

Enhancing Generalization of Depth Estimation Foundation Model via Weakly-Supervised Adaptation with Regularization

Yan Huang¹, Yongyi Su¹, Xin Lin², Le Zhang³, Xun Xu^{4*}

¹South China University of Technology

²Guangzhou University

³University of Electronic Science and Technology of China

⁴Institute for Infocomm Research (I²R), A*STAR

{eeyhuang, eesuyongyi}@mail.scut.edu.cn, linxin94@gzhu.edu.cn,
zhangleuestc@gmail.com, alex.xun.xu@gmail.com

Abstract

The emergence of foundation models has substantially advanced zero-shot generalization in monocular depth estimation (MDE), as exemplified by the Depth Anything series. However, given access to some data from downstream tasks, a natural question arises: can the performance of these models be further improved? To this end, we propose WeSTAR, a parameter-efficient framework that performs **Weakly supervised Self-Training Adaptation with Regularization**, designed to enhance the robustness of MDE foundation models in unseen and diverse domains. We first adopt a dense self-training objective as the primary source of structural self-supervision. To further improve robustness, we introduce semantically-aware hierarchical normalization, which exploits instance-level segmentation maps to perform more stable and multi-scale structural normalization. Beyond dense supervision, we introduce a cost-efficient weak supervision in the form of pairwise ordinal depth annotations to further guide the adaptation process, which enforces informative ordinal constraints to mitigate local topological errors. Finally, a weight regularization loss is employed to anchor the LoRA updates, ensuring training stability and preserving the model’s generalizable knowledge. Extensive experiments on both realistic and corrupted out-of-distribution datasets under diverse and challenging scenarios demonstrate that WeSTAR consistently improves generalization and achieves state-of-the-art performance across a wide range of benchmarks.

Introduction

Monocular depth estimation is crucial for applications such as stereo conversion (Xie, Girshick, and Farhadi 2016), augmented reality (Ganj et al. 2024), and 3D scene reconstruction (Yin et al. 2022). Recent progress has led to foundation models trained on large-scale combinations of labeled and unlabeled data, with models like those by (Yang et al. 2024a,b) showing strong generalization to diverse, unseen domains.

Despite these advancements, it remains unclear how much further performance can improve when data from downstream tasks becomes available. Figure 1 shows that zero-shot predictions remain imperfect especially under distribu-

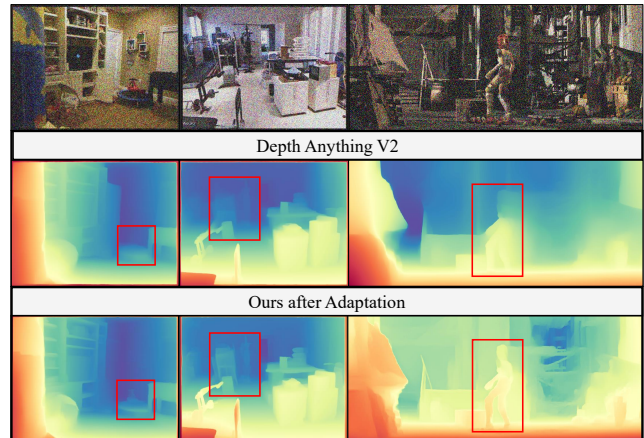


Figure 1: Illustration of MDE on unseen test samples with zero-shot results of model and model after adaptation.

tion shifts. Performance can be further improved via adaptation. In image segmentation, foundation models have been effectively adapted using unlabeled or weakly labeled task-specific data (Zhang et al. 2024; Zhao et al. 2025), often leveraging self-training (Sohn et al. 2020), self-supervised learning (Chen et al. 2020; Liu et al. 2021), and distribution alignment (Su, Xu, and Jia 2022).

However, extending such techniques to depth estimation, essentially a regression task, poses unique challenges. Unlike classification, regression complicates self-training due to the difficulty of generating reliable pseudo-labels. Inaccurate labels can mislead the model during adaptation, especially in unseen domains. As shown in Figure 4, naïve self-training can even degrade performance over time. Moreover, the dense self-training (ST) objective often provides only marginal gains on benchmarks, mainly due to the strong geometric understanding of the pre-trained model, which limits further improvement from ST alone. Finally, an overly aggressive adaptation process of self-training can be detrimental. It risks over-specializing the model to the specific domains, leading to overfitting and catastrophic forgetting, where the powerful, generalizable knowledge from the original pre-training is compromised.

Our framework addresses these limitations through a synergistic design. To mitigate confirmation bias and provide

*Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

targeted, non-redundant supervision, we complement dense self-training with a weakly supervised objective (Chen et al. 2016), incorporating a minimal set of relative depth annotations that correct fine-grained topological errors and introduce label signals independent of the model’s pseudo-labels. In other words, the weakly supervised loss introduces sparse but strong ordinal constraints, while the dense self-training objective provides global structural alignment.

To enhance adaptation efficiency and prevent catastrophic forgetting, motivated by robustness concerns in pre-trained models (Zhao et al. 2023), we introduce additive low-rank adapters (Hu et al. 2022) to constrain the model near the original. We further introduce a regularization that penalizes LoRA parameter updates, anchoring the model to the strong pre-trained depth prior and limiting the influence of noisy pseudo-labels and sparse ordinal constraints, suppressing the geometric distortions induced by weak supervision and mitigating the confirmation bias of self-training.

We also observe that self-training for MDE must address scale and shift ambiguity between teacher pseudo-labels and student predictions. While hierarchical depth normalization (HDN) (Zhang et al. 2022a; He et al. 2025) was designed to mitigate scale variation and depth discontinuities, its reliance on content-agnostic grids overlooks semantic context, often resulting in inferior performance. To address this, we propose using an external segmentation foundation model (Ravi et al. 2025) to inject semantic information as masks, enabling semantic-aware HDN.

We refer to our adaptation framework as **WeSTAR**. WeSTAR enables fine-tuning of pre-trained models using minimal data from the target domain, resulting in improved performance on downstream tasks. We validate the method on both unseen image domains and corrupted inputs (Kong et al. 2023), simulating real-world out-of-distribution scenarios. This makes it especially suitable for depth estimation tasks in specific contexts where limited data with weak annotations can be collected for adaptation.

Our contributions are summarized as follows:

- We address the challenge of generalizing depth estimation foundation models to unseen data distributions through a low-rank regularized self-training framework with semantic-aware hierarchical normalization.
- We incorporate low-cost weak labels, specifically pairwise ordinal depth annotations, to further enhance the robustness of the adaptation process.
- Our proposed method WeSTAR demonstrate robust generalization across diverse and challenging scenarios, including novel tasks and corrupted datasets.

