

Intrinsic Image Decomposition Using Point Cloud Representation

Xiaoyan Xing¹, Konrad Groh², Sezer Karaoglu¹, Theo Gevers¹

¹ UvA-Bosch Delta Lab, University of Amsterdam

² Bosch Center for Artificial Intelligence, Robert Bosch GmbH

Abstract. The purpose of intrinsic decomposition is to separate an image into its albedo (reflective properties) and shading components (illumination properties). This is challenging because it's an ill-posed problem. Conventional approaches primarily concentrate on 2D imagery and fail to fully exploit the capabilities of 3D data representation. 3D point clouds offer a more comprehensive format for representing scenes, as they combine geometric and color information effectively. To this end, in this paper, we introduce Point Intrinsic Net (PoInt-Net), which leverages 3D point cloud data to concurrently estimate albedo and shading maps. The merits of PoInt-Net include the following aspects. First, the model is efficient, achieving consistent performance across point clouds of any size with training only required on small-scale point clouds. Second, it exhibits remarkable robustness; even when trained exclusively on datasets comprising individual objects, PoInt-Net demonstrates strong generalization to unseen objects and scenes. Third, it delivers superior accuracy over conventional 2D approaches, demonstrating enhanced performance across various metrics on different datasets. [Code](#)

1 Introduction

The aim of intrinsic decomposition is to separate an image into its albedo (reflective properties) and shading components (illumination properties). Decomposing an image into these fundamental components is a challenging problem due to its ill-posed nature, requiring specific constraints. Geometric data, including depth and surface normals, are used in facilitating this process [16]. Typically, these geometric cues are estimated as part of the process of intrinsic decomposition [1, 3, 20, 30, 32, 44, 56]. However, the success of this approach heavily depends on the accuracy of surface normal estimation and is sensitive to the 2D nature of the data employed, indicating a lack of adaptability across various data types. Traditional methods mainly focus on 2D images, overlooking the full potential of 3D data representation. 3D point clouds provide a richer scene representation by effectively integrating geometric and color information.

This paper delves into utilizing 3D point clouds for the purpose of intrinsic decomposition. We focus on 3D point clouds obtained (1) directly from *RGB-D* cameras, or (2) from 2D *RGB* images where the depth *D* map is measured by a monocular depth estimation technique, see Figure 1. In this way, the proposed point cloud-based network, PoInt-Net, utilizes the 3D structure and appearance of objects or scenes to derive surface geometry and extract intrinsic features. Our method offers several benefits. First, point cloud representation naturally includes explicit 3D priors along with color details.

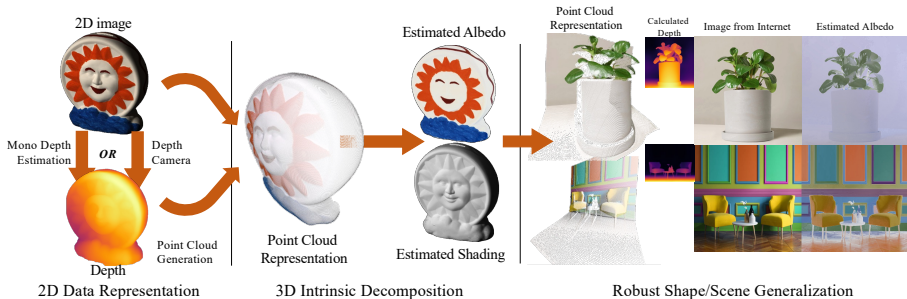


Fig. 1: Intrinsic image decomposition using a 3D point cloud representation. Our approach decomposes the intrinsic components of an object/scene based on the point cloud representation of its appearance from a particular viewing angle. Point clouds are generated from *RGB-D* images, where the depth maps are obtained by a depth camera (e.g. Lidar or ToF) or are estimated by a monocular depth estimation method such as [41].

Second, the intrinsic geometric information within the 3D point cloud is beneficial for more precise shading estimation, especially in areas where abrupt depth changes typically coincide with shifts in lighting [48]. Third, point clouds accurately capture the shape of a scene, providing superior generalization for low-level vision tasks as shown by [49]. Advancements in depth acquisition (such as Time-of-Flight cameras) and estimation technologies (for example, MiDaS [41]) have substantially lowered the costs associated with acquiring depth information necessary for constructing point clouds.

Experimental analysis demonstrates that PoInt-Net excels in both efficiency and generalization capabilities. It outperforms existing models in shading estimation across multiple datasets with a reduced number of parameters, and still achieves impressive albedo outcomes. Trained exclusively on datasets containing singular objects, PoInt-Net showcases remarkable ability for zero-shot intrinsic estimation in real-world scenarios by using point clouds obtained from depth estimations.

The contributions of the paper are:

- By applying intrinsic decomposition to a 3D point cloud framework, our approach innovatively merges geometric priors with sparse representations.
- Introducing PoInt-Net, a point-based intrinsic image decomposition network with specialized subnets for light direction, shading, and albedo tasks.
- PoInt-Net operates on sparse point clouds with far fewer parameters (1/10 to 1/100 of the existing methods), excelling on diverse datasets.
- PoInt-Net facilitates zero-shot intrinsic estimation in real-world settings through the use of point clouds derived from estimated depths.

2 Related Work

Intrinsic image decomposition (IID) is a complex, underdetermined problem that necessitates distinct constraints and priors, which can be sourced from images or supplementary data. We categorize prior IID approaches according to their input needs.

IID with image(s). One of the earliest work in 1970s [2] proposes to derive intrinsic characteristics from images. Numerous methods since then have employed simple network architectures for estimating these intrinsic characteristics [14, 28, 35, 40, 47, 54, 57]. [5, 7, 18, 33] exploring IID by employing perceptual priors, based on the assumption that pronounced gradient edges denote changes in reflectance, while subtle edges suggest variations in illumination. [13, 42, 45, 46] present IID techniques that utilize clustering or reflectance sparsity. Exploring shading involves geometric cues, where [1] introduces a framework that capitalizes on geometric information extracted from images. [3, 20, 30, 44, 56] divide shading estimation into two processes (estimating surface normals and lighting) and then applying a shader for the final shading calculation. However, the success of these approaches depends on their precision in estimating surface normals, which complicates their adaptability to different types of data, e.g., [44] is specifically designed for facial images and [30, 56] for indoor images.

IID with image(s) and additional input. *RGB-D* images are used for the purpose of IID. [1, 11, 19, 22] show that incorporating extra geometric cues (such as depth and surface normals) to refine the shading components results in enhanced decomposition outcomes. [26] proposes a model for IID using a sequence of *RGB-D* video frames. Recently, [43] leverages the LiDAR intensity to separate the reflectance from an image. In addition to *RGB-D* data type, [25] uses multi-view stereo to reconstruct 3D points and surface normals to estimate the intrinsic components. [52] employs a Neural Radiance Field (NeRF) [34] to obtain the intrinsic components in an implicit representation. [50] utilizes multi-view constraints and semantic labels to derive intrinsic properties from a NeRF. The inclusion of additional input aids models in understanding the geometric relationships between pixels. Nonetheless, the above approaches encounter two primary challenges: 1) it lacks computational efficiency for *RGB-D* image-based decomposition because of the expensive 3D operations involved; 2) it suffers from limited generalization capabilities in multi-view-based approaches as the Neural Radiance Fields (NeRF) are prone to overfitting specific scenes.

In contrast to previous *RGB-D* methods, our approach utilizes point cloud representation for estimating intrinsic components, without explicitly estimating surface normals. This strategy leads to enhanced robustness in decomposition, yielding a point-based operation network for intrinsic image decomposition, which improves accuracy and reduces computational costs.

3 Point Cloud Intrinsic Representation

In this section, we introduce a new intrinsic representation technique utilizing point clouds. Section 3.1 details the intrinsic decomposition process grounded in the rendering model. Section 3.2 revisits the intrinsic decomposition from the perspective of a point cloud representation.

