

# Intrinsic Image Decomposition for Robust Self-supervised Monocular Depth Estimation on Reflective Surfaces

Wonhyeok Choi\*, Kyumin Hwang\*, Minwoo Choi, Kiljoon Han,  
Wonjoon Choi, Mingyu Shin, Sunghoon Im†

Daegu Gyeongbuk Institute of Science and Technology, South Korea  
{smu06117, kyumin, subminu, kiljoon\_h, wjchoi, alsrb4446, sunghoonim}@dgist.ac.kr

## Abstract

Self-supervised monocular depth estimation (SSMDE) has gained attention in the field of deep learning as it estimates depth without requiring ground truth depth maps. This approach typically uses a photometric consistency loss between a synthesized image, generated from the estimated depth, and the original image, thereby reducing the need for extensive dataset acquisition. However, the conventional photometric consistency loss relies on the Lambertian assumption, which often leads to significant errors when dealing with reflective surfaces that deviate from this model. To address this limitation, we propose a novel framework that incorporates intrinsic image decomposition into SSMDE. Our method synergistically trains for both monocular depth estimation and intrinsic image decomposition. The accurate depth estimation facilitates multi-image consistency for intrinsic image decomposition by aligning different view coordinate systems, while the decomposition process identifies reflective areas and excludes corrupted gradients from the depth training process. Furthermore, our framework introduces a pseudo-depth generation and knowledge distillation technique to further enhance the performance of the student model across both reflective and non-reflective surfaces. Comprehensive evaluations on multiple datasets show that our approach significantly outperforms existing SSMDE baselines in depth prediction, especially on reflective surfaces.

## Introduction

Self-supervised monocular depth estimation (SSMDE) is a method to estimate depth using only sequential images during training, without the need for ground truth depth maps. This method has significantly expanded its applicability in commercial datasets, eliminating the need for expensive data acquisition methods like LiDAR. The key to the success of SSMDE lies in its ability to synthesize one image into another using estimated depth from sequential images and enforcing photometric consistency constraints. Over the years, SSMDE has consistently improved and is increasingly being applied in industries such as autonomous driving, robotics, and mixed reality.

\*These authors contributed equally.

†S. Im is the corresponding author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

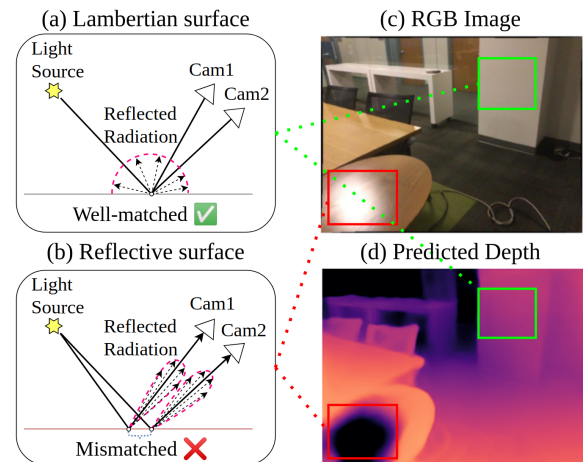


Figure 1: Illustrations of violations of the photometric consistency principle on non-Lambertian surfaces.

Recently, the focus of SSMDE research has shifted from merely enhancing depth estimation performance to investigating intrinsic network issues, particularly the generalization of trained models to changes in domain (Saunders, Vogiatzis, and Manso 2023; Gasperini et al. 2023; Bae, Hwang, and Im 2023). Some approaches have highlighted a critical issue: SSMDE often relies on the Lambertian assumption as a learning strategy, which does not hold true for non-Lambertian surfaces like mirrors and reflective materials commonly found in indoor scenes. These surfaces violate the photometric constancy principle, a cornerstone assumption of SSMDE, significantly degrading depth quality as shown in Figure 1. To address these issues, recent methods have been developed to localize the regions that violate these assumptions using additional segmentation masks (Costanzino et al. 2023), based on the uncertainty of multiple networks (Shi et al. 2023), and to selectively refine the depth in these areas using techniques such as inpainting and 3D mesh rendering. However, the computational intensity of 3D mesh rendering and ensemble uncertainty, along with the requirement for additional annotations, diminish the advantages of self-supervised approaches.

To address these challenges, we propose a novel end-to-end, plug-and-play method that leverages intrinsic im-

age decomposition, specifically utilizing the intrinsic residual model (Tominaga 1994; Maxwell, Friedhoff, and Smith 2008). This method decomposes an image into diffuse and residual components, the latter of which often varies with color or luminosity relative to the camera viewpoint—common complication in training SSMDEs. Our method specifically addresses these variations by localizing per-pixel reflective regions. It compares the photometric error between the diffuse image, where the residual component has been removed, and the original RGB image. Reflective regions are identified through this comparison and then excluded from the depth learning process. By systematically excluding these problematic reflective regions, our method significantly improves depth accuracy.

To demonstrate the effectiveness of the proposed method, we conduct evaluations using three representative SSMDE baselines (Godard et al. 2019; Lyu et al. 2021; Zhao et al. 2022) across three datasets (Dai et al. 2017; Shotton et al. 2013; Ramirez et al. 2023). These datasets include various realistic indoor scenes comprising numerous non-reflective and reflective surfaces. The results demonstrate significant improvements in depth prediction accuracy, particularly on reflective and non-Lambertian surfaces. Our contributions are threefold as follows:

- To the best of our knowledge, we are the first to introduce a novel SSMDE that employs the self-supervised intrinsic image decomposition for robust depth estimation.
- We present a method for identifying reflective regions using dynamic viewpoint monocular image sequences without requiring explicitly labeled data.
- We propose a reflection-aware depth training process that enhances the accuracy of depth prediction in our SSMDE framework.

