

Weather-Conditioned Depth Anything

Zhaoming Xu¹, Chan-Wei Hu¹, Kuan-Ru Huang¹, Zihao Zhu¹, Renjie Li¹, Yang Zhou¹, and Zhengzhong Tu¹*

Texas A&M University, College Station, TX 77843, USA
{zhaoming0812,tzz}@tamu.edu




Fig. 1: Depth estimation results under various real-weather degraded conditions. Our DA-W generates accurate details and well-preserved structure under snow, fog, rain, and low-light conditions. Compared with Depth Anything v2 [60] and Depth Anything-AC [48], our method produces finer-grained and more robust predictions.

Abstract. Monocular depth estimation foundation models, such as the Depth Anything series, have achieved remarkable performance across diverse domains. However, they still suffer from critical failures under adverse weather conditions, such as fog, rain, snow, or at night. To address this, we present Weather-Conditioned Depth Anything (DA-W), a framework that explicitly disentangles style from content for weather-robust depth estimation. Specifically, we introduce a *Style Filter* trained on a curated mix of real and synthetic degradation datasets to extract content-independent, degradation-aware weather embeddings. This style

* Corresponding Author

embedding is then injected into the Depth Anything backbone using a parameter-efficient, zero-initialized adapter. Such a lightweight modulation allows a single unified model to robustly adapt to diverse conditions—including fog, rain, snow, and low-light—while avoiding catastrophic forgetting of its core generalization abilities in normal conditions. We train the adapter using a pseudo-label distillation and alignment strategy. Our comprehensive experiments demonstrate that our proposed DA-W achieves state-of-the-art robust depth estimation, improving AbsRel by an average of 3.7% on our curated weather benchmarks, while matching or slightly outperforming performance on standard clean benchmarks.

Keywords: robust depth estimation · 3D from single images

1 Introduction

Monocular depth estimation (MDE) has become a fundamental component in 3D perception [15, 33, 57] for cost-effective applications, such as autonomous driving, robotics, and AR/VR, where it can replace relying on LiDAR or stereo cameras. Driven by large-scale datasets and model training, recent depth foundation models [2, 54] have demonstrated strong generalization in depth estimation across diverse environments. However, their performance degrades dramatically under adverse weather conditions—fog, snow, rain, and nighttime—due to atmospheric scattering, occlusion, specular artifacts, and low-light noise, which fundamentally alter the appearance (See Figure. 1).

This robustness gap persists even for state-of-the-art foundation models such as the Depth Anything [30, 59, 60]. Broadly, there are two strategies to improve robustness under adverse weather conditions: **(1)** pre-process the image to remove weather effects before depth estimation, and **(2)** adapt the depth model to the weather conditions.

The first strategy has an inherent limitation: restoration methods are typically optimized for perceptual quality rather than geometric accuracy, and often distort the very cues critical for reliable depth estimation [8, 9, 44].

Some studies train the restoration for task-oriented objectives [13, 27], yet they often overfit to a specific downstream model, making them non-generalizable to other estimators. The second strategy is more principled but requires large-scale annotated depth data under real adverse conditions, which is prohibitively expensive to collect. As a practical workaround, existing




Fig. 2: Performance $\delta_1 \uparrow$ comparison across three datasets: real adverse weather, synthetic adverse weather, and clean weather. Values are normalized per dataset (Baseline = 0.6).




Fig. 3: Feature space comparison for style characterization. Depth Anything v2 [60] *CLS* tokens fail to separate weather conditions, while our *style filter* features form distinct clusters.

methods predominantly rely on **synthetic** weather degradations applied to clean images [23, 42, 50].

Consequently, the key question is whether features learned from synthetic data generalize reliably to real-world scenarios. To probe this generalization gap, we analyze how a foundation depth model represents appearance variations across domains. We extract the **[CLS]** tokens in vision transformer (ViT) embeddings [51] as the global appearance vector on the Depth Anything v2 encoder [60] across diverse real and synthetic weather conditions, and visualize the feature embeddings via t-SNE [35] in Figure. 3 (left).

As may be seen, the features are fundamentally entangled. Representations from different weather types (e.g., fog, rain, snow) are heavily mixed, rather than forming separable clusters. The model conflates weather-specific style with scene geometry into a single, undifferentiated representation, instead of providing a pure weather code. This entanglement is the root cause of the model’s brittleness, as it cannot distinguish appearance variations caused by weather from those of the scene itself, preventing the learning of generalizable weather patterns and fatally hindering synthetic-to-real adaptation.

Motivated by these insights, we propose **Weather-Conditioned Depth Anything (DepthAnything-W or DA-W)**, a simple yet robust adaptation strategy to robustify depth estimation under adverse weather conditions. We first design a *Style Filter* network using contrastive learning on a curated collection of 7 real-world and 4 synthetic weather datasets. This module learns to align the style distributions of real and synthetic degradations, as shown in Figure 3(right), thereby producing a discriminative weather embedding. The embedding then conditions the Depth Anything foundation model using lightweight modulation techniques *AdaLN-Zero*. This lightweight conditioning allows the model to adjust its predictions based on the estimated weather while keeping the encoder frozen, preventing catastrophic forgetting and preserving clean-scene performance. Extensive experiments demonstrate that DA-W achieves state-of-the-art performance across both general and adverse-weather benchmarks Figure 2.

Our contributions are summarized as follows:

- We propose DepthAnything-W, a novel weather-robust depth estimation foundation model that learns to adapt to diverse real weather conditions.
- We introduce a *Style Filter* trained via contrastive loss on mixed datasets of real and synthetic data, learning domain-aligned weather embeddings that bridge the gap between synthetic and real data.
- We propose a parameter-efficient adaptation strategy that conditions the Depth Anything foundation model using the learned weather embedding via lightweight *AdaLN-Zero* modulation.
- Our comprehensive experiments show that DA-W achieves state-of-the-art performance on diverse adverse-weather benchmarks without degrading generalization performance on clean data.

2 Related Work

2.1 Monocular Depth Estimation (MDE).

Monocular depth estimation (MDE) predicts scene geometry from a single RGB image and serves as a foundational component for various 3D perception tasks in robotics [22, 46], augmented and virtual reality [2], and autonomous driving. The Initial deep learning approaches were either supervised using datasets such as KITTI [17] and NYUv2 [45], or self-supervised through stereo [18, 65] or video photometric consistency [55, 61, 64]. However, both strategies exhibited domain bias and limited generalization capabilities.

Recent research has addressed domain bias in early supervised and self-supervised pipelines by using more diverse datasets and larger models. These strategies help improve cross-domain robustness and facilitate foundational depth estimation. MiDaS [40] was one of the early methods to improve generalization by training on multiple datasets. Subsequent research advances these efforts by pretraining on larger and more diverse labeled and unlabeled datasets, as demonstrated by the Depth Anything series [31, 59, 60]. Additional methods, including ZoeDepth [1], UniDepth [38], and Metric3D(v2) [20], target generalization across various camera systems and domains. These strategies demonstrate strong performance on previously unseen tasks.

