






# Structure-Centric Robust Monocular Depth Estimation via Knowledge Distillation

Runze Chen<sup>1,2</sup>, Haiyong Luo<sup>2†</sup>, Fang Zhao<sup>1†</sup>, Jingze Yu<sup>3</sup>, Yupeng Jia<sup>1,2</sup>, Juan Wang<sup>4</sup>, and Xuepeng Ma<sup>4</sup>

<sup>1</sup> Beijing University of Posts and Telecommunications, Beijing, China

<sup>2</sup> Institute of Computing Technology Chinese Academy of Sciences, Beijing, China

<sup>3</sup> Guangdong Hong Kong Macao Greater Bay Area National Technology Innovation Center, Guangdong, China

<sup>4</sup> Shouguang Cheng Zhi Feng Xing Technology Co., Ltd, Shandong, China  
[yhluo@ict.ac.cn](mailto:yhluo@ict.ac.cn), [zfsse@bupt.edu.cn](mailto:zfsse@bupt.edu.cn)


**Abstract.** Monocular depth estimation, enabled by self-supervised learning, is a key technique for 3D perception in computer vision. However, it faces significant challenges in real-world scenarios, which encompass adverse weather variations, motion blur, as well as scenes with poor lighting conditions at night. Our research reveals that we can divide monocular depth estimation into three sub-problems: depth structure consistency, local texture disambiguation, and semantic-structural correlation. Our approach tackles the non-robustness of existing self-supervised monocular depth estimation models to interference textures by adopting a structure-centered perspective and utilizing the scene structure characteristics demonstrated by semantics and illumination. We devise a novel approach to reduce over-reliance on local textures, enhancing robustness against missing or interfering patterns. Additionally, we incorporate a semantic expert model as the teacher and construct inter-model feature dependencies via learnable isomorphic graphs to enable aggregation of semantic structural knowledge. Our approach achieves state-of-the-art out-of-distribution monocular depth estimation performance across a range of public adverse scenario datasets. It demonstrates notable scalability and compatibility, without necessitating extensive model engineering. This showcases the potential for customizing models for diverse industrial applications. The source code for our method is available at Scent-Depth.

**Keywords:** Monocular Depth Estimation · Robustness against Interference Patterns · 3D Perception in Computer Vision

## 1 Introduction

Monocular depth estimation (MDE) is a popular research topic in the field of computer vision due to its broad applications in fields such as autonomous driving, robotics, and augmented reality. However, accurately estimating objects'

<sup>†</sup> Joint corresponding authors: Haiyong Luo and Fang Zhao.



**Fig. 1:** Key Sub-Problems in Monocular Depth Estimation. This figure highlights three areas critical for enhancing depth estimation: *Depth Structure Consistency*, ensuring smooth depth transitions across frames; *Local Texture Disambiguation*, addressing the challenge of the model’s over-dependence on local textures that compromises depth estimation robustness, by improving performance in diverse or texture-sparse environments; and *Semantic and Structural Correlation*, leveraging object semantics and structure for continuous depth inference.

depth from a single 2D image remains a challenging task despite significant research efforts. Synthesizing new perspectives across multiple frames enables training a monocular affine-invariant depth estimation model in a self-supervised manner [38]. This approach imposes high demands on the image quality of adjacent frames, as even minor device or environmental disturbances can significantly degrade training quality [10].

Thus, it is crucial to investigate methods for enhancing the robustness of depth estimation models built through self-supervision based on clear visual sequences, particularly in adverse environments. Moreover, models that solely rely on pixel-level features fail to capture the critical structural information of objects, which negatively impacts their performance in complex and noisy real-world environments. The current approach constrains the model from learning reasonable affine-invariant depth through direct projection between pixels, reflecting the consistency of scene structure during motion. However, not all textures serve to represent scene structure effectively for the model. Over-reliance on structure-agnostic textures may lead to significant performance degradation in the presence of external disturbances, such as occlusion, adverse weather, equipment malfunction, and inconsistent lighting conditions [20]. Therefore, effectively incorporating structural information into the model becomes crucial in enhancing its depth estimation performance in various practical scenarios.

To distinguish the model’s understanding of effective scene structure, we propose to decouple the problem of representing structural information for robust monocular depth estimation into three sub-problems as displayed in Figure 1:

- *Depth structure consistency*, which ensures that the relative depth structure of the scene remains consistent across adjacent frames during camera motion, avoiding abrupt changes in estimated depths.

- *Local texture disambiguation* refers to resolving degradation issues in depth estimation models caused by overfitting of local detail textures, which can lead to estimation ambiguities in adverse scenarios due to interfering textures.
- *Semantic and structural correlation* illustrates the inherent link between an object’s semantics (e.g., shape, function) and its visual structure (e.g., size, position, orientation), where fixed semantic properties guide the estimation of continuous depth across different views.

Based on the above considerations, we propose solutions for each of the aforementioned sub-problems as illustrated in Figure 2. Adopting the retinex principle, we decouple ambient illumination from object surface reflectance in images. This reduces the monocular depth estimation model’s over-reliance on local textures of scene objects, enabling robustness against missing textures and distracting patterns. Further, we construct an isomorphic graph between the depth model and a semantic expert model [18] in a learnable manner. This achieves distillation of structural knowledge from the semantic expert into the monocular depth model. By employing the aforementioned methods, we substantially enhance the robustness of monocular depth estimation models trained on clear datasets in challenging scenarios, such as adverse weather variations, motion blur, and scenes with poor lighting conditions at night. This explicitly builds structure-centric monocular depth estimation.

In conclusion, our contributions are four-fold:

- We propose decoupling monocular depth estimation into three sub-problems: local texture disambiguation, semantic-structural correlation, and depth structure consistency. These highlight key challenges for monocular depth estimation methods in complex environments.
- We decouple ambient illumination from object surface reflectance to reduce the depth estimation model’s strong reliance on local textures. This mitigates performance degradation stemming from over-dependence on local patterns under adverse weather, intense motion, and poor lighting conditions.
- Our designed isomorphic graph distillation enables the aggregation of features across different tasks. It transforms inter-channel dependencies into a learnable loss, thereby distilling semantic structural knowledge into the depth model. This effectively learns mappings between teacher and student feature spaces, enabling effective knowledge transfer.
- This paper presents an effective benchmark for robustness and generalization under common real-world conditions. It reveals the challenges faced by current state-of-the-art models under these scenarios. It also demonstrates that our method is effective, achieving state-of-the-art performance across these challenging scenarios.

In this paper, we present a novel approach for monocular depth estimation that breaks down the problem into individual sub-problems. We address those sub-problems through the development of specific deep-learning frameworks designed to meet the demands of complex real-world scenes. Moving forward, we

describe the pipeline of our proposed method, present experimental results that validate its effectiveness, and provide our conclusions. Our method yields state-of-the-art performance on widely-used public benchmark datasets, even under various forms of corruption. This demonstrates the method’s strong potential for real-world applications in complex scenes.