Related Work

Monocular Depth Estimation: Monocular depth estimation (MDE) has advanced rapidly with deep learning, supported by architectural innovations and large-scale datasets (He et al. 2016; Dosovitskiy et al. 2020; Geiger et al. 2013). Early supervised (Eigen, Puhrsch, and Fergus 2014; Bhat, Alhashim, and Wonka 2021) and self-supervised methods (Godard, Mac Aodha, and Brostow 2017; Godard et al. 2019) achieved strong performance

but struggled to generalize due to dataset biases. Transfer learning with pretrained backbones (e.g., ImageNet (Deng et al. 2009)) improved generalization but still required task-specific tuning. Recent foundation models like MiDaS (Ranfil et al. 2020) and Depth Anything (Yang et al. 2024a,b) enhance robustness by training on large-scale data.

Source-Free Domain Adaptation: Source-Free Domain Adaptation (SFDA) (Liang, Hu, and Feng 2020) adapts models without access to source data. Our work aligns with SFDA, as zero-shot generalization often relies solely on pre-trained models. In classification, early SFDA methods re-construct source-like features (Ding et al. 2022) or augment target samples (Jing et al. 2022; Hwang et al. 2024). More recent approaches use contrastive learning (Zhang et al. 2022b; Chen et al. 2022), pseudo-labeling (Liang, Hu, and Feng 2020), and iterative self-training (Karim et al. 2023), along with diversity constraints (Mitsuzumi, Kimura, and Kashima 2024) and feature alignment (Lee and Lee 2023; Su, Xu, and Jia 2022). For regression, TASFAR (He et al. 2024) calibrates models using pseudo-label distributions. In depth estimation, some methods discretize depth into ordinal labels (Yi and Kim 2023), while others combine scale alignment and self-supervision (Li et al. 2023) to adapt without target supervision.

Robustness in Depth Estimation: MDE models often fail under severe domain shifts like weather or noise, as shown in RoboDepth (Kong et al. 2023). Our analysis reveals that even state-of-the-art models (e.g., Depth Anything V2) degrade significantly under compounded corruptions. While this vulnerability is known, systematic methods to address it remain limited. We present a principled solution to enhance robustness under such shifts.

Weakly Supervised Depth Estimation: Due to the difficulty of obtaining dense depth labels, alternatives like self-supervision from videos (Godard, Mac Aodha, and Brostow 2017; Zhao et al. 2022) and weak supervision via ordinal or semantic cues (Chen et al. 2016) have gained popularity. Inspired by this, we introduce a structured pairwise ranking loss that enforces transitive ordinal relations, promoting a globally coherent depth manifold.

Methodology

Problem Formulation

We begin by formally defining the task of adapting a pre-trained monocular depth estimation foundation model to a target domain. Let $\mathcal{D}_s = \{x_i, d_i\}$ denote a labeled source dataset, where x_i represents input RGB images and d_i their corresponding ground-truth depth maps. A depth estimation model $f(x; \Theta)$ is trained on this data using supervised or semi-supervised techniques (Yang et al. 2024a). Although the pre-trained model often generalizes well to many unseen domains, we hope to further enhance its effectiveness on the unseen domains. Our goal is to adapt $f(x; \Theta)$ to a new target domain dataset $\mathcal{D}_{te} = \{x_j\}$ by utilizing a small set of either unlabeled or weakly labeled training samples $\mathcal{D}_{tr} = \mathcal{D}_{tr}^w \cup \mathcal{D}_{tr}^u$, where $\mathcal{D}_{tr}^w = \{x_j, w_j\}$ is weakly labeled and $\mathcal{D}_{tr}^u = \{x_j\}$ is unlabeled. An overview of the framework is presented in Figure 2.

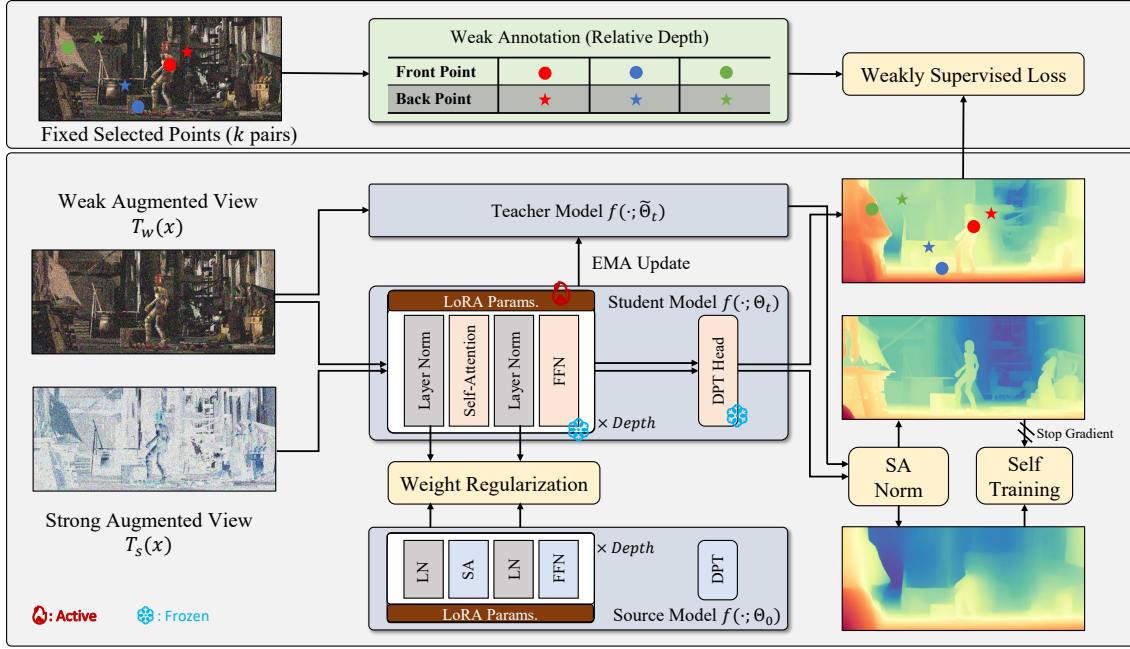


Figure 2: Illustration of the overall framework. Two augmentations are applied and the teacher model $\tilde{\Theta}$ generates pseudo labels. Self-training is regularized by model weights consistency and optionally weak labels to adapt pre-trained foundation model.