In addition to scaling discriminative models, another research direction investigates generative depth estimation, which leverages strong priors from large generative models. These methods formulate depth prediction as sampling from a conditional distribution. The framework inherently supports uncertainty-aware inference and enables tasks such as depth inpainting. For example, Marigold [24] employs diffusion-based image generators for monocular depth estimation. Flow matching models like DepthFM [19] formulates monocular depth estimation as distribution transport and enable rapid, high-quality sampling.

2.2 Robust Depth Estimation.

Although large-scale training improves average cross-domain performance, depth models remain brittle when input images violate standard assumptions of image

formation, such as rain streaks, fog scattering, or nighttime noise. These conditions corrupt or remove geometric cues [16, 42]. Among these factors, adverse weather is particularly challenging because of scattering, occlusions, and low illumination [16].

To address this, prior work largely falls into two categories: (1) *Two-stage pipelines*, which restore degraded images using diffusion or unified restoration models like UniRestore [7] and MWFormer [66] before depth estimation, but often alter geometric cues and degrade accuracy [42]; and (2) *Weather-aware training*, which incorporates robustness directly into the depth estimation model. WeatherDepth [53] employs curriculum contrastive learning from clear to severe weather, and md4all [16] enforces consistency between clean and degraded samples. Foundation model extensions like DepthAnything-AC [48] and ACDepth [21], incorporate physics-based degradations and knowledge distillation to achieve strong zero-shot robustness. However, although these methods enhance resilience, they often entangle weather appearance with scene geometry. Our approach explicitly disentangles them using a *Style Filter* that learns domain-invariant weather embeddings, thereby enabling robust depth estimation without sacrificing clean-image generalization.

3 Proposed Method

3.1 Preliminary: Depth Anything Model

We provide a concise overview of the Depth Anything [60] framework. As shown in Figure 4, it comprises two main components: the encoder employs a vision transformer (DINOv2) to extract hierarchical features, augmented with a multi-scale readout mechanism that produces pyramid feature representations, and the decoder adopts the DPT design [39], progressively fusing multi-scale features through refinement blocks with skip connections, followed by a lightweight prediction head for dense depth regression. The decoder performs progressive upsampling with skip-wise fusion to recover fine structure. Depth Anything is trained on large mixed-supervision data with an affine-invariant objective, yielding zero-shot generalization across diverse scenes. In our method, we keep this backbone intact and inject weather conditioning only in the decoder to preserve clean-scene behavior while improving robustness.

3.2 Weather Adaptation

Weather Embedding Formulation Given an RGB image $I \in \mathbb{R}^{H \times W \times 3}$ captured under a weather condition $c \in \mathcal{C}$ (e.g., fog, rain, snow, or low-light), our goal is to learn a *Style Filter* $F_\theta(\cdot)$ that maps I to a compact weather embedding $\mathbf{w} \in \mathbb{R}^d$:

$$\mathbf{w} = F_\theta(I). \quad (1)$$

The embedding $\mathbf{w}$ is designed to be both **weather-discriminative** and **domain-aligned**. It should encode the type of weather while remaining invariant to scene geometry and semantics, such that two distinct scenes under




Fig. 4: Overview of our proposed DepthAnything-W. (a) The training pipeline: clean and degraded images pass through the Depth Anything backbone, with weather embeddings from the *Style Filter* injected via *AdaLN-Zero* at multiple decoder scales to produce weather-conditioned depth maps supervised by $\mathcal{L}_{dis}$ and $\mathcal{L}_{align}$. (b) The *Style Filter* extracts weather embeddings from images via multi-scale Gram matrices, using contrastive learning to cluster by weather type while maintaining scene invariance and real-synthetic alignment.

the same weather yield similar embeddings, whereas the same scene under different weather conditions produces distinct ones. In addition, for a fixed weather category c , embeddings derived from real and synthetic sources are encouraged to occupy the same region of the latent manifold, promoting cross-domain consistency. We qualitatively verify this property using t-SNE [35] visualizations, where successful alignment manifests as overlapping clusters for real and synthetic samples of the same weather type.

Cross-domain Weather Alignment The goal of the *Style Filter* is to cluster images affected by similar weather degradations while remaining invariant to content and domain. We adopt a contrastive loss that encourages embeddings for the same weather condition to be close in the latent space, while pushing embeddings from different weather categories apart. The loss is defined as:

$$\mathcal{L}_{con} = \sum_{(a,b) \in \mathcal{P}} [\mathbb{I}_{ab}[m - d_{ab}]_+ + (1 - \mathbb{I}_{ab})d_{ab}], \quad (2)$$

where $\mathcal{P}$ denotes all possible image pairs in a batch, $d_{ab} = d(\mathbf{w}_a, \mathbf{w}_b)$ is the cosine similarity between weather embeddings $\mathbf{w}_a$ and $\mathbf{w}_b$ extracted by *Style Filter*, m is a positive margin, $\mathbb{I}_{ab} = 1$ when images (a, b) share the same weather type and 0 otherwise, and $[\cdot]_+ = \max(0, \cdot)$.

To promote cross-domain alignment between real and synthetic data, we further treat real-synthetic pairs that share the same weather condition as positive samples in Eq. 2. This modification differs from MWFormer [66], which considers only intra-domain pairs. By explicitly mixing real and synthetic samples during contrastive training, the *Style Filter* learns a unified embedding space in

which images from both domains cluster according to solely weather type. Consequently, $\mathbf{w}$ becomes both **weather-discriminative** and **domain-aligned**, facilitating robust conditioning of the depth estimation backbone across real-world and synthetic degradations.

Optional Designs Motivated by CLIP-style alignment, we also evaluated two variants. (i) *Prompt-driven only* [10] and (ii) *Image-image alignment* [11]. Despite extensive tuning, real and synthetic samples of the same weather still formed separable clusters. We conjecture this is because CLIP features are dominated by semantics and invariant to low-level, physics-driven cues, making them insufficient to close the real-synthetic domain gap. This motivates our Style Filter, which learns weather-specific statistics via an explicit domain-alignment objective. Additional details and negative results are in the supplementary material.

