4 segment:

$$L_{non-local}(\mathbf{x}) = \sum_{\mathbf{y} \in \mathcal{F}(\mathbf{x})} \omega_{cs}(\mathbf{x}, \mathbf{y}) \|r(\mathbf{x}) - r(\mathbf{y})\|_2^2, \quad (8)$$

where $\mathcal{F}(\mathbf{x})$ is the farhood of pixel \mathbf{x} . Note that the weight of this constraint is smaller than the reflectance sparsity's.

Shading Smoothness. Natural objects usually have smooth surfaces and the shading variance is expected to be smooth [44]. Moreover, neighboring pixels with different chromaticities, represent a reflectance edge, so we strongly enforce the shading smoothness:

$$L_{shade}(\mathbf{x}) = \sum_{\mathbf{y} \in \mathcal{N}(\mathbf{x})} \|c(\mathbf{x}) - c(\mathbf{y})\|_2^2 \|s(\mathbf{x}) - s(\mathbf{y})\|_2^2. \quad (9)$$

Intrinsic Residual Constraints. Since diffuse light generally dominates the scene, we want the image content to be recovered by reflectance and shading as much as possible. This prevents extreme cases when R and S both converge to zero and $Re = I$, which would destroy the efficacy of the previous constraints and fall into catastrophic results (see Fig. 4). We set this constraint as follows:

$$L_{residual}(\mathbf{x}) = \|re(\mathbf{x})\|_2^2. \quad (10)$$

The weight is set higher early, so $R(I) \odot S(I)$ is close to the target image I and then dropped later. As the output of R and S is stable, R_e can represent the view-dependent components, such as glossy reflections.




Figure 4: **Ablation Study of $L_{residual}$.** Without the residual constraint, the reflectance gets worse (see Tab. 2), and the shading and residual are exceptionally unfavorable.

Intensity Prior. The previous constraints on reflectance and shading only consider the relative relationship between two pixels. The absolute magnitude of R and S is required to prevent them from falling into certain extremes during optimization. The intensity of the unknown reflectance image R and the input image I should be close:

$$L_{intensity}(\mathbf{x}) = \|i_r(\mathbf{x}) - i(\mathbf{x})\|_2^2, \quad (11)$$

where i and i_r are the average intensities of the sampled points \mathbf{x} of the input image and reflectance r . The weight of this constraint is set higher early and then reduced.

3.3. Adaptive Reflectance Iterative Clustering

Although reflectance sparsity makes sense to some extent, there still remain inconsistencies of similar reflectance regions (see Fig. 8 Baseline+w/prior), therefore we propose an adaptive reflectance iterative clustering method by constructing a continuously updated clustering operation G , which maps similar reflectance colors $r(\mathbf{x})$ to the same target reflectance color $r_{cluster}(\mathbf{x})$ by adding a clustering constraint during the optimization of the network:

$$L_{cluster}(\mathbf{x}) = \|r_{cluster}(\mathbf{x}) - r(\mathbf{x})\|_2^2. \quad (12)$$

Next, we elucidate the detail of the clustering method.

RGB Transform. During the training of the network, we infer the reflectance r , shading s , and residual term re of multi-view posed images after every 10K iterations. Refer

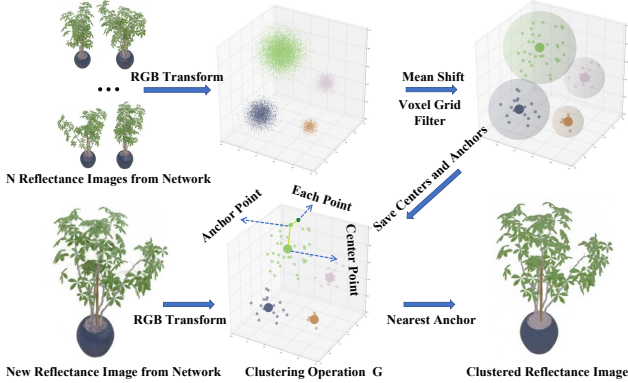


Figure 5: **Adaptive Reflectance Iterative Clustering Method.** The color of the reflectance pixels is first converted and then clustered with the Mean Shift algorithm. The voxel grid filter is performed to accelerate the processing of the fast approximation of clustering operation G .

to I1W [4], we take out all pixels of all r components and convert their colors to better cluster reflectances (pixel intensity, red chromaticity, green chromaticity [44]):

$$f([\mathbf{r}, \mathbf{g}, \mathbf{b}]) = \left[\beta \frac{\mathbf{r} + \mathbf{g} + \mathbf{b}}{\mathbf{3}}, \frac{\mathbf{r}}{\mathbf{r} + \mathbf{g} + \mathbf{b}}, \frac{\mathbf{g}}{\mathbf{r} + \mathbf{g} + \mathbf{b}} \right], \quad (13)$$

where β is set as 0.5 [4] in our experiment. The RGB transformation helps reduce the effect of intensity differences on the clustering, making the clustering more focused on the similarity of chromaticity between two pixels. The transformed RGB space is considered as f space.

Mean Shift. Unlike existing methods [44] using K-Means clustering to specify K clustering categories, we instead cluster all the pixel points P every 10K iterations with a Mean Shift clustering algorithm to adaptively determine the number of reflectance classes in the scene, for we do not know the reflectance class number.

Clustering Operation G . After Mean Shift clustering, we get a set of clustered centers, and a classification label for each pixel point in P . During each training iteration, it is unrealistic to cluster the reflectance of each rendered pixel because the clustering is time-consuming. So we define a fast approximation clustering operation G : for an RGB value of reflectance, it considers the category of the nearest point in P as its cluster category and set the value of the category center as the target clustered reflectance $r_{cluster}(\mathbf{x})$. When calculating the clustering loss, we only use the Clustering Operation G , shown in Fig. 5.