## 2 Related Works

### 2.1 Self-supervised Monocular Depth Estimation

The technique of using multi-view synthesis and pose estimation as proxy tasks proves to be an effective method for accomplishing self-supervised training of affine invariant depth estimation models. Employing a warping-based view synthesis approach with monocular image sequences [38] and stereo images [9] enables the model to learn from rudimentary 3D object structural details, providing an elegant solution for monocular depth estimation. The scale-invariant nature of object structure is a crucial aspect of its defining features. Furthermore, employing a multi-scale sampling approach [10] in depth estimation provides a strategy to counteract the restrictive impact of local texture on the model’s capacity to perceive the underlying structure of objects. In the context of image sequences, it is reasonable to posit that there exists a relative consistency in the structural organization of depth maps across different frames [2]. The fusion of Vision Transformers (ViT) and convolutional features [34, 36] effectively improves the modeling capacity for long-range structural features. Existing robust self-supervised monocular MDE works [6, 29, 32] mainly focus on learning models from huge datasets containing various scenes. Training a robust MDE model using clear, interference-free data is a significant challenge, as it must be capable of handling real-world conditions such as rain, snow, fog, motion blur, and night-time scenarios effectively.

### 2.2 Knowledge-Enhanced Monocular Depth Estimation

An effective approach to address challenges in the real world involves introducing more knowledge guidance during the self-supervised process, serving as an enhancement for self-supervised depth estimation training. Knowledge distillation is an effective approach, and many works [5, 23] accomplish the compression of the monocular depth estimation model through this approach. Constructing pseudo depth maps through supervised deep estimation models [27] to guide distillation in self-supervised training has to some extent enhanced the training effectiveness of self-supervised depth estimation models [1]. This approach essentially introduces implicit supervised information into self-supervised training. Semantic labels of objects can robustly represent the structure and contour of objects in the scene [19]. However, in common video sequence acquisition scenarios, the annotation of semantic labels typically consumes a substantial amount of manual effort and incurs high costs. Additional sensors, such as inertial sensors, can also assist in the scale consistency training of depth estimation [35].

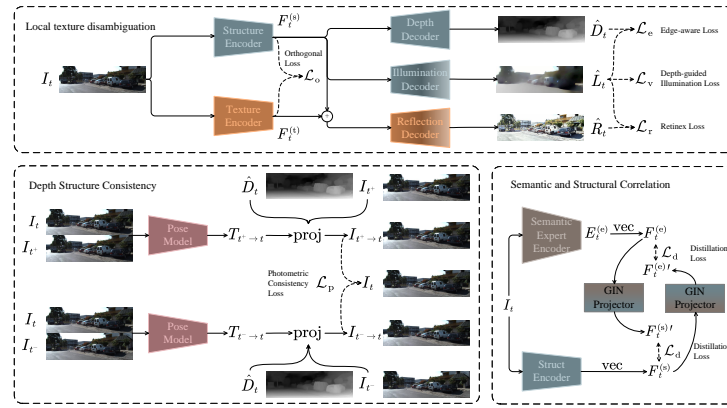
Integrating reliance on inertial sensors to some extent reduces the model’s usability for devices lacking inertial sensors. The use of affine invariant depth maps generated by unsupervised depth estimation models for dense reconstruction to create mesh grids, and the distillation of depth estimation models through the synthesis of pseudo depth maps from these mesh grids can provide unsupervised depth estimation of a higher surface quality [30]. The multi-stage training and its dense reconstruction process also imply more computational costs during the training process. Introducing additional knowledge requires consideration of various factors, including data diversity and training costs. Simultaneously, we believe robustness in adverse scenarios is crucial.

### 2.3 Structure-Centric Knowledge in Adverse Scenarios

Self-supervised monocular depth estimation shows promising prospects in affine invariant depth estimation and relative pose prediction. However, in real-world application scenarios, visual degradation caused by weather, lighting conditions, and sensor malfunctions poses a notable challenge [26]. Maintaining the robustness of neural networks towards visually degraded scenarios is always a focus in the academic field [15]. In self-supervising depth estimation, the out-of-distribution generalization ability for adverse scenarios [20] determines whether the depth estimation model can adapt to wide applications in the real world. How to generalize the monocular depth estimation model trained on clear visual sequence data to a wider and more complex application scenario still needs more attention. The semantic information of objects has consistency with the structure of the scene [16]. Guiding through the semantic prediction of the teacher model can bring significant performance improvement [13] in monocular depth estimation. The introduction of training constraints through semantic labels is also proven to enhance the domain transfer ability of depth estimation models [22]. Illumination in an image represents the interaction of the object structure in the scene with the natural light source [28]. Unlike the interference texture appearing in the image, the scene structure characteristics shown by the illumination remain unaffected in adverse weather and various light intensities [11]. Existing works [4] also pay attention to the connection between depth structure and illumination, guiding the learning of illumination through depth maps to provide low-light enhancement algorithms. Hence, we aim to enhance the out-of-distribution generalization ability of monocular depth estimation models in challenging scenarios never encountered before, guided by the structural semantic features of mature semantic expert models and the scene structure characteristics represented by illumination.

## 3 Structure-Centric Depth Estimation

We describe how our proposed structure-centric monocular depth estimation tackles the degradation problem under common camera corruptions in real-world scenarios. Specifically, guided by the three sub-problems including depth-



**Fig. 2:** Overview of the Structure-Centric Monocular Depth Estimation Approach. The solid lines in the figure indicate data flow, while the dashed lines represent constraints during the optimization process. The operation  $\text{vec}$  represents the process of flattening these features into vectors. The operation  $\text{proj}$  denotes the re-projection of views.

structure consistency, local texture disambiguation, and semantic-structural correlation, we demonstrate how our method mitigates performance degradation due to interference from adverse weather, motion, and poor lighting conditions at night. We accomplish this through knowledge distillation techniques to realize robust structure-centric monocular depth estimation.

### 3.1 Depth Structure Consistency

Given a monocular image sequence  $I = \{\dots, I_t \in \mathbb{R}^{3 \times W \times H}, \dots\}$  with camera intrinsics determined by  $K$ , the objective of monocular estimation is to establish a deep learning model for accurate estimation of the corresponding depth  $\hat{D}_t \in \mathbb{R}^{1 \times W \times H}$  from monocular image frame  $I_t$ . Given the difficulty and high cost of measuring depth ground truth  $D_t$ , we typically rely on unsupervised learning methods. This requires our monocular depth estimation approach to extract more sufficient scene structural information from a single image.