3.3 Weather-Conditioned Depth Anything

In this section, we seek a concise way to inject the compact *weather embedding* $\mathbf{w}$ into a depth foundation model without altering its core architecture. To achieve this, we propose a *weather-conditioned modulation* tailored for DPT-based Depth Anything backbones. As shown in Figure 4, a 64-dimensional weather embedding $\mathbf{w} \in \mathbb{R}^{64}$ is fused at four multi-scale feature maps $\{F_i \in \mathbb{R}^{C_i \times H_i \times W_i}\}_{i=1}^4$ after their per-scale reductions to $C_i=64$, and once on the final fused feature before the output head (five sites in total). A lightweight MLP maps $\mathbf{w}$ to per-channel shift, scale, and a gate, denoted by $[\beta_i(\mathbf{w}), \gamma_i(\mathbf{w}), g_i(\mathbf{w})]$. The features are then updated via an *AdaLN-Zero* style operation:

$$F_i \leftarrow F_i + g_i(\mathbf{w}) \left(\text{LN}(F_i) \odot (1 + \gamma_i(\mathbf{w})) + \beta_i(\mathbf{w}) \right), \quad (3)$$

where LN is channel-wise LayerNorm without affine parameters and broadcasting over spatial dimensions (H_i, W_i) is implicit. All modulation heads are *zero-initialized*, ensuring the initial network output matches the base model. The conditioned path then learns weather-specific increments during training.

The proposed modulation introduces only a negligible parameter yet yields substantial gains in adverse-weather robustness. Moreover, we initialize the model from a pretrained Depth Anything [60] and update only the decoder head and the small modulation heads, thereby preserving clean-domain behavior and mitigating catastrophic forgetting, effectively serving as a lightweight weather adapter for DPT-style Depth Anything backbones.

3.4 Overall Loss

Following Depth Anything [60], all supervision terms use the affine-invariant loss ℓ on disparity maps $y, \hat{y} \in \mathbb{R}^{H \times W}$:

$$\ell(y, \hat{y}) = \frac{1}{HW} \sum_{j=1}^{HW} \left| \frac{y_j - t(y)}{s(y)} - \frac{\hat{y}_j - t(\hat{y})}{s(\hat{y})} \right|, \quad (4)$$

where $t(\cdot)$ and $s(\cdot)$ denote per-image mean and mean absolute deviation, computed over valid pixels that are masked.

Data regime and objectives. We train with three data sources: (i) *clean* images $\mathcal{D}$, (ii) *synthetic weather* images $\mathcal{D}^s$ generated from $\mathcal{D}$, and (iii) *real weather* images $\mathcal{D}^r$ without clean counterparts.

Let $\mathcal{T}_{\text{syn}}$ denote a family of geometry-preserving synthetic weather operators (fog/rain/snow/low-light). For each clean image $x \in \mathcal{D}$ and operator $\tau \in \mathcal{T}_{\text{syn}}$, we generate its synthetic counterpart $x^s = \tau(x) \in \mathcal{D}^s$, forming a paired sample (x, x^s) .

A frozen teacher $\mathcal{F}_t$ provides pseudo labels; the student $\mathcal{F}_\phi(\cdot; \mathbf{w})$ is conditioned by the *Style Filter* embedding $\mathbf{w}(\cdot)$. All discrepancies are measured on *disparity* using the affine-invariant loss $\ell(\cdot, \cdot)$ in Eq. (4).

Distillation loss. For clean images $x \in \mathcal{D}$, synthetic weather images $x^s \in \mathcal{D}^s$, and real weather images $x \in \mathcal{D}^r$, we align the student model (ViT-S) to the teacher (ViT-L). For synthetic images that share geometry with clean counterparts, we use the teacher’s prediction on the corresponding clean image as the target:

$$\mathcal{L}_{\text{dist}} = \begin{cases} \ell(\mathcal{F}_\phi(x; \mathbf{w}(x)), \mathcal{F}_t(x)), & x \in \mathcal{D} \cup \mathcal{D}^r, \\ \ell(\mathcal{F}_\phi(x^s; \mathbf{w}(x^s)), \mathcal{F}_t(x)), & x^s = \tau(x), x \in \mathcal{D}. \end{cases} \quad (5)$$

Alignment losses. For *clean-synthetic* pairs (x, x^s) where $x \in \mathcal{D}$ and $x^s \in \mathcal{D}^s$ share geometry, we enforce agreement between the two student predictions:

$$\mathcal{L}_{\text{pair}} = \ell(\mathcal{F}_\phi(x; \mathbf{w}(x)), \mathcal{F}_\phi(x^s; \mathbf{w}(x^s))). \quad (6)$$

For *real weather* images $x \in \mathcal{D}^r$ without clean pairs, we apply a geometry-preserving appearance transform τ_{aug} (ColorJitter) and enforce augmentation consistency:

$$\mathcal{L}_{\text{aug}} = \ell(\mathcal{F}_\phi(x; \mathbf{w}(x)), \mathcal{F}_\phi(\tau_{\text{aug}}(x); \mathbf{w}(\tau_{\text{aug}}(x)))). \quad (7)$$

Total objective. The final loss combines teacher guidance with scene-pair alignment (*clean-synthetic*) and augmentation consistency (*real weather*):

$$\mathcal{L} = \mathcal{L}_{\text{dist}} + \mathcal{L}_{\text{pair}} + \mathcal{L}_{\text{aug}}, \quad (8)$$

4 Implementation

4.1 Dataset Curation

To learn a weather-embedding invariant across domains, we curate 7 real-world adverse-weather datasets and 4 clean datasets, along with their corresponding synthetic weather views. The adverse-weather collection comprises ACDC [41],

Table 1: Zero-shot relative depth estimation on real adverse-weather benchmarks. “–” denotes unavailable results (not reported or not supported by the released model, e.g., MWFormer on NuScenes-night and RobotCar-night).

Method	Encoder	NuScenes-night		RobotCar-night		DS-rain		DS-cloud		DS-fog	
		AbsRel↓	δ_1 ↑	AbsRel↓	δ_1 ↑	AbsRel↓	δ_1 ↑	AbsRel↓	δ_1 ↑	AbsRel↓	δ_1 ↑
DynaDepth [62]	ResNet	0.381	0.394	0.512	0.294	0.239	0.606	0.172	0.608	0.144	0.901
EC-Depth [47]	ViT-S	0.243	0.623	0.228	0.552	0.155	0.766	0.158	0.767	0.109	0.861
STEPS [63]	ResNet	0.252	0.588	0.350	0.367	0.301	0.480	0.252	0.588	0.216	0.641
robustdepth [42]	ViT-S	0.260	0.597	0.311	0.521	0.167	0.755	0.168	0.775	0.105	0.882
weather-depth [53]	ViT-S	–	–	–	–	0.158	0.764	0.160	0.767	0.105	0.879
Syn2Real [58]	ViT-S	–	–	–	–	0.171	0.729	–	–	0.128	0.845
DepthPro [2]	ViT-S	0.218	0.669	0.237	0.534	0.124	0.841	0.158	0.779	0.102	0.892
DepthAnything v1 [59]	ViT-S	0.232	0.679	0.239	0.518	0.133	0.819	0.150	0.801	0.098	0.891
DepthAnything v2 [60]	ViT-S	0.200	0.725	0.239	0.518	0.125	0.840	0.151	0.798	0.103	0.890
DepthAnything v3 [30]	ViT-S	0.275	0.586	0.169	0.755	0.134	0.822	0.154	0.785	0.105	0.887
DepthAnything-AC [48]	ViT-S	0.198	0.727	0.227	0.555	0.125	0.840	0.149	0.801	0.103	0.889
MWFormer [66]+DA v2 [60]	ViT-S	–	–	–	–	0.127	0.833	0.151	0.795	0.104	0.884
DarkIR [13]+DA v2 [60]	ViT-S	0.204	0.710	0.232	0.542	–	–	–	–	–	–
DepthAnything-W	ViT-S	0.194	0.737	0.239	0.513	0.123	0.842	0.151	0.795	0.101	0.896

Table 2: Zero-shot relative depth estimation on KITTI-C (synthetic corruptions). “–” denotes unavailable results.