Voxel Grid Filter. Since there are plenty of points P in the f space and most of them are clustered in very small regions due to reflectance sparsity, rather than finding the nearest neighbors in all points, we perform voxel grid filter (voxel size is 0.01) on the points P in the f space, and the filtered points are regarded as anchor points. The cluster-




Figure 6: **Hierarchical Clustering and Indexing Method.**

ing operation G therefore only needs to search the closest anchor point, and the anchor points are only been updated every 10K iterations by Mean-Shift.

Optimization. During the network optimization, the weight of the clustering loss $L_{cluster}(\mathbf{x})$ and the bandwidth parameter in the mean shift algorithm are gradually increased with the number of iterations (the larger the bandwidth is, the smaller the number of mean-shift clustering categories is). That is because, in the early stage of network optimization, the inferred reflectance r is not reliable and needs lower weight. While in the later stage, a higher weight can lead the output of the network to converge toward the effect of clustering, making the reflectance r before and after clustering indistinguishable.

3.4. Hierarchical Clustering and Indexing

The adaptive reflectance iterative clustering method can handle object-level scenes well, shown in Fig. 8 (Ours). However, when the reflectance in the scene is complex and similar, plenty of different instances with similar reflectance in room-scale scenes may be incorrectly clustered, shown in Fig. 9 (w/prior+cluster). So we propose a semantic-aware reflectance sparsity constraint, where only pixels with the same semantic label will be computed, thus significantly improving the quality of reflectance. Inspired by [79], we extend IntrinsicNeRF to jointly encode appearance, geometry, and semantics by adding a segmentation renderer to the original IntrinsicNeRF. Specially, we use a view-invariant MLP function $sl = F_{\Theta}(\mathbf{x})$ to map a spatial coordinate \mathbf{x} to semantic label and use the semantic loss L_{sem} in [79].

Depending on the semantic labels of each pixel, the pixel set P can be divided into multiple subsets $\{\mathbf{P}_1, \mathbf{P}_2, \dots, \mathbf{P}_N\}$, where N is the number of semantic categories. Then we can construct N clustering operations $\{\mathbf{G}_1, \mathbf{G}_2, \dots, \mathbf{G}_N\}$ as Sec. 3.3. The hierarchical clustering operation takes the reflectance RGB value of each pixel and the corresponding semantic label as input and outputs the result of the clustering operation for the pixel under the semantic label. Such a hierarchical clustering method allows the clustered information of each pixel to be stored in a tree structure, shown in Fig. 6, which can be indexed quickly.

Method	Reflectance (Invrender dataset)					View Synthesis (Invrender dataset)			Reflectance (our dataset)					View Synthesis (our dataset)		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	MSE \downarrow	LMSE \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	MSE \downarrow	LMSE \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
IIW[4]	22.0284	0.9307	<u>0.0847</u>	0.0099	0.0120	-	-	-	20.5299	0.9079	0.1131	0.0102	0.0727	-	-	-
CGIntrinsic[34]	20.1583	0.9209	0.0996	0.0129	<u>0.0141</u>	-	-	-	18.3542	0.8999	0.1229	0.0156	0.0659	-	-	-
USI3D [38]	20.7571	0.9267	0.0887	0.0079	0.0149	-	-	-	19.1489	0.9115	<u>0.1070</u>	0.0135	0.0524	-	-	-
NeRFactor[76]	19.9167	0.9156	0.1354	0.0059	0.0210	23.0133	0.9277	0.0822	21.4440	0.9170	0.1055	0.0063	0.0444	20.6880	0.8733	0.1185
PhySG[74]	23.3748	0.9231	0.1092	0.0034	0.0396	25.4225	0.9388	0.0804	-	-	-	-	-	-	-	-
Invrender [77]	26.3078	0.9380	0.0572	<u>0.0022</u>	0.0226	29.3870	0.9522	0.0505	-	-	-	-	-	-	-	-
Baseline	16.3209	0.8637	0.1301	0.0254	0.1955	34.0036	0.9670	0.0252	14.8572	0.8397	0.1738	0.0451	0.1849	28.2604	0.9383	0.0339
Baseline + w/ prior.	21.7370	0.9278	0.1086	0.0055	0.0186	<u>33.4909</u>	<u>0.9638</u>	<u>0.0304</u>	<u>20.9646</u>	<u>0.9140</u>	0.1216	0.0095	0.0538	<u>28.0633</u>	<u>0.9370</u>	<u>0.0369</u>
Ours	<u>24.2642</u>	<u>0.9371</u>	0.0880	0.0021	0.0173	33.4967	0.9630	0.0306	22.5677	0.9267	0.0975	<u>0.0066</u>	<u>0.0474</u>	27.9494	0.9357	0.0372

Table 1: **Quantitative Results of the Blender Object Dataset.** For reflectance estimation, IntrinsicNeRF achieved the best results on our dataset and ranked 2nd on the Invrender dataset. For novel view synthesis, IntrinsicNeRF achieved the best performance on both datasets, while Invrender [77] and PhySG [74] require good geometric prerequisites, which makes them fail on our dataset. Moreover, intrinsic decomposition methods can not perform novel view synthesis. - means failure.