## Related Work

### Self-supervised Monocular Depth Estimation

Monodepth (Godard, Mac Aodha, and Brostow 2017), a pioneering works in self-supervised monocular depth estimation, minimizes the photometric error by synthesizing a left image from a right image using stereo geometry. Following this work, interest in SSMDE has surged, leading to significant developments using CNN models (Zhou et al. 2017; Godard et al. 2019; Guizilini et al. 2020; Lyu et al. 2021; Zhou et al. 2021) and ViT models (Zhao et al. 2022; Bae, Moon, and Im 2023; Wang et al. 2024; Zhang et al. 2023).

In addition to structural advancements, substantial efforts have been made to address inherent challenges in SSMDE, such as scale-ambiguity (Guizilini et al. 2020, 2022), out-of-distribution generalization (Saunders, Vogiatzis, and Manso 2023; Gasperini et al. 2023; Bae, Hwang, and Im 2023). Recent works (Costanzino et al. 2023; Shi et al. 2023) have tackled the degradation of depth accuracy in non-Lambertian regions, a common issue in models trained under photometric consistency constraints. These works leverage additional resources such as ensemble uncertainty (Shi et al. 2023) or segmentation mask annotations (Costanzino et al. 2023) to identify the non-Lambertian regions. They refine the depth estimates of these regions using techniques

like inpainting-based color augmentation or 3D mesh rendering.

## Intrinsic Image Decomposition

Intrinsic image decomposition is a deeply researched field with a history spanning nearly five decades. Early approaches (Land and McCann 1971; Shen, Tan, and Lin 2008; Zhao et al. 2012; Rother et al. 2011; Shen and Yeo 2011; Garces et al. 2012) utilize prior knowledge, such as the assumption that large image gradients correspond to changes in albedo, to perform this task through non-learning-based method optimization. With the advent of labeled datasets, the focus has shifted towards learning-based methods (Baron and Malik 2014; Narihira, Maire, and Yu 2015; Shi et al. 2017; Kim et al. 2016), which primarily rely on ground truth intrinsic components. Some recent state-of-the-art learning-based intrinsic decomposition methods (Careaga and Aksoy 2023; Luo et al. 2024; Zeng et al. 2024) have demonstrated impressive performance, enabling natural and high-quality image rendering.

On the other hand, some studies (Weiss 2001; Sunkavalli et al. 2007; Laffont et al. 2012; Laffont and Bazin 2015; Duchêne et al. 2015; Ma et al. 2018; Yu and Smith 2019) have proposed self-supervised methods for training networks using only RGB image sequences without ground-truth information on intrinsic components. These studies primarily utilize image sequences with a fixed viewpoint and a static scene, but with varying illumination, such as time-lapse videos or synthetic image sequences. They exploit the assumption that the albedo remains constant in such static scenes to conduct the learning process. Notably, Yu and Smith (2019) is the first to use the multi-view image sequences but this method still relies on the depth annotations to obtain the pre-trained multi-view stereo network. To the best of our knowledge, our method is the first fully self-supervised framework that does not require any explicitly labeled data while leveraging dynamic viewpoint monocular image sequences.

## Method

Our framework consists of two primary branches: the intrinsic image decomposition branch and the depth estimation branch as shown in Figure 2. These branches collaboratively estimate intrinsic components and depth from RGB images, respectively. We replicate the architecture of the depth decoder in the intrinsic image decoder to additionally predict the intrinsic images. Our method specifically aims to enhance the depth prediction accuracy in the reflective regions, which are problematic for depth training, by leveraging the intrinsic components.

## Preliminary

**Basic assumption of general SSMDEs** To train the monocular depth estimation network in a self-supervised manner using photometric loss, conventional works typically assume that the following three conditions must be satisfied:




Figure 2: Overall structure of our method.

- a1** The scene being modeled is static while camera is in motion (Godard et al. 2019; Zhou et al. 2017).
- a2** Objects are perceived to maintain a consistent level of brightness across different views (Godard et al. 2019).
- a3** All surfaces in the scene are Lambertian (Godard et al. 2019; Zhou et al. 2017).

Nevertheless, outdoor scene datasets contain elements that violate these assumptions, such as moving vehicles or pedestrians, rapidly changing weather conditions. Some SSMDE methods use an auto-mask (Godard et al. 2019) to exclude pixel-level regions that violate the static scene assumption (**a1**). While the Lambertian surface assumption (**a3**) is not universally valid, it still provides a reasonable approximation for most outdoor scenes. However, this assumption (**a3**) fails notably in indoor environments where reflective surfaces, mirrors, and transparent objects can severely harm the depth accuracy due to the diverse object types and lighting conditions.

### General self-supervised intrinsic image decomposition

Self-supervised intrinsic image decomposition methods typically use time-lapse videos or synthetic datasets comprising image sequences under varying lighting conditions. These works aim to decompose an RGB image  $I$  into its albedo  $A$  and shading  $S$  components as follows:

$$I = A \odot S, \quad (1)$$

where  $\odot$  represents the element-wise multiplication operator. Given that albedo is inherently invariant to illumination changes, it remains consistent across all images within the same sequence. This consistency in albedo allows these methods to effectively decompose the intrinsic components of the images.

### Training of Intrinsic Image Decomposition Branch

**Our formulation of intrinsic decomposition** The intrinsic diffuse model in Equation 1 cannot be directly applied to a scenario of SSMDE due to these problems:

- p1** The lighting conditions in the image sequence are static ( $\therefore$  **a2**).
- p2** The camera coordinates within the image sequence are unaligned ( $\therefore$  **a1**).
- p3** The intrinsic diffuse model assumes all surfaces in the scene are Lambertian.

To address these challenges, we newly design the training method, facilitating the learning of intrinsic image decomposition in SSMDE scenarios, where only unaligned RGB image sequences are available.