Method	Encoder	Dark		Snow		Fog		Motion	
		AbsRel↓	δ_1 ↑	AbsRel↓	δ_1 ↑	AbsRel↓	δ_1 ↑	AbsRel↓	δ_1 ↑
DepthPro [2]	ViT-S	0.145	0.793	0.197	0.685	0.087	0.926	0.170	0.746
DepthAnything v1 [59]	ViT-S	0.129	0.826	0.113	0.870	0.092	0.912	0.133	0.822
DepthAnything v2 [60]	ViT-S	0.130	0.832	0.115	0.872	0.097	0.905	0.127	0.840
DepthAnything v3 [30]	ViT-S	0.180	0.737	0.182	0.715	0.119	0.858	0.196	0.698
DepthAnything-AC [48]	ViT-S	0.130	0.834	0.114	0.873	0.096	0.906	0.126	0.841
MWFormer [66] + DA v2 [60]	ViT-S	–	–	0.109	0.884	0.099	0.899	–	–
DarkIR [13] + DA v2 [60]	ViT-S	0.148	0.803	–	–	–	–	–	–
DepthAnything-W	ViT-S	0.126	0.837	0.107	0.884	0.093	0.910	0.125	0.841

RTTS [12], Snow100K [34], Muses [3], RID and RIS [28], and NightCity [49], totaling approximately 15K images covering fog, rain, snow, and low-light conditions. The clean datasets include COCO [32], MegaDepth [29], SA-1B [25], and HRWSI [14, 56], providing approximately 200K high-quality images.

During *Style Filter* training (Stage I), we uniformly subsample 20K images from the clean datasets to balance the scale with real-world adverse weather data, ensuring effective cross-domain learning without category imbalance. Following the synthetic degradation recipe in RobustSAM [8], we generate corresponding synthetic weather corruptions for each clean image, creating paired *clean-synthetic* views with shared underlying geometry.

During depth adaptation training (Stage II), we use the entire dataset, comprising all real adverse-weather images and synthetic corruptions. For depth supervision, we use validity masks when provided by the datasets; otherwise, we generate sky masks using Lang-SAM [37] and exclude these regions to avoid ill-posed supervision for unbounded depth. Detailed dataset statistics are provided in the supplementary material.




Fig. 5: Qualitative comparison on various synthetic weather degradations. Our method (DA-W) produces more accurate and stable depth predictions under challenging synthetic weather conditions. The red boxes highlight regions where DA-W maintains clearer structural details and smoother depth transitions, while other methods suffer from noisy artifacts and depth discontinuities.

4.2 Training Details

We adopt a parameter-efficient two-stage training protocol. **Stage I:** We train the *Style Filter* on the combined real and synthetic weather data to learn a compact 64-D weather embedding. The training procedure is described in MWFormer [66]. **Stage II:** We optimize only the DPT decoder head and the *AdaLN-Zero* modulation layers, which inject the weather embedding at five decoder stages. We train for 20 epochs using AdamW with a learning rate of $5e-6$, weight decay of 0.01, and a batch size of 16. All experiments are conducted on a single NVIDIA A100 GPU.

4.3 Evaluation Protocol

We follow the evaluation protocol of DepthAnything-AC [48], using standard relative depth metrics (AbsRel and δ_1) across two benchmark categories. **Adverse-weather benchmarks:** We evaluate on real world adverse weather datasets NuScenes-night [6], RobotCar-night [36], DrivingStereo-rain/cloud/fog [58], and the synthetic weather corruption datasets KITTI-C (Dark, Snow, Fog, Motion) [26]. **Mix-weather benchmarks:** To address the practical setting where multiple weather conditions may coexist (e.g., rainy nights or snowy nights), we introduce a *mix-weather* benchmark composed of both synthetic and real-world subsets. **Clean-domain benchmarks:** To verify that weather adaptation does not degrade general performance, we report results on KITTI [17], NYU-Depth v2 [45], Sintel [5], ETH3D [43], and DIODE [52].

5 Experiment

Comparison on real adverse weather. Table 1 presents a summary of zero-shot relative depth estimation performance across five real-world adverse-

Table 3: Relative depth estimation results on general benchmarks. Our method preserves DepthAnything’s performance on clean scenes. Best/second best are **bold**/underlined.

Method	Encoder	KITTI [17]		NYU-D [45]		Sintel [5]		ETH3D [43]		DIODE [52]	
		AbsRel↓	δ_1 ↑	AbsRel↓	δ_1 ↑	AbsRel↓	δ_1 ↑	AbsRel↓	δ_1 ↑	AbsRel↓	δ_1 ↑
DynaDepth [62]	ViT-S	0.085	0.916	0.212	0.630	0.458	0.391	0.192	0.703	0.180	0.737
EC-Depth [47]	ViT-S	<u>0.075</u>	0.931	0.220	0.606	0.447	0.405	0.181	0.779	0.172	0.760
STEPS [63]	ResNet	0.204	0.683	0.220	0.606	0.458	0.375	0.197	0.699	0.214	0.652
robustdepth [42]	ViT-S	0.076	0.929	0.174	0.718	0.411	0.436	0.185	0.776	0.182	0.741
weather-depth [53]	ViT-S	<u>0.075</u>	0.932	0.167	0.733	0.419	0.443	0.181	0.771	0.170	0.760
DepthPro [2]	ViT-S	0.068	0.947	<u>0.052</u>	<u>0.972</u>	0.240	0.682	0.057	0.963	0.060	0.952
DepthAnything v1 [59]	ViT-S	0.082	0.936	0.053	<u>0.972</u>	0.211	0.736	<u>0.064</u>	0.966	0.078	0.940
DepthAnything v2 [60]	ViT-S	0.083	0.934	0.051	0.973	0.240	0.712	<u>0.064</u>	0.969	0.073	<u>0.943</u>
DepthAnything-AC [48]	ViT-S	0.083	0.934	0.051	0.973	0.235	0.715	0.065	<u>0.968</u>	0.073	0.941
DepthAnything-W	ViT-S	0.080	<u>0.937</u>	0.051	0.973	<u>0.232</u>	<u>0.716</u>	<u>0.064</u>	<u>0.968</u>	<u>0.071</u>	0.942

weather tracks. **DepthAnything-W** demonstrates the highest overall robustness within the DepthAnything family on *NuScenes-night* and *DS-rain*, outperforming both DepthAnything v2 and DepthAnything-AC in terms of *both* AbsRel and δ_1 . On *DS-cloud* and *DS-fog*, our method remains competitive but does not achieve the top score. These conditions, characterized by reduced global contrast and partially suppressed fine-grained textures, may still require specialized approaches. The most challenging setting for our model is *RobotCar-night*, where methods such as DepthAnything v3 and DepthPro perform substantially better. This performance gap likely reflects domain characteristics, including camera response, exposure dynamics, and scene statistics, that are underrepresented in the training mixture. Finally, restoration-based two-stage pipelines (e.g., MWFormer + DepthAnything v2) can be competitive on some DrivingStereo subsets (e.g., *DS-fog*: AbsRel 0.104), but they require additional preprocessing. Overall, these results support our hypothesis that explicit weather embeddings with lightweight decoder-only AdaLN conditioning provide a practical alternative for improving adverse-weather robustness while preserving strong generalization.