Method	w/o L_{chrom}	w/o $L_{reflect}$	w/o $L_{non-local}$	w/o L_{shade}	w/o $L_{cluster}$	w/o $L_{residual}$	w/o $L_{intensity}$	w/o all prior	Ours
PSNR \uparrow	22.0243	22.4955	23.3032	22.9874	21.3508	21.1288	18.7466	15.5891	23.4160
MSE \downarrow	0.0067	0.0060	0.0044	0.0048	0.0075	0.0074	0.0172	0.0352	0.0043
LMSE \downarrow	0.0392	0.0378	0.0323	0.0338	0.0362	0.0387	0.0339	0.1902	0.0323

Table 2: **Ablation Studies of Each Loss Constraints for Reflectance Estimation on the Blender Object Dataset.**

3.5. Implementation Details

We implement IntrinsicNeRF on the top of SemanticNeRF [79] with additional three FC layers for intrinsic components which have 128 neurons. The network is optimized with photometric loss, semantic loss, unsupervised prior constraints, and clustering loss jointly. The final loss is:

$$\begin{aligned}
L_{final} = & \lambda_{pho}L_{pho} + \lambda_{sem}L_{sem} + \lambda_{chrom}L_{chrom} \\
& + \lambda_{reflect}L_{reflect} + \lambda_{non-local}L_{non-local} \\
& + \lambda_{shade}L_{shade} + \lambda_{cluster}L_{cluster} \\
& + \lambda_{residual}L_{residual} + \lambda_{intensity}L_{intensity}. \quad (14)
\end{aligned}$$

Here, $\lambda_{pho} = 1$, $\lambda_{sem} = 0.04$, $\lambda_{chrom} = 1$, $\lambda_{reflect} = 0.01$, $\lambda_{non-local} = 0.005$ and $\lambda_{shade} = 1$. While $\lambda_{cluster} = 10^{-2(1-iter/200K)}$, it exponentially increases from 0.01 to 1 every 10K iterations. We set $\lambda_{residual} = 1$ in the early 100K iterations and dropped to 0.02 in the later iterations. The $\lambda_{intensity}$ is set to 0.1 in the first 50K iterations and then set to 0.01. The batch size of the rays is 1024. The Adam [25] optimizer is used with a learning rate of $5e-4$ for 200K iterations. Tab. 3 shows the acceptable clustering and the total training time of our method.

4. Experiments

We first make qualitative and quantitative comparisons of IntrinsicNeRF with traditional optimization-based [4] and learning-based [34, 38] intrinsic decomposition methods, and neural rendering methods [76, 74, 77] combined with inverse rendering on synthetic object dataset in Sec. 4.2. Then we only compare qualitative results on synthetic scenes (e.g. Replica [57]) in Sec. 4.3, due to the lack

Dataset	Method	Training	Test	Clustering
Blender Object	NeRFactor [76]	5.7d	5.65s	-
Blender Object	PhySG [74]	5.2h	2.92s	-
Blender Object	InvRender [77]	11.4h	14.25s	-
Blender Object	NeRF [46]	5.4h	4.59s	-
Blender Object	Ours	6.5h	5.35s	39s
Replica Scene	Semantic-NeRF [79]	13.5h	2.50s	-
Replica Scene	Ours	17.5h	2.79s	220s

Table 3: **Time Comparison.** We show the total training time, the average synthesis time of each frame, and the average clustering time of our method, where **the clustering is performed every 10K training iterations.** All run on a single RTX3090.

of ground-truth labels. Finally, we perform ablation studies in Sec. 4.4 to analyze the design of our framework and demonstrate its applicability in Sec. 4.5 to both synthetic and real-world data.

4.1. Dataset

Synthetic Data. We collect 8 Blender Object dataset (4 from Invrender [77], and 4 from NeRF [46]) and 8 Replica Scene dataset. The Invrender dataset contains Hotdogs, Jugs, Chair, and Air balloons, and each dataset is rendered by Blender Cycles [15] with their masks, reflectance, and roughness maps. The NeRF dataset contains 4 objects (Lego, Drums, Ficus, and Chair2) that maintain complex geometry and realistic non-Lambertian materials. Note that some environment lighting maps in NeRF’s open-source blender model were missing, we search for some environment maps that look as realistic as possible and re-render the new image to match NeRF’s settings. We regard this dataset as our dataset. The image resolution is set as 400×400 .

Generated by Semantic-NeRF [79], each Replica

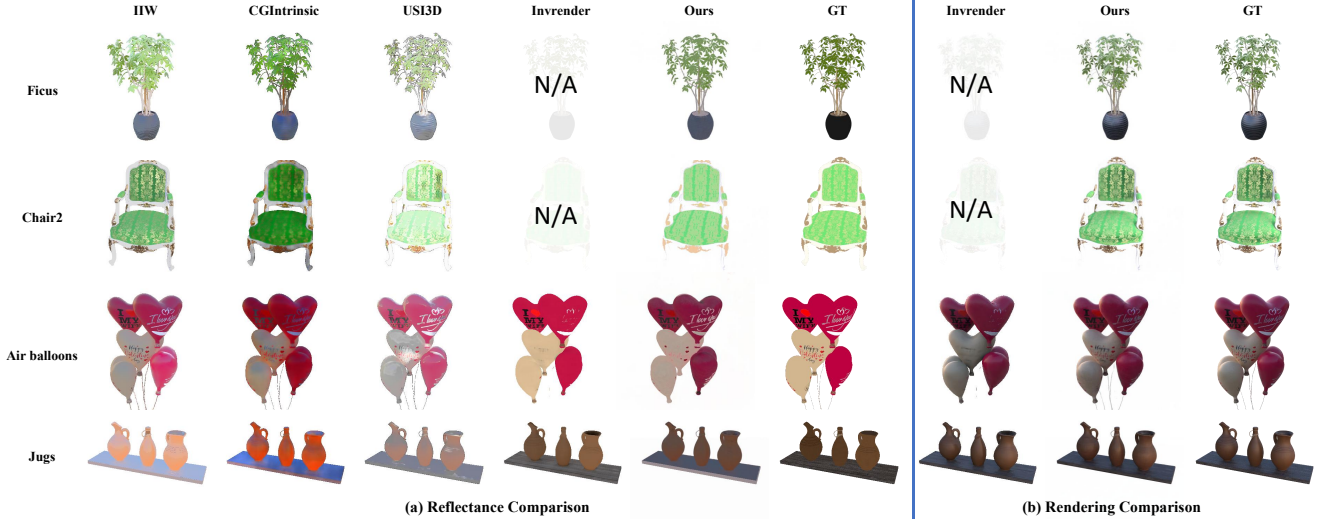


Figure 7: **Qualitative Comparison Results of Reflectance and Rendering on the Blender Object Dataset.** The top 2 rows represent our samples and the bottom 2 rows are the Invrender samples. Our method can perform reflectance estimation and novel view synthesis on both datasets well, while Invrender [77] fails to do that on our dataset. N/A means failure.