Monocular depth estimation typically involves synthesizing the view  $I_{t' \rightarrow t}$  by using the estimated relative pose  $\hat{T}_{t \rightarrow t'}$  and the estimated depth map  $\hat{D}_t$ , with respect to the source frame  $I_{t'}$  and the target frame  $I_t$ . We can express the synthesis operation as

$$I_{t' \rightarrow t} = I_{t'} \left\langle \text{proj} \left( \hat{D}_t, \hat{T}_{t \rightarrow t'}, K \right) \right\rangle, \quad (1)$$

where  $\text{proj}(\cdot)$  produces the two-dimensional positions resulting from the projection of depths  $D_t$  onto the image  $I_{t'}$ ,  $\langle \cdot \rangle$  upsamples the estimation to match the shape of  $I_{t'}$ , and  $I_{t' \rightarrow t}$  is an approximation of  $I_t$  obtained by projecting  $I_{t'}$ . Monocular depth estimation aims to model depth structure consistency, which

leverages the consistency of depth structure between adjacent frames to accomplish view synthesis tasks. Drawing on [37,38], we utilize  $\mathcal{L}_p$  to impose constraints on the quality of re-projected views,

$$\begin{aligned} \mathcal{L}_p^{u,v}(I_t, I_{t' \rightarrow t}) &= \frac{\alpha}{2} (1 - \text{ssim}(I_t, I_{t' \rightarrow t})) + (1 - \alpha) \|I_t - I_{t' \rightarrow t}\|_1, \\ \mathcal{L}_p &= \sum \mu \mathcal{L}_p(I_t, I_{t' \rightarrow t}) \end{aligned} \quad (2)$$

where  $\mu$  is a dynamic mask [10], and  $\text{ssim}$  measures the luminance consistency between two image frames [37].

The static environmental structure in the scene effectively provides structural information for depth estimation in the image sequence, which means we need to pay attention to the consistency of scene structure between adjacent frames. While maintaining the consistency between the synthesized view  $I_{t' \rightarrow t}$  and the target frame  $I_{t'}$ , the estimated depth  $\hat{D}_{t'}$  of the source frame also exhibits structural consistency with the estimated depth  $\hat{D}_t$  of the target frame.

### 3.2 Local Texture Disambiguation

Monocular depth estimation infers scene depth from camera motion cues. The shading patterns induced by ambient light reveal structural features of objects. Such light-dark contrasts in scenes facilitate important clues about scene structures that are critical for depth models. However, based on color constancy, abundant surface textures of objects appear unaffected by scene structures in camera-captured images. Conversely, excessively complex textures may impede depth judgment, necessitating the extraction of valid cues. By adopting the Retinex principle, we decouple structural information and surface textures in scenes. This reduces model sensitivity to surface patterns under complex imaging conditions and lowers susceptibility to noise interference.

For our camera observation  $I_t$ , there exist corresponding reflectance  $R_t$  and illumination  $L_t$  terms that satisfy the relation  $I_t = R_t * L_t$ . The reflectance term  $R_t$  primarily exhibits the surface texture patterns of objects in the scene, while the illumination term  $L_t$  mainly captures the ambient shading on structural surfaces within the scene. We leverage the estimated depth  $\hat{D}_t$  along with  $I_t$  to jointly optimize the illumination estimation  $\hat{L}_t$ . The optimization loss is

$$\mathcal{L}_v = \sum \max \left( \|I_t - \hat{L}_t\|, \alpha \right) + \sum \left\| \nabla \hat{L}_t \right\| \cdot e^{-\|\nabla \hat{D}_t\| - \|\nabla I_t\|}, \quad (3)$$

where the gradient of  $\hat{D}_t$  is not recorded during training. We apply a suitable threshold of  $\alpha = 0.2$  in  $\mathcal{L}_v$  to preclude absolute proximity between local pixels. This ensures relatively consistent luminance within regions, which facilitates reducing the influence of localized textures.

We leverage the complementary nature of ambient illumination and scene depth. Specifically, we utilize an edge-aware loss  $\mathcal{L}_e$  between the illumination component of ambient light and the depth map,

$$\mathcal{L}_e = \sum \left\| \nabla \hat{D}_t \right\| \cdot e^{-\|\nabla \hat{L}_t\|}, \quad (4)$$

where the gradient of  $\hat{L}_t$  is not recorded during training. The same encoder extracts features to generate the illumination and depth map. We guide the illumination estimation using the gradients of the depth map and original image. This encourages continuous illumination in regions with contiguous surface structures.

We derive the reflectance  $\hat{R}_t$  from the illumination  $\hat{L}_t$ . As reflectance mostly captures fine surface textures, we aim to reduce model sensitivity to superficial details without impairing the extraction of structural textures. We optimize the model’s computation of the reflectance  $\hat{R}_t$  through the loss  $\mathcal{L}_r$ ,

$$\mathcal{L}_r = \text{ssim} \left( \hat{R}_t * \hat{L}_t, I_t \right), \quad (5)$$

where  $\hat{L}_r$  does not participate in gradient calculation during training. We desire the reflectance model to adequately extract fine surface textures across all objects in the scene structure while minimizing the proportion of structural cues.

In Figure 2, the structure encoder generates the feature embedding  $E_t^{(s)} \in \mathbb{R}^{C \times H \times W}$ , while the texture encoder produces  $E_t^{(t)} \in \mathbb{R}^{C \times H \times W}$ . We expect  $E_t^{(s)}$  to represent all structural cues in the scene, and  $E_t^{(t)}$  to capture surface texture patterns across objects. We enforce orthogonality between the structural features  $E_t^{(s)}$  and texture features  $E_t^{(t)}$ . Specifically, we build an orthogonal loss  $\mathcal{L}_o$  using the vectorized  $F_t^{(s)} \in \mathbb{R}^{C \times (H \times W)}$ ,  $F_t^{(t)} \in \mathbb{R}^{C \times (H \times W)}$  and their Gram matrices [7],

$$\mathcal{L}_o = \text{cossim}(F_t^{(s)}, F_t^{(t)}) + \text{cossim} \left( \text{gram}(F_t^{(s)}), \text{gram}(F_t^{(t)}) \right), \quad (6)$$

where  $\text{cossim}$  calculates the cosine similarity between features, and  $\text{gram}$  computes the Gram matrix of the features and vectorizes them. The reflection decoder derives the reflection component  $\hat{R}_t$  from the sum of structural features  $E_t^{(s)}$  and texture features  $E_t^{(t)}$ . This prevents the structural features  $E_t^{(s)}$  from containing unrelated local texture information. Experiments validate that relying on local patterns reduces robustness when cameras encounter corruption.

### 3.3 Semantic-Structural Correlation

Decoupling surface patterns improves robustness against interfering textures from scenes or cameras. However, semantic cues aid structural understanding to provide depth clues. Semantic objects in the scene carry significant structural information, with depth continuity for the same object as shown in Figure 1.  $I_t$  passes through the structure encoder to produce  $F_t^{(s)}$ , and through the semantic expert encoder to generate  $F_t^{(e)} \in \mathbb{R}^{C \times (H' \times W')}$ . The structural embedding  $F_t^{(s)}$  and semantic embedding  $F_t^{(e)}$  of frame  $t$  have a structural correlation that exhibits graph-like characteristics in the feature embeddings. We aim to model the structural correlation between the structural embedding  $F_t^{(s)}$  and semantic embedding  $F_t^{(e)}$  of frame  $t$  as an isomorphic graph. This constructs feature mapping from semantic to structural cues, enabling efficient transfer of structure-relevant knowledge.