Domain Adaptation via Self-Training

We first consider a fully unsupervised adaptation scenario, where $\mathcal{D}_{tr}^w = \emptyset$. Self-training (ST) is a well-established technique for unsupervised adaptation of pre-trained models, particularly successful in classification tasks (Su, Xu, and Jia 2022; Liang, Hu, and Feng 2020). ST typically adopts a teacher-student architecture (Tarvainen and Valpola 2017), where the teacher model generates pseudo labels to guide the student model’s learning.

We employ an Exponential Moving Average (EMA) update rule for maintaining the teacher weights: $\tilde{\Theta}_{t+1} = \alpha\tilde{\Theta}_t + (1-\alpha)\Theta_{t+1}$, where α is a smoothing factor. Each unlabeled sample is augmented with weak (T_w) and strong (T_s) perturbations before being passed to the teacher and student models, respectively. The self-training loss is defined as:

$$\mathcal{L}_{st} = \frac{1}{|\mathcal{D}_{tr}^u|} \frac{1}{HW} \sum_{x_i \in \mathcal{D}_{tr}^u} \sum_{p \in x_i} \mathcal{L}_{st}(p), \quad (1)$$

where $\mathcal{L}_{st}(p)$ denotes the point-wise self-training loss, H and W are the height and width of the image.

Semantic-Aware Hierarchical Normalization: Depth normalization is a crucial component in our self-training framework, designed to resolve the inherent scale and shift ambiguity in both the teacher’s pseudo-labels d^* and the student’s predictions d . For each pixel p with depth d_p and associated context \mathcal{C}_p , normalization maps the depth into a canonical space using robust statistics:

$$\Phi(d_p, \mathcal{C}_p) = \frac{d_p - t(\mathcal{C}_p)}{s(\mathcal{C}_p) + \epsilon} \quad (2)$$

where t and s denote the median and the median absolute deviation (MAD) of depths within \mathcal{C}_p .

Traditional normalization approaches differ primarily in how they define \mathcal{C}_p . Global Normalization (Ranftl et al. 2020; Yang et al. 2024a) uses a single context $\mathcal{C}_p = \mathcal{C}_{global}$ containing all valid pixels. In contrast, Hierarchical Depth Normalization (HDN) (Zhang et al. 2022a; He et al. 2025) constructs a multi-scale hierarchy $\mathcal{C}_p = \{\mathcal{C}_{global}, \mathcal{C}_{l-1}, \mathcal{C}_{l-2}, \dots\}$ based on fixed grids or depth bins. However, these content-agnostic groupings often fragment semantic objects, leading to unstable statistics.

We propose Semantically-Aware Hierarchical Depth Normalization (SA-HDN) to address these limitations by defining the hierarchy using semantic instance masks. SA-HDN includes two levels:

- The **Global Context** \mathcal{C}_{global} provides a scene-level normalization using all valid pixels.
- The **Instance Contexts** $\{\mathcal{C}_{ins}^k\}$ capture fine-grained object-level structure, where each \mathcal{C}_{ins}^k contains pixels in a specific object instance M_k .

We employ SAM2 (Ravi et al. 2025) to automatically generate M_k on the fly, allowing flexible domain-adaptive context construction without retraining. For each pixel p assigned to instance k , the context is defined as $\mathcal{C}_p = \{\mathcal{C}_{global}, \mathcal{C}_{ins}^k\}$, enabling dual supervision at global and instance levels. The self-training loss at pixel p is then defined as the average MAE between the normalized pseudo-label and prediction over all associated contexts:

$$\mathcal{L}_{st}(p) = \frac{1}{|\mathcal{C}_p|} \sum_{c \in \mathcal{C}_p} |\mathbf{sg}(\Phi(d_p^*, c)) - \Phi(d_p, c)| \quad (3)$$

$$d_p^* = f(T_w(x_p); \tilde{\Theta}_t), \quad d_p = f(T_s(x_p); \Theta_t)$$

where $\text{sg}(\cdot)$ indicates stop-gradient to block teacher back-propagation. Φ applies the context-dependent normalization to the predicted relative depth.

Weakly Supervised Adaptation

While self-training strategy described above provides consistent performance gains even without labeled data, its effectiveness is hampered by confirmation bias from pseudo-labels, which can reinforce structural errors. Furthermore, it becomes redundant and offers only marginal gains when the baseline model is already proficient. To overcome these issues, we complement the dense self-training with a weakly supervised objective (Chen et al. 2016), incorporating a minimal and inexpensive set of relative depth annotations such that $\mathcal{D}_{tr}^w \neq \emptyset$, to guide a more efficient adaptation.

Each weak label $w_j = \{p_{jn}^+, p_{jn}^-, l_{jn}\}$ represents the ordinal relationship between two pixels. The label $l_{jn} \in \{-1, 0, 1\}$ indicates whether p_{jn}^+ is farther, equal in depth, or closer than p_{jn}^- , respectively. We use a margin ranking loss to enforce these relations:

$$\mathcal{L}_{weak} = \sum_{(p_{jn}^+, p_{jn}^-, l_{jn}) \in w_j} \ell(\hat{d}_{jn}^+, \hat{d}_{jn}^-, l_{jn}) \quad (4)$$

where the loss for a single pair is defined as

$$\begin{aligned} \ell(\hat{d}_{jn}^+, \hat{d}_{jn}^-, l_{jn}) &= \mathbf{1}(l_{jn} \neq 0) \cdot \max(0, -l_{jn} \Delta d_{jn} + \delta) \\ &\quad + \mathbf{1}(l_{jn} = 0) \cdot |\Delta d_{jn}| \\ \text{s.t. } \Delta d_{jn} &= \hat{d}_{jn}^+ - \hat{d}_{jn}^-, \quad d_{jn}^\pm = f(T_w(x_j); \Theta_t)_{p_{jn}^\pm} \end{aligned} \quad (5)$$

where δ is a slack variable controlling the margin. This loss encourages depth estimates that are consistent with the weak supervision, even with sparse annotations. Our experiments demonstrate that a small amount of such supervision significantly enhances training stability and generalization.