Synthetic corruptions (KITTI-C). Table 2 evaluates robustness under four synthetic corruptions. **DepthAnything-W** achieves the best performance on *Dark*, *Snow*, and *Motion*, and remains competitive on *Fog*. Compared with DepthAnything v2, our weather-conditioned model consistently improves both metrics across all settings. The largest gain appears on *Snow*, suggesting that domain-aligned weather embeddings help preserve cues related to scattering and partial occlusions that are easily suppressed by generic representations. On *Fog*, we observe a smaller margin and slightly trail methods that explicitly target this corruption, indicating that dense fog remains challenging when appearance cues are heavily washed out. Notably, our conditioning approach is comparable to restoration-based or two-stage pipelines while avoiding additional preprocessing and the potential for error accumulation, offering a simpler and more efficient path to corruption robustness. Figure 5 qualitatively illustrates these improvements, with additional results provided in the Appendix.

Table 4: Mix-weather benchmark on KITTI Eigen split with composed corruptions (RoboDepth [26]). We report $\delta_1 \uparrow$ and AbsRel $\downarrow$. Best/second best are **bold**/underlined.

Corruption	DA v1 [59]		DA v2 [60]		DA v3 [30]		DA-AC [48]		DA-W (Ours)	
	$\delta_1 \uparrow$	AbsRel $\downarrow$	$\delta_1 \uparrow$	AbsRel $\downarrow$	$\delta_1 \uparrow$	AbsRel $\downarrow$	$\delta_1 \uparrow$	AbsRel $\downarrow$	$\delta_1 \uparrow$	AbsRel $\downarrow$
Clean	0.937	<u>0.082</u>	<u>0.934</u>	0.083	0.912	0.094	0.925	0.088	0.937	0.080
Fog+Rain	0.708	0.192	0.726	0.185	0.636	0.225	<u>0.731</u>	<u>0.181</u>	0.750	0.173
Fog+Snow	0.681	0.205	0.712	0.192	0.481	0.303	<u>0.731</u>	<u>0.179</u>	0.734	0.177
Fog+Night	0.715	0.185	0.738	0.178	0.640	0.233	<u>0.742</u>	<u>0.175</u>	0.747	0.171
Rain+Snow	0.671	0.209	<u>0.754</u>	0.171	0.627	0.228	<u>0.754</u>	<u>0.167</u>	0.771	0.158
Rain+Night	0.658	0.217	<u>0.680</u>	0.210	0.657	0.224	<u>0.680</u>	<u>0.207</u>	0.714	0.192
Snow+Night	0.724	0.182	0.742	0.174	0.639	0.230	0.756	<u>0.166</u>	<u>0.754</u>	0.165
Fog+Rain+Snow	0.562	0.269	0.596	0.251	0.451	0.322	<u>0.634</u>	<u>0.230</u>	0.670	0.214
Fog+Rain+Night	0.661	0.222	<u>0.668</u>	<u>0.218</u>	0.546	0.283	0.638	0.232	0.673	0.217
Fog+Snow+Night	0.529	0.287	0.577	<u>0.262</u>	0.496	0.307	<u>0.617</u>	0.237	0.623	0.237
Rain+Snow+Night	0.616	0.238	0.645	0.225	0.611	0.250	<u>0.679</u>	<u>0.207</u>	0.696	0.197
Fog+Rain+Snow+Night	0.545	0.278	0.562	0.268	0.501	0.309	<u>0.608</u>	<u>0.245</u>	0.630	0.236

General benchmarks (clean scenes). To assess catastrophic forgetting, we evaluate on standard clean datasets (Table 3). **DepthAnything-W** maintains the backbone’s performance, with small but consistent gains on *KITTI* and *Sintel*, parity on *ETH3D*. This indicates that decoder-side AdaLN with zero-initialized heads preserves clean-scene behavior and improves robustness under adverse weather conditions.

Mix-weather benchmark (coexisting conditions). Real driving often exhibits co-occurring degradations, such as rainy nights or snowy nights, which are not covered by single-factor benchmarks. To evaluate this setting, we construct a *mix-weather* benchmark with both synthetic and real-world subsets. For the *synthetic* subset, we follow RoboDepth [26] to compose weather corruptions on the KITTI Eigen split, testing 11 mixed corruption combinations across 5 severity levels. As summarized in Table 4, **DepthAnything-W** consistently improves both δ_1 and AbsRel over DepthAnything v1/v2/v3 and DepthAnything-AC across most mixed-condition tracks, indicating that explicit weather embeddings and decoder-only AdaLN conditioning remain effective even when multiple condition factors overlap. For the *real-world* subset, we qualitatively evaluate on Boreas [4] *night+snow* sequences, where our predictions exhibit fewer structural artifacts and more stable geometry under simultaneous low-illumination and snowfall (see Figure. 6). These results indicate that adaptation to individual weather conditions can transfer to *coexisting* degradations: our conditioning design remains effective on mixed corruptions without requiring restoration-stage preprocessing.

5.1 Ablation Study

Table 5 shows three design choices: fine-tuning strategy (full or decoder-only), data composition (with or without real-world degradations), and AdaLN-based

weather conditioning. Although no single configuration dominates all metrics, the complete decoder-only variant with real-world degradations and weather injection offers the best overall trade-off.

Fine-tuning strategy. Transitioning from full fine-tuning to decoder-only tuning gives a better overall trade-off across adverse and clean benchmarks. Specifically, decoder-only tuning results in higher δ_1 on *NuScenes-night* (0.741 versus 0.727) and enhances clean-scene generalization on *NYU-D* and *KITTI*. These findings indicate that the pretrained encoder is sufficiently general and that freezing it helps prevent representation drift.

Data composition. Incorporating real-world degradations during training primarily benefits real adverse-weather domains, as demonstrated by clear gains on *RobotCar-night* and a modest improvement on *DS-rain*, while clean benchmarks remain largely unaffected. A minor trade-off is observed on *NuScenes-night*, indicating a residual domain mismatch between the real-night data characteristics and this benchmark.