The semantic expert model extracts features  $E_t^{(e)}$  from  $I_t$ . To align with  $F_t^{(s)}$  from our structure encoder, we map  $E_t^{(e)}$  to  $F_t^{(e)}$  via bilinear interpolation, convolution, and vectorization. The correlation between  $F_t^{(s)}$  and  $F_t^{(e)}$  reflects cross-model feature relationships. We represent this via the inter-channel correlation matrix  $A(F_t^{(s)}, F_t^{(e)})$ , where

$$A(F_t^{(s)}, F_t^{(e)}) = \frac{F_t^{(s)} \cdot F_t^{(e)}}{\|F_t^{(s)}\| \cdot \|F_t^{(e)}\|}. \quad (7)$$

We formulate  $A(F_t^{(s)}, F_t^{(e)})$  as the adjacency matrix of a graph. By performing feature aggregation via GIN (Graph Isomorphism Network) projector [33], this effectively constructs an isomorphic graph between the feature nodes across the structure encoder and the semantic expert encoder. It enables the distillation of structural knowledge from the expert to the student model.

The expert encoder and structure encoder aggregate structural knowledge between the embedding nodes  $F_t^{(e)}$  and  $F_t^{(s)}$  to construct the virtual embedding nodes  $F_t^{(s)'}$  and  $F_t^{(e)'}$ ,

$$\begin{aligned} F_t^{(s)'} &= \text{gin}\left(\theta^{(e \rightarrow s)}, F_t^{(e)}, A(F_t^{(e)}, F_t^{(s)})\right) \\ F_t^{(e)'} &= \text{gin}\left(\theta^{(s \rightarrow e)}, F_t^{(s)}, A(F_t^{(s)}, F_t^{(e)})\right) \end{aligned} \quad (8)$$

where gin performs feature aggregation based on parameters  $\theta$ , embedding nodes  $F_t$ , and adjacency matrix  $A(\cdot, \cdot)$ . We update the feature nodes via two-dimensional convolutional layers in the aggregation function of GIN. The virtual nodes  $F_t^{(s)'}$  and  $F_t^{(e)'}$  progressively converge to the feature nodes  $F_t^{(e)}$  and  $F_t^{(s)}$  respectively. We define the loss function for this convergence as

$$\mathcal{L}_d = 2 - \text{cossim}(F_t^{(s)'}, F_t^{(s)}) - \text{cossim}(F_t^{(e)'}, F_t^{(e)}) \quad (9)$$

By optimizing the cosine distance between the virtual and feature nodes, the structure encoder obtains semantic structural knowledge from the expert encoder through bidirectional feature aggregation.

### 3.4 Total Loss

Modeling depth-structure consistency, local texture disambiguation, and semantic-structural correlation constitutes multiple optimization objectives. We consolidate these goals into one overall loss  $\mathcal{L}$  for joint optimization,

$$\mathcal{L} = \lambda_p \mathcal{L}_p + \lambda_v \mathcal{L}_v + \lambda_e \mathcal{L}_e + \lambda_r \mathcal{L}_r + \lambda_o \mathcal{L}_o + \lambda_d \mathcal{L}_d, \quad (10)$$

where  $\lambda_p$ ,  $\lambda_v$ ,  $\lambda_e$ ,  $\lambda_r$ ,  $\lambda_o$ ,  $\lambda_d$  balance the weights of different loss terms. The structure-centric self-supervised monocular depth model is trained with  $\mathcal{L}$  as the overall optimization objective. We provide a comprehensive list of all the weights  $\lambda$  used during the training process in the supplementary materials.

## 4 Experiments

To validate the efficacy of our structure-centric monocular depth estimation, we experimentally verify that: our method effectively incorporates scene structure information from semantics and illumination, enhancing the estimation accuracy of self-supervised depth estimation models in a wider range of real-world scenes; our proposed approach substantially improves depth estimation accuracy under common outdoor adverse weather, motion blur and in poor illumination conditions like night-time environments, realizing robust monocular depth estimation.

### 4.1 Implementation Details

We train our model on the clear KITTI dataset [8] to avoid cheating under corrupted monocular depth estimation. No corrupted samples are seen during training. The original KITTI samples have a resolution of  $1242 \times 375$ . We scale them to  $640 \times 192$  for faster training. We train our model for 20 epochs using the Adam optimizer on the `eigen_zhou` split [8] with static sequences removed. We implement the model in PyTorch on a server with NVIDIA V100S GPUs and Intel Xeon Gold 6132 CPUs. We utilize the advanced SAM-Huge model [18] as our semantic expert model in the knowledge distillation process. Throughout the distillation process, we exclusively rely on commonly used semantic segmentation teacher models and do not introduce any prior depth knowledge. In our experiments, we present two model configurations: one built upon MonoDepth2, referred to as Ours<sub>R</sub>, and another based on MonoViT, denoted as Ours<sub>T</sub> (distilled from SAM-Huge [18]). These configurations illustrate the broad applicability of our proposed training approach. The learning rates are  $5 \times 10^{-5}$  for the structure and texture encoders of Ours<sub>T</sub>, and  $1 \times 10^{-4}$  for other components. We use exponential decay with a ratio of 0.9 for learning rate scheduling.

### 4.2 Depth Estimation Performance under Corruptions

Adverse weather and motion blur are common camera degradation scenarios in outdoor driving or sports. For adverse weather, we select the snow, frost, fog, and rain conditions that interfere most with camera observations for validation. We validate robustness and out-of-distribution generalization on corrupted KITTI, DENSE [12] and nuScenes [3] nighttime datasets using models trained on clear KITTI. We compare against a series of state-of-the-art models with innovations in training approaches, e.g. SCDepthV1 [1], SCDepthV3 [31], PackNet-SfM [13], DynaDepth [35], and SGDepth [19], and model architectures, e.g. MonoViT [36], LiteMono [34], HRDepth [24]. All of the above works using different training approaches adopt the model architecture configuration used in MonoDepth2. We favor lower Abs Rel, Sq Rel, RMSE, and RMSE log, while higher  $\delta_1$ ,  $\delta_2$ , and  $\delta_3$  indicate better performance for all metrics in this section. The presentation of representative indicators takes place in Table 1 and Table 2. For more detailed indicators, please refer to the supplementary materials. In the following comparisons, to ensure a fair comparison, we separately compare the models of

**Table 1:** Performance comparison of monocular depth estimation on the KITTI Snow, KITTI Frost, and KITTI Motion Blur test sets. To ensure fairness in comparisons and avoid introducing any corrupted samples, all the models mentioned above are trained on the KITTI Clear training dataset. Upper Section: ResNet-based Architectures. Lower Section: Hybrid Architectures.  $\downarrow$ : Lower is better.  $\uparrow$ : Higher is better.

