

# DEFOM-Stereo: Depth Foundation Model Based Stereo Matching

Hualie Jiang<sup>1</sup> Zhiqiang Lou<sup>1</sup> Laiyan Ding<sup>2</sup>

Rui Xu<sup>1\*</sup> Minglang Tan<sup>1</sup> Wenjie Jiang<sup>1</sup> Rui Huang<sup>2</sup>

<sup>1</sup>Insta360 Research <sup>2</sup>The Chinese University of Hong Kong, Shenzhen

{jianghualie, johnlou, jerry1}@insta360.com, laiyanding@link.cuhk.edu.cn

<https://insta360-research-team.github.io/DEFOM-Stereo>

## Abstract

Stereo matching is a key technique for metric depth estimation in computer vision and robotics. Real-world challenges like occlusion and non-texture hinder accurate disparity estimation from binocular matching cues. Recently, monocular relative depth estimation has shown remarkable generalization using vision foundation models. Thus, to facilitate robust stereo matching with monocular depth cues, we incorporate a robust monocular relative depth model into the recurrent stereo-matching framework, building a new framework for depth foundation model-based stereo-matching, DEFOM-Stereo. In the feature extraction stage, we construct the combined context and matching feature encoder by integrating features from conventional CNNs and DEFOM. In the update stage, we use the depth predicted by DEFOM to initialize the recurrent disparity and introduce a scale update module to refine the disparity at the correct scale. DEFOM-Stereo is verified to have much stronger zero-shot generalization compared with SOTA methods. Moreover, DEFOM-Stereo achieves top performance on the KITTI 2012, KITTI 2015, Middlebury, and ETH3D benchmarks, ranking 1<sup>st</sup> on many metrics. In the joint evaluation under the robust vision challenge, our model simultaneously outperforms previous models on the individual benchmarks, further demonstrating its outstanding capabilities.

## 1. Introduction

Stereo matching involves recovering the disparity between two views and has been a fundamental topic in computer vision for decades. Stereo matching has obtained steady improvement with the consistently proposed algorithms, which range from hand-designed ones [40] to end-to-end deep architectures [24, 48]. The deep matching frameworks have mainly shifted from cost-volume filtering ones [6, 17, 23, 58, 66] to recent recurrent refinement

\*corresponding author

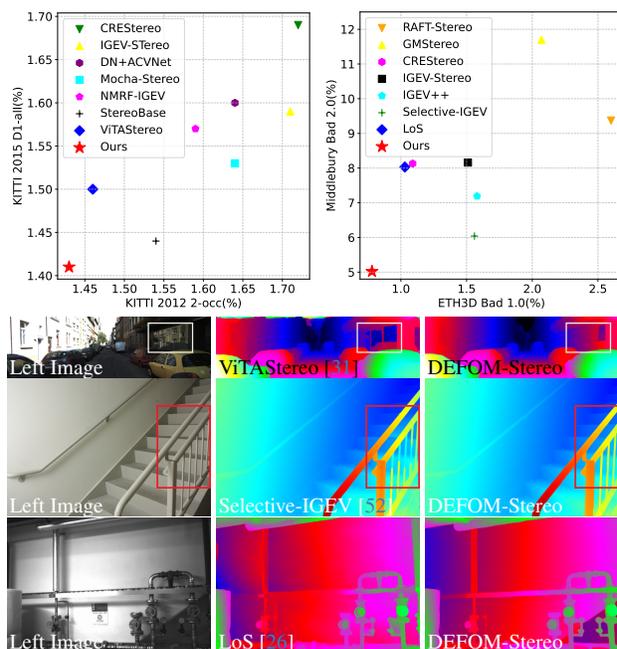


Figure 1. **Row 1:** Comparisons with SOTA stereo methods on KITTI 2012 [15] and 2015 [33], Middlebury [41] and ETH3D [42] leaderboards. **Row 2-4:** Visual comparison with ViTAStereo [31] on KITTI 2012, Selective-IGEV [52] on Middlebury, and LoS [26] on ETH3D. Our method performs robustly on reflective, textureless, and ill-exposure areas and recovers more structural details.

ones [8, 25, 29, 52, 56, 69]. According to [48], despite RAFT-Stereo [29] has made good progress, the generalization and robustness remain challenging. The critical challenges lie in occlusion, non-texture, image blur, high-resolution, *etc. etc.*

On the other hand, a closely related task, monocular relative depth estimation [3, 22, 36, 37, 62, 63] has achieved advancing zero-shot generalization ability across diverse scenes. Among them, Depth Anything V2 [63], is a model with a pre-trained vision transformer (ViT) [11] backbone of the general-purpose vision foundation model DINOv2 [35] and a head of the dense prediction transformer (DPT) [37]. It is trained on synthetic labeled images and distilled on

large realistic unlabeled images to conduct robust and efficient depth prediction and recover impressive fine-grained details. These properties allow it to be a depth foundation model (DEFOM). Nevertheless, DEFOM does not provide metric depth, which is crucial for robotic applications. To achieve more robust stereo depth estimation, we incorporate Depth Anything V2 into RAFT-Stereo [29] to build a DEFOM-based stereo matching (DEFOM-Stereo) model. To be specific, we perform the incorporation with two aspects, 1) utilizing DEFOM’s pre-trained features to improve the feature extraction for stereo matching, and 2) facilitating disparity updating with the predicted monocular depth.

Utilizing the pre-trained ViT to improve stereo matching has recently been explored in [31, 68] via developing attention-based adapters to convert the ViT backbone features into matching features. In contrast, the DEFOM provides dense feature maps with a DPT head, and the feature is more relevant to the depth task. We demonstrate that a simple convolutional fusion of DPT features and plain CNN features is effective. As the existing DPT should be fixed to predict depth, we initialize a trainable DPT head to provide more flexible features for fusion. Apart from the matching feature, there is a context encoder of the recurrent stereo framework to provide monocular cues controlling the recurrent disparity update process. Thus, we propose a combined matching feature and context encoder with CNN and DEFOM.

Besides its features, the DEFOM depth can intuitively be used to improve disparity recovery. The DEFOM depth is an affine disparity with unknown scale and shift. However, finding a single scale and shift values does help to recover accurate disparity. A preliminary evaluation of Depth Anything V2 [63] on typical stereo datasets using the least-square affine alignment illustrates the disparity error is significantly large, indicating the scale inconsistency for depth prediction within the image is serious. The phenomenon is more serious for synthetic stereo datasets, for example of the synthetic Scene Flow [32] shown in Fig. 2. We conjecture that the scale inconsistency is caused by the training in different FoV images with affine invariant loss. As the stereo model is usually pre-trained on the synthetic stereo dataset, the scale inconsistency of DEFOM poses challenges for recovering disparity from its depth estimate.