Weather conditioning. Enabling AdaLN-based weather injection further enhances nighttime robustness, particularly on *RobotCar-night* and to a lesser extent on *NuScenes-night*, while maintaining strong performance on clean benchmarks (*KITTI*: 0.080/0.937). For *DrivingStereo* corruptions, the effect is mixed, as performance is slightly reduced on *DS-fog*, suggesting that not all synthetic corruptions benefit equally from explicit weather conditioning.

Loss and restoration ablations. Table 6 isolates the objective terms requested by the reviewers. Removing L_{dist} causes the largest degradation, increasing Global AbsRel from 0.117 to 0.352, which confirms that teacher distillation anchors the adapted model to geometry. L_{pair} provides a small but consistent gain, while L_{aug} shows no measurable independent gain under this protocol and is best interpreted as a real-weather consistency regularizer. The RGB restoration baseline disables Style Filter/AdaLN conditioning and adds a lightweight degraded-to-clean RGB head trained with masked L1 loss on synthetic pairs, weighted by $\lambda_{\text{rgb}} = 0.1$. It is competitive but does not improve the Global score over DA-W, suggesting that weather-conditioned modulation is not merely image restoration. Full fine-tuning improves the real-night average but is not better globally, supporting decoder-only adaptation as the better robustness-generalization trade-off.

5.2 Analysis and Discussion

Feature modulation. To validate the effectiveness of our weather conditioning, we analyze the decoder features before and after AdaLN modulation. Figure 7 visualizes the intermediate features for nighttime and foggy scenarios with different maps. Specifically, the spatial distribution of differences correlates with degradation characteristics: in nighttime scenes, darker regions exhibit markedly

Table 5: Ablation across training and weather-conditioning choices. Metrics are AbsRel $\downarrow$ / δ_1 $\uparrow$. Avg. rank $\downarrow$ is computed over all 14 reported metrics with ties averaged. Bold and underlined denote the best and second-best values for each metric, respectively.

Config ($\checkmark$ = enabled)				Zero-shot (AbsRel $\downarrow$ / δ_1 $\uparrow$)							Overall
Encoder FT	Decoder FT	Real Deg.	Weather Inject.	NS-N	RC-N	DS-Rain	DS-Fog	NYU-D	KITTI	KITTI-Snow	Avg. rank $\downarrow$
$\checkmark$	$\checkmark$	—	—	<u>0.194</u> /0.727	0.245/0.493	<u>0.122</u> / 0.842	0.098 /0.898	0.056/0.968	0.084/0.932	<u>0.109</u> / 0.888	2.89
—	$\checkmark$	—	—	0.190 / 0.741	0.248/0.495	0.123/0.840	0.102/ 0.902	0.050 / 0.974	<u>0.081</u> / <u>0.936</u>	0.113/0.877	2.68
—	$\checkmark$	$\checkmark$	—	0.195/0.731	<u>0.243</u> / <u>0.507</u>	0.121 / 0.842	<u>0.100</u> / <u>0.901</u>	0.050 / 0.974	<u>0.081</u> / <u>0.936</u>	0.112/0.878	<u>2.29</u>
—	$\checkmark$	$\checkmark$	$\checkmark$	<u>0.194</u> / <u>0.737</u>	0.239 / 0.513	0.123/ 0.842	0.101/0.896	<u>0.051</u> / <u>0.973</u>	0.080 / 0.937	0.107 / <u>0.884</u>	2.14




Fig. 6: Qualitative comparison on mixed degradations. Each row shows real-world images of snow at night.

stronger modulation than illuminated areas, whereas in foggy scenes, texture-rich regions on the right side show more pronounced differences than the homogeneous regions on the left. This spatially varying behavior indicates that our weather-aware conditioning adaptively focuses on regions most affected by degradation, without requiring explicit spatial supervision. The feature-level differences propagate to depth predictions, confirming that our learned weather representations tangibly impact the final outputs.

Limitations. Our controlled adaptation experiments focus on DA-v2-S. Extending the same weather-conditioned interface to larger Depth Anything variants and other depth foundation models requires matched adaptation heads and optimization settings. Real nighttime benchmarks such as RobotCar also include camera response, exposure dynamics, motion blur, and scene-distribution shifts beyond weather appearance alone, which remain challenging for the current weather-conditioned modulation.

6 Conclusion

This paper introduces Weather-Conditioned Depth Anything (**DepthAnything-W**), a model designed for robust monocular depth estimation across diverse adverse weather conditions. The model’s effectiveness stems from its key components, notably the *Style Filter*, which learns domain-aligned weather embeddings, and the *AdaLN-Zero*-based decoder, which injects these embeddings into a frozen Depth Anything backbone. To evaluate the proposed architecture, a

Table 6: Controlled loss ablations. Results are AbsRel↓. All rows use the same 20% clean/synthetic subset, full real-weather data, seed, and evaluation protocol unless noted. Real night averages RC-N and NS-N; DrivingStereo averages DS-Rain and DS-Fog; KITTI-C averages Fog, Snow, and Dark. Global is the pooled per-image mean over evaluated weather sets.

Variant	Real night	DrivingStereo	KITTI-C	Global
Full objective	0.217	0.108	0.110	0.117
w/o L_{dist}	0.535	0.409	0.337	0.352
w/o L_{pair}	0.222	0.109	0.114	0.120
w/o L_{aug}	0.215	0.109	0.109	0.116
L_{dist} only	0.223	0.110	0.111	0.118
RGB restore ($\lambda_{\text{rgb}} = 0.1$)	0.214	0.106	0.111	0.118
Full FT + Real Deg. + Weather Inject.	0.209	0.108	0.113	0.119




Fig. 7: Visualization of decoder feature modulation by weather-aware conditioning. From left to right: input image, decoder features before modulation, decoder features after modulation, and the absolute difference map.

mixed real-and-synthetic weather dataset is curated, and training incorporates distillation, pairwise alignment, and augmentation consistency to address the synthetic-to-real domain gap while maintaining generalization to clean domains. Extensive experiments on both real and synthetic weather benchmarks demonstrate that our method consistently outperforms strong depth foundation baselines and restoration-based pipelines, without sacrificing performance on standard clean datasets. Its performance sets a new standard for weather robustness in zero-shot depth estimation, offering a promising direction for future research on condition-aware depth foundation models.

Acknowledgements

This work was supported in part by the GPU hardware provided to the Texas A&M University through the NVIDIA Academia Grant Program.