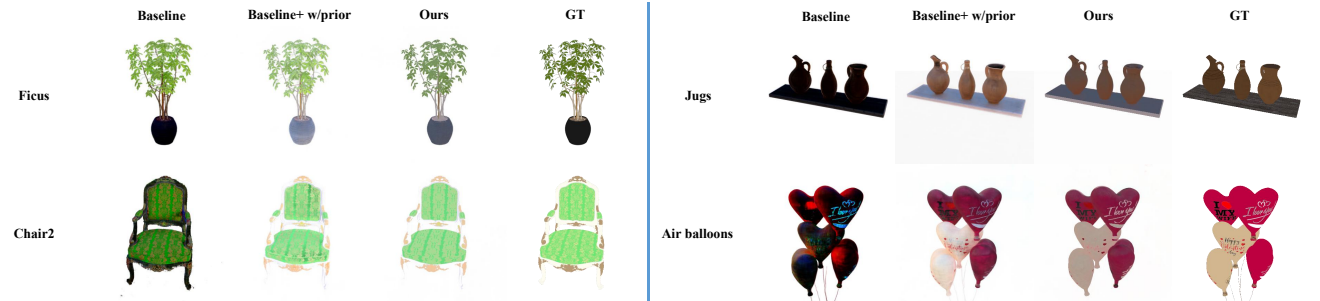


Figure 8: **Ablation study of Reflectance Estimation Sample on the Blender Object Dataset.** Left: our dataset, right: Invrender dataset. The reflectance estimation of the baseline method is stochastic and unstable, while the intrinsic prior makes the optimization of the network traceable. Our final model achieves more plausible reflectance results.

Scene [57] of rooms and offices consists of RGB images, depth maps, and semantic labels at resolution 320x240 from randomly generated 6-DOF trajectories. It contains challenging illumination effects, such as glossy reflections.

Real-world Data. We selected 4 real data of natural scenes (Orchids, Flowers, Horns, and Ferns) at 504x378 resolution from LLFF [45] to demonstrate the generalization ability of our method in real-world lighting and reflection and its applicability such as recoloring, illumination variation.

4.2. Comparison on the Blender Object Dataset

We exploit Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index Measure (SSIM), Learned Perceptual Image Patch Similarity (LPIPS) [75], Mean Squared Error (MSE) and Local Mean Squared Error (LMSE) as reflectance evaluation metrics. We do not evaluate the shading quantitatively because different methods model lighting differently, and we cannot get the ground-truth shading of our model in Blender (Eq. 2). In contrast, reflectance is a

Method	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
NeRF [46]	31.0838	0.9525	0.0302
Ours	30.7230	0.9494	0.0339

Table 4: **Comparable Quantitative Results for Novel View Synthesis on the Blender Object Dataset.**

Method	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	mIoU \uparrow
[79]	30.9770	0.8955	0.1066	0.9725
Ours	30.7044	0.8908	0.1140	0.9702

Table 5: **Comparable Results for View Synthesis and Semantic Segmentation on the Replica Scene Dataset.**

common output and has ground-truth values, so we focus on the evaluation of reflectance. The view synthesis evaluation metrics are PSNR, SSIM, and LPIPS.

We compare IntrinsicNeRF with the following methods: IIW [4] is a classic intrinsic decomposition method




Figure 9: **Qualitative Reflectance Comparison Samples with Previous Methods on the Replica Scene Dataset.** Experiments demonstrate the progressive facilitation effect of our different variants. Compared with previous methods, our final method achieves more plausible and consistent reflectance estimation results, retaining the boundaries of objects.

that does not require training. CGIntrinsic [34] is a learning method with good generalization trained on large-scale datasets, and USI3D [38] is another state-of-the-art unsupervised learning method, and we use their pre-trained models. We do not choose IRISformer [81] and intrinsic video decomposition methods [44, 43], because their codes are not available. NeRFactor [76], PhySG [74], and InvRender [77] are the neural rendering methods, and we have retrained them in the same setting for fair comparisons. Tab. 1 shows our method achieves the best results on our dataset and ranked 2nd on the Invrender dataset for reflectance estimation. Compared with single-view intrinsic decomposition methods, our method yields more consistent and plausible decomposition results, even in challenging object scenes, such as Chair2, and Ficus. As for view synthesis, IntrinsicNeRF achieves the best performance on both datasets, while Invrender [77] and PhySG [74] require good geometric prerequisites using IDR method [66], which makes them fail on our dataset, as shown in Fig. 7. Moreover, traditional intrinsic decomposition methods can not perform novel view synthesis. Tab. 4 shows IntrinsicNeRF achieves comparable novel view synthesis results with NeRF [46] while giving the power of modeling the intrinsic components of scenes.

4.3. Comparison on the Replica Scene Dataset

We only compare qualitative results with intrinsic decomposition methods [38, 4, 34] on the Replica Scene dataset in reflectance estimation, because we cannot obtain the ground truth of reflectance. Fig. 9 shows that we can obtain more plausible results than other intrinsic decomposition methods, and maintain consistent reflectance estimation for multi-view images in the supplementary material. Moreover, our method obtains comparable results with Semantic-NeRF [79] in novel view synthesis and semantic segmentation (the metric is mIOU), shown in Tab. 5, and we give Semantic-NeRF the ability to model the intrinsic properties of the scene (Fig. 1). While PhySG [74] and Invrender [77] fail to do that in room-scale scenes.