3.1 Intrinsic Decomposition

For a given point x , the reflected radiance L is defined by [21]:

$$L(x, \omega_o) = \int_{\omega_i \in \Omega^+} f_r(x, \omega_i, \omega_o) L_i(x, \omega_i) (N \cdot \omega_i) d\omega_i, \quad (1)$$




Fig. 2: Our proposed framework for intrinsic point cloud decomposition starts by transforming the *RGB-D* representation into a point cloud representation. (a) The point cloud representation is used as input to train two separate components: the shading and the albedo estimations. The shading estimation is supported by the DirectionNet (Light Direction Estimation Net), which takes (a) as input and outputs surface light direction estimates (c). Surface normals (d) are calculated using local neighborhoods within (a). The Shader (Learnable Shader) then uses the concatenated vectors of (c) and (d) to generate the final shading estimation (e). The albedo estimation is obtained by the AlbedoNet (Point-Albedo Net) which extracts invariant reflectance (b) from (a) based on the Lambertian assumption. Finally, by multiplying (b) and (e), the reconstructed image (f) is generated. Please refer to the supplementary for a more detailed explanation.

where ω_i is the light angle from the upper hemisphere Ω_+ , ω_o is the viewing angle, N is the surface normal, $L_i(x, \omega_i)$ is the position of the lighting angle and its direction, and f_r is the surface reflectance, modeled by a Bidirectional Reflectance Distribution Function (BRDF) [37].

Given a viewing angle, assuming that the surface is Lambertian, the diffuse appearance $I_{diffuse}$ is given by:

$$\mathbf{I}_{diffuse} = \int_{\omega_i \in \Omega_+} f_r(\omega_i) L_i(\omega_i) (N \cdot \omega_i) d\omega_i. \quad (2)$$

Conventionally, $\frac{\rho_d}{\pi}$ denotes the reflectivity of the surface (albedo), where $f_r(\omega_i) = \frac{\rho_d}{2\pi}$. Therefore, if the illumination is uniform, the model is defined by:

$$\mathbf{I}_{diffuse} = \frac{\rho_d}{\pi} \cdot (\mathbf{N} \cdot \mathbf{L}_{in}), \quad (3)$$

where, \mathbf{L}_{in} represents the visible incident light. The aim of intrinsic decomposition is to disentangle the albedo $\mathbf{A} = \frac{\rho_d}{\pi}$ and shading $\mathbf{S} = (\mathbf{N} \cdot \mathbf{L}_{in})$ from the appearance $\mathbf{I}_{diffuse}$, where (\cdot) is the dot product.

3.2 Intrinsic Appearance of Point Clouds

According to Equation (2), the appearance of an object under a given lighting condition is represented by a *RGB* image $\mathbf{I} = [\mathbf{I}_r, \mathbf{I}_g, \mathbf{I}_b] \in \mathbb{R}^{U \times V \times 3}$. Additionally, its corresponding depth map is given by $D \in \mathbb{R}^{U \times V \times 1}$. Depth data can be acquired from a depth camera, such as LiDAR, or it can be estimated from 2D images using a monocular depth estimation method, e.g., [41].

The *RGB* image and its associated depth map are converted into a colored point cloud representation, $\mathbf{P} = \{\mathbf{p}_i | i \in 1, \dots, n\}$. Each point \mathbf{p}_i is represented as a vector of $[x, y, d, r, g, b]$ values:

$$\mathbf{p}_i = \left(\frac{(u - c_x)d}{f_x}, \frac{(v - c_y)d}{f_y}, d, r, g, b \right), \quad (4)$$

where, f_x and f_y are the focal lengths, and (c_x, c_y) is the principal point.

Given a dataset of M point clouds, $\mathcal{P} = \{\mathbf{P}_1, \mathbf{P}_2, \dots, \mathbf{P}_M\}$, its intrinsic components are defined by: 1) Albedo $\mathcal{A} = \{\mathbf{A}_1, \mathbf{A}_2, \dots, \mathbf{A}_M\}$, 2) Shading $\mathcal{S} = \{\mathbf{S}_1, \mathbf{S}_2, \dots, \mathbf{S}_M\}$, 3) Surface normal $\mathcal{N} = \{\mathbf{N}_1, \mathbf{N}_2, \dots, \mathbf{N}_M\}$, and 4) Light source position $\mathcal{L} = \{\mathbf{L}_1, \mathbf{L}_2, \dots, \mathbf{L}_M\}$.

Albedo contains the reflectance information. Hence, a direct point based mapping ($f_\alpha : \mathcal{P} \rightarrow \mathcal{A}$) is employed to decompose the reflectance from input point cloud.

Shading depends on the object geometry, viewing and lighting conditions. Thus, instead of directly learning the shading, a point-light direction net ($f_\theta : \mathcal{P} \rightarrow \mathcal{L}$) is used to estimate the light direction from the point cloud representation. Then, a point-learnable shader ($f_\sigma : \mathcal{L}, \mathcal{N} \rightarrow \mathcal{S}$) is trained to generate the rendering effects based on the surface normals (from input point cloud) and light direction estimation (from point-light direction net).

Adopting point cloud representation for intrinsic decomposition is beneficial due to several reasons: 1) It naturally integrates depth and *RGB* information into a cohesive structure. 2) It facilitates the derivation of surface normal data, ascertainable through local neighborhood analysis. 3) It provides a resilient framework capable of withstanding errors in depth measurements, as our additional ablation study indicates that minor inaccuracies in certain points have a negligible impact on the overall representation.

4 Point Based Intrinsic Decomposition

In this section, a novel point based intrinsic decomposition technique is proposed. Section 4.1 provides the technical details of the network architecture and Section 4.2 introduces the learning strategy to train the network.

4.1 Point Intrinsic Net

Our proposed network, PoInt-Net, consists of three key components: 1) the Point Albedo-Net, which is designed to capture the properties of surface materials, 2) the Light Direction Estimation Net, tasked with identifying lighting conditions to support the albedo estimation of the point cloud, and 3) the Learnable Shader, a module that combines the deduced light direction with surface normals to generate the shading map. The approach most similar to ours is [20]. However, it needs ground-truth normal information to initialize and is limited by its generalization capabilities to individual objects.

Figure 2 shows the architecture of the proposed network and the specifics of its forward connections. The design of the three sub-networks is largely consistent, with minor variations such as the activation functions used. Specifically, all three sub-nets are adopted from [38], and employ Multi-Layer Perceptrons (MLPs) for point-feature extraction and decoding, with the aim of solving the point-to-point relationship.

Point albedo-net inputs a 6-D point cloud, which includes color data and spatial coordinates, and generates surface reflectance estimates. The Rectified Linear Unit (ReLU) is employed as the activation function to yield scaled output colors.

Light direction estimation Net takes the same input as *Point Albedo-Net*, and predicts point-wise surface light directions. ReLU is used as the activation function in most of the layers. The final two layers use the hyperbolic tangent function (Tanh) to ensure that all light directions are estimated.

Surface normal calculation computes the surface normal from the given point cloud through several steps: 1) identifying neighboring points and calculating the covariance matrix, 2) computing the eigenvectors of the covariance matrix, and 3) selecting the normal vector associated with the smallest eigenvalue.

Learnable shader takes concatenated vectors of surface normal information (derived from the input point cloud) and surface light direction estimation as input, and produces a point-wise shading map as output.

Our choice of module configuration is based on empirical evidence: 1) *Module Specialization.*: Given PointNet’s limitations as an relatively weaker encoder, we use three specialized modules for separate estimation of albedo, lighting, and shading. Which also guarantees the robustness in the generalization. 2) *Input Integration.* We found that concatenating surface normals with lighting, rather than multiplying them, yields improved outcomes.