Robust Adaptation with LoRA

Full fine-tuning of large-scale MDE foundation models is often impractical due to excessive computational demands and limited batch sizes, which lead to inefficiencies and potential overfitting on small target datasets. Additionally, full fine-tuning risks catastrophic forgetting of pre-trained knowledge, leading to potentially poorer performance on the target domain. To address this challenge, we adopt a low-rank adaptation (LoRA) strategy (Hu et al. 2022), injecting trainable low-rank matrices into each attention layer $\Theta_a \in \mathbb{R}^{d_1 \times d_2}$ of the encoder. The adapted weights are formulated as:

$$\Theta_a + UV \quad (6)$$

where $U \in \mathbb{R}^{d_1 \times r}$ and $V \in \mathbb{R}^{r \times d_2}$ with $r \ll \min(d_1, d_2)$. Only U and V are updated during adaptation, substantially reducing memory and computational overhead. This lightweight tuning mitigates overfitting and catastrophic forgetting while retaining the model’s generalization capabilities.

Weight Regularization: Despite LoRA’s benefits, adaptation under severe domain shifts or noisy pseudo-labels remains susceptible to confirmation bias, where erroneous

pseudo-labels may reinforce model errors. To stabilize adaptation, we introduce a weight regularization term that constrains the student weights to remain close to the pre-trained initialization.

Let U_0, V_0 be the initial LoRA weights (typically $U_0 V_0 = 0$), and U_t, V_t be the student’s current weights. We define a regularization loss that penalizes large deviations:

$$\mathcal{L}_{reg} = \sum_{U_{tk} \in U_t, V_{tk} \in V_t} \left\| \frac{\alpha}{r} U_{tk} V_{tk} \right\|_2^2 \quad (7)$$

where α controls the regularization strength. This encourages parameter updates only when strongly justified by new evidence from the target domain, helping to mitigate overfitting and stabilize adaptation.

Overall Loss

Starting from a public depth foundation model, we perform adaptation using a self-training framework enhanced with weight regularization and weak supervision. The total loss is a weighted sum of three components.

$$\mathcal{L} = \lambda_{st} \mathcal{L}_{st} + \lambda_w \mathcal{L}_{weak} + \lambda_r \mathcal{L}_{reg} \quad (8)$$

Experiment

Datasets and Metrics

We evaluate our method WeSTAR on several widely-used datasets across diverse scenarios. These datasets exhibit significant domain shifts from the training data, providing a robust testbed for assessing the generalization ability of our method under diverse distribution shifts. The evaluated datasets include: **NYU-V2** (Silberman et al. 2012), **KITTI** (Geiger et al. 2013), **Sintel** (Butler et al. 2012), **DIODE** (Vasiljevic et al. 2019), and their corresponding corrupted benchmarks **NYU-C**, **KITTI-C**, **DIODE-C**, and **Sintel-C**, following (Kong et al. 2023). These benchmarks introduce 6 types of corruptions at the highest severity 5 to the original datasets. Moreover, to evaluate our method on more realistic datasets under distribution shift, we choose the night version of **NuScenes** (Caesar et al. 2020) and the **DrivingStereo** (Yang et al. 2019) dataset with different weathers (Sunny, Cloudy, Foggy and Rainy). More dataset details are provided in the Appendix.

For each dataset, we construct non-overlapping training and test splits. Adaptation is performed only on the training split, while evaluation is carried out on the unseen test set.

Evaluation Metrics: Following (Yang et al. 2024a,b), we adopt two standard metrics to evaluate depth estimation performance on all datasets: δ_1 (the percentage of pixels satisfying $\max(d_i^*/d_i, d_i/d_i^*) < 1.25$) and AbsRel (the mean absolute relative error $1/N \sum_{i=1}^N |d_i^* - d_i|/d_i$).

Experimental Settings

MDE Model: We focus on Relative Depth Estimation and adopt two Relative MDE foundation models, Depth Anything v2 (Yang et al. 2024b) and MiDas v3.1 (Birkel, Wofk, and Müller 2023) as our base model, known for their strong zero-shot generalization.

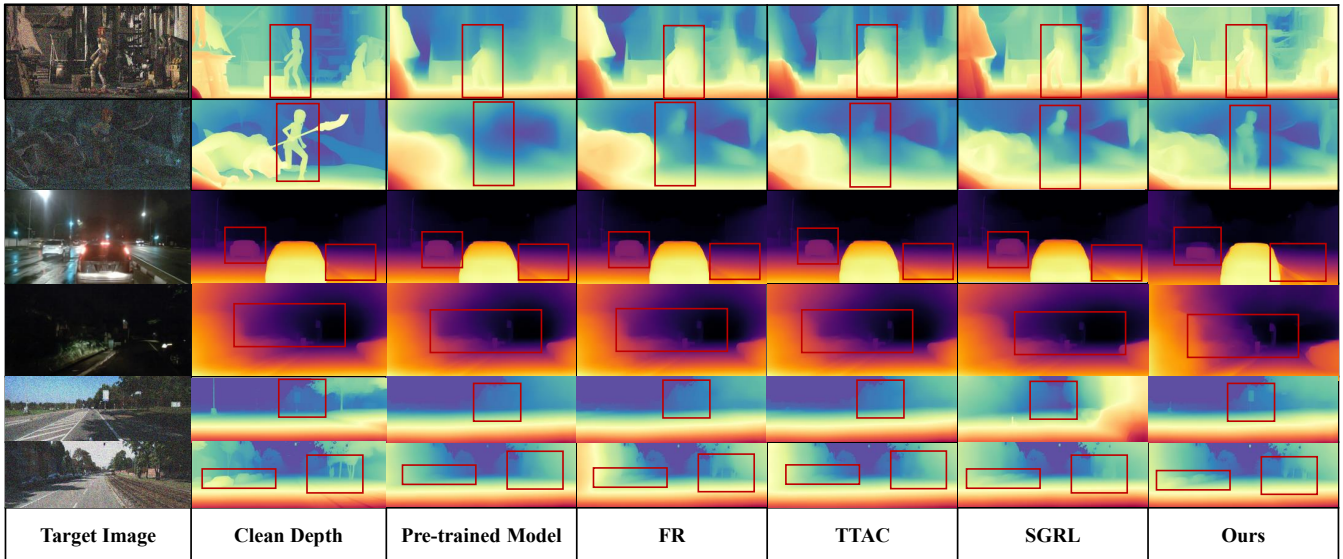


Figure 3: Qualitative results on some selected examples.