To cope with the scale inconsistency, we propose a scale update (SU) module that performs recurrent dense scale updates on the initialized disparity from the depth estimate. We found that the depth amplitude by Depth Anything V2 varies across different image samples with different backbone sizes. Considering that the disparity range is usually proportional to image resolution, we use the image width to normalize the depth estimate in initialization. To facilitate SU, we design a scale lookup (SL) on the correlation volume on which the current disparity estimate is multiplied with a set of scale values to generate sampling indexes. Compared with the con-



Figure 2. Illustration of the significant scale inconsistency in the depth estimate of Depth Anything V2 [63] compared to the ground-truth disparity on the synthetic Scene Flow dataset.

ventional pyramid lookup (PL), SL has a complete search range on the whole image, enabling SU with global matching. Therefore, we insert the SU module before the conventional delta update (DU) module in RAFT-Stereo [29], which can continue to recover local details upon the result of SU.

The overall framework of DEFOM-Stereo is illustrated in Fig. 3. We have conducted extensive experiments to verify DEFOM-Stereo’s effectiveness. When pre-training only on Scene Flow, DEFOM-Stereo achieves comparable performance with SOTA methods and obtains considerable progress in the zero-shot generation evaluation on KITTI 2012 and 2015, Middlebury, and ETH3D, for example, reducing the 2-pixel thresholding error rate by about a third on Middlebury. When submitting to the official benchmarks of these datasets, DEFOM-Stereo achieves top performance, ranking 1<sup>st</sup> on many metrics on the four leaderboards among all submissions while writing this paper. A brief comparison with published SOTA methods is demonstrated in Fig. 1, where DEFOM-Stereo exhibits remarkable advantages. Furthermore, we also examine our model under the joint evaluation of the Robust Vision Challenge (RVC), and our model simultaneously outperforms all previous RVC models on the three benchmarks, particularly on KITTI 2015 and Middlebury with a large margin.

The contributions can be summarized as follows:

- We propose a novel recurrent stereo-matching framework incorporating monocular depth cues from a depth foundation model to improve robustness.
- We develop a simple technique that utilizes pre-trained DEFOM features to construct stronger combined feature and context encoders.
- We present a recurrent scale update module empowered with the scale lookup, serving to recover accurate pixel-wise scales for the coarse DEFOM depth.

## 2. Related Work

### 2.1. Learning-Based Stereo Matching

Deep learning methods have dominated stereo matching for past decades. Early methods [6, 12, 17, 23, 58, 65] mainly construct a cost volume on defined disparities with the deep feature maps of left and right images and use 3D convolutions to perform cost aggregation. The final disparity estimate is regressed from the processed cost volume. In recent years, the recurrent refinement-based methods [8, 14, 21, 25, 26, 29, 52, 56, 69] became the mainstream,

as well as the transformer-based ones [27, 53, 59]. The transformer-based methods usually use the cross-attention between two views to correlate them for the final disparity prediction. Unlike traditional cost-volume construction approaches, these methods circumvent the constraints of a pre-defined disparity range and effectively capture long-range pixel associations. More recently, a unique attempt to leverage the power of deep learning to create a neural Markov random field (NMRF-Stereo) model was presented [16], establishing an accurate and highly efficient model.

The recurrent stereo-matching stems from the framework of recurrent all-pairs field transforms (RAFT) [47] for optical flow estimation. RAFT-Stereo [29] is a pioneering work that adapts the RAFT framework for stereo matching mainly by limiting the 2D flow field to the 1D disparity field and upgrading the convolutional Gated Recurrent Unit (ConvGRU) for recurrent updates from a single level to multiple levels to enlarge the receptive field. By careful modification, RAFT-Stereo inherits strong cross-dataset generalization from RAFT. Subsequently, there are a series of improved models from different aspects [8, 14, 21, 25, 26, 52, 56, 69]. For example, CREStereo [25] proposed a hierarchical recurrent refinement structure with an adaptive group correlation layer. IGEV-Stereo [56] uses an aggregated cost volume to estimate the initial disparity and combines it with the all-pairs correlations into a combined geometry encoding volume. Selective-Stereo [52] preposes a selective recurrent unit which enables the recurrent update to adaptively capture and fuse multi-frequency information. Most of the improved models have achieved a better in-domain fitting than RAFT-Stereo. However, with comprehensive zero-shot generalization evaluation, we found that RAFT-Stereo is still more advantageous. Therefore, we use the based RAFT-Stereo framework as a baseline model and construct our model by integrating powerful monocular cues from DEFOM.

## 2.2. Zero-shot Generalized Stereo Matching

Zero-shot generalization is a crucial challenge in stereo matching, particularly for synthetic-to-real transfer, as labeled real-world datasets are difficult to obtain at scale. Earlier cost-volume-based methods [6, 17, 65] exhibit a significant performance drop when trained on synthetic datasets (e.g., Scene Flow) and tested on real-world data. This decline is primarily due to the poor transferability of the matching features learned from synthetic data to real scenes. Consequently, substantial efforts have been devoted to learning domain-agnostic features. For instance, DSMNet [66] and FCStereo [67] aim to learn domain-invariant features. DSMNet introduces domain normalization to regularize feature distributions, while FCStereo applies a contrastive loss to ensure feature consistency between matched pixels and a stereo whitening loss to preserve feature consistency across domains. Other approaches, such as ITSA [10] and HVT [7],

tackle the issue of shortcut learning—where superficial cues are exploited instead of transferable representations—by using information-theoretic losses to constrain the encoding of shortcut-related information or by transforming synthetic training images to diversify the training domain.

In contrast to these works, which focus on strategies tailored to improve generalization, RAFT-Stereo achieves significant improvements in zero-shot performance via architectural innovation. We further enhance zero-shot generalization by integrating a robust monocular depth estimation model to adapt RAFT-Stereo into a new framework.

Note that there are some contemporary works [2, 9, 54] that incorporate the depth foundation model into the recurrent stereo matching framework. They both use 3D convolutions upon a 4D cost volume to obtain initial disparity following IGEV-Stereo [56] and [2, 9] further get another refined initial disparity via alignment with the DEFOM depth. Our model differs in directly replacing the 0 initial disparity of the original RAFT-stereo with the DEFOM depth and reforming the recurrent update process accordingly.

## 2.3. Foundation Model for Depth Estimation

MiDaS [36] pioneers zero-shot relative depth estimation (ZS-RDE), using a scale-shift-invariant loss to handle depth distribution inconsistencies in large-scale mixed datasets, achieving robust depth estimation across diverse scenes with CNNs. Later, DPT [37] introduces a dense prediction head for the vision transformers, yielding finer-grained and more globally coherent depth predictions on large-scale training datasets. MiDaS v3.1 [3] conducts an in-depth and comprehensive evaluation of vision transformer backbones with varying architectures and pre-training methods, demonstrating that enhanced backbones consistently yield performance improvements. More recently, studies [22, 62, 63] leverage powerful vision foundation models for ZS-RDE task, resulting in a significant enhancement of zero-shot generalization capabilities and ushering the field into a new era.

Among the public foundational models, Stable Diffusion (SD) [39], trained on vast datasets for high-quality image generation, and DINOv2 [35], a general-purpose vision foundation model, are proven to have strong zero-shot generalization capabilities in both monocular and multi-view geometric estimation [13]. Marigold [22] fine-tunes pre-trained SD on limited synthetic depth data, achieving strong cross-dataset accuracy and capturing fine details. However, the diffusion process results in high inference costs and slow processing. Depth Anything [62] employs DINOv2 as the encoder in the DPT architecture and proposes training a large powerful teacher model on labeled data, followed by distilling a student model on vast unlabeled real-world images to enhance its generalization performance. While being much faster, Depth Anything recovers fewer details than Marigold.