4.4. Ablation Studies

We perform ablation studies to analyze three components of our method that primarily affect the intrinsic decomposition quality. The baseline method is the NeRF [46] variant with intrinsic neural radiance fields, using the proposed distance-aware point sampling policy. Tab. 1 shows that the introduction of the intrinsic prior and iterative clustering leads to more accurate reflectance estimation, with a slight decrease in the accuracy of the novel view synthesis. Fig. 8 show that the reflectance estimated by the baseline method is more stochastic and unstable. While adding the intrinsic prior, the network output is plausible. The adaptive reflectance iterative clustering method can make the reflectance regions of the same material cluster together. The average results of the Blender Object dataset in Tab. 2 show the effectiveness of our method and the necessity of each loss. However, reflectance clustering may lose some distinguishable boundaries in room-scale scenes such as Replica Scene, for complex and similar reflectance may be clustered incorrectly. Whereas the hierarchical clustering method with semantic constraints can retain the boundaries and still yields more plausible results, as shown in Fig. 9. See more qualitative comparison results of IntrinsicNeRF variants in different scenarios in the supplementary material.

Method	IIW [4]	CGIntrinsic [34]	USI3D [38]	Ours
PSNR \uparrow	16.1219	17.9740	12.8545	15.4692
MSE \downarrow	0.0263	0.0159	0.0655	0.0345
LMSE \downarrow	0.1498	0.1342	0.1675	0.1415

Table 6: Reflectance Comparison on the EDEN dataset.




Figure 10: Reflectance and Recoloring on the EDEN dataset.

4.5. Applications

We demonstrate the applicability of IntrinsicNeRF with its decomposed components and the novel view synthesis results on both synthetic and real-world data.

Scene Recoloring. In IntrinsicNeRF, the predicted reflectance is saved as [Semantic category, reflectance category] in the hierarchical iterative clustering and indexing method. We can simply modify the color of a certain reflectance category, the reflectance values of all pixels belonging to the selected category can be modified at the same time, and then the edited images can be reconstructed using the modified reflectance with the original shading and residual through Eq. 2. Fig. 11 shows some recoloring examples on both synthetic and real-world data.

Illumination Variation. The decomposed residual term beyond the Lambertian assumption, can represent the properties such as glossy illumination, we can adjust the overall brightness directly by a multiplicative factor. Fig. 12 shows the effect of different light intensities after enhancing or diminishing the light. Please see more edited samples and the novel view synthesis results in the supplementary material.

5. Limitations and Future Work

The main limitation is that when the scenario does not conform to unsupervised intrinsic prior, it will struggle to obtain the correct decomposition results. A refinement method based on intrinsic decomposition prediction is required. Clustering errors may occur when the reflectance in a scene is complex and similar, especially in the real-world lacking semantic constraints. This can be solved by unsupervised semantic segmentation [20]. Estimating the reflectance requires a trade-off between preserving the texture and modeling the shadows correctly. Although our method performs well on room-scale scenes, such as Replica Scene, when applied to outdoor scenes with large scene sizes and fewer images, e.g. EDEN [31] under various lighting (clear and sunset), IntrinsicNeRF may lead to detail loss in rendering results (Fig. 10). Meanwhile, the reflectance estimation may fall into local optimality due to limited observations, shown in Tab. 6 (Note that [74, 77] all fail). This can be addressed by combining our method with outdoor NeRF works [54, 60]. Although IntrinsicNeRF gives NeRF the ability to model the basic properties of scenes, it retains other shortcomings of NeRF. Given the high degree of integration of our approach with NeRF, NeRF extensions can be seamlessly incorporated into our IntrinsicNeRF, such as NeRF in the wild [12, 46, 59], NeRF in dynamic environments [33, 51, 52, 69], fast NeRF [48, 18, 10, 71], NeRF with generalization [11, 64, 72, 27], generative NeRF [55, 62], NeRF with panoptic segmentation [26, 68], NeRF-based SLAM [47, 58, 82], Geometry and Texture Editing with NeRF [1, 14] etc, which will be helpful to the com-




Figure 11: Recoloring on Synthetic/Real-World Data.




Figure 12: Illumination Variation on Replica Scene.

munity. Another interesting direction is to unify intrinsic decomposition and inverse rendering to construct a hierarchical representation of the intrinsic properties of the scene. Since our approach yields multi-view consistent intrinsic decomposition results, IntrinsicNeRF can improve the performance of the intrinsic decomposition method by providing more datasets with pseudo-Ground-Truth labels for the intrinsic decomposition task. We leave this as future work.

6. Conclusion

We introduce intrinsic decomposition into neural rendering and propose intrinsic neural radiance fields that can decompose the images into reflectance, shading, and residual layers. Several techniques are proposed to make decomposition learning feasible and support online augmented applications such as recoloring, illumination variation, and editable novel view synthesis. We believe our method is the step toward the intrinsic decomposition (beyond Lambertian assumption) of more general scenes with neural rendering and will inspire follow-up work.

Acknowledgments. The authors thank Yuanqing Zhang for providing us with the pre-trained model of InvRender [77], Jiarun Liu for reproducing the results of NeRFactor [76] and PhysSG [74]. We thank Hai Li and Jundan Luo for proof-reading the paper. This work was partially supported by NSF of China (No. 61932003) and ZJU-SenseTime Joint Lab of 3D Vision. Weicai Ye was partially supported by China Scholarship Council (No. 202206320316).