4.2 Joint-learning Strategy

A two-step training strategy is employed to arrive at an end-to-end intrinsic decomposition learning pipeline. First, in terms of shading estimation, the *Light Direction Estimation Net* and *Learnable Shader* are trained using the ground-truth light position \mathbf{L} and shading map \mathbf{S} . Then, for albedo estimation, the parameters in these two sub-nets are preserved and frozen, while the *Point Albedo-Net* is constrained by the ground-truth albedo \mathbf{A} and the final reconstructed image $\hat{\mathbf{I}}$ (multiplied by the estimated albedo map $\hat{\mathbf{A}}$ and the estimated shading map $\hat{\mathbf{S}}$). During training, the mean square error is used. The loss function³, for stage one, is:

$$\mathcal{L}_{shading} = \frac{1}{M} \sum^M (|\mathbf{L} - \hat{\mathbf{L}}|^2 + |\mathbf{S} - \hat{\mathbf{S}}|^2). \quad (5)$$

For stage two, a number of loss functions are used. To address reflectance changes, a color cross ratio loss inspired by [17] is taken, formulated as follows:

$$\mathcal{L}_{ccr} = |M_{RG} - M_{\hat{R}\hat{G}}| + |M_{RB} - M_{\hat{R}\hat{B}}| + |M_{GB} - M_{\hat{G}\hat{B}}|, \quad (6)$$

where $\{M_{RG}, M_{RB}, M_{GB}\}$, $\{M_{\hat{R}\hat{G}}, M_{\hat{R}\hat{B}}, M_{\hat{G}\hat{B}}\}$ are the cross color ratios from the ground-truth albedo and the estimated albedo respectively. Please refer to supplemental for the details of the cross color ratios calculation. Similar to [14], the gradient difference is considered and is formulated by:

$$\mathcal{L}_{grad} = |\nabla \mathbf{A} - \nabla \hat{\mathbf{A}}|_2^2. \quad (7)$$

³ For datasets without light source labels, the loss function only contains the shading map $\hat{\mathbf{S}}$.

Hence, the reconstruction loss is applied to constrain the estimated albedo:

$$\mathcal{L}_{rec} = \frac{1}{M} \sum^M (|\mathbf{A} - \hat{\mathbf{A}}|^2 + |\mathbf{I} - \hat{\mathbf{I}}|^2). \quad (8)$$

The final loss function is given by:

$$\mathcal{L}_{albedo} = \mathcal{L}_{rec} + \mathcal{L}_{grad} + \mathcal{L}_{ccr}, \quad (9)$$

where $\{\hat{\cdot}\}$ represents the estimated values, and M is the number of input point clouds in the mini-batch. Adam [23] is employed as the optimizer.

5 Experiments

This section outlines the experimental framework used to assess the effectiveness of our proposed approach.

Dataset. Five publicly accessible datasets are employed in the assessment:

- *ShapeNet-Intrinsic* [20]: Based on ShapeNet [10], albedo and shading are generated by the Blender-cycle. The dataset contains ground-truth depth, normal, and light position information. We follow the same dataset split as Liu *et al.* [31].
- *MIT-Intrinsic* [18]: For this real-world dataset, albedo and shading under different illumination conditions are provided. Depth information, calculated by [1], is used. The train and test split are kept the same as proposed by [1].
- *MPI-Sintel* [8]: This synthetic dataset provides albedo, shading, and depth information. We use the same training and test splits as those employed by existing methods to assess our approach.
- *Inverender* [53]: This synthetic dataset includes ground-truth albedo and normals. To evaluate the performance of our method, we adhere to the same training and test splits established by previous methods [50, 52].
- *IIW* [6]: A real-world image dataset comprising a wide variety of scenes and lighting conditions, along with the weighted human disagreement rate (WHDR) label.

Moreover, we employ various images sourced from the internet to demonstrate our capability to generalize to real-world scenarios.

Depth acquisition. For depth information, we utilize ground-truth depth from the datasets where available. In cases without ground-truth depth, including certain datasets [6, 52] and real-world images, we apply the mono-depth estimation method [41] to generate relative depth information for point cloud construction.

Metrics and visual comparison For quantitative evaluation, we employ three commonly used metrics: Mean Square Error (MSE), Local Mean Squared Error (LMSE), and Structural Dissimilarity (DSSIM). These metrics are our standard evaluation criteria, unless specified otherwise. For fair qualitative comparison, visualization images for all benchmarks are objectively selected, sourced directly from the publications of the methods being compared or acquired from their official implementations.




Fig. 3: Comparison to state-of-the-art method USI3D (fine-tuned version) [31] and ablation study on the ShapeNet-intrinsic dataset. **Zoom to see details.**

Pretrained on ShapeNet-Intrinsic dataset. Due to the limited size of the intrinsic decomposition datasets, such as MIT-intrinsic [18] with only 20 objects and MPI-Sintel [8] with hundreds of images, employing pre-trained models is a common practice. For instance, [12, 40] use pre-trained weights from ImageNet [24]. Similarly, [13] pre-train on the NED dataset [4], while [20], [31], and [5] use ShapeNet for pre-training. In our comparative tables, we indicate which models have undergone pre-training and specify the datasets utilized for this pre-training. In this way, fairness and clarity is ensured in our evaluation.

5.1 Evaluation on IID datasets

	MSE(10^{-2})			LMSE(10^{-2})	DSSIM(10^{-2})
	Albedo	Shading	Average	Total	Total
CGIntrinsics [29]	3.38	2.96	3.17	6.23	-
Fan <i>et al.</i> [14]	3.02	3.15	3.09	7.17	-
Ma <i>et al.</i> * [33]	2.84	2.62	2.73	5.44	-
USI3D* [31]	1.85	1.08	1.47	4.65	-
Ours (w/o. shader)	0.48	0.57	0.53	1.15	4.93
Ours	0.46	0.38	0.42	1.00	4.15

* Unsupervised methods but finetune on the dataset.

Table 1: Results and ablation study for ShapeNet-Intrinsic [20].

ShapeNet-Intrinsic Table 1 shows a quantitative comparison of PoInt-Net with the latest open source methods [14, 29, 31, 33]. Our results clearly indicate that our approach significantly surpasses existing methods across all three metrics. To provide specific values, PoInt-Net achieves an MSE of 0.0046 for albedo and 0.0038 for shading, with an LMSE of 1.00 and a DSSIM of 0.0415. This exceptional performance is attributed to our method’s capacity to capture and utilize intricate relationships among intrinsic properties, leading to a more robust and dependable estimation.

Qualitative results of the method are shown in Figure 3, where USI3D [31] (fine-tuned version) is used as a reference. PoInt-Net generates a shading map by utilizing surface light direction and normal information, effectively separating shading from the composite image. This process results in realistic and consistent shading in the output images, closely reflecting the underlying surface geometry. Particularly noticeable is the clear distinction between shading and surface in darker areas of objects, underscoring the method’s efficacy and robustness in generating high-quality, visually appealing outputs that faithfully represent intrinsic properties of objects.




Fig. 4: Qualitative results on the MIT-intrinsic benchmark [18]. Comparison to the state-of-the-art method PIE-Net [13]. Ablation study on shader is conducted.

MIT-intrinsic In addition to the synthetic dataset, this section encompasses an evaluation on the MIT-intrinsic dataset to assess the proposed method’s ability to generalize to real-world scenarios. The results obtained from the MIT-intrinsic dataset are consistent with those from the synthetic dataset, validating the method’s efficacy and robustness across varied datasets.

Table 2 reports the quantitative results. PoInt-Net produces state-of-the-art shading results on the MIT-intrinsic dataset for all metrics. Moreover, our approach obtains the best LMSE and the second-best performance for the albedo output in terms of MSE and DSSIM metrics. Note that [13] employs an extra input to provide additional information for albedo estimation.