Method	Knowledge	Snow		Frost		Motion Blur		Clear	
		Abs Rel $\downarrow$	$\delta_1$ $\uparrow$	Abs Rel $\downarrow$	$\delta_1$ $\uparrow$	Abs Rel $\downarrow$	$\delta_1$ $\uparrow$	Abs Rel $\downarrow$	$\delta_1$ $\uparrow$
MonoDepth2	M	0.462	0.330	0.289	0.534	<u>0.242</u>	<u>0.613</u>	0.117	0.875
SCDepthV1	M	0.435	0.335	<u>0.259</u>	0.572	0.334	0.492	0.118	0.860
HRDepth	M	0.520	<b>0.478</b>	0.492	0.504	0.635	0.415	<b>0.106</b>	<b>0.891</b>
SCDepthV3	M+D	<u>0.410</u>	0.330	0.321	0.436	0.282	0.537	0.117	0.866
DynaDepth	M+I	0.484	0.301	0.320	0.493	0.345	0.487	<u>0.114</u>	<u>0.876</u>
SGDepth	M+S	0.473	0.322	0.265	<u>0.575</u>	0.259	0.578	0.117	0.874
PackNet-SfM	M+S	0.426	0.349	0.315	0.491	0.291	0.519	0.116	0.868
<b>Ours<sub>R</sub></b>	M+S	<b>0.404</b>	<u>0.358</u>	<b>0.236</b>	<b>0.608</b>	<b>0.238</b>	<b>0.617</b>	0.118	0.869
LiteMono	M	0.381	0.355	0.258	0.567	0.252	0.582	0.113	0.882
MonoViT	M	<u>0.230</u>	<u>0.621</u>	<u>0.212</u>	<u>0.663</u>	<u>0.189</u>	<u>0.699</u>	<b>0.100</b>	<b>0.898</b>
<b>Ours<sub>T</sub></b>	M+S	<b>0.224</b>	<b>0.627</b>	<b>0.202</b>	<b>0.678</b>	<b>0.183</b>	<b>0.710</b>	<u>0.105</u>	<u>0.894</u>

the MonoDepth2 architecture (including Ours<sub>R</sub>) and the models of the hybrid architecture of Convolutional Networks and Transformer (including Ours<sub>T</sub>).

**Robustness and Generalization.** Obtaining accurate depth data under corrupted conditions remains challenging. In the adverse real world, the robustness to corruptions reflects the model’s out-of-distribution generalization ability. To reliably benchmark monocular depth methods, we build corrupted versions of the KITTI test set following the approaches in [25] and [17]. As shown in Table 1, under the benchmarks of Snow, Frost, and Motion Blur scenes constructed based on KITTI, both Ours<sub>T</sub> and Ours<sub>R</sub> have leading performances in their respective sections. Notably, Ours<sub>R</sub> achieves significant gains over MonoDepth2 under the same model architecture, demonstrating the efficacy of our training scheme for robustness under corruptions. Monocular depth estimation also faces challenges in handling out-of-distribution images due to camera specifics. We conduct benchmarks on the DENSE dataset and the nuScenes nighttime dataset, which have camera configurations completely different from the KITTI dataset. These tests cover scenarios with rain, fog, clear weather, and poor nighttime lighting conditions. As shown in Table 2, under the rain, fog, and clear scenarios, both Ours<sub>R</sub> and Ours<sub>T</sub> significantly surpass other methods. This demonstrates the effectiveness of our approach in improving out-of-distribution monocular depth estimation performance in real adverse weather scenarios. The DENSE dataset also covers daytime and nighttime testing scenarios, each including testing samples of rain, fog, and clear scenarios during both daytime and nighttime. To highlight the validation in nighttime environments with poor lighting conditions and uneven light distribution, we compared benchmarks of the nighttime

**Table 2:** Generalization performance comparison of monocular depth estimation on the DENSE benchmark. The benchmark comprises real rain and fog scenarios under both daytime and night conditions. Upper Section: ResNet-based Architectures. Lower Section: Hybrid Architectures. ↓: Lower is better. ↑: Higher is better.

Method	Know-ledge	Fog			Rain			Clear		
		Abs Rel ↓	RMSE ↓	$\delta_1$ ↑	Abs Rel ↓	RMSE ↓	$\delta_1$ ↑	Abs Rel ↓	RMSE ↓	$\delta_1$ ↑
MonoDepth2	M	0.266	4.865	0.526	0.274	4.962	0.517	0.281	4.995	0.572
SCDepthV1	M	0.269	5.125	0.539	0.266	5.118	0.537	0.297	5.147	0.507
HRDepth	M	0.262	4.707	0.533	0.262	4.730	0.528	0.272	4.681	0.532
SCDepthV3	M+D	0.257	5.087	0.578	0.253	5.087	0.581	0.289	5.152	0.531
DynaDepth	M+I	0.271	5.191	0.519	0.281	5.311	0.504	0.273	4.912	0.493
SGDepth	M+S	0.260	4.967	0.548	0.252	4.912	0.563	0.261	4.856	0.547
PackNet-SfM	M+S	0.254	4.697	0.532	0.256	4.765	0.533	0.258	4.585	0.550
<b>Ours<sub>R</sub></b>	M+S	<b>0.240</b>	<b>4.567</b>	<b>0.580</b>	<b>0.247</b>	<b>4.642</b>	<b>0.587</b>	<b>0.256</b>	<b>4.645</b>	<b>0.575</b>
LiteMono	M	0.325	4.863	0.422	0.322	4.787	0.414	0.310	5.095	0.481
MonoViT	M	0.239	4.600	0.556	0.237	4.767	0.564	0.242	4.320	0.555
<b>Ours<sub>T</sub></b>	M+S	<b>0.224</b>	<b>3.907</b>	<b>0.583</b>	<b>0.224</b>	<b>4.046</b>	<b>0.571</b>	<b>0.218</b>	<b>3.725</b>	<b>0.606</b>

nuScenes dataset in Table 3. Here, Ours<sub>R</sub> has a significant advantage over other works, and Ours<sub>T</sub> also has certain advantages for hybrid architecture tasks.

**Ablation Studies.** As shown in Table 4, we conduct ablation studies on the generalizability of the training scheme, as well as the effects of local texture disambiguation and semantic-structural correlation on monocular depth robustness. Both Ours<sub>R</sub> and Ours<sub>T</sub> outperform their respective baselines, demonstrating the general applicability of our framework. Removing local texture disambiguation and semantic-structural correlation from Ours<sub>T</sub> reduces performance, but it still surpasses MonoViT. Removing the GIN Projector during semantic distillation led to further performance degradation, validating its role in constructing isomorphic mappings between feature channels. These results highlight the positive impact of addressing these sub-problems and the versatility of our approach in supporting semantic knowledge distillation with different ViT expert models.