Building upon Depth Anything, Depth Anything V2 [63]

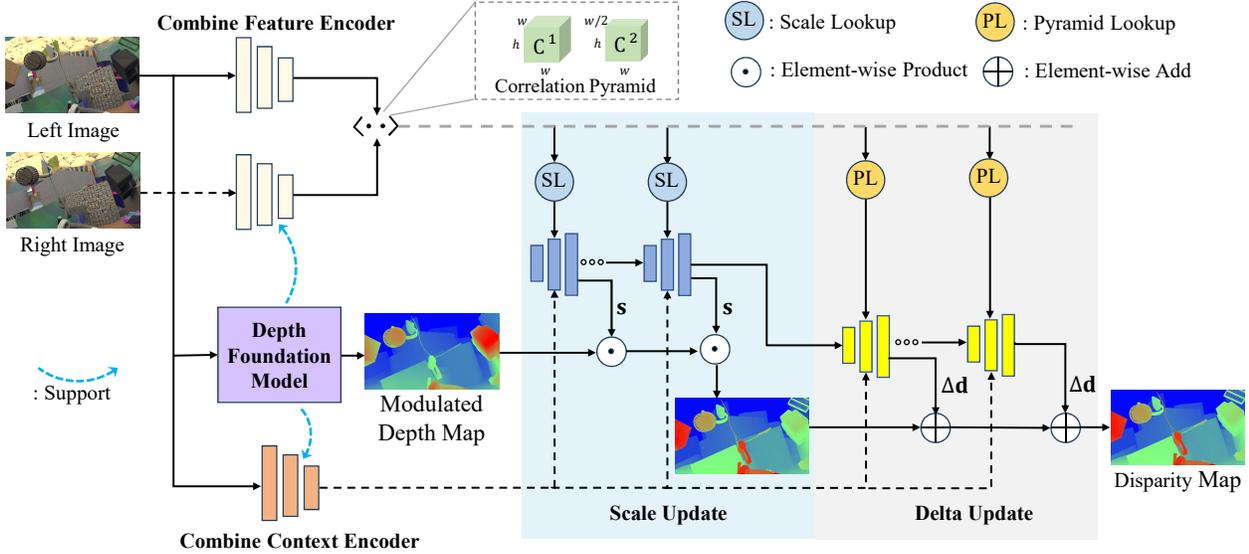


Figure 3. Overview of our proposed DEFOM-Stereo. At first, we use a depth foundation model to augment the plain CNN encoders, thus obtaining more powerful feature and context encoders. Next, we insert the scale update module before the conventional delta update module. The SU starts from the modulated depth estimate from DEFOM and recurrently performs dense scale recovery with scaling lookup from correlation volume, while the DU continues to refine the result of SU with the pyramid lookup to obtain the final disparity map.

reevaluates training data construction and notes that depth label noise from real captured datasets hinders the network’s ability to generate detailed predictions. It turns to use large-scale precise synthetic datasets to train a giant teacher model which subsequently distills student models on extensive unlabeled real data, yielding accurate and efficient depth foundation models. Since Depth Anything V2 can rapidly infer detailed and accurate relative depth across various scenes, even including challenging areas such as reflective and transparent regions, we propose integrating its monocular depth cues for robust disparity estimation from stereo images.

### 3. Methodology

#### 3.1. Overall Framework

Fig. 3 shows the overall framework of DEFOM-Stereo, which incorporates a depth foundation model, Depth Anything V2 [63] into the recurrent stereo model, RAFT-Stereo [29]. The framework consists of two main stages. The first stage mainly includes feature extraction, correlation pyramid construction, and recurrent disparity initialization. In this stage, we use the pre-trained DEFOM to enhance feature extraction and its predicted depth map to initialize disparity. In the next disparity update stage, we propose a scale update module to recover the accurate repeatedly dense scales for the disparities and then apply the conventional delta update to obtain finer details.

#### 3.2. Combined Feature Extraction

There are two encoders in feature extraction. The first encoder is the feature encoder, which extracts the matching

feature maps at  $1/4$  resolution from the left and right images. Given the image size  $H \times W$ , we define  $h = H/4$  and  $w = W/4$  as the feature spatial size. The matching feature maps are used to construct the all-pair correlation volume, which provides matching cues for the recurrent update.

Another is the context encoder, which is only applied on the left images to obtain multi-level context features at  $1/4$ ,  $1/8$ , and  $1/16$  resolutions. The context features are employed to initialize the update operator’s hidden state and integrated into the GRU at each iteration of the update process, providing the monocular cue for the update process. In RAFT-Stereo [29], the feature and context encoders are implemented with plain CNNs.

We propose a simple but effective way to combine the pre-trained DEFOM and plain CNNs to construct improved feature extractors for stereo matching. We use the pre-trained backbone from the DEFOM, and initialize a new shared dense prediction transformer (DPT) to predict the ViT features for both the matching feature and context, as it is trainable compared with the existing DPT in the DEFOM, which has to be fixed for predicting depth.

**Combined Feature Encoder.** We take the final fused feature map of the DPT and further, use a convolutional block to align it with the CNN feature map in the channel. After that, we simply add them as a combined feature map.

**Combined Context Encoder.** We take the feature maps from the  $Reassemble_4$ ,  $Reassemble_8$ , and  $Reassemble_{16}$  of the DPT as the ViT context maps. Similarly, we use another three convolutional blocks for channel alignment. Finally, we add the aligned ViT map and the CNN map to get the combined context map at different resolutions.

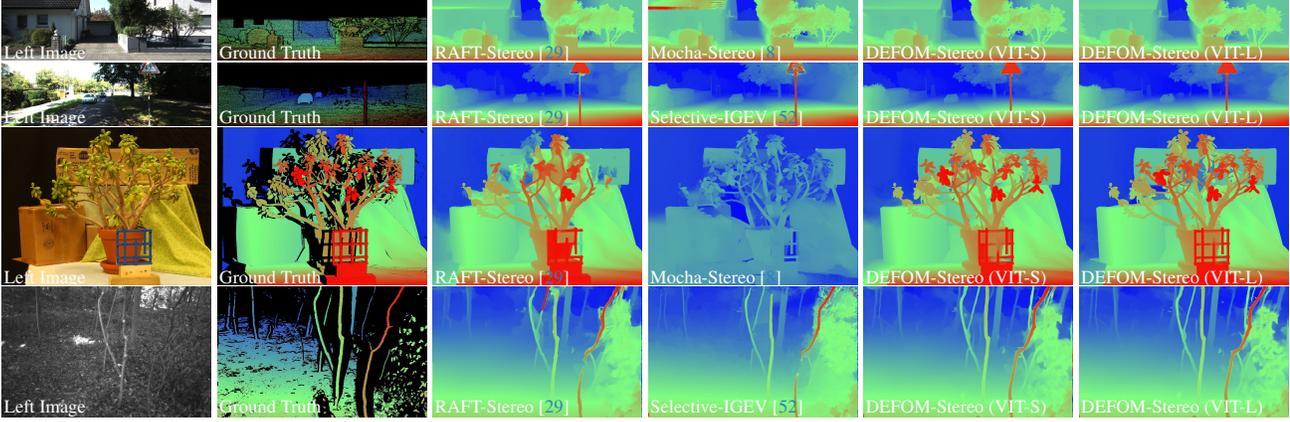


Figure 4. Zero-Shot Qualitative Comparison with RAFT-Stereo [29], Selective-IGEV [52] and Mocha-Stereo [8] on the four stereo datasets. **Row 1:** KITTI 2012. **Row 2:** KITTI 2015. **Row 3:** Middlebury-full. **Row 4:** ETH3D. Best viewed in color and by zooming in.