References

- [1] Chong Bao, Yinda Zhang, Bangbang Yang, Tianxing Fan, Zesong Yang, Hujun Bao, Guofeng Zhang, and Zhaopeng Cui. SINE: Semantic-driven Image-based NeRF Editing with Prior-guided Editing Field. In *Proceedings of the IEEE/CVF Computer Vision and Pattern Recognition Conference*, 2023.
- [2] Harry Barrow, J Tenenbaum, A Hanson, and E Riseman. Recovering Intrinsic Scene Characteristics. *Computer Vision Systems*, 2(3-26):2, 1978.
- [3] Anil S Baslamisli, Thomas T Groenestegge, Partha Das, Hoang-An Le, Sezer Karaoglu, and Theo Gevers. Joint Learning of Intrinsic Images and Semantic Segmentation. In *Proceedings of the European Conference on Computer Vision*, pages 286–302, 2018.
- [4] Sean Bell, Kavita Bala, and Noah Snavely. Intrinsic Images in the Wild. *ACM Transactions on Graphics*, 33(4):1–12, 2014.
- [5] Volker Blanz and Thomas Vetter. A Morphable Model for The Synthesis of 3D Faces. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 187–194, 1999.
- [6] Boming Zhao and Bangbang Yang, Zhenyang Li, Zuoyue Li, Guofeng Zhang, Jiashu Zhao, Dawei Yin, Zhaopeng Cui, and Hujun Bao. Factorized and Controllable Neural Re-Rendering of Outdoor Scene for Photo Extrapolation. In *Proceedings of the 30th ACM International Conference on Multimedia*, 2022.
- [7] Nicolas Bonneel, Kalyan Sunkavalli, James Tompkin, Deqing Sun, Sylvain Paris, and Hanspeter Pfister. Interactive Intrinsic Video Editing. *ACM Transactions on Graphics*, 33(6):1–10, 2014.
- [8] Nicolas Bonneel, James Tompkin, Kalyan Sunkavalli, Deqing Sun, Sylvain Paris, and Hanspeter Pfister. Blind Video Temporal Consistency. *ACM Transactions on Graphics*, 34(6):1–9, 2015.
- [9] 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.
- [10] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. TensorRF: Tensorial Radiance Fields. *arXiv preprint arXiv:2203.09517*, 2022.
- [11] Anpei Chen, Zexiang Xu, Fuqiang Zhao, Xiaoshuai Zhang, Fanbo Xiang, Jingyi Yu, and Hao Su. MVSNeRF: Fast Generalizable Radiance Field Reconstruction from Multi-View Stereo. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14124–14133, 2021.
- [12] Xingyu Chen, Qi Zhang, Xiaoyu Li, Yue Chen, Ying Feng, Xuan Wang, and Jue Wang. Hallucinated Neural Radiance Fields In The Wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12943–12952, 2022.
- [13] Yizong Cheng. Mean Shift, Mode Seeking, and Clustering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(8):790–799, 1995.
- [14] Chong Bao and Bangbang Yang, Zeng Junyi, Bao Hujun, Zhang Yinda, Cui Zhaopeng, and Zhang Guofeng. NeuMesh: Learning Disentangled Neural Mesh-based Implicit Field for Geometry and Texture Editing. In *Proceedings of the European Conference on Computer Vision*, 2022.
- [15] Blender Online Community. *Blender - a 3D modelling and rendering package*. Blender Foundation, Stichting Blender Foundation, Amsterdam, 2018.
- [16] Sylvain Duchêne, Clement Riant, Gaurav Chaurasia, Jorge Lopez-Moreno, Pierre-Yves Laffont, Stefan Popov, Adrien Bousseau, and George Drettakis. Multi-view Intrinsic Images of Outdoors Scenes with An Application to Relighting. *ACM Transactions on Graphics*, page 16, 2015.
- [17] Qingnan Fan, Jiaolong Yang, Gang Hua, Baoquan Chen, and David Wipf. Revisiting Deep Intrinsic Image Decompositions. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8944–8952, 2018.
- [18] Stephan J Garbin, Marek Kowalski, Matthew Johnson, Jamie Shotton, and Julien Valentin. FastNeRF: High-fidelity Neural Rendering at 200FPS. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14346–14355, 2021.
- [19] 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.
- [20] Mark Hamilton, Zhoutong Zhang, Bharath Hariharan, Noah Snavely, and William T. Freeman. Unsupervised Semantic Segmentation by Distilling Feature Correspondences. In *International Conference on Learning Representations*, 2022.
- [21] Berthold KP Horn. Determining Lightness from An Image. *Computer graphics and image processing*, 3(4):277–299, 1974.
- [22] Michael Janner, Jiajun Wu, Tejas D Kulkarni, Ilker Yildirim, and Josh Tenenbaum. Self-supervised Intrinsic Image Decomposition. *Advances in Neural Information Processing Systems*, 30, 2017.
- [23] Alen Joy and Charalambos Poullis. Multi-view Gradient Consistency for SVBRDF Estimation of Complex Scenes under Natural Illumination. *arXiv preprint arXiv:2202.13017*, 2022.
- [24] Yoshihiro Kanamori and Yuki Endo. Relighting Humans: Occlusion-aware Inverse Rendering for Full-body Human Images. *arXiv preprint arXiv:1908.02714*, 2019.
- [25] Diederik P Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [26] Abhijit Kundu, Kyle Genova, Xiaoqi Yin, Alireza Fathi, Caroline Pantofaru, Leonidas Guibas, Andrea Tagliasacchi, Frank Dellaert, and Thomas Funkhouser. Panoptic Neural Fields: A Semantic Object-Aware Neural Scene Representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022.
- [27] Youngjoong Kwon, Dahun Kim, Duygu Ceylan, and Henry Fuchs. Neural Human Performer: Learning Generalizable Radiance Fields for Human Performance Rendering. *Advances in Neural Information Processing Systems*, 34, 2021.