### 4.3 Feature Map Visualization

To enable the structure encoder to focus on scene structure and avoid over-reliance on local textures that degrades depth estimation under interference, we aim to solve local texture disambiguation and semantic-structural correlation. By introducing a texture encoder to model fine textures in the scene, we expect mutual suppression between the two encoders to improve the structure encoder’s robustness to interfering textures. To validate this, we visualize the intermediate features  $F_t^{(s)}$  and  $F_t^{(t)}$  of the structure encoder and texture decoder. We upsample the multi-scale features of the encoder and merge the channels. Additionally, we perform PCA (Principal Component Analysis) dimensionality reduction on the channel dimensions to form visualizations of  $F_t^{(s)}$  and  $F_t^{(t)}$ . As seen in Figure 3, the reflection encoder successfully extracts interfering textures, which manifests

**Table 3:** Generalization performance comparison of monocular depth estimation on the nuScenes nighttime benchmark. This benchmark has poor lighting conditions and is subject to a large amount of uneven lighting interference. Upper Section: ResNet-based Architectures. Lower Section: Hybrid Architectures. ↓: Lower is better. ↑: Higher is better.

Method	Knowledge	Abs	Sq	RMSE	RMSE	$\delta_1 \uparrow$	$\delta_2 \uparrow$	$\delta_3 \uparrow$
		Rel ↓	Rel ↓	↓	Log ↓			
MonoDepth2	M	<u>0.386</u>	<u>5.030</u>	<u>13.090</u>	<u>0.548</u>	<u>0.386</u>	<u>0.671</u>	<b>0.820</b>
SCDepthV1	M	0.468	6.760	14.773	0.662	0.322	0.570	0.733
HRDepth	M	0.458	6.131	14.439	0.632	0.305	0.556	0.739
SCDepthV3	M+D	0.414	5.560	14.488	0.627	0.321	0.613	0.774
DynaDepth	M+I	0.483	6.820	14.724	0.643	0.293	0.559	0.744
SGDepth	M+S	0.414	5.585	13.888	0.591	0.350	0.621	0.785
PackNet-SfM	M+S	0.442	6.095	13.867	0.588	0.319	0.609	0.791
<b>Ours<sub>R</sub></b>	M+S	<b>0.364</b>	<b>4.831</b>	<b>12.998</b>	<b>0.535</b>	<b>0.412</b>	<b>0.686</b>	<u>0.811</u>
LiteMono	M	0.282	3.794	12.074	0.457	0.531	0.777	0.861
MonoViT	M	<u>0.272</u>	<b>3.268</b>	<b>10.775</b>	<b>0.411</b>	<u>0.557</u>	<b>0.783</b>	<b>0.882</b>
<b>Ours<sub>T</sub></b>	M+S	<b>0.269</b>	<u>3.416</u>	<u>11.282</u>	<u>0.420</u>	<b>0.569</b>	<u>0.781</u>	<u>0.872</u>


**Table 4:** Ablation studies of monocular depth estimation on the KITTI Snow, KITTI Frost, KITTI Motion Blur, and KITTI Clear test sets. TI: Local texture disambiguation. SD: Semantic distillation. GP: GIN Projector. R: The architecture of MonoDepth2 based on ResNet-18 [14]. T: The architecture of MonoViT based on MPViT-Tiny [21].

Ablation				Snow		Frost		Motion Blur		Clear	
Arch	TI	SD	GP	Abs Rel ↓	$\delta_1 \uparrow$	Abs Rel ↓	$\delta_1 \uparrow$	Abs Rel ↓	$\delta_1 \uparrow$	Abs Rel ↓	$\delta_1 \uparrow$
R	×	×	×	0.462	0.330	0.289	0.534	0.242	0.613	0.117	0.875
R	✓	✓	✓	0.404	0.358	0.236	0.608	0.238	0.617	0.118	0.869
T	×	×	×	0.230	<u>0.621</u>	0.212	0.663	0.189	0.699	<b>0.100</b>	<b>0.898</b>
T	×	✓	×	<u>0.227</u>	0.617	0.209	<b>0.681</b>	0.186	0.704	<u>0.103</u>	0.893
T	×	✓	✓	0.228	0.590	0.205	<u>0.678</u>	0.188	0.677	0.107	0.889
T	✓	×	×	0.232	0.623	<u>0.204</u>	0.677	0.185	0.692	0.109	0.893
T	✓	✓	✓	<b>0.224</b>	<b>0.627</b>	<b>0.203</b>	<u>0.678</u>	<b>0.183</b>	<b>0.710</b>	0.105	<u>0.894</u>

in  $F_t^{(t)}$  as clearer snow tracks and frost textures, and even accurately extracts radial blurring caused by motion. Given that the texture feature  $F_t^{(t)}$  is decoupled from the structural feature  $F_t^{(s)}$ ,  $F_t^{(s)}$  clearly filters out a large amount of interfering textures. This corroborates the effectiveness of our method.

## 5 Conclusion

We start from the principles of imaging in nature and the perspective of object textures, exploring three potential sub-problems in monocular depth estimation: local texture disambiguation and semantic-structural correlation. Current meth-



**Fig. 3:** Visualization of intermediate feature maps. For features  $F_t^{(s)}$  and  $F_t^{(t)}$ , yellow indicates high values. For predicted depth  $\hat{D}_t$ , purple denotes high values. We provide a detailed description of the computation method for visual feature maps in the supplementary materials.

ods do not adequately model these factors, leading to significant degradation in challenging scenarios like adverse weather, motion blur, and nighttime conditions.

We approach these sub-problems by modeling them as optimizable self-supervised objectives, greatly enhancing the robustness of self-supervised monocular depth estimation under interference through knowledge distillation. We disentangle image features into scene structure and local texture, promoting a mutual suppression between the structure encoder and texture encoder to minimize over-reliance on local textures. Furthermore, learnable isomorphic graphs, coupled with a semantic expert model, facilitate the integration of structural knowledge into the depth model.

Our method achieves state-of-the-art performance under adverse weather, poor nighttime lighting, and motion blur scenarios on corrupted KITTI, DENSE, and nighttime nuScenes datasets, demonstrating its effectiveness. Visualizations show successful separation of scene structure from interfering textures, highlighting the robustness of our self-supervised approach for monocular depth estimation. We establish that not all texture details are beneficial for depth estimation; starting from global object semantics with effective guidance enhances estimation performance.

**Acknowledgments.** This work was supported in part by the Strategic Priority Research Program of Chinese Academy of Sciences under Grant XDA28040500, the National Natural Science Foundation of China under Grant 62261042, the Key Research Projects of the Joint Research Fund for Beijing Natural Science Foundation and the Fengtai Rail Transit Frontier Research Joint Fund under Grant L221003, and the Beijing Natural Science Foundation under Grant 4232035 and 4222034.