To align the camera coordinate of multiple images in a sequence, we use the view-synthesized method (Godard et al. 2019), which is generally adopted in SSMDEs. We also reformulate the intrinsic image decomposition to decompose the diffuse and residual, which are the view-independent and view-dependent components, respectively. Instead of using the intrinsic diffuse model, we adopt the intrinsic residual model (Tominaga 1994; Maxwell, Friedhoff, and Smith 2008), which accommodates the specular reflections that vary with the camera view direction. We formulate the residual model in logarithmic space, where the RGB image exhibits shadow and illumination-invariant chromacity (Maxwell, Friedhoff, and Smith 2008). This allows for the decomposition of the diffuse image and residual image from the original image as follows:

$$\log I = \log L + \log R, \text{ where } L = A \odot S, \quad (2)$$

where  $L$  is the diffuse image, defined as the element-wise multiplication of albedo and shading images. The diffuse component remains constant irrespective of the camera view direction.  $R$  is the residual component varying depending on the camera view direction.

**Loss term of intrinsic network** The objective of the intrinsic image decomposition network is to appropriately decompose the diffuse images and residual images from a given reference image  $I_r$ , a source image  $I_s$  (or source images), relative camera pose between reference view and source view  $[R|t]$ , and their camera intrinsic  $K$ . The network is designed to predict the diffuse component  $L_r$  and the residual component  $R_r$  from the reference image  $I_r$ . Similarly, it predicts the source diffuse  $L_s$  and residual  $R_s$  from the source image  $I_s$ . The predicted intrinsic components are required to flawlessly reconstruct the original RGB image by Equation 2. Accordingly, we design the reconstruction loss  $\mathcal{L}_{\text{recon}}$  based on this formulation as follows:

$$\mathcal{L}_{\text{recon}} = |\log(I_r) - \log(L_r) - \log(R_r)|_1. \quad (3)$$

Corresponding surface points in diffuse images must have the same color and luminance because diffuse components do not vary with the camera view direction. To enable the model to leverage the diffuse consistency with static camera coordinates, the model aligns the camera coordinate of the source RGB image and diffuse image  $\Gamma \in \{I, L\}$  to the reference camera coordinate by using view synthesis (Godard et al. 2019) as follows:

$$\Gamma_{s_{2r}}[u, v] = \Gamma_s[\hat{u}, \hat{v}], \Gamma \in \{I, L\},$$

$$(\hat{u}, \hat{v}) = \text{proj} \left( \mathbf{K}[\mathbf{R}|\mathbf{t}]_{r_{2s}} D_r[u, v] \mathbf{K}^{-1} \begin{pmatrix} u \\ v \\ 1 \end{pmatrix} \right), \quad (4)$$

where  $(u, v)$  and  $(\hat{u}, \hat{v})$  represent the corresponding indices of image coordinates in the warping process. The function  $\text{proj}(\cdot)$  serves as the projection function that maps homogeneous coordinates to image coordinates. Additionally,  $D_r$  is the reference depth prediction from the depth branch. If the depth prediction  $D_r$  is sufficiently accurate, the aligned diffuse image  $L_r$  and  $L_{s_{2r}}$  should be identical, reflecting diffuse consistency. To leverage this property, we design the cross-reconstruction loss  $\mathcal{L}_{\text{cross}}$  as follows:

$$\mathcal{L}_{\text{cross}} = |\log(I_r) - \log(L_{s_{2r}}) - \log(R_r)|_1. \quad (5)$$

The proposed loss terms  $\mathcal{L}_{\text{recon}}$  and  $\mathcal{L}_{\text{cross}}$  are designed to ensure that the outputs of intrinsic model adhere to the diffuse-residual model formulation and diffuse consistency.

However, these constraints allow for numerous optimal solutions, including some that are unsuitable. For example, an outcome where the diffuse image is a plain white image and the residual image matches the original RGB image would reduce these loss terms to zero, yielding nonsensical intrinsic components. To ensure that the model does not provide such abnormal intrinsic images, we design additional contrastive loss term  $\mathcal{L}_{\text{cts}}$  that encourages the model to decompose a different diffuse image for each image sequence as follows:

$$\mathcal{L}_{\text{cts}} = \sum_{i=0}^b \sum_{\substack{j=0 \\ i \neq j}}^b \max(\delta - \|L_{s_{2r}}^i - L_r^j\|_2, 0), \quad (6)$$

where  $b$  is the batch size and  $\delta$  is the proper margin to regulate the distance between diffuse images from different images. The total loss for intrinsic image decomposition  $\mathcal{L}_{\text{itr}}$  is formed as follows:

$$\mathcal{L}_{\text{itr}} = \lambda_{\text{recon}} \mathcal{L}_{\text{recon}} + \lambda_{\text{cross}} \mathcal{L}_{\text{cross}} + \lambda_{\text{cts}} \mathcal{L}_{\text{cts}}. \quad (7)$$

### Training of Depth Estimation Branch

Most SSMDEs use a combination of an L1 and single scale SSIM (Wang et al. 2004) as the photometric error, which is proposed by Godard, Mac Aodha, and Brostow (2017):

$$\mathcal{P}(\Gamma_r, \Gamma_{s_{2r}}) = M \odot \left( \alpha(1 - \text{SSIM}(\Gamma_r, \Gamma_{s_{2r}}))/2 \right. \\ \left. + (1 - \alpha)\|\Gamma_r - \Gamma_{s_{2r}}\|_1 \right), \Gamma \in \{I, L\}, \quad (8)$$

where  $\Gamma_{s_{2r}}$  is a set that includes warped image or diffuse component  $\{I_{s_{2r}}, L_{s_{2r}}\}$  according to Equation 4. Here,  $M$

is the auto-mask (Godard et al. 2019), and  $\alpha$  is a parameter within  $[0, 1]$  balancing the contribution of the two loss terms.

While general SSMDEs minimize the photometric error  $\mathcal{P}(I_r, I_{s_{2r}})$ , the mask  $M$  is employed to selectively filter out the gradients from regions affected by moving objects or static camera movements, addressing the assumption **a1**. However, this method does not account for the Lambertian assumption **a3**. To identify for regions that substantially violate the assumption **a3**, we compute the photometric error for both aligned RGB pairs and diffuse pairs as follows:

$$E_I = \mathcal{P}(I_r, I_{s_{2r}}), E_L = \mathcal{P}(L'_r, L'_{s_{2r}}), \\ L' = \exp(\log(I) - \log(R)) \quad (\cdot: \text{Equation 2}), \quad (9)$$

where  $E_I$  and  $E_L$  are the photometric errors between RGB pairs and diffuse pairs, respectively. Note that we use the pseudo-diffuse  $L'$  computed from the RGB image and the predicted residual image rather than the predicted diffuse image  $L$  due to the leak of high-frequency details during image-to-image translation.