Hyper-parameters: We use AdamW optimizer with weight decay of 0.0001. For Depth Anything v2 backbone, loss weights $\{\lambda_{st}, \lambda_w, \lambda_r\}$ are set to $\{1.0, 0.001, 1.0\}$. Initial learning rate is set to 0.1 and scaled linearly with batch size (Goyal et al. 2017), i.e., $0.1 \cdot (\text{BS}/256)$. A cosine annealing schedule is applied to decay learning rate. We train for up to 100 epochs with early stopping after 30 epochs. The EMA decay factor is 0.996. We use a LoRA with the rank of 8 and alpha of 16. Hyper-Parameters for Midas v3.1 backbone follow a similar configuration, as detailed in the appendix. All experiments are conducted on single NVIDIA RTX 3090 GPU with a batch size of 4.

Pixel-Pairs Sampling: To construct the ranking loss in weakly-supervised component, we employ a structured pixel-pair sampling strategy instead of sampling pairs independently. For each image, we perform 5 sampling iterations. In the k -th iteration, we first randomly select an anchor pixel p_k , then randomly sample a farther point p_k^+ and a nearer point p_k^- . This results in two structured pairs, (p_k^+, p_k) and (p_k, p_k^-) for supervision. Unlike random sampling, this approach enforces the geometric property of transitivity, providing more coherent and consistent supervision.

Competing Methods: We compare our method WeSTAR against approaches from source-free domain adaptation and weakly supervised depth estimation. As a baseline, we include the zero-shot performance of the original pretrained model without fine-tuning (**Source**). Among adaptation methods, **TTT++**(Liu et al. 2021) applies contrastive learning and online feature alignment for improved test-time robustness. **FR**(Eastwood et al. 2022) uses softly-binned histograms for feature alignment, while **TTAC**(Su, Xu, and Jia 2022) aligns Gaussian feature distributions across domains. **SSA**(Adachi et al. 2025) proposes a regression-specific strategy by aligning features in a dominant subspace. Note that FR, TTAC, and SSA rely on source-domain statistics, which are unavailable in true zero-shot settings,

but we include them here for comparison. We also consider recent ViT-based self-supervised learning methods, which have proven effective for unsupervised feature representation. In particular, we use iBOT(Zhou et al. 2021) as a representative baseline due to its strong performance. We further fine-tune our encoder module using iBOT and denote this variant as **iBOT***. For weak supervision, **SGRL**(Xian et al. 2020) introduces a structure-guided relative ranking loss. To ensure fairness, all methods are evaluated using the same backbone.

Adaptation Results on Diverse Benchmarks

Adaptation to Corrupted Datasets: We present the comparison results on four corrupted benchmarks and report the average performance over all 6 corruptions in Table 2 (detailed results on each corruption are deferred to the Appendix). When the target images are affected by data corruptions, sensor failures, or adverse weather conditions, the performance of the base model degrades significantly. Specifically, the average performance of NYU-C across 6 corruption types drops from δ_1 : 97.7 to 87.4 (a 10.5% decrease) and AbsRel: 4.6 to 10.6 (a 56.6% increase). These results underscore the substantial challenge that domain shifts pose to the zero-shot generalization capability of monocular depth estimation (MDE) foundation models.

When SFDA methods such as TTAC, FR, and SSA are applied, they consistently improve performance compared to the baseline model. In contrast, applying TTT++ results in a noticeable degradation in performance in KITTI-C. This suggests that simple contrastive loss and feature alignment mechanisms are insufficient to effectively handle the distribution shifts caused by the corruptions.

Meanwhile, the self-supervised method iBOT* demonstrates competitive performance without access to the ground truth, highlighting the strength of robust feature learning under severe domain shifts. Similarly, the weakly

Method	NYU		KITTI		Sintel		DIODE		NuScenes		D-Sunny		D-Foggy		D-Cloudy		D-Rainy	
	δ_1	AbsRel	δ_1	AbsRel	δ_1	AbsRel	δ_1	AbsRel	δ_1	AbsRel	δ_1	AbsRel	δ_1	AbsRel	δ_1	AbsRel	δ_1	AbsRel
Source	97.7	4.6	93.4	8.4	74.8	20.3	95.0	7.0	74.4	18.5	78.4	15.7	91.3	9.0	81.5	14.5	84.8	12.1
TTT++	97.7	4.6	93.4	8.3	74.8	20.2	95.2	7.0	74.5	18.4	78.4	15.7	91.3	9.0	81.5	14.4	84.9	12.1
TTAC	97.7	4.6	93.4	8.5	75.0	20.5	95.0	7.0	74.4	18.5	79.5	15.4	91.1	9.1	82.0	14.4	84.5	12.2
FR	97.6	4.7	93.8	8.2	74.6	20.6	95.0	7.4	74.4	18.5	79.1	15.5	91.3	9.0	82.3	14.3	84.8	12.1
SSA	97.8	4.5	93.5	8.4	75.5	18.9	95.1	6.9	74.3	18.4	79.6	15.3	90.8	9.1	82.1	14.4	84.7	11.9
iBOT*	97.7	4.6	93.2	8.7	74.1	20.5	95.0	7.1	74.5	18.4	78.4	15.6	91.5	8.9	81.6	14.4	84.7	12.1
SGRL	97.6	4.8	94.1	7.7	76.9	21.9	95.1	6.9	75.8	17.6	83.4	13.5	92.0	8.5	84.4	12.9	85.3	12.0
WeSTAR	98.2	4.3	95.1	7.2	82.2	16.9	95.2	6.5	78.1	16.2	82.8	13.6	92.5	8.2	85.4	12.5	87.4	10.6

Table 1: Adaptation results on unseen realistic datasets. For δ_1 , higher is better, while for AbsRel, lower is better.

supervised method SGRL also achieves strong results, illustrating the potential of weak supervision for improving robustness in the presence of corruptions. Notably, WeSTAR consistently achieves state-of-the-art performance across nearly all corruption types, clearly demonstrating superior generalization under challenging distribution shifts.

We observe a similar results on other datasets, but TTT++ shows better performance than the unadapted baseline model. WeSTAR maintains a consistent advantage across the majority of corruption types, particularly under severe distortions. The performance gains are especially pronounced on the more challenging synthetic scenes in the Sintel-C dataset, further validating the robustness and generalization capability of our WeSTAR. Notably, with the MiDaS backbone, TTT++ outperforms other SFDA methods, while WeSTAR maintains a clear advantage, further validating its effectiveness in various backbones. We excluded NYU-C and KITTI-C since their clean versions are part of MiDaS’ training data, ensuring fair evaluation.