## References

1. Bian, J., Zhan, H., Wang, N., Li, Z., Zhang, L., Shen, C., Cheng, M., Reid, I.: Unsupervised scale-consistent depth learning from video. *Int. J. Comput. Vis.* **129**(9), 2548–2564 (2021). <https://doi.org/10.1007/s11263-021-01484-6>, <https://doi.org/10.1007/s11263-021-01484-6>
2. Bian, J., Li, Z., Wang, N., Zhan, H., Shen, C., Cheng, M., Reid, I.D.: Unsupervised scale-consistent depth and ego-motion learning from monocular video. In: Wallach, H.M., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E.B., Garnett, R. (eds.) *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*. pp. 35–45 (2019), <https://proceedings.neurips.cc/paper/2019/hash/6364d3f0f495b6ab9dcf8d3b5c6e0b01-Abstract.html>
3. Caesar, H., Bankiti, V., Lang, A.H., Vora, S., Liong, V.E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., Beijbom, O.: nuscenes: A multimodal dataset for autonomous driving. In: *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*. pp. 11618–11628. Computer Vision Foundation / IEEE (2020). <https://doi.org/10.1109/CVPR42600.2020.01164>, [https://openaccess.thecvf.com/content\\_CVPR\\_2020/html/Caesar\\_nuScenes\\_A\\_Multimodal\\_Dataset\\_for\\_Autonomous\\_Driving\\_CVPR\\_2020\\_paper.html](https://openaccess.thecvf.com/content_CVPR_2020/html/Caesar_nuScenes_A_Multimodal_Dataset_for_Autonomous_Driving_CVPR_2020_paper.html)
4. Chen, Y., Wen, C., Liu, W., He, W.: A depth iterative illumination estimation network for low-light image enhancement based on retinex theory. *Scientific Reports* **13**(1), 19709 (Nov 2023). <https://doi.org/10.1038/s41598-023-46693-w>, <https://doi.org/10.1038/s41598-023-46693-w>
5. Gao, W., Rao, D., Yang, Y., Chen, J.: Edge devices friendly self-supervised monocular depth estimation via knowledge distillation. *IEEE Robotics and Automation Letters* **8**(12), 8470–8477 (2023). <https://doi.org/10.1109/LRA.2023.3330054>
6. Gasperini, S., Morbitzer, N., Jung, H., Navab, N., Tombari, F.: Robust monocular depth estimation under challenging conditions. In: *IEEE/CVF International Conference on Computer Vision, ICCV 2023, Paris, France, October 1-6, 2023*. pp. 8143–8152. IEEE (2023). <https://doi.org/10.1109/ICCV51070.2023.00751>, <https://doi.org/10.1109/ICCV51070.2023.00751>
7. Gatys, L.A., Ecker, A.S., Bethge, M.: A neural algorithm of artistic style. *CoRR* **abs/1508.06576** (2015), <http://arxiv.org/abs/1508.06576>
8. Geiger, A., Lenz, P., Stiller, C., Urtasun, R.: Vision meets robotics: The KITTI dataset. *Int. J. Robotics Res.* **32**(11), 1231–1237 (2013). <https://doi.org/10.1177/0278364913491297>, <https://doi.org/10.1177/0278364913491297>
9. Godard, C., Aodha, O.M., Brostow, G.J.: Unsupervised monocular depth estimation with left-right consistency. In: *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*. pp. 6602–6611. IEEE Computer Society (2017). <https://doi.org/10.1109/CVPR.2017.699>, <https://doi.org/10.1109/CVPR.2017.699>
10. Godard, C., Aodha, O.M., Firman, M., Brostow, G.J.: Digging into self-supervised monocular depth estimation. In: *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*. pp. 3827–3837. IEEE (2019). <https://doi.org/10.1109/ICCV.2019.00393>, <https://doi.org/10.1109/ICCV.2019.00393>
11. Goyal, B., Dogra, A., Lepcha, D.C., Goyal, V., Alkhayyat, A., Chohan, J.S., Kukreja, V.: Recent advances in image dehazing: Formal analysis to automated

- approaches. *Information Fusion* **104**, 102151 (2024). <https://doi.org/https://doi.org/10.1016/j.inffus.2023.102151>, <https://www.sciencedirect.com/science/article/pii/S1566253523004670>
12. Gruber, T., Bijelic, M., Heide, F., Ritter, W., Dietmayer, K.: Pixel-accurate depth evaluation in realistic driving scenarios. In: 2019 International Conference on 3D Vision (3DV). pp. 95–105 (2019). <https://doi.org/10.1109/3DV.2019.00020>
  13. Guizilini, V., Hou, R., Li, J., Ambrus, R., Gaidon, A.: Semantically-guided representation learning for self-supervised monocular depth. In: 8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020. OpenReview.net (2020), <https://openreview.net/forum?id=ByxT7TNFvH>
  14. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016. pp. 770–778. IEEE Computer Society (2016). <https://doi.org/10.1109/CVPR.2016.90>, <https://doi.org/10.1109/CVPR.2016.90>
  15. Hendrycks, D., Dietterich, T.G.: Benchmarking neural network robustness to common corruptions and perturbations. In: 7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019. OpenReview.net (2019), <https://openreview.net/forum?id=HJz6tiCqYm>
  16. Hoyer, L., Dai, D., Chen, Y., Köring, A., Saha, S., Van Gool, L.: Three ways to improve semantic segmentation with self-supervised depth estimation. In: 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 11125–11135 (2021). <https://doi.org/10.1109/CVPR46437.2021.01098>
  17. Kar, O.F., Yeo, T., Atanov, A., Zamir, A.: 3d common corruptions and data augmentation. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022. pp. 18941–18952. IEEE (2022). <https://doi.org/10.1109/CVPR52688.2022.01839>, <https://doi.org/10.1109/CVPR52688.2022.01839>
  18. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W., Dollár, P., Girshick, R.B.: Segment anything. In: IEEE/CVF International Conference on Computer Vision, ICCV 2023, Paris, France, October 1-6, 2023. pp. 3992–4003. IEEE (2023). <https://doi.org/10.1109/ICCV51070.2023.00371>, <https://doi.org/10.1109/ICCV51070.2023.00371>
  19. Klingner, M., Termöhlen, J., Mikolajczyk, J., Fingscheidt, T.: Self-supervised monocular depth estimation: Solving the dynamic object problem by semantic guidance. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J. (eds.) *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part XX. Lecture Notes in Computer Science*, vol. 12365, pp. 582–600. Springer (2020). [https://doi.org/10.1007/978-3-030-58565-5\\_35](https://doi.org/10.1007/978-3-030-58565-5_35), [https://doi.org/10.1007/978-3-030-58565-5\\_35](https://doi.org/10.1007/978-3-030-58565-5_35)
  20. Kong, L., Xie, S., Hu, H., Ng, L.X., Cottureau, B., Ooi, W.T.: Robodepth: Robust out-of-distribution depth estimation under corruptions. In: Oh, A., Naumann, T., Globerson, A., Saenko, K., Hardt, M., Levine, S. (eds.) *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023* (2023), [http://papers.nips.cc/paper\\_files/paper/2023/hash/43119db5d59f07cc08fca7ba6820179a-Abstract-Datasets\\_and\\_Benchmarks.html](http://papers.nips.cc/paper_files/paper/2023/hash/43119db5d59f07cc08fca7ba6820179a-Abstract-Datasets_and_Benchmarks.html)