The visualization results are given in Figure 4. Our method outperforms the latest state-of-the-art method [13] in accurately distinguishing the markings on the frog’s back. This success is due to PoInt-Net’s incorporation of surface light direction estimation and surface normal calculation, which contribute to the production of high-quality shading outcomes.

Note: Ground-truth depth information is not included in the MIT-intrinsic dataset. To this end, depth information is used and computed by the method of [1]. Although depth estimations frequently contain noise, such as outliers and invalid points, PoInt-Net consistently learns intrinsic features even when the input data includes noise. This underscores its robustness to imperfect depth information.

MPI-Sintel Unlike object-wise datasets, the MPI dataset provides color information for each pixel. To ensure a fair evaluation, we adopt the approach outlined in [14], which employs scale-invariant MSE (si-MSE) and local scale-invariant MSE (si-LMSE).

We conduct a comparative analysis of our approach against several state-of-the-art methods (Table 3). Our method surpasses others in terms of si-LMSE for albedo and shading, as well as si-MSE for albedo. Notably, methods like [1, 11, 22] also require

	MSE(10^{-2})		LMSE(10^{-2})		DSSIM(10^{-2})	
	Albedo	Shading	Albedo	Shading	Albedo	Shading
SIRFS [1]	1.47	1.83	4.16	<u>1.68</u>	12.38	9.85
Zhou <i>et al.</i> [54]	2.52	2.29	-	-	-	-
Shi <i>et al.</i> [47]	2.78	1.26	5.03	2.40	14.65	12.00
DI [35]	2.77	1.54	5.86	2.95	15.26	13.28
Ma <i>et al.</i> * [33]	3.13	2.07	1.16	0.95	-	-
Janner <i>et al.</i> * [20]	3.36	1.95	2.10	1.03	-	-
CGIntrinsics [29]	1.67	1.27	3.19	2.21	12.87	13.76
US3D*† [31]	1.57	1.35	1.46	2.31	-	-
FFI-Net† [40]	1.11	0.93	2.91	3.19	10.14	11.39
PIE-Net† [13]	0.28	0.35	<u>1.36</u>	1.83	3.40	<u>4.93</u>
Ours†	<u>0.89</u>	0.34	0.97	0.37	<u>4.39</u>	3.02

* Unsupervised methods but finetuned on the dataset.

† Using pre-trained parameters.

Table 2: Results for MIT Intrinsic.



Fig. 5: Visual results (*image split*) for the MPI Sintel dataset [8]. Comparison to state-of-art methods Fan *et al.* [14] and LapPyrNet [12]. The method of Chen *et al.* [11] is provided as reference, since it uses *RGB-D* images as input.

	si-MSE(10^{-2})			si-LMSE(10^{-2})			DSSIM(10^{-2})		
	Albedo	Shading	Average	Albedo	Shading	Average	Albedo	Shading	Average
Retinex [18]	6.06	7.27	6.67	3.66	4.19	3.93	22.70	24.00	23.35
Lee <i>et al.</i> [26]	4.63	5.07	4.85	2.24	1.92	2.08	19.90	17.70	18.80
SIRFS [1]	4.20	4.36	4.28	2.98	2.64	2.81	21.00	20.60	20.80
Chenet <i>et al.</i> [11]	3.07	2.77	2.92	1.85	1.90	1.88	19.60	16.50	18.05
DI [35]	1.00	0.92	0.96	0.83	0.85	0.84	20.14	15.05	17.60
DARN [27]	1.24	1.28	1.26	0.69	0.70	0.70	12.63	12.13	12.38
Kimet <i>et al.</i> [22]	0.70	0.90	0.70	0.60	0.70	0.70	9.20	10.10	9.70
Fanet <i>et al.</i> † [14]	0.69	0.59	<u>0.64</u>	<u>0.44</u>	<u>0.42</u>	<u>0.43</u>	11.94	8.22	10.08
LapPyrNet† [12]	<u>0.66</u>	<u>0.60</u>	0.63	<u>0.44</u>	<u>0.42</u>	<u>0.43</u>	6.56	6.37	6.47
US3D*† [31]	1.59	1.48	1.54	0.87	0.81	0.84	17.97	14.74	16.35
Ours	0.57	0.71	<u>0.64</u>	0.29	0.38	0.34	<u>8.74</u>	<u>8.83</u>	<u>8.79</u>

* Unsupervised methods but finetuned on the dataset.

† Using pre-trained parameters.

Table 3: Numerical results for MPI-Sintel (image split).

RGB-D data to train their model. Furthermore, PoInt-Net demonstrates competitive performance for si-MSE shading and ranks second-best in DSSIM, highlighting its effectiveness in handling complex scenes with diverse lighting conditions. Notably, we do not compare with PIE-Net [13], due to it only reports results of scene split.

The qualitative results, presented in Figure 5, highlight PoInt-Net’s ability to generate high-quality results, particularly in terms of sharpness. This characteristic is advantageous for a point-based intrinsic network, where intrinsic features are processed on a point-by-point basis. In comparison to [11], which also utilizes depth information for intrinsic decomposition, PoInt-Net consistently demonstrates improved reflectance and shading estimation. This improvement can be attributed mainly to the adoption of point cloud representation. For more results, please refer to the supplementary material.

Inverender Recent advancements in NeRF methods have extended their functionalities to encompass scene intrinsic decomposition. In our evaluation, we contrast our approach with contemporary NeRF methodologies that similarly offer intrinsic decomposition capabilities (Table 4).

By leveraging a pre-trained PoInt-Net, our approach effectively estimates reflectance for the Inverender dataset, achieving improved intrinsic image decomposition results comparable to NeRFactor [52]. This highlights our model’s zero-shot estimation capability. Additionally, when fine-tuned on the Inverender dataset, PoInt-Net outperforms all competing methods, including NeRF-based techniques trained to over-fit a single scene (Figure 6). For more visual comparisons, please refer to the supplemental.




Fig. 6: Visual comparison with the NeRF based methods.

Method	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	MSE \downarrow	LMSE \downarrow
NeRFactor [52]	19.9167	0.9156	0.1354	0.0059	0.0210
PhySG [51]	23.3748	0.9231	0.1092	0.0034	0.0396
Inverender [53]	<u>26.3078</u>	<u>0.9380</u>	<u>0.0572</u>	0.0022	0.0226
Intrinsic-NeRF [50]	24.2642	0.9371	0.0880	<u>0.0021</u>	0.0173
IIW [6]	22.0284	0.9307	0.0847	0.0099	0.0120
CGIntrinsic [28]	20.1583	0.9209	0.0996	0.0129	0.0141
USI3D [31]	20.7571	0.9267	0.0887	0.0079	0.0149
Li <i>et al.</i> [30]	16.8167	0.8224	0.1661	0.0075	0.0089
PIE-Net [13]	18.8119	0.7870	0.2319	0.0163	0.0194
Ours (zero-shot)	22.4548	0.8986	0.1052	0.0059	<u>0.0071</u>
Ours (finetuned)	29.0404	0.9426	0.0543	0.0014	0.0015

Table 4: Numerical results for Inverender dataset (reflectance).




Fig. 7: Visual comparison on the IIW dataset [6]. Our results are generated using the same model that we released for review (only trained on single-level objects). Results of [14] are taken from their paper (trained on IIW dataset with edge guidance).

5.2 Real-world Generalization Performance

The real-world decomposition is conducted to further evaluate the generalization capability of our method.

IIW dataset. Figure 7 provides visual results on the IIW dataset, using our model (only trained on a single object level dataset). The comparison is conducted with the visual results provided by [14] and estimated by [31]. Our method shows comparable results on the scene-level dataset without prior learning. Notably, our method effectively reconstructs the shape of objects, such as the red door in the background, which previous work [14] struggled to accomplish.