- [28] Pierre-Yves Laffont, Adrien Bousseau, Sylvain Paris, Frédéric Durand, and George Drettakis. Coherent Intrinsic Images from Photo Collections. *ACM Transactions of Graphics*, 2012.
- [29] Wei-Sheng Lai, Jia-Bin Huang, Oliver Wang, Eli Shechtman, Ersin Yumer, and Ming-Hsuan Yang. Learning Blind Video Temporal Consistency. In *Proceedings of the European Conference on Computer Vision*, pages 170–185, 2018.
- [30] Edwin H Land and John J McCann. Lightness and Retinex Theory. *Josa*, 61(1):1–11, 1971.
- [31] Hoang-An Le, Thomas Mensink, Partha Das, Sezer Karaoglu, and Theo Gevers. EDEN: Multimodal Synthetic Dataset of Enclosed Garden Scenes. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1579–1589, 2021.
- [32] Tzu-Mao Li, Miika Aittala, Frédo Durand, and Jaakko Lehtinen. Differentiable Monte Carlo Ray Tracing through Edge Sampling. *ACM Transactions on Graphics*, 37(6):1–11, 2018.
- [33] Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. Neural Scene Flow Fields for Space-time View Synthesis of Dynamic Scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6498–6508, 2021.
- [34] Zhengqi Li and Noah Snavely. CGIntrinsics: Better Intrinsic Image Decomposition through Physically-based Rendering. In *Proceedings of the European Conference on Computer Vision*, pages 371–387, 2018.
- [35] Zhengqi Li and Noah Snavely. Learning Intrinsic Image Decomposition from Watching The World. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9039–9048, 2018.
- [36] Andrew Liu, Shiry Ginosar, Tinghui Zhou, Alexei A Efros, and Noah Snavely. Learning to Factorize and Relight A City. In *Proceedings of the European Conference on Computer Vision*, pages 544–561. Springer, 2020.
- [37] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural Sparse Voxel Fields. *Advances in Neural Information Processing Systems*, 33:15651–15663, 2020.
- [38] 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.
- [39] Guillaume Loubet, Nicolas Holzschuch, and Wenzel Jakob. Reparameterizing Discontinuous Integrands for Differentiable Rendering. *ACM Transactions on Graphics*, 38(6):1–14, 2019.
- [40] Jundan Luo, Zhaoyang Huang, Yijin Li, Xiaowei Zhou, Guofeng Zhang, and Hujun Bao. NIID-Net: Adapting Surface Normal Knowledge for Intrinsic Image Decomposition in Indoor Scenes. *IEEE Transactions on Visualization and Computer Graphics*, 26(12):3434–3445, 2020.
- [41] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7210–7219, 2021.
- [42] Bruce A Maxwell, Richard M Friedhoff, and Casey A Smith. A Bi-illuminant Dichromatic Reflection Model for Understanding Images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2008.
- [43] Abhimitra Meka, Mohammad Shafiei, Michael Zollhöfer, Christian Richardt, and Christian Theobalt. Real-time Global Illumination Decomposition of Videos. *ACM Transactions on Graphics*, 40(3):1–16, 2021.
- [44] Abhimitra Meka, Michael Zollhöfer, Christian Richardt, and Christian Theobalt. Live Intrinsic Video. *ACM Transactions on Graphics*, 35(4):1–14, 2016.
- [45] Ben Mildenhall, Pratul P. Srinivasan, Rodrigo Ortiz-Cayon, Nima Khademi Kalantari, Ravi Ramamoorthi, Ren Ng, and Abhishek Kar. Local Light Field Fusion: Practical View Synthesis with Prescriptive Sampling Guidelines. *ACM Transactions on Graphics*, 2019.
- [46] 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 *Proceedings of the European conference on computer vision*, pages 405–421. Springer, 2020.
- [47] Yuhang Ming, Weicai Ye, and Andrew Calway. iDF-SLAM: End-to-End RGB-D SLAM with Neural Implicit Mapping and Deep Feature Tracking. *arXiv preprint arXiv:2209.07919*, 2022.
- [48] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant Neural Graphics Primitives with A Multiresolution Hash Encoding. *ACM Transactions on Graphics*, 41(4):102:1–102:15, July 2022.
- [49] Merlin Nimier-David, Delio Vicini, Tizian Zeltner, and Wenzel Jakob. Mitsuba 2: A Retargetable Forward and Inverse Renderer. *ACM Transactions on Graphics*, 38(6):1–17, 2019.
- [50] Byong Mok Oh, Max Chen, Julie Dorsey, and Frédo Durand. Image-based Modeling and Photo Editing. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 433–442, 2001.
- [51] Keunhong Park, Utkarsh Sinha, Jonathan T. Barron, Sofien Bouaziz, Dan B Goldman, Steven M. Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable Neural Radiance Fields. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021.
- [52] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-NeRF: Neural Radiance Fields for Dynamic Scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020.
- [53] Mike Roberts, Jason Ramapuram, Anurag Ranjan, Atulit Kumar, Miguel Angel Bautista, Nathan Paczan, Russ Webb, and Joshua M Susskind. Hypersim: A Photorealistic Synthetic Dataset for Holistic Indoor Scene Understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10912–10922, 2021.
- [54] Viktor Rudnev, Mohamed Elgharib, William Smith, Lingjie Liu, Vladislav Golyanik, and Christian Theobalt. NeRF for