21. Lee, Y., Kim, J., Willette, J., Hwang, S.J.: Mpvit: Multi-path vision transformer for dense prediction. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022. pp. 7277–7286. IEEE (2022). <https://doi.org/10.1109/CVPR52688.2022.00714>, <https://doi.org/10.1109/CVPR52688.2022.00714>
22. Lopez-Rodriguez, A., Mikolajczyk, K.: Desc: Domain adaptation for depth estimation via semantic consistency. *International Journal of Computer Vision* **131**(3), 752–771 (Mar 2023). <https://doi.org/10.1007/s11263-022-01718-1>, <https://doi.org/10.1007/s11263-022-01718-1>
23. Luginov, A., Makarov, I.: Swiftdepth: An efficient hybrid cnn-transformer model for self-supervised monocular depth estimation on mobile devices. In: 2023 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct). pp. 642–647 (2023). <https://doi.org/10.1109/ISMAR-Adjunct60411.2023.00137>
24. Lyu, X., Liu, L., Wang, M., Kong, X., Liu, L., Liu, Y., Chen, X., Yuan, Y.: Hr-depth: High resolution self-supervised monocular depth estimation. In: Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021. pp. 2294–2301. AAAI Press (2021), <https://ojs.aaai.org/index.php/AAAI/article/view/16329>
25. Michaelis, C., Mitzkus, B., Geirhos, R., Rusak, E., Bringmann, O., Ecker, A.S., Bethge, M., Brendel, W.: Benchmarking robustness in object detection: Autonomous driving when winter is coming. *CoRR* **abs/1907.07484** (2019), <http://arxiv.org/abs/1907.07484>
26. Ming, Y., Meng, X., Fan, C., Yu, H.: Deep learning for monocular depth estimation: A review. *Neurocomputing* **438**, 14–33 (2021). <https://doi.org/https://doi.org/10.1016/j.neucom.2020.12.089>, <https://www.sciencedirect.com/science/article/pii/S0925231220320014>
27. Ranftl, R., Bochkovskiy, A., Koltun, V.: Vision transformers for dense prediction. In: 2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021. pp. 12159–12168. IEEE (2021). <https://doi.org/10.1109/ICCV48922.2021.01196>, <https://doi.org/10.1109/ICCV48922.2021.01196>
28. Rematas, K., Ferrari, V.: Neural voxel renderer: Learning an accurate and controllable rendering tool. In: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020. pp. 5416–5426. Computer Vision Foundation / IEEE (2020). <https://doi.org/10.1109/CVPR42600.2020.00546>, [https://openaccess.thecvf.com/content\\_CVPR\\_2020/html/Rematas\\_Neural\\_Voxel\\_Renderer\\_Learning\\_an\\_Accurate\\_and\\_Controllable\\_Rendering\\_Tool\\_CVPR\\_2020\\_paper.html](https://openaccess.thecvf.com/content_CVPR_2020/html/Rematas_Neural_Voxel_Renderer_Learning_an_Accurate_and_Controllable_Rendering_Tool_CVPR_2020_paper.html)
29. Saunders, K., Vogiatzis, G., Manso, L.J.: Self-supervised monocular depth estimation: Let’s talk about the weather. In: IEEE/CVF International Conference on Computer Vision, ICCV 2023, Paris, France, October 1-6, 2023. pp. 8873–8883. IEEE (2023). <https://doi.org/10.1109/ICCV51070.2023.00818>, <https://doi.org/10.1109/ICCV51070.2023.00818>
30. Shi, X., Dikov, G., Reitmayr, G., Kim, T.K., Ghafoorian, M.: 3d distillation: Improving self-supervised monocular depth estimation on reflective surfaces. In: 2023 IEEE/CVF International Conference on Computer Vision (ICCV). pp. 9099–9109 (2023). <https://doi.org/10.1109/ICCV51070.2023.00838>

31. Sun, L., Bian, J., Zhan, H., Yin, W., Reid, I.D., Shen, C.: Sc-depthv3: Robust self-supervised monocular depth estimation for dynamic scenes. *IEEE Trans. Pattern Anal. Mach. Intell.* **46**(1), 497–508 (2024). <https://doi.org/10.1109/TPAMI.2023.3322549>, <https://doi.org/10.1109/TPAMI.2023.3322549>
32. Wang, K., Zhang, Z., Yan, Z., Li, X., Xu, B., Li, J., Yang, J.: Regularizing nighttime weirdness: Efficient self-supervised monocular depth estimation in the dark. In: 2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021. pp. 16035–16044. IEEE (2021). <https://doi.org/10.1109/ICCV48922.2021.01575>, <https://doi.org/10.1109/ICCV48922.2021.01575>
33. Xu, K., Hu, W., Leskovec, J., Jegelka, S.: How powerful are graph neural networks? In: 7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019. OpenReview.net (2019), <https://openreview.net/forum?id=ryGs6iA5Km>
34. Zhang, N., Nex, F., Vosselman, G., Kerle, N.: Lite-mono: A lightweight cnn and transformer architecture for self-supervised monocular depth estimation. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, Canada, June 18-22, 2023. IEEE (2023)
35. Zhang, S., Zhang, J., Tao, D.: Towards scale-aware, robust, and generalizable unsupervised monocular depth estimation by integrating imu motion dynamics. In: Computer Vision – ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXVIII. p. 143–160. Springer-Verlag, Berlin, Heidelberg (2022). [https://doi.org/10.1007/978-3-031-19839-7\\_9](https://doi.org/10.1007/978-3-031-19839-7_9), [https://doi.org/10.1007/978-3-031-19839-7\\_9](https://doi.org/10.1007/978-3-031-19839-7_9)
36. Zhao, C., Zhang, Y., Poggi, M., Tosi, F., Guo, X., Zhu, Z., Huang, G., Tang, Y., Mattoccia, S.: Monovit: Self-supervised monocular depth estimation with a vision transformer. In: International Conference on 3D Vision, 3DV 2022, Prague, Czech Republic, September 12-16, 2022. pp. 668–678. IEEE (2022). <https://doi.org/10.1109/3DV57658.2022.00077>, <https://doi.org/10.1109/3DV57658.2022.00077>
37. Zhao, H., Gallo, O., Frosio, I., Kautz, J.: Loss functions for image restoration with neural networks. *IEEE Trans. Computational Imaging* **3**(1), 47–57 (2017). <https://doi.org/10.1109/TCI.2016.2644865>, <https://doi.org/10.1109/TCI.2016.2644865>
38. Zhou, T., Brown, M., Snavely, N., Lowe, D.G.: Unsupervised learning of depth and ego-motion from video. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017. pp. 6612–6619. IEEE Computer Society (2017). <https://doi.org/10.1109/CVPR.2017.700>, <https://doi.org/10.1109/CVPR.2017.700>