Since the residual image varies with color or luminosity relative to the camera viewpoint, photometric errors can arise from the residual component. As a result, the photometric error between a pair of diffuse images, where the residual components have been removed from the color images, is lower than the photometric error between the original color images. We use the Mahalanobis distance (Mahalanobis 1936) to pinpoint and exclude problematic pixel-level regions from the training where the difference between the image error and diffuse error in the photometric error space is significant as follows:

$$M_R[u, v] = \begin{cases} 0, & \text{if } z_L[u, v] < z_I[u, v] \\ 1, & \text{else} \end{cases}, \text{ where}$$

$$z_I = \sum_{u, v} \sqrt{(E_I[u, v] - \mu_I) \sigma_I^{-1} (E_I[u, v] - \mu_I)^T}, \quad (10)$$

$$z_L = \sum_{u, v} \sqrt{(E_L[u, v] - \mu_L) \sigma_L^{-1} (E_L[u, v] - \mu_L)^T},$$

where  $\{\mu_I, \mu_L\}$  represent the mean vectors and  $\{\sigma_I, \sigma_L\}$  denote the covariance vector of per-pixel vectors from the photometric error of RGB and diffuse images, respectively.  $M_R$  is a per-pixel binary mask, similar in function to the auto-mask  $M$ , designed to exclude non-diffuse regions. Importantly, the Mahalanobis distance is employed as a metric to assess the divergence between a sample point and a distribution, allowing the network to selectively identify and filter out specular regions. This is achieved by recognizing areas where the photometric error significantly decreases following the removal of the residual component in specific pixels.

Similar to the auto-masking scheme, our method mitigates the influence of the attenuated gradients from reflective regions during training. This is achieved by performing element-wise multiplication of the photometric error with the obtained mask  $M_R$  as follows:

$$\mathcal{L}_{\text{depth}} = M_R \odot \mathcal{P}(I_r, I_{s_{2r}}). \quad (11)$$

Dataset	Method	Abs Rel ↓	Sq Rel ↓	RMSE ↓	RMSE log ↓	$\delta < 1.25 \uparrow$	$\delta < 1.25^2 \uparrow$	$\delta < 1.25^3 \uparrow$
ScanNet-Ref. Test set	Monodepth2	0.181	0.160	0.521	0.221	<b>0.758</b>	0.932	0.976
	Monodepth2 + <i>Ours</i>	<b>0.163</b>	<b>0.103</b>	<b>0.472</b>	<b>0.205</b>	0.757	<b>0.940</b>	<b>0.986</b>
	HRDepth	0.182	0.168	0.530	0.225	0.749	<b>0.937</b>	0.979
	HRDepth + <i>Ours</i>	<b>0.165</b>	<b>0.104</b>	<b>0.471</b>	<b>0.206</b>	<b>0.752</b>	0.934	<b>0.984</b>
	MonoViT	0.154	0.129	0.458	0.197	<b>0.822</b>	0.955	0.979
	MonoViT + <i>Ours</i>	<b>0.142</b>	<b>0.093</b>	<b>0.421</b>	<b>0.180</b>	0.815	<b>0.957</b>	<b>0.985</b>
ScanNet-Ref. Val. set	Monodepth2	0.206	0.227	0.584	0.246	0.750	0.912	0.961
	Monodepth2 + <i>Ours</i>	<b>0.158</b>	<b>0.100</b>	<b>0.462</b>	<b>0.200</b>	<b>0.769</b>	<b>0.939</b>	<b>0.986</b>
	HRDepth	0.213	0.244	0.605	0.255	0.741	0.906	0.961
	HRDepth + <i>Ours</i>	<b>0.160</b>	<b>0.102</b>	<b>0.463</b>	<b>0.201</b>	<b>0.773</b>	<b>0.939</b>	<b>0.986</b>
	MonoViT	0.179	0.206	0.557	0.227	0.819	0.930	0.963
	MonoViT + <i>Ours</i>	<b>0.139</b>	<b>0.106</b>	<b>0.446</b>	<b>0.179</b>	<b>0.840</b>	<b>0.953</b>	<b>0.984</b>

Table 1: Main results on the ScanNet-Reflection Test and Validation sets.

Lastly, our framework is designed to simultaneously learn intrinsic image decomposition and depth estimation by summing the losses  $\mathcal{L}_{\text{itr}}$  and  $\mathcal{L}_{\text{depth}}$  as follows:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{itr}} + \mathcal{L}_{\text{depth}}. \quad (12)$$

## Experiments

### Baselines and Training Setups

We incorporate experiments with three leading architectures that have demonstrated outstanding performance in previous research: Monodepth2 (Godard et al. 2019), HRDepth (Lyu et al. 2021), and MonoViT (Zhao et al. 2022). Each baseline was trained in a self-supervised manner using the photometric error minimization strategy initially proposed by Monodepth2, and they are referred to using the general label [Baseline]. To demonstrate the effectiveness of our method relative to these baselines, we adapt the decoder of each baseline architecture to our training protocol, incorporating intrinsic networks. This modified training approach is denoted as [Baseline + Ours].

### Implementation Details

The batch size for training is configured differently across the architecture: 12 for Monodepth2 and HRDepth, and 8 for MonoViT. All networks start with a learning rate of  $1e-4$ , which is decreased by a factor of 10 following the 26th and 36th epochs, within a total training duration of 41 epochs. The balancing weights for photometric error and smoothness loss follow the policy established by Monodepth2. For resolution, a training dimension of  $384 \times 288$  is utilized, and no post-processing techniques, such as median-scaling, are employed in any of the quantitative evaluations. Additionally, similar to 3D Distillation (Shi et al. 2023), camera poses are assumed to be known during training. Depth boundaries were set between a minimum of 0.1m and a maximum of 10m. Weight balancing parameters for intrinsic branch  $\{\lambda_{\text{recon}}, \lambda_{\text{cross}}, \lambda_{\text{cts}}\}$  are set at  $\{1.0, 1.0, 0.01\}$  respectively, with  $\delta$  for  $\lambda_{\text{cts}}$  fix at 5.0. For the quantitative evaluation of depth estimation, we utilized metrics that are widely recognized and standard within the field (Eigen, Puhrsch, and Fergus 2014; Geiger, Lenz, and Urtasun 2012).