References

1. Bhat, S.F., Birkel, R., Wofk, D., Wonka, P., Müller, M.: Zoedepth: Zero-shot transfer by combining relative and metric depth. arXiv preprint arXiv:2302.12288 (2023)
2. Bochkovskii, A., Delaunoy, A., Germain, H., Santos, M., Zhou, Y., Richter, S.R., Koltun, V.: Depth pro: Sharp monocular metric depth in less than a second. arXiv preprint arXiv:2410.02073 (2024)
3. Brödermann, T., Bruggemann, D., Sakaridis, C., Ta, K., Liagouris, O., Corkill, J., Van Gool, L.: Muses: The multi-sensor semantic perception dataset for driving under uncertainty. In: European Conference on Computer Vision. pp. 21–38. Springer (2024)
4. Burnett, K., Yoon, D.J., Wu, Y., Li, A.Z., Zhang, H., Lu, S., Qian, J., Tseng, W.K., Lambert, A., Leung, K.Y., Schoellig, A.P., Barfoot, T.D.: Boreas: A multi-season autonomous driving dataset. The International Journal of Robotics Research **42**(1–2), 33–42 (2023). <https://doi.org/10.1177/02783649231160195>
5. Butler, D.J., Wulff, J., Stanley, G.B., Black, M.J.: A naturalistic open source movie for optical flow evaluation. In: European conference on computer vision. pp. 611–625. Springer (2012)
6. 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: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 11621–11631 (2020)
7. Chen, I., Chen, W.T., Liu, Y.W., Chiang, Y.C., Kuo, S.Y., Yang, M.H., et al.: Unirestore: Unified perceptual and task-oriented image restoration model using diffusion prior. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 17969–17979 (2025)
8. Chen, W.T., Vong, Y.J., Kuo, S.Y., Ma, S., Wang, J.: Robustsam: Segment anything robustly on degraded images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4081–4091 (2024)
9. Danier, D., Aygün, M., Li, C., Bilen, H., Mac Aodha, O.: Depthcues: Evaluating monocular depth perception in large vision models. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 20049–20059 (2025)
10. Fahes, M., Vu, T.H., Bursuc, A., Pérez, P., de Charette, R.: P {O} da: Prompt-driven zero-shot domain adaptation. arXiv preprint arXiv:2212.03241 (2022)
11. Fahes, M., Vu, T.H., Bursuc, A., Pérez, P., de Charette, R.: Domain adaptation with a single vision-language embedding. arXiv preprint arXiv:2410.21361 (2024)
12. Fang, C., He, C., Xiao, F., Zhang, Y., Tang, L., Zhang, Y., Li, K., Li, X.: Real-world image dehazing with coherence-based pseudo labeling and cooperative unfolding network. Advances in Neural Information Processing Systems **37**, 97859–97883 (2024)
13. Feijoo, D., Benito, J.C., Garcia, A., Conde, M.V.: Darkir: Robust low-light image restoration. In: Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR). pp. 10879–10889 (June 2025)
14. Gaidon, A., Wang, Q., Cabon, Y., Vig, E.: Virtual worlds as proxy for multi-object tracking analysis. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4340–4349 (2016)
15. Ganj, A., Zhao, Y., Su, H., Guo, T.: Mobile ar depth estimation: Challenges & prospects. In: Proceedings of the 25th International Workshop on Mobile Computing Systems and Applications. pp. 21–26 (2024)

16. Gasperini, S., Morbitzer, N., Jung, H., Navab, N., Tombari, F.: Robust monocular depth estimation under challenging conditions. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 8177–8186 (2023)
17. Geiger, A., Lenz, P., Stiller, C., Urtasun, R.: Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research* **32**(11), 1231–1237 (2013)
18. Godard, C., Mac Aodha, O., Firman, M., Brostow, G.J.: Digging into self-supervised monocular depth estimation. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 3828–3838 (2019)
19. Gui, M., Schusterbauer, J., Prestel, U., Ma, P., Kotovenko, D., Grebenkova, O., Baumann, S.A., Hu, V.T., Ommer, B.: Depthfm: Fast monocular depth estimation with flow matching (2024)
20. Hu, M., Yin, W., Zhang, C., Cai, Z., Long, X., Chen, H., Wang, K., Yu, G., Shen, C., Shen, S.: Metric3d v2: A versatile monocular geometric foundation model for zero-shot metric depth and surface normal estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2024)
21. Jiang, K., Cao, J., Yu, Z., Jiang, J., Zhou, J.: Always clear depth: Robust monocular depth estimation under adverse weather. In: Proc. Int. Joint Conf. Artif. Intell. pp. 1251–1259 (2025)
22. Job, M., Stastny, T., Kazik, T., Siegwart, R., Pantic, M.: Radar meets vision: robustifying monocular metric depth prediction for mobile robotics. *arXiv preprint arXiv:2410.00736* (2024)
23. Karvat, M., Givigi, S.: Adver-city: Open-source multi-modal dataset for collaborative perception under adverse weather conditions. *arXiv preprint arXiv:2410.06380* (2024)
24. Ke, B., Obukhov, A., Huang, S., Metzger, N., Daudt, R.C., Schindler, K.: Repurposing diffusion-based image generators for monocular depth estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9492–9502 (2024)
25. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment anything. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 4015–4026 (2023)
26. Kong, L., Xie, S., Hu, H., Ng, L.X., Cottureau, B., Ooi, W.T.: Robodepth: Robust out-of-distribution depth estimation under corruptions. *Advances in Neural Information Processing Systems* **36**, 21298–21342 (2023)
27. Li, J., Li, B., Tu, Z., Liu, X., Guo, Q., Juefei-Xu, F., Xu, R., Yu, H.: Light the night: A multi-condition diffusion framework for unpaired low-light enhancement in autonomous driving. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15205–15215 (2024)
28. Li, S., Araujo, I.B., Ren, W., Wang, Z., Tokuda, E.K., Junior, R.H., Cesar-Junior, R., Zhang, J., Guo, X., Cao, X.: Single image deraining: A comprehensive benchmark analysis. In: 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 3833–3842 (2019). <https://doi.org/10.1109/CVPR.2019.00396>
29. Li, Z., Snavely, N.: Megadepth: Learning single-view depth prediction from internet photos. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2041–2050 (2018)
30. Lin, H., Chen, S., Liew, J.H., Chen, D.Y., Li, Z., Shi, G., Feng, J., Kang, B.: Depth anything 3: Recovering the visual space from any views. *arXiv preprint arXiv:2511.10647* (2025)