### 3.3. Correlation Construction and Lookup

The section shows how to construct the all-pair correlation volume pyramid and perform correlation pyramid lookup. Given the extracted feature maps,  $\mathbf{f}^l, \mathbf{f}^r$  ( $\mathbf{f} \in \mathbb{R}^{c \times h \times w}$ ), the first level all-pairs correlation volume is computed as,

$$\mathbf{C}_{ijk}^1 = \sum_h \mathbf{f}_{hij}^l \cdot \mathbf{f}_{hik}^r, \mathbf{C}^1 \in \mathbb{R}^{h \times w \times w}. \quad (1)$$

Then, the multi-level correlation pyramid  $\{\mathbf{C}^l\}$  can be obtained by repeatedly applying 1D average pooling at the last dimension to the upper correlation volume, and the last dimension will be reduced by half each time.

With the correlation pyramid created, the pyramid lookup (PL) is performed as follows. Given the current disparity estimate, the correlation values of it and its neighbors within a radius  $r$  will be sampled from the correlation pyramid and concatenated into the retrieved correlation feature.

RAFT-Stereo [29] uses a 4-level pyramid and sets the radius  $r$  as 4, so the maximum search range is  $2^2 \times 2^3 \times 4 = 128$ , which might be insufficient for the large disparity. In contrast, the proposed scale lookup does not have this limitation. As our scale update can provide approximated results, the delta update should focus on local details and we set the pyramid level as 2 following IGEV-Stereo [56].

### 3.4. Delta Update

In the iterative motion estimation framework [29, 47], the motion field is initialized with  $\mathbf{0}$  and updated with a residual. For disparity, it is updated in an additive manner,

$$\mathbf{d}_n = \mathbf{d}_{n-1} + \Delta \mathbf{d}, \quad (2)$$

where  $\Delta d$  is a delta disparity map estimated by a multi-level convolutional Gated Recurrent Unit (ConvGRU).

The ConvGRU at resolutions of  $1/8$  or  $1/16$  takes the context map at the same resolution and the reshaped hidden state at adjacent resolutions as input and updates its hidden

state. These two ConvGRUs help to increase the receptive field of the update operation. The finest ConvGRU at  $1/4$  resolution takes not only the context map at  $1/4$  resolution and the reshaped hidden state at  $1/8$  resolution, but also the encoded feature of the current disparity estimate and retrieved correlation feature map to update its hidden state. The hidden state  $h_n$  at  $1/4$  resolution is updated as follows,

$$\begin{aligned} x_n &= [\text{Encoder}_c(\mathbf{c}), \text{Encoder}_d(\mathbf{d}_{n-1}), \text{Up}^2(h_n^{1/8})] \\ z_n &= \sigma(\text{Conv}([h_{n-1}, x_n], W_z) + c_z), \\ r_n &= \sigma(\text{Conv}([h_{n-1}, x_n], W_r) + c_r), \\ \tilde{h}_n &= \tanh(\text{Conv}([r_n \odot h_{n-1}, x_n], W_h) + c_h), \\ h_n &= (1 - z_n) \odot h_{n-1} + z_n \odot \tilde{h}_n, \end{aligned} \quad (3)$$

where  $c_z, c_r, c_h$  are context features generated from the context network.  $h_n^{1/8}$  is the hidden state at  $1/8$  resolution.  $\mathbf{c}$  is the retrieved correlation feature map from the correlation volume pyramid with the previous disparity estimate  $\mathbf{d}_{n-1}$ .

Finally, the updated  $h_n$  will be used to predict the delta disparity map  $\Delta \mathbf{d}$  in Eqn. 2 with two convolutional layers. As the disparity estimate is at  $1/4$  resolution,  $h_n$  also predicts a convex weight map, so that  $\mathbf{d}_n$  can be upsampled to the original resolution with the convex combination [47].

### 3.5. Monocular Depth Initialization

We first propose to utilize the predicted depth map from the depth foundation model to initialize the recurrent disparity map. The depth foundation model usually outputs the relative depth, which has unknown scales and shifts to the disparity. To make reasonable initialization, we should convert the predicted depth proportional to the true disparity. Consider that the disparity amplitude of an image is usually proportional to its width. Suppose that the depth map predicted by the depth model is  $\mathbf{z}$ , we obtain the initialized disparity map with the following formula,

$$\mathbf{d}_0 = \frac{\eta w * \mathbf{z}}{\max(\mathbf{z})} + \epsilon. \quad (4)$$

where  $\eta$  is the controlling ratio and  $\epsilon$  is a tiny positive bias to make scale update work when the predicted depth is zero.

### 3.6. Scale Update

As the initialized disparity map is scale-uncertain and inconsistent, we first use the scale update module to obtain an overall scale-certain and consistent disparity map recurrently. Supposed the disparity map in the last iteration is  $\mathbf{d}_{n-1}$ , the ConvGRU for SU would predict a dense scale map  $\mathbf{s}$ , the next disparity map is updated by,

$$\mathbf{d}_n = \mathbf{s} \cdot \mathbf{d}_{n-1}. \quad (5)$$

It is insufficient to feed the ConvGRU with the sampled correlation feature from the pyramid lookup to recover the scale, as its search range is limited. To this end, we design the scale lookup (SL) that retrieves correlation in a scaling manner. The SL is only applied to the finest correlation volume. Given the disparity estimation  $\mathbf{d}_n$ , we multiply it with a series of pre-defined scale factors  $s^m$ , then the correlation value corresponding to the scaled disparity  $s^m \mathbf{d}_n$ ,  $\mathbf{C}^1(s^m \mathbf{d}_n)$  will be sampled. To, make the sampled correlation map more robust to local noise, the correlation of two nearby pixels,  $\mathbf{C}^1(s^m \mathbf{d}_n - 1)$  and  $\mathbf{C}^1(s^m \mathbf{d}_n + 1)$  will also be retrieved. To enable the SL to reach the whole image pixels, the ratio  $\eta$  in Eqn. 4 is set to  $1/2$ , and the scale factors are set to  $\{1, 2, 4, 6, 8, 10, 12, 16\}/8$ . In total, 24 correlation values are retrieved in the scale lookup.