### Datasets

- **ScanNetv2** (Dai et al. 2017) consists of indoor scenes captured by an RGB-D scanner, including 1,201 training scenes, 312 validation scenes, and 100 test scenes. To simulate training scenarios dominated by non-Lambertian surfaces and to ensure fair comparison with 3D Distillation (Shi et al. 2023), we utilize the training triplet splits provided by the authors, encompassing 45,539 training images, 439 validation images, and 121 test images.
- **7-Scenes** (Shotton et al. 2013) comprises seven indoor scenes and has been extensively utilized in 3D reconstruction and camera localization studies. To evaluate the cross-dataset generalization of the proposed method, we follow the evaluation splits specified by Long et al. (2021); Bae, Budvytis, and Cipolla (2022).
- **Booster** (Ramirez et al. 2023) comprises indoor scenes with a variety of non-Lambertian surfaces, including transparent, mirrored, and specular surfaces. In line with the work of Costanzino et al. (2023), who evaluate the performance of monocular and stereo depth networks on non-Lambertian surfaces, we adopt their training split as our test split to evaluate our method.

### Performance Evaluation

**ScanNet-Reflection dataset** To demonstrate the effectiveness of the proposed method on reflective surfaces, we measure depth accuracy using the Scannet-Reflection test and validation sets, which specifically include scenes with reflective surfaces. The results indicate that all baselines demonstrate consistent performance improvements across nearly all evaluated metrics. Notably, the *Abs rel* metric display significant gains, with an average improvement of 9.02% on the test set and an impressive 23.51% on the validation set, as shown in Table 1. This indicates the proposed method encourages the network to accurately learn depth in reflective surfaces, as depicted in Figure 3.

**ScanNet-NoReflection dataset** Since the proposed method pinpoints the disturbed regions from the intrinsic image and excludes them from training process, there is

Method	ScanNet-NoReflection Validation Set						
	Abs Rel ↓	Sq Rel ↓	RMSE ↓	RMSE log ↓	$\delta < 1.25 \uparrow$	$\delta < 1.25^2 \uparrow$	$\delta < 1.25^3 \uparrow$
Monodepth2	<b>0.169</b>	<b>0.100</b>	<b>0.395</b>	<b>0.206</b>	<b>0.759</b>	<b>0.932</b>	<b>0.979</b>
Monodepth2 + <i>Ours</i>	0.174	0.103	0.398	0.210	0.752	0.927	0.977
HRDepth	<b>0.169</b>	0.102	<b>0.388</b>	<b>0.202</b>	<b>0.766</b>	<b>0.933</b>	<b>0.980</b>
HRDepth + <i>Ours</i>	<b>0.169</b>	<b>0.097</b>	<b>0.388</b>	0.205	0.763	0.932	<b>0.980</b>
MonoViT	<b>0.140</b>	0.074	<b>0.333</b>	<b>0.171</b>	<b>0.829</b>	<b>0.952</b>	0.984
MonoViT + <i>Ours</i>	0.144	<b>0.073</b>	0.339	0.176	0.816	0.951	<b>0.986</b>

Table 2: Main results on the ScanNet-NoReflection Validation set.

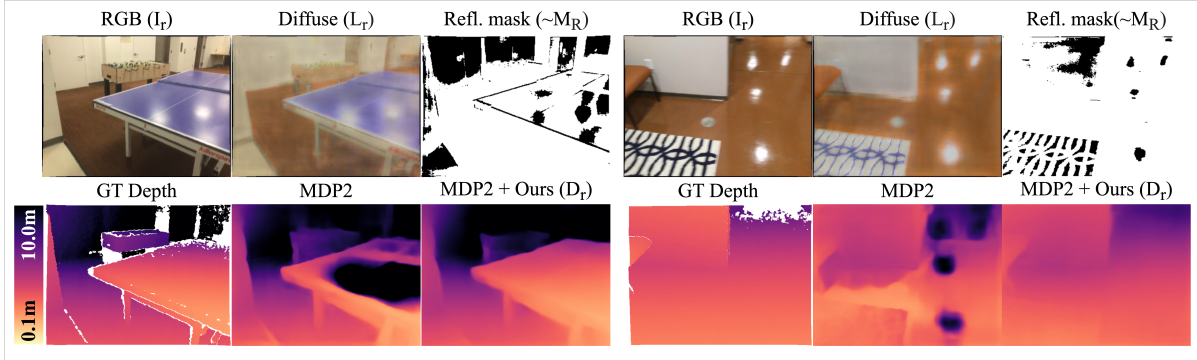


Figure 3: Qualitative results between monodepth2 (denoted as MDP2) and our method.

room for high-frequency detail to be lost depending on the quality of the intrinsic image output. Therefore, to evaluate the impact of these potential sources of instability on the performance, we measure the generalization performance of the proposed method through evaluation on the ScanNet-NoReflection set, where there are few reflective surfaces in the scene. As shown in Table 2, on average, across all metrics, Monodepth2 has a 1.47% performance drop, HRDepth shows a 0.42% performance improvement, and MonoViT records a 1.10% performance drop. However, compared to the performance gains achieved by the proposed method on the reflection set, the performance drops for Monodepth2 and MonoViT are negligible, while HRDepth even shows performance gains. This suggests that the proposed method accurately localizes reflective surfaces and achieves strong generalization performance by preserving features that are conducive to training.