The IIW dataset introduces WHDR as a performance metric, relying on human judgment scores to assess the accuracy of intrinsic image decompositions. Nevertheless, recent literature has raised doubts regarding the WHDR score’s reliability in accurately gauging the performance of intrinsic decomposition methods [9, 15, 16, 36], e.g., [36] achieves a WHDR of 25.7% by simply re-scaling the value of the original image into [0.55, 1]. In line with [9], we adjust our estimated reflectance by 0.5, yielding a comparable outcome, even when compared to models trained directly on the IIW dataset (Table 5).

	Bell <i>et al.</i> [6]	PIE-Net [13]	Careaga <i>et al.</i> [9]	USI3D [31]	Fan <i>et al.</i> [14]	IRISFormer [56]	Ours	Ours'
WHDR(%)	20.6	21.3	24.9	18.6	14.7	12.0	26.7	17.3
Trained on IIW	×	×	×	✓	✓	✓	×	×
Type of the trained data	S	S	S	S	S	S	O	O

Table 5: Results for the IIW dataset [6]. The letter 'S' and 'O' represent the scene-level and object-level datasets, respectively. Ours ' indicates the rescaled estimation, in line with [9].



Fig. 8: Real-world intrinsic estimation comparison. USI3D [31] and PIE-Net [13] are trained or finetuned on large scene level datasets. Although our model is exclusively trained on datasets featuring single objects, our approach demonstrates the capability to accurately compute intrinsic properties for previously unseen objects and scenes. Notably, red arrows emphasize our exceptional estimation performance.

Figure 8 demonstrates PoInt-Net’s generalization to real-world images, randomly sourced from the internet. The point clouds are generated based on estimated depth maps obtained from [41]. We employ the state-of-the-art unsupervised model [31] (finetuned on scene level datasets, e.g., CGI [28]& IIW [6]) as a reference. Even though PoInt-Net is trained on a dataset focused on single objects (as described in Section 5.1), it shows the capability to accurately estimate surface reflectance and shading not only for single objects but also for complex scenes, as evidenced by its performance in regions such as the shadow area between the sofa and the wall. For further details and results, please refer to the supplementary materials.

5.3 Point Cloud Matters

To demonstrate the effectiveness of point cloud representation in the task of intrinsic decomposition, we use deep CNNs (used in [30]) and Vision Transformers (ViT, used in [13, 56]) as backbones and carefully design two alternative frameworks with a similar structure as our PoInt-Net. The same training strategy is applied, i.e., pre-training on the

Backbone	MSE (10^{-2})		Model Size (MB)	Data Format
	Albedo	Shading		
DeepCNN [30]	1.36	0.67	~200	RGB
ViT [13, 56]	1.15	0.66	~1450	RGB
DeepCNN [30]	1.14	0.59	~200	RGB-D
ViT [13, 56]	1.58	1.06	~1450	RGB-D
Ours	0.89	0.34	~20	PC

Table 6: Evaluation of different data modalities on MIT-Intrinsic dataset, all pre-trained on ShapeNet-Intrinsic dataset.

Method	MSE (10^{-2})		Data Format
	Albedo	Shading	
Janner <i>et al.</i> [20]	3.03 (0.30)↓	1.77 (0.18)↓	RGB-D
LapPyrNet [12]	1.72 (0.05)↓	1.03 (0.21)↓	RGB-D
Fan <i>et al.</i> [14]	1.53 (0.04)↓	0.81 (0.07)↓	RGB-D
Ours	0.89	0.34	PC

Table 7: Evaluation of adding depth as extra input on MIT-Intrinsic dataset, all pre-trained on ShapeNet-Intrinsic dataset. (·) ↓ indicates error decreasing.

ShapeNet-Intrinsic dataset [20] and fine-tuning on MIT-Intrinsic dataset [18]. Table 6 shows the point cloud representation superiority in intrinsic estimation with even fewer parameters. Interestingly, a few weeks before our paper submission, a parallel work in robotic learning [55] showcases a notable trend, "Point cloud-based methods, even those with the simplest designs, frequently surpass their RGB and RGB-D counterparts in performance." Although we are working on different subjects, a similar trend is also observed in our case. We argue that the community should not overlook the uniqueness of point cloud representation.

5.4 Ablation Study

Learnable shader. As discussed in Section 3.2, shading depends on the surface geometry. Table 1 illustrates the significant impact of the shader in improving shading quality numerically. Additionally, Figures 3 and 4 demonstrate how the shader aids PoInt-Net in differentiating between invariant and ambient colors. This distinction is notably visible in areas such as the face and cloud segments of the "Sun" from MIT-Intrinsic, which have historically presented difficulties for numerous learning methods [4, 13, 14].

Using depth information. We recognize the necessity of depth information in our methodology, which brings up issues regarding the equitable comparison with earlier intrinsic image decomposition (IID) techniques, especially given that only a select few incorporate depth and their implementations are not openly accessible. To ensure fairness in our assessment, we included depth information as an extra input in three publicly accessible IID frameworks [12, 14, 20] (excluding PIE-Net [13] due to its unavailability of training code). Table 7 demonstrates that, with depth information, performance of existing methods increased, yet not to the same level of performance as ours. This result further emphasizes the uniqueness and effectiveness of our proposed PoInt-Net.

Depth quality. We assess the impact of depth quality in controlled conditions. PoIntNet demonstrates robust performance even with moderately noisy depth information. However, in highly noisy scenarios, there may be a decline in shading performance of up to 50%, while albedo performance remains relatively stable, with a decrease of less than 10%. These outcomes underscore the effectiveness of our method's design and are consistent with the fundamental principle of image formation, where shading is influenced by geometry. Please refer to supplemental for more details.

Backbone selection. PoInt-Net processes the point cloud by using multiple MLP layers. We acknowledge that, there are many powerful point-based backbones, such as Point++ [39]. We chose the relatively sampler backbone as it provides a good trade-off between computational efficiency and the ability to perform low-level vision tasks effectively, as evidenced in [48]. Please refer to supplemental for more details.

Model size of networks. As depicted in Figure 9, our method excels in intrinsic estimation performance despite having a compact model size. The reported model size is in accordance with its official pre-trained model specifications. Overall, our approach maintains a smaller model size compared to others that incorporate additional components like a mapping module [31], an adversarial network [57], a multi-scale CNN [12], and a transformer [13] in their network architecture.

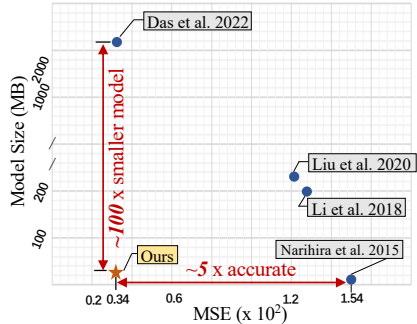


Fig. 9: Model size (MB) vs. MSE ($\times 10^2$) on shading for the MIT-intrinsic dataset. Our method is very efficient, outperforming the state-of-the-art model with just 1/100 of the model size and achieving 5 times the accuracy of a model with a similar size.

6 Conclusion

We introduced point intrinsic representation and PoInt-Net for 3D-based intrinsic decomposition. PoInt-Net employs a point cloud representation to efficiently decompose surface light direction, reflectance, and shading

maps, outperforming larger models on the MIT-intrinsic dataset while being highly efficient. Our experiments across different scenarios highlight its robustness and zero-shot generalization ability. Extra evaluation on the data modality evidenced point cloud as a valuable data format for intrinsic decomposition.

Limitations and Future Work. Although PoInt-Net demonstrates stability with non-Lambertian scenes and multiple light sources, explicitly investigating such complex scenarios in the future is essential to extend the versatility of our approach. Additionally, our method works with a 3D point cloud derived from a 2D image, such conversion is not able to perform spatial relations like occlusion. Creating 3D dataset containing point clouds and intrinsic components will benefit the research in the future. Finally, some of our generalization results are only evaluated visually, due to the unavailability of either ground-truth or proper evaluation metrics. A quantitative evaluation process is needed in our future research.