### 3.7. Loss Function

Suppose that the total number of update iterations is  $N$  and the set of predicted disparity maps are  $\{\mathbf{d}_1, \dots, \mathbf{d}_N\}$ . We follow [29, 47] to supervise those predictions in exponentially increasing weights with the ground truth disparity  $\mathbf{d}_{gt}$ . The loss is computed as,

$$\mathcal{L} = \sum_{i=n}^N \gamma^{N-n} \|\mathbf{d}_{gt} - \mathbf{d}_n\|_1, \quad \text{where } \gamma = 0.9. \quad (6)$$

## 4. Experiments

In this section, we verify the effectiveness of our model via synthetic-to-real zero-shot generalization and popular online stereo benchmarks. We put the ablation study on the supplementary, where the effectiveness and complexities of the proposed components, the combined encoders' design choices, and the iterations of scale updates will be discussed.

### 4.1. Implementation Details

We implement our DEFOM-Stereo using the PyTorch framework on NVIDIA RTX 4090 GPUs. We use a simplified REAF-Stereo that reduces the correlation pyramid level from 4 to 2 as our baseline model and add the proposed components to construct DEFOM-Stereo. For all training, we use

Methods	KITTI		Middlebury			ETH3D
	2012	2015	full	half	quarter	
DSMNet [66]	6.2	6.5	21.8	13.8	8.1	6.2
RAFT-Stereo [29]	4.35	5.74	18.33	12.59	9.36	3.28
DLNR [69]	9.09	16.05	16.78	9.82	7.82	22.99
IGEV-Stereo [56]	5.14	6.03	36.41	13.36	8.82	4.05
Selective-RAFT [52]	5.19	6.68	33.98	13.25	8.66	4.36
Selective-IGEV [52]	5.65	6.05	35.82	13.27	9.82	6.07
NMRF-Stereo [16]	4.23	5.52	44.30	15.88	7.49	3.80
Mocha-Stereo [8]	4.85	5.93	35.34	11.49	7.39	3.89
DEFOM-Stereo (ViT-S)	4.29	5.29	14.70	6.76	6.38	2.61
DEFOM-Stereo (ViT-L)	3.76	4.99	11.95	5.91	5.65	2.35

Table 1. **Synthetic to real generalization.** All models are trained on Scene Flow. Standard thresholding error rates are used: 3-pixel for KITTI, 2-pixel for Middlebury, and 1-pixel for ETH3D.

the AdamW optimizer, a one-cycle learning rate schedule with a learning rate of  $2e-4$  and a batch size of 8. The total number of update iterations (SU+DU) is set to 18 in training and 32 in evaluation. The number of SU is fixed to 8 during both training and evaluation.  $\epsilon$  in Eqn. 4 is set to 0.05.

### 4.2. Zero-shot Generalization

We first pre-train our model on Scene Flow for 200k with a crop size of  $320 \times 736$ . We follow [29, 66] to examine the zero-shot generalization ability of our model on the train-sets of four realistic datasets, and for Middlebury, we consider the three available resolutions. For Middlebury and ETH3D, we count all valid pixels instead of only those in non-occluded regions, as this allows for a more comprehensive test. For other methods, we perform such evaluation using the provided pre-trained model on Scene Flow. The results are shown in Tab. 1. We found that despite later recurrent optimization-based methods [8, 52, 56, 69] have improved on pre-training fitting on Scene Flow compared with RAFT-Stereo, fail to make comprehensive progress on zero-shot evaluation. For example, DLNR [69], shows a clear performance gain on Middlebury with different resolutions, but significantly increases the error rate on KITTI and ETH3D. For the recent neural Markov random field-based model, NMRF-Stereo [16], progress can be seen only in driving datasets KITTI 2012 and 2015, and Middlebury-quarter.

For DEFOM-Stereo, we present two variants with two different ViT backbone sizes, large (ViT-L) and small (ViT-S). Both model variants achieve improvement in all these realistic datasets. In particular, compared with RAFT-Stereo, DEFOM-Stereo (ViT-L) reduces the error rates, by about 13% on KITTI, by 28% on ETH3D, and by over a half on Middlebury. Furthermore, on the high-resolution Middlebury-full, DEFOM-Stereo can even reduce the error rate by over 29% compared with the previous best model (DLNR), demonstrating the strong generalization ability with high-resolution even pre-trained at low resolution.

Method	KITTI-2012						KITTI-2015						Run-time (s)
	non-occluded			all			non-occluded			all			
	Bad 2.0	Bad 3.0	AvgErr	Bad 2.0	Bad 3.0	AvgErr	D1-bg	D1-fg	D1-all	D1-bg	D1-fg	D1-all	
ACVNet [55]	1.83	1.13	0.4	2.34	1.47	0.5	1.26	2.84	1.52	1.37	3.07	1.65	0.2
PCWNet [44]	1.69	1.04	0.4	2.18	1.37	0.5	1.26	2.93	1.53	1.37	3.16	1.67	0.44
LaC + GANet [30]	1.72	1.05	0.4	2.26	1.42	0.5	1.26	2.64	1.49	1.44	2.83	1.67	1.8
CREStereo [25]	1.72	1.14	0.4	2.18	1.46	0.5	1.33	2.60	1.54	1.45	2.86	1.69	0.40
IGEV-Stereo [56]	1.71	1.12	0.4	2.17	1.44	0.4	1.27	2.62	1.49	1.38	2.67	1.59	0.18
MC-Stereo [14]	1.55	1.04	0.4	1.99	1.34	0.4	1.24	2.55	1.46	1.36	2.51	1.55	0.40
DN+ACVNet [68]	1.64	1.02	0.4	2.20	1.41	0.5	1.21	2.62	1.44	1.32	2.95	1.60	0.24
Selective-IGEV [52]	1.59	1.07	0.4	2.05	1.38	0.4	1.22	2.55	1.44	1.33	2.61	1.55	0.24
NMRF-Stereo [16]	1.59	1.01	0.4	2.07	1.35	0.4	1.18	2.90	1.46	1.28	3.07	1.57	0.09
LoS [26]	1.69	1.10	0.4	2.12	1.38	0.4	1.29	2.66	1.52	1.42	2.81	1.65	0.19
GANet+ADL [60]	1.52	0.98	0.4	2.01	1.29	0.5	1.24	<b>2.18</b>	1.40	1.38	2.38	1.55	0.67
MoCha-Stereo [8]	1.64	1.06	0.4	2.07	1.36	0.4	1.24	2.42	1.44	1.36	2.43	1.53	0.33
IGEV++ [57]	1.56	1.04	0.4	2.03	1.36	0.4	1.20	2.54	1.42	1.31	2.54	1.51	0.28
StereoBase [18]	1.54	1.00	0.4	1.95	1.26	0.4	1.17	<u>2.23</u>	<u>1.35</u>	1.28	<u>2.26</u>	<u>1.44</u>	0.29
ViTAStereo [31]	<u>1.46</u>	<b>0.93</b>	0.4	<u>1.80</u>	<b>1.16</b>	0.4	<b>1.12</b>	2.90	1.41	<b>1.21</b>	2.99	1.50	0.22
DEFOM-Stereo (Ours)	<b>1.43</b>	<u>0.94</u>	<b>0.3<sup>1</sup></b>	<b>1.79</b>	<u>1.18</u>	0.4	<u>1.15</u>	2.24	<b>1.33<sup>1</sup></b>	<u>1.25</u>	<b>2.23<sup>1</sup></b>	<b>1.41<sup>1</sup></b>	0.30

Table 2. **Quantitative evaluation on KITTI 2012 and KITTI 2015.** <sup>1</sup> 1<sup>st</sup>-ranking metric on the leaderboards at the time of writing.