**7-Scenes & Booster dataset** To further validate the effectiveness of the proposed method across various environments, we evaluate its cross-dataset generalization performance on 7-Scenes and Booster datasets. These datasets are characterized by scenes dominated by non-reflective and reflective objects (*i.e.*, specular, transparent, and, mirrors), respectively. As shown in Table 3, models trained with the proposed method demonstrate consistent performance gains on the Booster dataset, which includes a significant presence of non-Lambertian surfaces. Specifically, Monodepth2 shows an improvement of 9.96%, HRDepth improved by 8.09%, and MonoViT by 4.39%, averaged across all metrics. These results indicate that the models trained with the proposed method are robust in estimating various non-Lambertian sur-

faces effectively.

On the other hand, on the 7-Scenes dataset, predominantly featuring opaque objects, our method achieves average performance improvements of 1.38% on Monodepth2 and 1.80% on MonoViT across all evaluation metrics. HRDepth shows an average performance degradation of 2.15% across all metrics. However, the decrease is considered negligible compared to the significant 8.09% improvement observed on non-Lambertian surfaces of the Booster dataset.

**Comparison with multi-stage training methods** To more fairly evaluate the scalability of our proposed end-to-end training method, we propose a simple yet efficient knowledge distillation technique. We designate the model trained with the proposed method (designated as [Baseline + Ours]) and those trained using the self-supervised method with conventional photometric error (designated as [Baseline]) as the teacher models. Distillation is performed using a mask generated from the photometric error between each RGB and diffuse image, as defined in Equation 10. Specifically, if the photometric error for diffuse images exceeds that for RGB images by a certain margin, the depth learned by the [Baseline] method under conventional photometric error conditions is treated as pseudo-ground truth. Conversely, if the photometric error for RGB images is significantly higher, the depth from [Baseline + Ours] is used as pseudo ground truth. (Detailed distillation process is provided in the appendix). As shown in Table 4, our proposed distillation model outperforms the self-teaching approach (Poggi et al. 2020), which uses uncertainty-based fusion for generating pseudo ground truth, with a significant enhancement across all metrics. Compared to 3D Distillation (Shi et al. 2023),

Dataset	Method	Abs Rel ↓	Sq Rel ↓	RMSE ↓	RMSE log ↓	$\delta < 1.25 \uparrow$	$\delta < 1.25^2 \uparrow$	$\delta < 1.25^3 \uparrow$
7-Scenes	Monodepth2	0.210	0.130	0.445	0.248	0.656	<b>0.906</b>	<b>0.974</b>
	Monodepth2 + <i>Ours</i>	<b>0.209</b>	<b>0.126</b>	<b>0.431</b>	<b>0.244</b>	<b>0.666</b>	<b>0.906</b>	0.972
	HRDepth	<b>0.193</b>	<b>0.115</b>	<b>0.421</b>	<b>0.231</b>	<b>0.682</b>	<b>0.921</b>	<b>0.982</b>
	HRDepth + <i>Ours</i>	0.204	0.116	<b>0.421</b>	0.239	0.661	0.907	0.978
	MonoViT	0.173	0.093	0.365	0.201	0.752	0.945	0.988
	MonoViT + <i>Ours</i>	<b>0.171</b>	<b>0.086</b>	<b>0.355</b>	<b>0.200</b>	<b>0.754</b>	<b>0.949</b>	<b>0.988</b>
Booster	Monodepth2	0.520	0.429	0.601	0.444	0.305	0.591	0.827
	Monodepth2 + <i>Ours</i>	<b>0.471</b>	<b>0.341</b>	<b>0.536</b>	<b>0.412</b>	<b>0.334</b>	<b>0.639</b>	<b>0.861</b>
	HRDepth	0.495	0.391	0.559	0.426	0.307	0.611	0.852
	HRDepth + <i>Ours</i>	<b>0.462</b>	<b>0.315</b>	<b>0.506</b>	<b>0.404</b>	<b>0.325</b>	<b>0.645</b>	<b>0.890</b>
	MonoViT	0.418	0.327	0.504	0.374	0.425	0.679	0.888
	MonoViT + <i>Ours</i>	<b>0.396</b>	<b>0.293</b>	<b>0.482</b>	<b>0.364</b>	<b>0.426</b>	<b>0.716</b>	<b>0.909</b>

Table 3: Main results on the 7-Scenes and Booster dataset.

Backbone	Method	ScanNet-Reflection Validation Set						
		Abs Rel ↓	Sq Rel ↓	RMSE ↓	RMSE log ↓	$\delta < 1.25 \uparrow$	$\delta < 1.25^2 \uparrow$	$\delta < 1.25^3 \uparrow$
Monodepth2	Self-teaching	0.192	0.188	0.548	0.233	0.764	0.920	0.967
	3D Distillation	0.156	0.093	0.442	0.191	0.786	0.943	0.987
	<i>Ours (distillation)</i>	<b>0.148</b>	<b>0.089</b>	<b>0.429</b>	<b>0.185</b>	<b>0.798</b>	<b>0.950</b>	<b>0.989</b>
HRDepth	Self-teaching	0.202	0.208	0.565	0.243	0.756	0.914	0.964
	3D Distillation	0.153	<b>0.090</b>	<b>0.430</b>	0.188	0.789	0.948	<b>0.989</b>
	<i>Ours (distillation)</i>	<b>0.150</b>	<b>0.090</b>	<b>0.430</b>	<b>0.187</b>	<b>0.798</b>	<b>0.949</b>	0.988
MonoViT	Self-teaching	0.176	0.195	0.537	0.224	0.823	0.930	0.963
	3D Distillation	<b>0.126</b>	<b>0.068</b>	<b>0.367</b>	<b>0.159</b>	0.851	<b>0.965</b>	<b>0.991</b>
	<i>Ours (distillation)</i>	0.129	0.083	0.406	0.166	<b>0.856</b>	0.960	0.989

Table 4: Comparison between the proposed methods and Multi-stage methods on ScanNet-Reflection Validation set.

our method achieves an average performance improvement of 2.57% for Monodepth2 across all metrics, with HRDepth showing an improvement of 0.52%, while MonoViT exhibits a performance decrease of 5.66%. Notably, the proposed method achieves comparable performance to 3D Distillation but achieves this without the extensive computational demands associated with 3D Distillation’s approach of using uncertainty-based multi-model ensembles and 3D mesh reconstruction for pseudo ground truth generation.