31. Lin, H., Chen, S., Liew, J., Chen, D.Y., Li, Z., Shi, G., Feng, J., Kang, B.: Depth anything 3: Recovering the visual space from any views. arXiv preprint arXiv:2511.10647 (2025)
32. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: European conference on computer vision. pp. 740–755. Springer (2014)
33. Liu, Y., Kwon, H.: Efficient depth estimation for unstable stereo camera systems on ar glasses. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 6252–6261 (2025)
34. Liu, Y.F., Jaw, D.W., Huang, S.C., Hwang, J.N.: Desnownet: Context-aware deep network for snow removal. *IEEE Transactions on Image Processing* **27**(6), 3064–3073 (2018)
35. van der Maaten, L., Hinton, G.: Visualizing data using t-sne. *Journal of Machine Learning Research* **9**(86), 2579–2605 (2008), <http://jmlr.org/papers/v9/vandermaten08a.html>
36. Maddern, W., Pascoe, G., Linegar, C., Newman, P.: 1 year, 1000 km: The oxford robotcar dataset. *The International Journal of Robotics Research* **36**(1), 3–15 (2017)
37. Medeiros, L.: Lang segment anything. <https://github.com/luca-medeiros/lang-segment-anything> (2023), gitHub repository
38. Piccinelli, L., Yang, Y.H., Sakaridis, C., Segu, M., Li, S., Van Gool, L., Yu, F.: Unidepth: Universal monocular metric depth estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10106–10116 (2024)
39. Ranftl, R., Bochkovskiy, A., Koltun, V.: Vision transformers for dense prediction. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 12179–12188 (2021)
40. Ranftl, R., Lasinger, K., Hafner, D., Schindler, K., Koltun, V.: Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE transactions on pattern analysis and machine intelligence* **44**(3), 1623–1637 (2020)
41. Sakaridis, C., Wang, H., Li, K., Jadon, A., Abbeloos, W., Reino, D.O., Van Gool, L., Dai, D., et al.: Acdc: The adverse conditions dataset with correspondences for robust semantic driving scene perception. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2025)
42. Saunders, K., Vogiatzis, G., Manso, L.J.: Self-supervised monocular depth estimation: Let’s talk about the weather. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 8907–8917 (2023)
43. Schops, T., Schonberger, J.L., Galliani, S., Sattler, T., Schindler, K., Pollefeys, M., Geiger, A.: A multi-view stereo benchmark with high-resolution images and multi-camera videos. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3260–3269 (2017)
44. Shyam, P., Yoo, H.: Pair: Perception aided image restoration for natural driving conditions. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). pp. 7459–7470 (January 2024)
45. Silberman, N., Hoiem, D., Kohli, P., Fergus, R.: Indoor segmentation and support inference from rgb-d images. In: European Conference on Computer Vision. pp. 746–760 (2012)
46. Simon, N., Majumdar, A.: Mononav: Mav navigation via monocular depth estimation and reconstruction. In: International Symposium on Experimental Robotics. pp. 415–426. Springer (2023)

47. Song, Z., Zhu, R., Wang, C., Deng, J., He, J., Zhang, T.: Ec-depth: Exploring the consistency of self-supervised monocular depth estimation in challenging scenes. arXiv preprint arXiv:2310.08044 (2023)
48. Sun, B., Jin, M., Yin, B., Hou, Q.: Depth anything at any condition. arXiv preprint arXiv:2507.01634 (2025)
49. Tan, X., Xu, K., Cao, Y., Zhang, Y., Ma, L., Lau, R.W.H.: Night-time scene parsing with a large real dataset. *IEEE Transactions on Image Processing* **30**, 9085–9098 (2021). <https://doi.org/10.1109/TIP.2021.3122004>
50. Tosi, F., Ramirez, P.Z., Poggi, M.: Diffusion models for monocular depth estimation: Overcoming challenging conditions. In: *European Conference on Computer Vision*. pp. 236–257. Springer (2024)
51. Tumanyan, N., Bar-Tal, O., Bagon, S., Dekel, T.: Splicing vit features for semantic appearance transfer. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 10748–10757 (2022)
52. Vasiljevic, I., Kolkin, N., Zhang, S., Luo, R., Wang, H., Dai, F.Z., Daniele, A.F., Mostajabi, M., Basart, S., Walter, M.R., et al.: Diode: A dense indoor and outdoor depth dataset. arXiv preprint arXiv:1908.00463 (2019)
53. Wang, J., Lin, C., Nie, L., Huang, S., Zhao, Y., Pan, X., Ai, R.: Weatherdepth: Curriculum contrastive learning for self-supervised depth estimation under adverse weather conditions. In: *2024 IEEE International Conference on Robotics and Automation (ICRA)*. pp. 4976–4982. IEEE (2024)
54. Wang, R., Xu, S., Dai, C., Xiang, J., Deng, Y., Tong, X., Yang, J.: Moge: Unlocking accurate monocular geometry estimation for open-domain images with optimal training supervision. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 5261–5271 (2025)
55. Watson, J., Mac Aodha, O., Prisacariu, V., Brostow, G., Firman, M.: The temporal opportunist: Self-supervised multi-frame monocular depth. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 1164–1174 (2021)
56. Xian, K., Zhang, J., Wang, O., Mai, L., Lin, Z., Cao, Z.: Structure-guided ranking loss for single image depth prediction. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 611–620 (2020)
57. Xue, F., Zhuo, G., Huang, Z., Fu, W., Wu, Z., Ang, M.H.: Toward hierarchical self-supervised monocular absolute depth estimation for autonomous driving applications. In: *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. pp. 2330–2337. IEEE (2020)
58. Yan, W., Li, M., Li, H., Shao, S., Tan, R.T.: Synthetic-to-real self-supervised robust depth estimation via learning with motion and structure priors. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 21880–21890 (2025)
59. Yang, L., Kang, B., Huang, Z., Xu, X., Feng, J., Zhao, H.: Depth anything: Unleashing the power of large-scale unlabeled data. In: *CVPR* (2024)
60. Yang, L., Kang, B., Huang, Z., Zhao, Z., Xu, X., Feng, J., Zhao, H.: Depth anything v2. *Advances in Neural Information Processing Systems* **37**, 21875–21911 (2024)
61. Yin, Z., Shi, J.: Geonet: Unsupervised learning of dense depth, optical flow and camera pose. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 1983–1992 (2018)
62. Zhang, S., Zhang, J., Tao, D.: Towards scale-aware, robust, and generalizable unsupervised monocular depth estimation by integrating imu motion dynamics. In: *European Conference on Computer Vision*. pp. 143–160. Springer (2022)

- 63. Zheng, Y., Zhong, C., Li, P., Gao, H.a., Zheng, Y., Jin, B., Wang, L., Zhao, H., Zhou, G., Zhang, Q., et al.: Steps: Joint self-supervised nighttime image enhancement and depth estimation. arXiv preprint arXiv:2302.01334 (2023)
- 64. Zhou, T., Brown, M., Snavely, N., Lowe, D.G.: Unsupervised learning of depth and ego-motion from video. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1851–1858 (2017)
- 65. Zhou, Z., Dong, Q.: Learning occlusion-aware coarse-to-fine depth map for self-supervised monocular depth estimation. In: Proceedings of the 30th ACM International Conference on Multimedia. pp. 6386–6395 (2022)
- 66. Zhu, R., Tu, Z., Liu, J., Bovik, A.C., Fan, Y.: Mwformer: Multi-weather image restoration using degradation-aware transformers. IEEE Transactions on Image Processing pp. 1–1 (2024). <https://doi.org/10.1109/TIP.2024.3501855>