Method	NYU-C		KITTI-C		Sintel-C		DIODE-C	
	δ_1	AbsRel	δ_1	AbsRel	δ_1	AbsRel	δ_1	AbsRel
Source	87.4	10.6	83.2	13.2	60.3	30.6	88.0	11.1
TTT++	90.1	9.7	70.2	19.4	62.7	27.1	88.7	10.5
TTAC	91.2	8.9	84.7	12.7	62.3	29.7	89.2	10.6
FR	91.2	8.8	84.0	12.9	62.6	29.5	89.2	10.5
SSA	91.3	8.8	84.2	12.7	63.7	28.5	88.9	10.6
iBOT*	92.1	8.5	85.6	12.5	62.7	29.0	89.6	10.3
SGRL	92.4	8.4	87.4	11.3	66.5	29.9	90.0	10.0
WeSTAR	94.6	7.1	88.7	10.5	71.8	24.1	91.3	9.3

Table 2: Adaptation results averaged over all 6 types of corruptions on corrupted datasets.

Results on Realistic Datasets With Diverse Scenarios:

As shown in Table 1, the powerful Depth Anything V2 model already exhibits a high performance ceiling on standard benchmarks such as NYU, KITTI, and DIODE, leaving limited room for further improvements, where most adaptation methods struggle to yield further gains and may even degrade performance. In contrast, our proposed WeSTAR consistently maintains or improves upon this strong baseline. Notably, the benefits of WeSTAR are especially evident on the synthetic Sintel dataset and in realistic scenarios involving complex environmental changes, such as varying

weather and lighting conditions. This highlights the superior generalization ability of our WeSTAR, even when adapting to clean yet distributionally distinct target domains. We also reach similar conclusions on Midas v3.1 backbone in Table 3. To ensure fair evaluation, we excluded NYU and KITTI as they are included in the training data.

Qualitative Studies: We present qualitative results on corrupted and realistic test samples in Figure 3. The comparison includes predictions on the corresponding clean images (Clean Depth), predictions from the pre-trained model without adaptation (Pre-trained Model), and results from various adaptation methods. Note that when the target image is not the corrupted version, "clean depth" and "Pretrained Model" is the same result. Overall, our proposed method WeSTAR consistently produces more accurate depth estimates, often surpassing even the predictions on clean images. More visualization results are provided in appendix.

Ablation Study and Additional Analysis

Effects of Individual Components: We conduct ablation studies on three datasets as shown in Table 4. Starting from the zero-shot baseline, applying self-training (ST) consistently improves performance, especially on the corrupted Sintel-C dataset (δ_1 : 60.3 \rightarrow 63.3, AbsRel: 30.6 \rightarrow 27.6), indicating enhanced robustness under domain shifts. However, ST brings only marginal gains on clean datasets like Sintel and NuScenes, due to confirmation bias from inaccurate pseudo-labels and the strong geometric prior of the pre-trained model, which limits further improvement from ST alone. Weak supervision (WS) moderately improves δ_1 (e.g., 74.8 \rightarrow 77.5 on Sintel), but its scale insensitivity introduces geometric distortions and yields inconsistent AbsRel results, sometimes even degrading performance (e.g., 20.3 \rightarrow 24.1 on Sintel). Weight regularization (WR) further stabilizes training by anchoring to initial parameters, reducing overfitting to noisy pseudo-labels or ordinal relations. Combining all components, WeSTAR significantly outperforms all variants, highlighting their complementary effects.

Effects of Tuning Different Modules: Table 5 shows that LoRA-based adaptation consistently outperforms full fine-tuning. Although full fine-tuning slightly outperforms LoRA on Sintel-C, this likely stems from its tendency to overfit corruptions. In contrast, LoRA constrains updates to a low-rank subspace, mitigating overfitting and preserving pre-trained knowledge. Tuning only the decoder brings no im-

Method	Sintel-C		DIODE-C		Sintel		DIODE		NuScenes		D-Sunny		D-Foggy		D-Cloudy		D-Rainy	
	δ_1	AbsRel	δ_1	AbsRel	δ_1	AbsRel	δ_1	AbsRel	δ_1	AbsRel	δ_1	AbsRel	δ_1	AbsRel	δ_1	AbsRel	δ_1	AbsRel
Source	43.2	40.5	74.6	17.7	58.6	30.4	91.4	9.3	52.3	32.5	73.7	17.8	84.2	11.9	77.2	16.4	62.2	22.4
TTAC	48.4	36.6	81.7	14.3	61.5	26.6	91.4	9.3	53.0	32.3	75.9	16.8	84.5	11.9	78.1	15.9	67.5	18.6
TTAC	44.8	38.3	80.2	15.0	58.7	30.4	91.4	9.3	52.3	32.5	73.6	17.8	84.2	11.9	77.2	16.4	62.1	22.5
FR	45.2	38.0	81.4	14.7	58.6	30.4	91.4	9.3	53.6	32.0	73.5	17.9	84.5	11.7	78.5	15.7	69.4	18.6
SSA	45.5	38.8	80.3	15.2	58.6	30.4	91.4	9.3	53.1	31.8	75.2	17.2	83.7	12.0	78.1	15.8	70.7	18.0
SGRL	52.0	35.6	85.9	12.2	65.5	26.2	91.6	9.1	63.4	24.4	79.8	15.2	85.3	11.2	79.1	15.0	74.5	16.2
WeSTAR	56.6	34.8	86.6	12.0	67.9	25.8	91.7	9.0	67.5	22.0	78.3	15.8	87.1	10.1	80.4	14.5	75.3	15.5

Table 3: Adaptation results for MiDas v3.1 backbone.

ST	WR	WS	Sintel-C		Sintel		NuScenes	
			δ_1	AbsRel	δ_1	AbsRel	δ_1	AbsRel
-	-	-	60.3	30.6	74.8	20.3	74.4	18.5
✓	-	-	63.3	27.6	74.9	20.3	74.4	18.4
-	-	✓	68.3	28.2	77.5	24.1	76.9	16.9
✓	-	✓	70.1	27.1	80.2	21.3	76.8	16.7
✓	✓	✓	71.8	24.1	82.2	16.9	78.1	16.2

Table 4: Ablation Studies of Individual Component. "ST" is the self-training loss, "WS" is the weakly-supervised loss and "WR" is the weight regularization loss.

provement, and updating both encoder and decoder offers no advantage over encoder-only tuning, emphasizing the encoder's key role in adaptation.

Component	Sintel-C		Sintel		NuScenes	
	δ_1	AbsRel	δ_1	AbsRel	δ_1	AbsRel
All Params	72.4	24.8	80.8	19.5	77.1	16.4
Encoder	72.2	24.7	80.7	19.4	77.2	16.6
Decoder	60.3	30.7	74.8	20.5	74.7	18.3
LoRA	71.8	24.1	82.2	16.9	78.1	16.2

Table 5: Effects of tuning different modules.