Method	ScanNet-Reflection Test Set	
	Abs Rel ↓	RMSE ↓
Monodepth2	0.181	0.521
Monodepth2 + <i>Ours</i> (w/o. $\mathcal{L}_{cts}$ )	0.176	0.510
Monodepth2 + <i>Ours</i>	<b>0.163</b>	<b>0.472</b>

Table 5: Ablation study on the ScanNet-Reflection Test set.

**Ablation study of intrinsic loss term** In the loss terms of the intrinsic network,  $\mathcal{L}_{recon}$  and  $\mathcal{L}_{cross}$  serve as fundamental conditions necessary to derive meaningful intrinsic images. To investigate the impact of the proposed contrastive loss,  $\mathcal{L}_{cts}$ , we integrate it into Monodepth2 and compare the performance changes with and without this loss. As shown

in Table 5, the experimental results show that the omitting  $\mathcal{L}_{cts}$  causes the network to learn the diffuse image that approaches to a non-valid image, adversely affecting the accuracy of reflective region localization. This misalignment results in marginal performance improvements.

## Conclusion

This paper presents the first end-to-end approach that leverages intrinsic image decomposition to address the challenges posed by reflective surfaces, which are known to compromise depth accuracy in SSMDE. We introduce a new method for localizing the pixel-level reflective regions without the need for additional labeled annotation by using estimated intrinsic components. Building on this, our reflection-aware SSMDE method leverages the localized information to significantly improve depth accuracy. It prevents erroneous depth learning caused by these regions violating the Lambertian assumption. Experimental results demonstrate that our method significantly outperforms conventional self-supervised approaches across the various indoor datasets consisting of numerous reflective surfaces. Moreover, the proposed method achieves competitive performance compared to the method leveraging 3D knowledge distillation, which requires excessive training costs.

## Acknowledgements

This work was supported by Korea Research Institute for defense Technology planning and advancement through Defense Innovation Vanguard Enterprise Project, funded by Defense Acquisition Program Administration (R230206) and the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. RS-2023-00210908).

## References

- Bae, G.; Budvytis, I.; and Cipolla, R. 2022. Multi-view depth estimation by fusing single-view depth probability with multi-view geometry. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2842–2851.
- Bae, J.; Hwang, K.; and Im, S. 2023. A study on the generality of neural network structures for monocular depth estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Bae, J.; Moon, S.; and Im, S. 2023. Deep digging into the generalization of self-supervised monocular depth estimation. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, 187–196.
- Barron, J. T.; and Malik, J. 2014. Shape, illumination, and reflectance from shading. *IEEE transactions on pattern analysis and machine intelligence*, 37(8): 1670–1687.
- Careaga, C.; and Aksoy, Y. 2023. Intrinsic image decomposition via ordinal shading. *ACM Transactions on Graphics*, 43(1): 1–24.
- Costanzino, A.; Ramirez, P. Z.; Poggi, M.; Tosi, F.; Mattocchia, S.; and Di Stefano, L. 2023. Learning Depth Estimation for Transparent and Mirror Surfaces. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9244–9255.
- Dai, A.; Chang, A. X.; Savva, M.; Halber, M.; Funkhouser, T.; and Nießner, M. 2017. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5828–5839.
- Duchêne, S.; Riant, C.; Chaurasia, G.; Lopez-Moreno, J.; Laffont, P.-Y.; Popov, S.; Bousseau, A.; and Drettakis, G. 2015. Multi-view intrinsic images of outdoors scenes with an application to relighting. *ACM Transactions on Graphics*, 16.
- Eigen, D.; Puhrsch, C.; and Fergus, R. 2014. Depth map prediction from a single image using a multi-scale deep network. *Advances in neural information processing systems*, 27.
- Garces, E.; Munoz, A.; Lopez-Moreno, J.; and Gutierrez, D. 2012. Intrinsic images by clustering. In *Computer graphics forum*, volume 31, 1415–1424. Wiley Online Library.
- Gasperini, S.; Morbitzer, N.; Jung, H.; Navab, N.; and Tombari, F. 2023. Robust monocular depth estimation under challenging conditions. In *Proceedings of the IEEE/CVF international conference on computer vision*, 8177–8186.
- Geiger, A.; Lenz, P.; and Urtasun, R. 2012. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE conference on computer vision and pattern recognition*, 3354–3361. IEEE.
- Godard, C.; Mac Aodha, O.; and Brostow, G. J. 2017. Unsupervised monocular depth estimation with left-right consistency. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 270–279.
- Godard, C.; Mac Aodha, O.; Firman, M.; and Brostow, G. J. 2019. Digging into self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF international conference on computer vision*, 3828–3838.
- Guizilini, V.; Ambrus, R.; Pillai, S.; Raventos, A.; and Gaidon, A. 2020. 3d packing for self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2485–2494.
- Guizilini, V.; Vasiljevic, I.; Ambrus, R.; Shakhnarovich, G.; and Gaidon, A. 2022. Full surround monodepth from multiple cameras. *IEEE Robotics and Automation Letters*, 7(2): 5397–5404.
- Kim, S.; Park, K.; Sohn, K.; and Lin, S. 2016. Unified depth prediction and intrinsic image decomposition from a single image via joint convolutional neural fields. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part VIII 14*, 143–159. Springer.
- Laffont, P.-Y.; and Bazin, J.-C. 2015. Intrinsic decomposition of image sequences from local temporal variations. In *Proceedings of the IEEE international conference on computer vision*, 433–441.
- Laffont, P.-Y.; Bousseau, A.; Paris, S.; Durand, F.; and Drettakis, G. 2012. Coherent intrinsic images from photo collections. *ACM Transactions on Graphics*, 31(6).
- Land, E. H.; and McCann, J. J. 1971. Lightness and retinex theory. *Josa*, 61(1): 1–11.
- Long, X.; Liu, L.; Li, W.; Theobalt, C.; and Wang, W. 2021. Multi-view depth estimation using epipolar spatio-temporal networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8258–8267.
- Luo, J.; Ceylan, D.; Yoon, J. S.; Zhao, N.; Philip, J.; Frühstück, A.; Li, W.; Richardt, C.; and Wang, T. 2024. Intrinsicdiffusion: joint intrinsic layers from latent diffusion models. In *ACM SIGGRAPH 2024 Conference Papers*, 1–11.
- Lyu, X.; Liu, L.; Wang, M.; Kong, X.; Liu, L.; Liu, Y.; Chen, X.; and Yuan, Y. 2021. Hr-depth: High resolution self-supervised monocular depth estimation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 2294–2301.
- Ma, W.-C.; Chu, H.; Zhou, B.; Urtasun, R.; and Torralba, A. 2018. Single image intrinsic decomposition without a single intrinsic image. In *Proceedings of the European conference on computer vision (ECCV)*, 201–217.
- Mahalanobis, P. C. 1936. Mahalanobis distance. In *Proceedings National Institute of Science of India*, volume 49, 234–256.