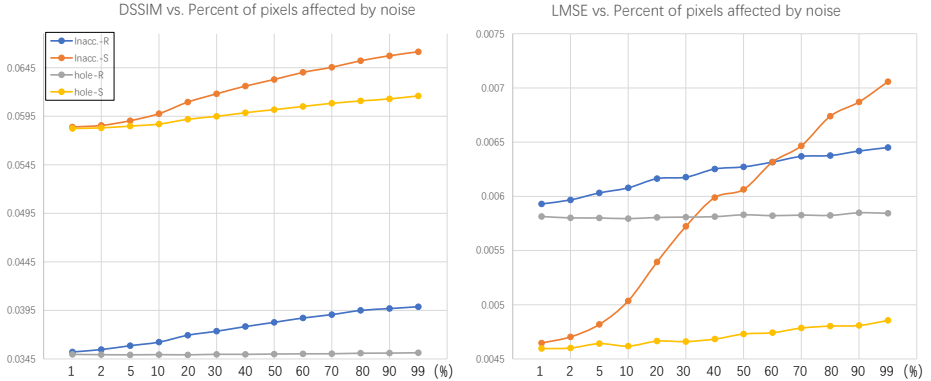


Fig. 10: DSSIM (left plot) and LMSE (right plot) on Shapenet-Intrinsic dataset vs. percent of pixels affected by noise. "Inacc." denotes inaccurate depth distribution, and "Hole" refers to the deletion in the depth. "R" represents reflectance, and "S" is for shading.

A Ablation Study

A.1 Depth quality

Depth information plays a critical role in our method, particularly in shading estimation. To assess how our method performs under varying depth quality, we conducted a quantitative study focusing on two common types of depth inconsistencies:

- **Inaccurate Depth:** Here, we simulate depth inaccuracy by adding random values to a number of pixels. This introduces a type of noise that mimics the effect of depth imprecision.
- **Holes in Depth:** For this scenario, we randomly set a number of pixels to zero. This action creates 'holes' in the depth data, replicating instances where depth information is missing or unreliable.

The number of pixels affected can be adjusted to various scales, simulating different levels of noise intensity. This allows for a more comprehensive evaluation of how our method is effected by depth inaccuracies ranging from minor to severe.

Figure 10 demonstrates albedo and shading performance under varying noise levels. Our method maintains stability with incomplete depth (e.g., holes). Yet, inaccurate depth significantly disrupts intrinsic estimation. Specifically, Shading estimation performance (LMSE) may decrease by 50% under inaccurate depth.

Discussion Since our method is robust to noisy depth. The inclusion of depth data does not pose a considerable limitation to our method. Our approach requires only an RGB image alongside its estimated depth, providing adaptability across a diverse range of application scenarios. Please refer to Section F to check our method's performance on complex scenes.

	CNN [3, 5, 14, 30]	Transformer [13, 56]	NeRF [50, 52]	Ours
Real-world generalize	some of them	✓	×	✓
Model size	>100MB	>1GB	-	20MB
Training Time	> 5 h	> 1 d	> 1 d	< 1.5 h
Diverse input size	×	×	-	✓

Table 8: Comparison with various IID methods using different backbones.

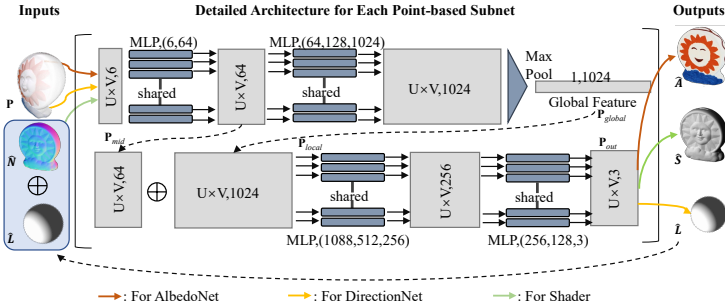


Fig. 11: The sub-net of PoInt-Net is based on a modified PointNet [38] architecture. For internal connections, refer to Figure 2 in the main paper. Three colors represent input/output for three subnets.

A.2 Backbone selection

Technically, several backbones can process RGB-D images. Table 8 summarizes their capabilities for IID. Our point cloud-based IID estimates real-world intrinsics with a smaller model size and shorter training time. Transformers and CNNs are not able to handle multiple size of images and are limited by the training time. NeRF based methods show good results yet lack of generalization ability.

B Details of Point Cloud Generation

Building Point Cloud. We generate the point cloud based on a pair of RGB image and its corresponding depth information. Normal information is pre-computed from the point cloud, which can be performed online or in advance for faster training and inference.

Camera Intrinsic Matrix. In Equation (4), the camera intrinsic matrix is employed to translate the point cloud from RGB-D image, i.e., images $\mathbf{I} = [\mathbf{I}_r, \mathbf{I}_g, \mathbf{I}_b] \in \mathbb{R}^{U \times V \times 3}$ and depth map $\mathbf{D} \in \mathbb{R}^{U \times V \times 1}$. For the MPI-Sintel dataset, we use the intrinsic matrix provide from the dataset. For other datasets, we use a default setting for the intrinsic matrix, where as the focal lengths and principal point are set as $(\frac{1}{2}U, \frac{1}{2}V)$, respectively.

C Point Intrinsic Net: Architecture

Figure 11 demonstrates the detailed architecture of the subnet of PoInt-Net. The input feature $\mathbf{P}_{input} \in \mathbb{R}^{U \times V \times 6}$ is first extracted by MLPs, and Max Pooling is employed

to extract the significant global feature $\mathbf{P}_{global} \in \mathbb{R}^{1 \times 1024}$. To maintain the ability in solving local information, the global feature is then repeated to the same number as the input points, and then concatenate with the inter-middle feature $\mathbf{P}_{mid} \in \mathbb{R}^{U \times V \times 64}$ to get the local feature $\mathbf{P}_{local} \in \mathbb{R}^{U \times V \times 1088}$. Finally, the output $\mathbf{P}_{output} \in \mathbb{R}^{U \times V \times 3}$ is obtained by MLPs which act as channel wise decoder. Since the three subnets share the similar network architecture, we use different color of the arrow to indicate the different input for the different nets, the color refer to the same color used in the Figure 2 of the main paper.

Discussion Our choice of module configuration is based on empirical evidence: 1) *Module Specialization.*: Given PointNet’s limitations as anrelatively weaker encoder, we use three specialized modules for separate estimation of albedo, lighting, and shading. 2) *Input Integration.* We found that concatenating surface normals with lighting, rather than multiplying them, yields improved outcomes. These insights will be added into the revision.

D Cross Color Ratios Loss Calculation

In paper, we propose the cross color ratios (CCR) loss to address the reflectance changes. Where the CCR $\{M_{RG}, M_{RB}, M_{GB}\}$ is calculated as:

$$M_{RG} = \frac{R_{p1}G_{p2}}{R_{p2}G_{p1}}, M_{RB} = \frac{R_{p1}B_{p2}}{R_{p2}B_{p1}}, M_{GB} = \frac{G_{p1}B_{p2}}{G_{p2}B_{p1}}, \quad (10)$$

where the $\{R_{p1}, G_{p1}, B_{p1}\}$ and $\{R_{p2}, G_{p2}, B_{p2}\}$ represent the R, G, B value at the two adjacent points p_1 and p_2 , respectively.

E Implementation Details

Environment Setting. PoInt-Net is developed using the PyTorch framework, enabling it to run on both CPU and GPU platforms. During evaluation, our method runs on a single NVIDIA 1080Ti GPU, with inference times dependent on the image size. For a 512×512 image, the average inference time is ~ 10 frames per second. It’s important to note that this inference time is achieved without any pruning techniques applied.