Effects of Normalization: Table 6 validates SA-HDN against Global Norm and HDN. Our method clearly outperforms both, consistently surpassing the strong Global Norm baseline and the grid-based HDN. This is due to its ability to enforce geometric consistency at multiple semantic scales. By normalizing within instances, it preserves fine-grained structure often lost in global alignment.

Norm	Sintel-C		Sintel		NuScenes	
	δ_1	AbsRel	δ_1	AbsRel	δ_1	AbsRel
Global	70.9	24.2	80.8	16.4	77.2	16.7
HDN	68.9	28.1	78.6	22.4	77.2	16.6
SA-HDN	71.8	24.1	82.2	16.9	78.1	16.2

Table 6: Effects of different Normalization methods.

Training Stability: We evaluate training stability by examining performance trajectories on Sintel in Figure 4. Feature alignment and contrastive methods (e.g., TTAC, FR)

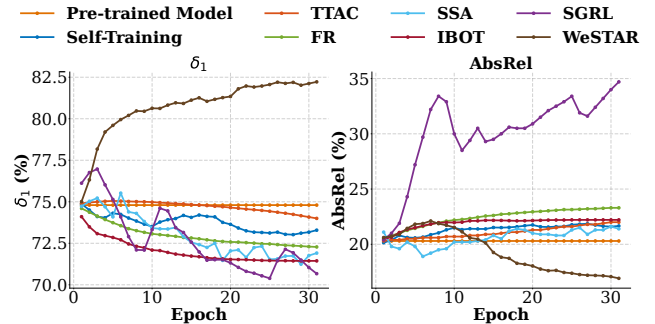


Figure 4: Comparative Performance over Training Epochs on Sintel.

show modest, inconsistent gains due to sensitivity to low-quality pseudo-labels and feature drift. SSA yields smoother but conservative improvements, avoiding abrupt changes yet limiting progress. The self-supervised method iBOT achieves rapid early gains but quickly plateaus, suggesting pretrained features alone are insufficient for sustained adaptation under corruption. SGRL, a weakly supervised approach, improves sharply initially but degrades over time, especially on DIODE Indoor-C, due to limited and noisy supervision causing divergence toward suboptimal solutions. By contrast, WeSTAR delivers stable, consistent improvements on both δ_1 and AbsRel through training, benefiting from reliable task-specific supervision and effective regularization that prevent performance drops and model drift.

Conclusion

This work investigates the generalization capabilities of pre-trained depth estimation foundation models under unseen and corrupted conditions. When limited unlabeled and weakly labeled target domain data is available, we propose an adaptation strategy that combines self-training with weight regularization and optional ordinal supervision. This design effectively addresses challenges posed by unreliable pseudo-labels and facilitates stable adaptation. Our comprehensive experiments across diverse benchmark datasets validate the efficacy of our method, showing substantial improvements in both realistic and challenging corrupted domains. The results underscore the potential of combining weak supervision and regularized adaptation to enhance the robustness and applicability of depth estimation foundation models in real-world scenarios.

Acknowledgments

This work is supported by the Guangdong Basic and Applied Basic Research Foundation (No. 2023A1515110077) and the Agency for Science, Technology and Research (A*STAR) under its MTC Programmatic Funds (Grant No. M23L7b0021).

References

- Adachi, K.; Yamaguchi, S.; Kumagai, A.; and Hamagami, T. 2025. Test-time Adaptation for Regression by Subspace Alignment. In *International Conference on Learning Representations*.
- Bhat, S. F.; Alhashim, I.; and Wonka, P. 2021. Adabins: Depth estimation using adaptive bins. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*.
- Birkel, R.; Wofk, D.; and Müller, M. 2023. Midas v3. 1—a model zoo for robust monocular relative depth estimation. *arXiv preprint arXiv:2307.14460*.
- Butler, D. J.; Wulff, J.; Stanley, G. B.; and Black, M. J. 2012. A naturalistic open source movie for optical flow evaluation. In *European Conference on Computer Vision (ECCV)*. Springer.
- Caesar, H.; Bankiti, V.; Lang, A. H.; Vora, S.; Liong, V. E.; Xu, Q.; Krishnan, A.; Pan, Y.; Baldan, G.; and Beijbom, O. 2020. nuScenes: A multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Chen, D.; Wang, D.; Darrell, T.; and Ebrahimi, S. 2022. Contrastive test-time adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*. PMLR.
- Chen, W.; Fu, Z.; Yang, D.; and Deng, J. 2016. Single-image depth perception in the wild. In *Advances in neural information processing systems*.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*.
- Ding, N.; Xu, Y.; Tang, Y.; Xu, C.; Wang, Y.; and Tao, D. 2022. Source-free domain adaptation via distribution estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Eastwood, C.; Mason, I.; Williams, C. K. I.; and Schölkopf, B. 2022. Source-Free Adaptation to Measurement Shift via Bottom-Up Feature Restoration. In *International Conference on Learning Representations*.
- Eigen, D.; Puhrsch, C.; and Fergus, R. 2014. Depth map prediction from a single image using a multi-scale deep network. In *Advances in neural information processing systems*.
- Ganj, A.; Zhao, Y.; Su, H.; and Guo, T. 2024. Mobile AR depth estimation: Challenges & prospects. In *Proceedings of the 25th International Workshop on Mobile Computing Systems and Applications*.
- Geiger, A.; Lenz, P.; Stiller, C.; and Urtasun, R. 2013. Vision meets robotics: The kitti dataset. *The international journal of robotics research*, 1231–1237.
- Godard, C.; Mac Aodha, O.; and Brostow, G. J. 2017. Unsupervised monocular depth estimation with left-right consistency. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- Godard, C.; Mac Aodha, O.; Firman, M.; and Brostow, G. J. 2019. Digging into self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF international conference on computer vision*.
- Goyal, P.; Dollár, P.; Girshick, R.; Noordhuis, P.; Wesolowski, L.; Kyrola, A.; Tulloch, A.; Jia, Y.; and He, K. 2017. Accurate, large minibatch sgd: Training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- He, T.; Xia, Z.; Chen, J.; Li, H.; and Chan, S.-H. G. 2024. Target-agnostic source-free domain adaptation for regression tasks. In *2024 IEEE 40th International Conference on Data Engineering (ICDE)*. IEEE.
- He, X.; Guo, D.; Li, H.; Li, R.; Cui, Y.; and Zhang, C. 2025. Distill Any Depth: Distillation Creates a Stronger Monocular Depth Estimator. *arXiv preprint arXiv: 2502.19204*.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2022. LoRA: Low-Rank Adaptation of Large Language Models. In *International Conference on Learning Representations*.
- Hwang, U.; Lee, J.; Shin, J.; and Yoon, S. 2024. SF(DA)²: Source-free Domain Adaptation Through the Lens of Data Augmentation. In *International Conference on Learning Representations*.
- Jing, M.; Zhen, X.; Li, J.; and Snoek, C. 2022. Variational model perturbation for source-free domain adaptation. In *Advances in Neural Information Processing Systems*.
- Karim, N.; Mithun, N. C.; Rajvanshi, A.; Chiu, H.-p.; Samarasekera, S.; and Rahnavard, N. 2023. C-sfda: A curriculum learning aided self-training framework for efficient source free domain adaptation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*.
- Kong, L.; Xie, S.; Hu, H.; Ng, L. X.; Cottureau, B.; and Ooi, W. T. 2023. Robodepth: Robust out-of-distribution depth estimation under corruptions. In *Advances in Neural Information Processing Systems*.
- Lee, J.; and Lee, G. 2023. Feature alignment by uncertainty and self-training for source-free unsupervised domain adaptation. *Neural Networks*.