- Maxwell, B. A.; Friedhoff, R. M.; and Smith, C. A. 2008. A bi-illuminant dichromatic reflection model for understanding images. In *2008 IEEE conference on computer vision and pattern recognition*, 1–8. IEEE.
- Narihira, T.; Maire, M.; and Yu, S. X. 2015. Direct intrinsics: Learning albedo-shading decomposition by convolutional regression. In *Proceedings of the IEEE international conference on computer vision*, 2992–2992.
- Poggi, M.; Aleotti, F.; Tosi, F.; and Mattocchia, S. 2020. On the uncertainty of self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3227–3237.
- Ramirez, P. Z.; Costanzino, A.; Tosi, F.; Poggi, M.; Salti, S.; Mattocchia, S.; and Di Stefano, L. 2023. Booster: a benchmark for depth from images of specular and transparent surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Rother, C.; Kiefel, M.; Zhang, L.; Schölkopf, B.; and Gehler, P. 2011. Recovering intrinsic images with a global sparsity prior on reflectance. *Advances in neural information processing systems*, 24.
- Saunders, K.; Vogiatzis, G.; and Manso, L. J. 2023. Self-supervised Monocular Depth Estimation: Let’s Talk About The Weather. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 8907–8917.
- Shen, L.; Tan, P.; and Lin, S. 2008. Intrinsic image decomposition with non-local texture cues. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, 1–7. IEEE.
- Shen, L.; and Yeo, C. 2011. Intrinsic images decomposition using a local and global sparse representation of reflectance. In *CVPR 2011*, 697–704. IEEE.
- Shi, J.; Dong, Y.; Su, H.; and Yu, S. X. 2017. Learning non-lambertian object intrinsics across shapenet categories. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1685–1694.
- Shi, X.; Dikov, G.; Reitmayr, G.; Kim, T.-K.; and Ghafoorian, M. 2023. 3D Distillation: Improving Self-Supervised Monocular Depth Estimation on Reflective Surfaces. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9133–9143.
- Shotton, J.; Glocker, B.; Zach, C.; Izadi, S.; Criminisi, A.; and Fitzgibbon, A. 2013. Scene coordinate regression forests for camera relocalization in RGB-D images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2930–2937.
- Sunkavalli, K.; Matusik, W.; Pfister, H.; and Rusinkiewicz, S. 2007. Factored time-lapse video. In *ACM SIGGRAPH 2007 papers*, 101–es.
- Tominaga, S. 1994. Dichromatic reflection models for a variety of materials. *Color Research & Application*, 19(4): 277–285.
- Wang, Y.; Liang, Y.; Xu, H.; Jiao, S.; and Yu, H. 2024. Sqldepth: Generalizable self-supervised fine-structured monocular depth estimation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 5713–5721.
- Wang, Z.; Bovik, A. C.; Sheikh, H. R.; and Simoncelli, E. P. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4): 600–612.
- Weiss, Y. 2001. Deriving intrinsic images from image sequences. In *Proceedings Eighth IEEE International Conference on Computer Vision. ICCV 2001*, volume 2, 68–75. IEEE.
- Yu, Y.; and Smith, W. A. 2019. Inverserendernet: Learning single image inverse rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3155–3164.
- Zeng, Z.; Deschaintre, V.; Georgiev, I.; Hold-Geoffroy, Y.; Hu, Y.; Luan, F.; Yan, L.-Q.; and Hašan, M. 2024. Rgb $\leftrightarrow$ x: Image decomposition and synthesis using material-and lighting-aware diffusion models. In *ACM SIGGRAPH 2024 Conference Papers*, 1–11.
- Zhang, N.; Nex, F.; Vosselman, G.; and Kerle, N. 2023. Lite-mono: A lightweight cnn and transformer architecture for self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18537–18546.
- Zhao, C.; Zhang, Y.; Poggi, M.; Tosi, F.; Guo, X.; Zhu, Z.; Huang, G.; Tang, Y.; and Mattocchia, S. 2022. Monovit: Self-supervised monocular depth estimation with a vision transformer. In *2022 International Conference on 3D Vision (3DV)*, 668–678. IEEE.
- Zhao, Q.; Tan, P.; Dai, Q.; Shen, L.; Wu, E.; and Lin, S. 2012. A closed-form solution to retinex with nonlocal texture constraints. *IEEE transactions on pattern analysis and machine intelligence*, 34(7): 1437–1444.
- Zhou, T.; Brown, M.; Snavely, N.; and Lowe, D. G. 2017. Unsupervised learning of depth and ego-motion from video. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1851–1858.
- Zhou, Z.; Fan, X.; Shi, P.; and Xin, Y. 2021. R-msfm: Recurrent multi-scale feature modulation for monocular depth estimating. In *Proceedings of the IEEE/CVF international conference on computer vision*, 12777–12786.