Method	Middlebury						ETH3D					
	non-occluded			all			non-occluded			all		
	Bad 2.0	AvgErr	RMS	Bad 2.0	AvgErr	RMS	Bad 1.0	AvgErr	RMS	Bad 1.0	AvgErr	RMS
HITNet [46]	6.46	1.71	9.97	12.8	3.29	14.5	2.79	0.20	0.46	3.11	0.22	0.55
RAFT-Stereo [29]	4.74	1.27	8.41	9.37	2.71	12.6	2.44	0.18	0.36	2.60	0.19	0.42
CREStereo [25]	3.71	1.15	7.70	8.13	2.10	10.5	0.98	0.13	0.28	1.09	0.14	0.31
CroCo-Stereo [53]	7.29	1.76	8.91	11.1	2.36	10.6	0.99	0.14	0.30	1.14	0.15	0.35
GMStereo [59]	7.14	1.31	<u>6.45</u>	11.7	1.89	<u>8.03</u>	1.83	0.19	0.38	2.07	0.21	0.44
IGEV-Stereo [56]	4.83	2.89	12.8	8.16	3.64	15.1	1.12	0.14	0.34	1.51	0.20	0.86
DLNR [69]	3.20	1.06	7.78	6.98	1.91	10.2	-	-	-	-	-	-
LoS [26]	4.20	1.12	6.99	8.03	1.75	8.78	<u>0.91</u>	0.14	0.31	<u>1.03</u>	<u>0.15</u>	<u>0.34</u>
IGEV++ [57]	3.23	0.97	7.23	7.19	1.83	10.2	1.14	0.13	0.34	1.58	0.19	0.74
Selective-IGEV [52]	<u>2.51</u>	<u>0.91</u>	7.26	<u>6.04</u>	<u>1.54</u>	9.26	1.23	<u>0.12</u>	<u>0.29</u>	1.56	<u>0.15</u>	0.57
DEFOM-Stereo (Ours)	<b>2.39<sup>2</sup></b>	<b>0.79<sup>1</sup></b>	<b>5.81<sup>1</sup></b>	<b>5.02<sup>1</sup></b>	<b>1.27<sup>1</sup></b>	<b>7.73<sup>1</sup></b>	<b>0.70<sup>4</sup></b>	<b>0.11<sup>3</sup></b>	<b>0.22<sup>2</sup></b>	<b>0.78<sup>2</sup></b>	<b>0.11<sup>1</sup></b>	<b>0.26<sup>1</sup></b>

Table 3. **Quantitative evaluation on ETH3D and Middlebury benchmarks.** The superscript number on the metric of our model represents its ranking on the leaderboards at the time of writing.

### 4.3. Benchmark Comparisons

In this section, we present the results of our models against the previous SOTA methods on the typical stereo leaderboards of KITTI 2012 [15], KITTI 2015 [33], Middlebury [41], and ETH3D [42].

**KITTI.** For submitting the benchmarks of KITTI, we first finetune the Scene Flow pre-trained model 50k on the mixed dataset of KITTI 2012 [15], KITTI 2015 [33] and virtual KITTI 2 [5], where the samples of KITTI 2012 and KITTI 2015 are augmented to take account for 50% of the mixed dataset. On the official leaderboards, our model obtains many 1<sup>st</sup> ranking metrics among all the submissions at the time of paper writing, including Avg-Noc, Avg-All(Reflective), and 2-Out-All(Reflective) for KITTI 2012 and D1-fg, D1-all and D1-all(Noc) for KITTI. The results compared with public SOTA methods are shown in Tab. 2. Compared with StereoBase [18] and ViTAStereo [31], which perform best among the public methods, our model still shows clear advantages. For ViTAStereo, while its performance on KITTI 2012 is comparable, it underperforms our model overall on KITTI 2015, particularly in the foreground region. For StereoBase,

our model is slightly better than it on KITTI 2015 but outperforms it on KITTI 2012 by a clear margin, for example, 8.2% reduction of Bad 2.0 (all).

**Middlebury.** Following previous work [25, 52, 59], we first finetune the Scene Flow pre-trained model on a mixed dataset of Tartan Air [51], CREStereo Dataset [25], Scene Flow [32], Falling Things [49], InStereo2k [1], CARLA HR-VS [61] and Middlebury [41] datasets with a crop size of 384 × 512 for 200k steps. Next, we perform finetuning on a combination of CREStereo Dataset [25], Falling Things [49], InStereo2k [1], CARLA HR-VS [61] and Middlebury [41] datasets with a crop size of 512 × 768 for 100k steps. Our model delivers outstanding performance on the Middlebury benchmark: DEFOM-Stereo ranks first in 7 out of 10 accuracy metrics when evaluated on all valid pixels, and in 3 out of 10 metrics when evaluated on non-occluded valid pixels. The comparison with public methods is shown on the left of Tab. 3. Quantitatively, our model outperforms the best-published method, Selective-IGEV [52] by 4.8% on Bad 2.0 (noc) and 16.9% on Bad 2.0 (all). The gap in performance gain between evaluation on non-occluded and all valid pixels

Method	KITTI-2015						Middlebury				ETH3D			
	non-occluded			all			non-occluded		all		non-occluded		all	
	D1-bg	D1-fg	D1-all	D1-bg	D1-fg	D1-all	Bad 2.0	AvgErr	Bad 2.0	AvgErr	Bad 1.0	AvgErr	Bad 1.0	AvgErr
<b>Robust Vision Challenge 2018</b>														
DN-CSS_ROB [19]	2.23	4.96	2.68	2.39	5.71	2.94	22.8	4.04	28.3	5.48	2.69	0.22	3.00	0.24
iResNet_ROB [28]	2.10	3.96	2.40	2.27	4.89	2.71	24.8	4.51	31.7	6.56	4.23	0.25	4.67	0.27
<b>Robust Vision Challenge 2020</b>														
NLCA_NET_v2_RVC [38]	1.36	3.49	1.71	<u>1.51</u>	3.97	1.92	10.4	3.89	16.4	5.60	3.84	0.27	4.11	0.29
CFNet_RVC [43]	1.50	3.03	1.76	1.65	3.53	1.96	10.1	3.49	16.1	5.07	3.31	0.24	3.70	0.26
<b>Robust Vision Challenge 2022</b>														
iRaftStereo_RVC [20, 29]	1.76	2.94	1.95	1.88	<u>3.03</u>	2.07	8.07	1.71	13.3	2.90	1.62	0.16	1.88	0.17
CREStereo++_RVC [21]	1.43	3.36	1.75	1.55	3.53	1.88	<u>4.68</u>	1.28	9.46	<u>2.20</u>	1.59	0.15	1.70	0.16
More Recently														
UCFNet_RVC [45]	<u>1.41</u>	<u>2.93</u>	<u>1.66</u>	1.57	3.33	1.86	10.7	3.74	16.7	5.96	3.09	0.24	3.37	0.25
LoS_RVC [26]	1.46	2.95	1.71	1.58	3.08	<u>1.83</u>	5.14	<u>1.57</u>	<u>9.30</u>	2.36	<u>1.26</u>	<b>0.13</b>	<u>1.47</u>	<u>0.14</u>
DEFOM-Stereo_RVC	<b>1.32</b>	<b>2.67</b>	<b>1.54</b>	<b>1.42</b>	<b>2.68</b>	<b>1.63</b>	<b>3.28</b>	<b>0.97</b>	<b>6.90</b>	<b>1.61</b>	<b>0.98</b>	<b>0.13</b>	<b>1.09</b>	<b>0.13</b>

Table 4. **Robust Vision Challenge.** Comparison with winners and runner-ups of the Robust Vision Challenges 2018, 2020, and 2022, and more recent RVC models, UCFNet\_RVC [45] and LoS\_RVC [26].

suggests our model brings more improvement on occluded regions than non-occluded ones.