- Outdoor Scene Relighting. In *Proceedings of the European Conference on Computer Vision*, 2022.
- [55] Katja Schwarz, Yiyi Liao, Michael Niemeyer, and Andreas Geiger. GRAF: Generative Radiance Fields for 3D-aware Image Synthesis. *Advances in Neural Information Processing Systems*, 33:20154–20166, 2020.
- [56] Li Shen, Ping Tan, and Stephen Lin. Intrinsic Image Decomposition with Non-local Texture Cues. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1–7. IEEE, 2008.
- [57] Julian Straub, Thomas Whelan, Lingni Ma, Yufan Chen, Erik Wilmans, Simon Green, Jakob J. Engel, Raul Mur-Artal, Carl Ren, Shobhit Verma, Anton Clarkson, Mingfei Yan, Brian Budge, Yajie Yan, Xiaqing Pan, June Yon, Yuyang Zou, Kimberly Leon, Nigel Carter, Jesus Briales, Tyler Gillingham, Elias Mueggler, Luis Pesqueira, Manolis Savva, Dhruv Batra, Hauke M. Strasdat, Renzo De Nardi, Michael Goesele, Steven Lovegrove, and Richard Newcombe. The Replica Dataset: A Digital Replica of Indoor Spaces. *arXiv preprint arXiv:1906.05797*, 2019.
- [58] Edgar Sucar, Shikun Liu, Joseph Ortiz, and Andrew J Davison. iMAP: Implicit Mapping and Positioning in Real-time. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6229–6238, 2021.
- [59] Jiaming Sun, Xi Chen, Qianqian Wang, Zhengqi Li, Hadar Averbuch-Elor, Xiaowei Zhou, and Noah Snavely. Neural 3D Reconstruction in the Wild. In *SIGGRAPH Conference Proceedings*, 2022.
- [60] Matthew Tancik, Vincent Casser, Xinchen Yan, Sabeek Pradhan, Ben Mildenhall, Pratul P Srinivasan, Jonathan T Barron, and Henrik Kretschmar. Block-NeRF: Scalable Large Scene Neural View Synthesis. *arXiv preprint arXiv:2202.05263*, 2022.
- [61] Shoji Tominaga. Dichromatic Reflection Models for A Variety of Materials. *Color Research & Application*, 19(4):277–285, 1994.
- [62] Alex Trevithick and Bo Yang. GRF: Learning A General Radiance Field for 3d Representation and Rendering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15182–15192, 2021.
- [63] 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. *Advances in Neural Information Processing Systems*, 2021.
- [64] Qianqian Wang, Zhicheng Wang, Kyle Genova, Pratul P Srinivasan, Howard Zhou, Jonathan T Barron, Ricardo Martin-Brualla, Noah Snavely, and Thomas Funkhouser. IBRNet: Learning Multi-view Image-based Rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4690–4699, 2021.
- [65] Bangbang Yang, Yinda Zhang, Yijin Li, Zhaopeng Cui, Sean Fanello, Hujun Bao, and Guofeng Zhang. Neural Rendering in a Room: Amodal 3D Understanding and Free-Viewpoint Rendering for the Closed Scene Composed of Pre-Captured Objects. *ACM Transactions on Graphics*, 41(4):101:1–101:10, July 2022.
- [66] Lior Yariv, Yoni Kasten, Dror Moran, Meirav Galun, Matan Atzmon, Basri Ronen, and Yaron Lipman. Multiview Neural Surface Reconstruction by Disentangling Geometry and Appearance. *Advances in Neural Information Processing Systems*, 33, 2020.
- [67] Genzhi Ye, Elena Garces, Yebin Liu, Qionghai Dai, and Diego Gutierrez. Intrinsic Video and Applications. *ACM Transactions on Graphics*, 33(4):1–11, 2014.
- [68] Weicai Ye, Xinyue Lan, Shuo Chen, Yuhang Ming, Xinyuan Yu, Chong Bao, Hujun Bao, Zhaopeng Cui, and Guofeng Zhang. PVO: Panoptic Visual Odometry. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023.
- [69] Weicai Ye, Xinyue Lan, Shuo Chen, Yuhang Ming, Xinyuan Yu, Jinyu Li, Hujun Bao, Zhaopeng Cui, and Guofeng Zhang. DeFlowSLAM: Self-Supervised Scene Motion Decomposition for Dynamic Dense SLAM. 2022.
- [70] Renjiao Yi, Ping Tan, and Stephen Lin. Leveraging Multi-view Image Sets for Unsupervised Intrinsic Image Decomposition and Highlight Separation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020.
- [71] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. PlenOctrees for Real-time Rendering of Neural Radiance Fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5752–5761, 2021.
- [72] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. PixelNeRF: Neural Radiance Fields from One or Few Images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4578–4587, 2021.
- [73] Ye Yu and William AP Smith. Inverserendernet: Learning Single Image Inverse Rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3155–3164, 2019.
- [74] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5453–5462, 2021.
- [75] 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/CVF Conference on Computer Vision and Pattern Recognition*, 2018.
- [76] 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*, 40(6):1–18, 2021.
- [77] Yuanqing Zhang, Jiaming Sun, Xingyi He, Huan Fu, Rongfei Jia, and Xiaowei Zhou. Modeling Indirect Illumination for Inverse Rendering. *arXiv preprint arXiv:2204.06837*, 2022.
- [78] Shuang Zhao, Wenzel Jakob, and Tzu-Mao Li. Physics-based Differentiable Rendering: from Theory to Implementation. In *ACM SIGGRAPH 2020 Courses*, pages 1–30. ACM SIGGRAPH 2020 Courses, 2020.

- [79] Shuaifeng Zhi, Tristan Laidlow, Stefan Leutenegger, and Andrew J Davison. In-place Scene Labelling and Understanding with Implicit Scene Representation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15838–15847, 2021.
- [80] Jingsen Zhu, Yuchi Huo, Qi Ye, Fujun Luan, Jifan Li, Dianbing Xi, Lisha Wang, Rui Tang, Wei Hua, Hujun Bao, et al. I2-sdf: Intrinsic indoor scene reconstruction and editing via raytracing in neural sdf. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12489–12498, 2023.
- [81] Rui Zhu, Zhengqin Li, Janarbek Matai, Fatih Porikli, and Manmohan Chandraker. IRISformer: Dense Vision Transformers for Single-Image Inverse Rendering in Indoor Scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2822–2831, 2022.
- [82] Zihan Zhu, Songyou Peng, Viktor Larsson, Weiwei Xu, Hujun Bao, Zhaopeng Cui, Martin R Oswald, and Marc Pollefeys. NICE-SLAM: Neural Implicit Scalable Encoding for SLAM. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12786–12796, 2022.