Training Details. For single object datasets, the point cloud is sampled by voxel down-sampling where the voxel size is set to 0.03. For scene datasets, the point cloud is resized to 64×64 points by average downsampling. The batch size is set based on the GPU memory accordingly. Adam [23] is employed as the optimizer. The learning rate is 3×10^{-4} . All networks are trained till convergences.

F More Results: Intrinsic Decomposition

Real World. The use of monocular depth estimation in PoInt-Net allows for effective decomposition of intrinsic properties in real-world scenes. Figure 12 showcases additional visual results across various environments, both indoor and outdoor. PoInt-Net



Fig. 12: Visualization results across different shapes, scenes, and domains in the real world. Our method consistently estimates the reasonable results. Note: our method is only trained on single-object level datasets.

consistently delivers reliable reflectance and shading estimates. Notably, it performs well even in challenging conditions like dark or low-resolution outdoor scenes, accurately decomposing invariant colors and shading.

As discussed in the paper, PoInt-Net is trained on single-object level datasets, i.e., ShapeNet-Intrinsic and MIT-Intrinsic. However, our method still achieves impressive intrinsic decomposition results across various shapes, scenes, and domains. These findings suggest that PoInt-Net is capable of solving intrinsic decomposition tasks at a very low-level, highlighting its robustness and versatility.

ShapeNet-Intrinsic. We present additional results on the ShapeNet dataset in Figure 13. These results showcase PoInt-Net’s ability to estimate surface light direction, with direct shading results displayed in the second column of the figure.

MIT-Intrinsic. Figure 14 presents additional results on the MIT-Intrinsic dataset, including ablation study results for the shader. The findings demonstrate that the shader significantly enhances PoInt-Net’s ability to distinguish between invariant color and illumination.

MPI-Sintel. Figure 15 presents extensional results that provide evidence for PoInt-Net’s ability to handle complex scenes during intrinsic decomposition training. The findings suggest that the architecture of PoInt-Net is capable of undertaking such tasks effectively.

Inverender. Figure 16 shows the additional results of the reflectance estimation on the Inverender dataset.




Fig. 13: Qualitative results on ShapeNet-Intrinsic dataset. The Direct shading is computed by surface normal and surface light direction.

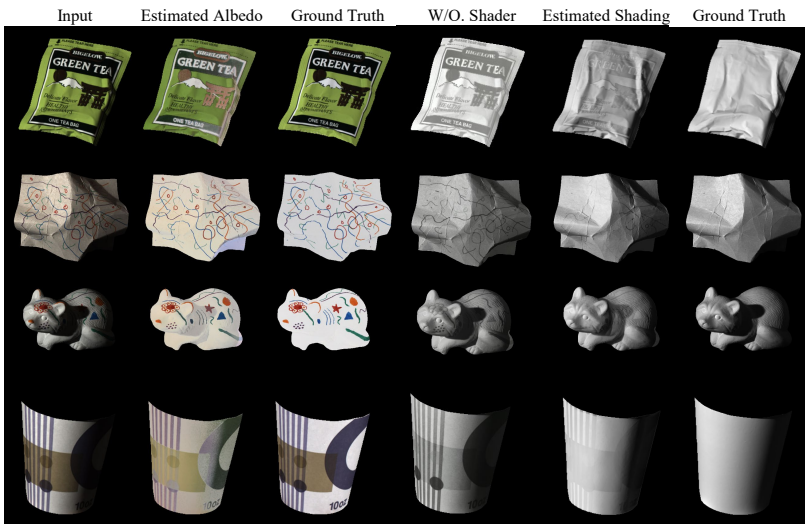


Fig. 14: Qualitative results on MIT-Intrinsic dataset. The ablation study on the shader (W/O. Shader) is applied.

IIW-Dataset. Figure 17 provides more visual results of the intrinsic estimation on IIW dataset [6].



Fig. 15: Visual results on MPI-Sintel dataset.




Fig. 16: Qualitative results on Inverender dataset.



Fig. 17: Qualitative results on IIW dataset. Note: Our method is not trained on IIW dataset.

References

1. Barron, J.T., Malik, J.: Shape, illumination, and reflectance from shading. *IEEE TPAMI* **37**(8), 1670–1687 (2015) [1](#), [3](#), [7](#), [9](#), [10](#)
2. Barrow, H., Tenenbaum, J., Hanson, A., Riseman, E.: Recovering intrinsic scene characteristics. *Comput. Vis. Syst* **2**, 3–26 (1978) [3](#)
3. Baslamisli, A.S., Das, P., Le, H.A., Karaoglu, S., Gevers, T.: Shadingnet: Image intrinsics by fine-grained shading decomposition. *IJCV* **129**(8), 2445–2473 (2021) [1](#), [3](#), [16](#)
4. Baslamisli, A.S., Groenestege, T.T., Das, P., Le, H.A., Karaoglu, S., Gevers, T.: Joint learning of intrinsic images and semantic segmentation. In: *ECCV* (2018) [8](#), [13](#)
5. Baslamisli, A.S., Le, H.A., Gevers, T.: Cnn based learning using reflection and retinex models for intrinsic image decomposition. In: *CVPR* (2018) [3](#), [8](#), [16](#)
6. Bell, S., Bala, K., Snavely, N.: Intrinsic images in the wild. *ACM Transactions on Graphics (TOG)* **33**(4), 159 (2014) [7](#), [11](#), [12](#), [19](#)
7. Buchsbaum, G.: A spatial processor model for object colour perception. *Journal of the Franklin Institute* **310**(1), 1–26 (1980) [3](#)
8. Butler, D.J., Wulff, J., Stanley, G.B., Black, M.J.: A naturalistic open source movie for optical flow evaluation. In: *ECCV* (2012) [7](#), [8](#), [10](#)
9. Careaga, C., Aksoy, Y.: Intrinsic image decomposition via ordinal shading. *ACM Transactions on Graphics* (2023) [11](#), [12](#)
10. Chang, A.X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., Xiao, J., Yi, L., Yu, F.: ShapeNet: An Information-Rich 3D Model Repository. Tech. Rep. arXiv:1512.03012 [cs.GR], Stanford University — Princeton University — Toyota Technological Institute at Chicago (2015) [7](#)
11. Chen, Q., Koltun, V.: A simple model for intrinsic image decomposition with depth cues. In: *ICCV* (2013) [3](#), [9](#), [10](#)
12. Cheng, L., Zhang, C., Liao, Z.: Intrinsic image transformation via scale space decomposition. In: *CVPR* (2018) [8](#), [10](#), [13](#), [14](#)
13. Das, P., Karaoglu, S., Gevers, T.: Pie-net: Photometric invariant edge guided network for intrinsic image decomposition. In: *CVPR* (2022) [3](#), [8](#), [9](#), [10](#), [11](#), [12](#), [13](#), [14](#), [16](#)
14. Fan, Q., Yang, J., Hua, G., Chen, B., Wipf, D.: Revisiting deep intrinsic image decompositions. In: *CVPR* (2018) [3](#), [6](#), [8](#), [9](#), [10](#), [11](#), [12](#), [13](#), [16](#)
15. Forsyth, D., Rock, J.J.: Intrinsic image decomposition using paradigms. *IEEE TPAMI* **44**(11), 7624–7637 (2021) [11](#)
16. Garces, E., Rodriguez-Pardo, C., Casas, D., Lopez-Moreno, J.: A survey on intrinsic images: Delving deep into lambert and beyond. *IJCV* **130**(3), 836–868 (2022) [1](#), [11](#)
17. Gevers, T., Smeulders, A.W.: Color-based object recognition. *Pattern recognition* **32**(3), 453–464 (1999) [6](#)
18. Grosse, R., Johnson, M.K., Adelson, E.H., Freeman, W.T.: Ground truth dataset and baseline evaluations for intrinsic image algorithms. In: *ICCV* (2009) [3](#), [7](#), [8](#), [9](#), [10](#), [13](#)
19. Hachama, M., Ghanem, B., Wonka, P.: Intrinsic scene decomposition from rgb-d images. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 810–818 (2015) [3](#)
20. Janner, M., Wu, J., Kulkarni, T.D., Yildirim, I., Tenenbaum, J.: Self-supervised intrinsic image decomposition. In: *NIPS* (2017) [1](#), [3](#), [5](#), [7](#), [8](#), [9](#), [13](#)
21. Kajiya, J.T.: The rendering equation. In: *SIGGRAPH*. pp. 143–150 (1986) [3](#)
22. Kim, S., Park, K., Sohn, K., Lin, S.: Unified depth prediction and intrinsic image decomposition from a single image via joint convolutional neural fields. In: *ECCV* (2016) [3](#), [9](#), [10](#)

23. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014) [7](#), [17](#)
24. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: NIPS (2012) [8](#)
25. Laffont, P.Y., Bousseau, A., Paris, S., Durand, F., Drettakis, G.: Coherent intrinsic images from photo collections. ACM ToG **31**(6) (2012) [3](#)
26. Lee, K.J., Zhao, Q., Tong, X., Gong, M., Izadi, S., Lee, S.U., Tan, P., Lin, S.: Estimation of intrinsic image sequences from image+ depth video. In: ECCV (2012) [3](#), [10](#)
27. Lettry, L., Vanhoey, K., Van Gool, L.: Darn: a deep adversarial residual network for intrinsic image decomposition. In: WACV (2018) [10](#)
28. Li, Z., Snaveley, N.: Cgintrinsics: Better intrinsic image decomposition through physically-based rendering. In: ECCV (2018) [3](#), [11](#), [12](#)
29. Li, Z., Snaveley, N.: Learning intrinsic image decomposition from watching the world. CVPR (2018) [8](#), [9](#)
30. Li, Z., Yu, T.W., Sang, S., Wang, S., Song, M., Liu, Y., Yeh, Y.Y., Zhu, R., Gundavarapu, N., Shi, J., et al.: Openrooms: An open framework for photorealistic indoor scene datasets. In: CVPR. pp. 7190–7199 (2021) [1](#), [3](#), [11](#), [12](#), [13](#), [16](#)
31. Liu, Y., Li, Y., You, S., Lu, F.: Unsupervised learning for intrinsic image decomposition from a single image. In: CVPR (2020) [7](#), [8](#), [9](#), [10](#), [11](#), [12](#), [14](#)
32. Luo, J., Huang, Z., Li, Y., Zhou, X., Zhang, G., Bao, H.: Niid-net: adapting surface normal knowledge for intrinsic image decomposition in indoor scenes. IEEE TVCG **26**(12), 3434–3445 (2020) [1](#)
33. Ma, W.C., Chu, H., Zhou, B., Urtasun, R., Torralba, A.: Single image intrinsic decomposition without a single intrinsic image. In: ECCV (2018) [3](#), [8](#), [9](#)
34. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: ECCV. pp. 405–421. Springer (2020) [3](#)
35. Narihira, T., Maire, M., Yu, S.X.: Direct intrinsics: Learning albedo-shading decomposition by convolutional regression. In: ICCV (2015) [3](#), [9](#), [10](#)
36. Nestmeyer, T., Gehler, P.V.: Reflectance adaptive filtering improves intrinsic image estimation. In: CVPR (2017) [11](#)
37. Nicodemus, F.E.: Directional reflectance and emissivity of an opaque surface. Applied optics **4**(7), 767–775 (1965) [4](#)
38. Qi, C.R., Su, H., Mo, K., Guibas, L.J.: Pointnet: Deep learning on point sets for 3d classification and segmentation. In: CVPR. pp. 652–660 (2017) [5](#), [16](#)
39. Qi, C.R., Yi, L., Su, H., Guibas, L.J.: Pointnet++: Deep hierarchical feature learning on point sets in a metric space. arXiv preprint arXiv:1706.02413 (2017) [14](#)
40. Qian, Y., Shi, M., Kamarainen, J.K., Matas, J.: Fast fourier intrinsic network. In: WACV (2021) [3](#), [8](#), [9](#)
41. Ranftl, R., Lasinger, K., Hafner, D., Schindler, K., Koltun, V.: Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. IEEE TPAMI **44**(3) (2022) [2](#), [4](#), [7](#), [12](#)
42. Rother, C., Kiefel, M., Zhang, L., Schölkopf, B., Gehler, P.V.: Recovering intrinsic images with a global sparsity prior on reflectance. In: NIPS (2011) [3](#)
43. Sato, S., Yao, Y., Yoshida, T., Kaneko, T., Ando, S., Shimamura, J.: Unsupervised intrinsic image decomposition with lidar intensity. In: CVPR. pp. 13466–13475 (2023) [3](#)
44. Sengupta, S., Kanazawa, A., Castillo, C.D., Jacobs, D.W.: Sfsnet: Learning shape, reflectance and illuminance of faces in the wild. In: CVPR. pp. 6296–6305 (2018) [1](#), [3](#)
45. Shen, L., Yeo, C.: Intrinsic images decomposition using a local and global sparse representation of reflectance (2011) [3](#)

46. Shen, L., Yeo, C., Hua, B.S.: Intrinsic image decomposition using a sparse representation of reflectance. *IEEE TPAMI* **35**(12), 2904–2915 (2013) [3](#)
47. Shi, J., Dong, Y., Su, H., Stella, X.Y.: Learning non-lambertian object intrinsics across shapenet categories. In: *CVPR* (2017) [3](#), [9](#)
48. Xing, X., Qian, Y., Feng, S., Dong, Y., Matas, J.: Point cloud color constancy. In: *CVPR* (2022) [2](#), [14](#)
49. Xu, Q., Xu, Z., Philip, J., Bi, S., Shu, Z., Sunkavalli, K., Neumann, U.: Point-nerf: Point-based neural radiance fields. In: *CVPR* (2022) [2](#)
50. Ye, W., Chen, S., Bao, C., Bao, H., Pollefeys, M., Cui, Z., Zhang, G.: Intrinsicnerf: Learning intrinsic neural radiance fields for editable novel view synthesis. In: *ICCV*. pp. 339–351 (2023) [3](#), [7](#), [11](#), [16](#)
51. Zhang, K., Luan, F., Wang, Q., Bala, K., Snavely, N.: Physg: Inverse rendering with spherical gaussians for physics-based material editing and relighting. In: *CVPR*. pp. 5453–5462 (2021) [11](#)
52. Zhang, X., Srinivasan, P.P., Deng, B., Debevec, P., Freeman, W.T., Barron, J.T.: Nerfactor: Neural factorization of shape and reflectance under an unknown illumination. *ACM TOG* **40**(6), 1–18 (2021) [3](#), [7](#), [10](#), [11](#), [16](#)
53. Zhang, Y., Sun, J., He, X., Fu, H., Jia, R., Zhou, X.: Modeling indirect illumination for inverse rendering. In: *CVPR*. pp. 18643–18652 (2022) [7](#), [11](#)
54. Zhou, T., Krahenbuhl, P., Efros, A.A.: Learning data-driven reflectance priors for intrinsic image decomposition. In: *CVPR* (2015) [3](#), [9](#)
55. Zhu, H., Wang, Y., Huang, D., Ye, W., Ouyang, W., He, T.: Point cloud matters: Re-thinking the impact of different observation spaces on robot learning. *arXiv preprint arXiv:2402.02500* (2024) [13](#)
56. Zhu, R., Li, Z., Matai, J., Porikli, F., Chandraker, M.: Irisformer: Dense vision transformers for single-image inverse rendering in indoor scenes. In: *CVPR*. pp. 2822–2831 (2022) [1](#), [3](#), [12](#), [13](#), [16](#)
57. Zoran, D., Isola, P., Krishnan, D., Freeman, W.T.: Learning ordinal relationships for mid-level vision. In: *CVPR* (2015) [3](#), [14](#)