**ETH3D.** Similarly, in finetuning for ETH3D benchmark submission, we follow CREStereo [25], GMStereo [59] and Selective-Stereo [52]. The first finetuning was performed on the mixture of Tartan Air [51], CREStereo Dataset [25], Scene Flow [32], Sintel Stereo [4], InStereo2k [1] and ETH3D [42] with a crop size of  $384 \times 512$  for 300k steps. Then, another finetuning was only performed on CREStereo Dataset [25], InStereo2k [1] and ETH3D [42] datasets for 90k steps. Still, our model obtains the top performance on the ETH3D leaderboard and has several 1<sup>st</sup>-ranking metrics when evaluated with all valid masks. The results of previous methods and ours are listed on the right of Tab. 3. Compared with LoS [26], which has the best Bad 1.0 among the published methods, our model reduces Bad 1.0 (noc) by 23.1%, and Bad 1.0 (all) by 24.3%, a noteworthy progress.

#### 4.4. Robust Vision Challenge

In this section, we examine our model under the setting of the **Robust Vision Challenge**. Different from the previous section which focuses on the performance of only a specific domain, this section will compare models’ performance across many domains with the same model parameters. The three stereo benchmarks in RVC are KITTI 2015 [33], Middlebury [41], and ETH3D [42]. The RVC evaluation reveals the joint generalization capability of models.

To provide a lightweight RVC model, we use the small ViT backbone, *i.e.*, ViT-S, instead of the ViT-L for individual benchmarks. Following the training strategy for individual benchmarks in Sec. 4.3, we use a mixture of synthetic and realistic datasets for finetuning the pre-trained model on Scene Flow [32]. Specifically, we perform the first finetuning for 200k on a mixture of synthetic datasets, including Tartan Air [51], CREStereo Dataset [25], Scene Flow [32], Falling Things [49], CARLA HR-VS [61], Sintel Stereo [4], virtual KITTI 2 [5], IRS [50] and 3D Ken Burns [34]. Next, we

add the training split of realistic datasets, including KITTI 2012 [15], KITTI 2015 [33], Middlebury [41], ETH3D [42], InStereo2k [1] and Booster dataset [64] for additional 100k finetuning. Following LoS\_RVC [26], we augment the KITTI 2012 and KITTI 2015 to make up half of the mixed dataset in the final 20k finetuning. During all the finetuning, the batch size and input crop size are set to 8 and  $384 \times 768$ .

We compare our RVC model with previous RVC models in Tab 4. The previous models include the winners and runners-up of the Robust Vision Challenges in 2018, 2020, and 2022, and two more recent RVC models. Our DEFOM-Stereo\_RVC achieves the best results among all RVC models on three benchmarks and shows a significant performance gap over previous models on the KITTI 2015 and Middlebury. For example, DEFOM-Stereo\_RVC outperforms the next-best CREStereo++\_RVC and LoS\_RVC, by around 25% on the listed four typical metrics on Middlebury.

#### 5. Conclusion

In this paper, we have presented a novel recurrent stereo-matching framework by integrating a foundation model, Depth Anything V2. We developed simple and effective techniques to harness its powerful pre-retrained feature representations, enhancing both the matching feature and context extraction in recurrent stereo matching. Despite its universal robustness, we still observed that when evaluating its depth estimate on standard stereo datasets, scale inconsistency across different image regions exists, particularly for the synthetic ones. To address this in utilizing the depth estimate to facilitate the recurrent disparity, we introduced a scale update module designed to iteratively recover dense scale maps. Our model has been verified to have significant effectiveness on both in-domain fitting and zero-shot generation and achieves leading performance in standard stereo benchmarks. We believe that the work presented in this paper contributes to advancing the exploration of the foundational model for stereo matching.