- Li, Z.; Shi, S.; Schiele, B.; and Dai, D. 2023. Test-time domain adaptation for monocular depth estimation. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE.
- Liang, J.; Hu, D.; and Feng, J. 2020. Do we really need to access the source data? source hypothesis transfer for unsupervised domain adaptation. In *International conference on machine learning*. PMLR.
- Liu, Y.; Kothari, P.; Van Delft, B.; Bellot-Gurlet, B.; Mordan, T.; and Alahi, A. 2021. Ttt++: When does self-supervised test-time training fail or thrive? In *Advances in Neural Information Processing Systems*.
- Mitsuzumi, Y.; Kimura, A.; and Kashima, H. 2024. Understanding and improving source-free domain adaptation from a theoretical perspective. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Ranftl, R.; Lasinger, K.; Hafner, D.; Schindler, K.; and Koltun, V. 2020. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE transactions on pattern analysis and machine intelligence*.
- Ravi, N.; Gabeur, V.; Hu, Y.-T.; Hu, R.; Ryali, C.; Ma, T.; Khedr, H.; Rädle, R.; Rolland, C.; Gustafson, L.; et al. 2025. SAM 2: Segment Anything in Images and Videos. In *International Conference on Learning Representations*.
- Silberman, N.; Hoiem, D.; Kohli, P.; and Fergus, R. 2012. Indoor segmentation and support inference from rgb-d images. In *European Conference on Computer Vision (ECCV)*.
- Sohn, K.; Berthelot, D.; Carlini, N.; Zhang, Z.; Zhang, H.; Raffel, C. A.; Cubuk, E. D.; Kurakin, A.; and Li, C.-L. 2020. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. In *Advances in neural information processing systems*.
- Su, Y.; Xu, X.; and Jia, K. 2022. Revisiting realistic test-time training: Sequential inference and adaptation by anchored clustering. In *Advances in Neural Information Processing Systems*.
- Tarvainen, A.; and Valpola, H. 2017. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *Advances in neural information processing systems*.
- 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. 2019. Diode: A dense indoor and outdoor depth dataset. *arXiv preprint arXiv:1908.00463*.
- Xian, K.; Zhang, J.; Wang, O.; Mai, L.; Lin, Z.; and Cao, Z. 2020. Structure-guided ranking loss for single image depth prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Xie, J.; Girshick, R.; and Farhadi, A. 2016. Deep3d: Fully automatic 2d-to-3d video conversion with deep convolutional neural networks. In *European conference on computer vision*.
- Yang, G.; Song, X.; Huang, C.; Deng, Z.; Shi, J.; and Zhou, B. 2019. DrivingStereo: A Large-Scale Dataset for Stereo Matching in Autonomous Driving Scenarios. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Yang, L.; Kang, B.; Huang, Z.; Xu, X.; Feng, J.; and Zhao, H. 2024a. Depth anything: Unleashing the power of large-scale unlabeled data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Yang, L.; Kang, B.; Huang, Z.; Zhao, Z.; Xu, X.; Feng, J.; and Zhao, H. 2024b. Depth anything v2. In *Advances in Neural Information Processing Systems*.
- Yi, E.; and Kim, J. 2023. Test-time synthetic-to-real adaptive depth estimation. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE.
- Yin, W.; Zhang, J.; Wang, O.; Niklaus, S.; Chen, S.; Liu, Y.; and Shen, C. 2022. Towards accurate reconstruction of 3d scene shape from a single monocular image. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Zhang, C.; Yin, W.; Wang, B.; Yu, G.; Fu, B.; and Shen, C. 2022a. Hierarchical normalization for robust monocular depth estimation. In *Advances in Neural Information Processing Systems*.
- Zhang, H.; Su, Y.; Xu, X.; and Jia, K. 2024. Improving the generalization of segmentation foundation model under distribution shift via weakly supervised adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Zhang, Z.; Chen, W.; Cheng, H.; Li, Z.; Li, S.; Lin, L.; and Li, G. 2022b. Divide and contrast: Source-free domain adaptation via adaptive contrastive learning. In *Advances in Neural Information Processing Systems*.
- Zhao, C.; Zhang, Y.; Poggi, M.; Tosi, F.; Guo, X.; Zhu, Z.; Huang, G.; Tang, Y.; and Mattocchia, S. 2022. Monovit: Self-supervised monocular depth estimation with a vision transformer. In *2022 international conference on 3D vision (3DV)*. IEEE.
- Zhao, D.; Li, J.; Wang, S.; Wu, M.; Zang, Q.; Sebe, N.; and Zhong, Z. 2025. FisherTune: Fisher-Guided Robust Tuning of Vision Foundation Models for Domain Generalized Segmentation. *arXiv preprint arXiv:2503.17940*.
- Zhao, H.; Liu, Y.; Alahi, A.; and Lin, T. 2023. On pitfalls of test-time adaptation. In *Proceedings of the 40th International Conference on Machine Learning*.
- Zhou, J.; Wei, C.; Wang, H.; Shen, W.; Xie, C.; Yuille, A.; and Kong, T. 2021. ibot: Image bert pre-training with online tokenizer. *arXiv preprint arXiv:2111.07832*.