## References

- [1] Wei Bao, Wei Wang, Yuhua Xu, Yulan Guo, Siyu Hong, and Xiaohu Zhang. Instereo2k: a large real dataset for stereo matching in indoor scenes. *Science China Information Sciences*, 63:1–11, 2020. [7](#), [8](#)
- [2] Luca Bartolomei, Fabio Tosi, Matteo Poggi, and Stefano Mattoccia. Stereo anywhere: Robust zero-shot deep stereo matching even where either stereo or mono fail. *arXiv preprint arXiv:2412.04472*, 2024. [3](#)
- [3] Reiner Birkel, Diana Wofk, and Matthias Müller. Midas v3.1 – a model zoo for robust monocular relative depth estimation. *arXiv preprint arXiv:2307.14460*, 2023. [1](#), [3](#)
- [4] Daniel J Butler, Jonas Wulff, Garrett B Stanley, and Michael J Black. A naturalistic open source movie for optical flow evaluation. In *Computer Vision—ECCV 2012: 12th European Conference on Computer Vision, Florence, Italy, October 7–13, 2012, Proceedings, Part VI 12*, pages 611–625. Springer, 2012. [8](#)
- [5] Yohann Cabon, Naila Murray, and Martin Humenberger. Virtual kitti 2. *arXiv preprint arXiv:2001.10773*, 2020. [7](#), [8](#)
- [6] Jia-Ren Chang and Yong-Sheng Chen. Pyramid stereo matching network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5410–5418, 2018. [1](#), [2](#), [3](#)
- [7] Tianyu Chang, Xun Yang, Tianzhu Zhang, and Meng Wang. Domain generalized stereo matching via hierarchical visual transformation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9559–9568, 2023. [3](#)
- [8] Ziyang Chen, Wei Long, He Yao, Yongjun Zhang, Bingshu Wang, Yongbin Qin, and Jia Wu. Mocha-stereo: Motif channel attention network for stereo matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 27768–27777, 2024. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#)
- [9] Junda Cheng, Longliang Liu, Gangwei Xu, Xianqi Wang, Zhaoxing Zhang, Yong Deng, Jinliang Zang, Yurui Chen, Zhipeng Cai, and Xin Yang. Monster: Marry monodepth to stereo unleashes power. *arXiv preprint arXiv:2501.08643*, 2025. [3](#)
- [10] WeiQin Chuah, Ruwan Tennakoon, Reza Hoseinnezhad, Alireza Bab-Hadiashar, and David Suter. Itsa: An information-theoretic approach to automatic shortcut avoidance and domain generalization in stereo matching networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13022–13032, 2022. [3](#)
- [11] Alexey Dosovitskiy. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. [1](#)
- [12] Sébastien Drouyer, Serge Beucher, Michel Bilodeau, Maxime Moreaud, and Loïc Sorbier. Sparse stereo disparity map densification using hierarchical image segmentation. In *International symposium on mathematical morphology and its applications to signal and image processing*, pages 172–184. Springer, 2017. [2](#)
- [13] Mohamed El Banani, Amit Raj, Kevis-Kokitsi Maninis, Abhishek Kar, Yuanzhen Li, Michael Rubinstein, Deqing Sun, Leonidas Guibas, Justin Johnson, and Varun Jampani. Probing the 3d awareness of visual foundation models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21795–21806, 2024. [3](#)
- [14] Miaojie Feng, Junda Cheng, Hao Jia, Longliang Liu, Gangwei Xu, and Xin Yang. Mc-stereo: Multi-peak lookup and cascade search range for stereo matching. In *2024 International Conference on 3D Vision (3DV)*, pages 344–353. IEEE, 2024. [2](#), [3](#), [7](#)
- [15] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE conference on computer vision and pattern recognition*, pages 3354–3361. IEEE, 2012. [1](#), [7](#), [8](#)
- [16] Tongfan Guan, Chen Wang, and Yun-Hui Liu. Neural markov random field for stereo matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5459–5469, 2024. [3](#), [6](#), [7](#)
- [17] Xiaoyang Guo, Kai Yang, Wukui Yang, Xiaogang Wang, and Hongsheng Li. Group-wise correlation stereo network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3273–3282, 2019. [1](#), [2](#), [3](#)
- [18] Xianda Guo, Chenming Zhang, Juntao Lu, Yiqi Wang, Yiqun Duan, Tian Yang, Zheng Zhu, and Long Chen. Openstereo: A comprehensive benchmark for stereo matching and strong baseline. *arXiv preprint arXiv:2312.00343*, 2023. [7](#)
- [19] Eddy Ilg, Tonmoy Saikia, Margret Keuper, and Thomas Brox. Occlusions, motion and depth boundaries with a generic network for disparity, optical flow or scene flow estimation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 614–630, 2018. [8](#)
- [20] Hualie Jiang, Rui Xu, and Wenjie Jiang. An improved raft-stereo trained with a mixed dataset for the robust vision challenge 2022. *arXiv preprint arXiv:2210.12785*, 2022. [8](#)
- [21] Junpeng Jing, Jiankun Li, Pengfei Xiong, Jiangyu Liu, Shuaicheng Liu, Yichen Guo, Xin Deng, Mai Xu, Lai Jiang, and Leonid Sigal. Uncertainty guided adaptive warping for robust and efficient stereo matching. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3318–3327, 2023. [2](#), [3](#), [8](#)
- [22] Bingxin Ke, Anton Obukhov, Shengyu Huang, Nando Metzger, Rodrigo Caye Daudt, and Konrad Schindler. Repurposing diffusion-based image generators for monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9492–9502, 2024. [1](#), [3](#)
- [23] Alex Kendall, Hayk Martirosyan, Saumitro Dasgupta, Peter Henry, Ryan Kennedy, Abraham Bachrach, and Adam Bry. End-to-end learning of geometry and context for deep stereo regression. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 66–75, 2017. [1](#), [2](#)
- [24] Hamid Laga, Laurent Valentin Jospin, Farid Boussaid, and Mohammed Bennamoun. A survey on deep learning techniques for stereo-based depth estimation. *IEEE transactions on pattern analysis and machine intelligence*, 44(4):1738–1764, 2020. [1](#)
- [25] Jiankun Li, Peisen Wang, Pengfei Xiong, Tao Cai, Ziwei Yan, Lei Yang, Jiangyu Liu, Haoqiang Fan, and Shuaicheng Liu. Practical stereo matching via cascaded recurrent network

- with adaptive correlation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16263–16272, 2022. 1, 2, 3, 7, 8
- [26] Kunhong Li, Longguang Wang, Ye Zhang, Kaiwen Xue, Shunbo Zhou, and Yulan Guo. Los: Local structure-guided stereo matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19746–19756, 2024. 1, 2, 3, 7, 8
- [27] Zhaoshuo Li, Xingtong Liu, Nathan Drenkow, Andy Ding, Francis X Creighton, Russell H Taylor, and Mathias Unberath. Revisiting stereo depth estimation from a sequence-to-sequence perspective with transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6197–6206, 2021. 3
- [28] Zhengfa Liang, Yiliu Feng, Yulan Guo, Hengzhu Liu, Wei Chen, Linbo Qiao, Li Zhou, and Jianfeng Zhang. Learning for disparity estimation through feature constancy. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2811–2820, 2018. 8
- [29] Lahav Lipson, Zachary Teed, and Jia Deng. Raft-stereo: Multilevel recurrent field transforms for stereo matching. In *2021 International Conference on 3D Vision (3DV)*, pages 218–227. IEEE, 2021. 1, 2, 3, 4, 5, 6, 7, 8
- [30] Biyang Liu, Huimin Yu, and Yangqi Long. Local similarity pattern and cost self-reassembling for deep stereo matching networks. In *Proceedings of the AAAI conference on artificial intelligence*, pages 1647–1655, 2022. 7
- [31] Chuang-Wei Liu, Qijun Chen, and Rui Fan. Playing to vision foundation model’s strengths in stereo matching. *IEEE Transactions on Intelligent Vehicles*, 2024. DOI:10.1109/TIV.2024.3467287. 1, 2, 7
- [32] N. Mayer, E. Ilg, P. Häusser, P. Fischer, D. Cremers, A. Dosovitskiy, and T. Brox. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. arXiv:1512.02134. 2, 7, 8
- [33] Moritz Menze and Andreas Geiger. Object scene flow for autonomous vehicles. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3061–3070, 2015. 1, 7, 8
- [34] Simon Niklaus, Long Mai, Jimei Yang, and Feng Liu. 3d ken burns effect from a single image. *ACM Transactions on Graphics (ToG)*, 38(6):1–15, 2019. 8
- [35] Maxime Oquab, Timothée Darcet, Theo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Russell Howes, Po-Yao Huang, Hu Xu, Vasu Sharma, Shang-Wen Li, Wojciech Galuba, Mike Rabbat, Mido Assran, Nicolas Ballas, Gabriel Synnaeve, Ishan Misra, Herve Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. DINOv2: Learning robust visual features without supervision, 2023. 1, 3
- [36] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. 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. 1, 3
- [37] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 12179–12188, 2021. 1, 3
- [38] Zhibo Rao, Yuchao Dai, Zhelun Shen, and Renjie He. Rethinking training strategy in stereo matching. *IEEE Transactions on Neural Networks and Learning Systems*, 2022. 8
- [39] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 3
- [40] Daniel Scharstein and Richard Szeliski. A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *International journal of computer vision*, 47:7–42, 2002. 1
- [41] Daniel Scharstein, Heiko Hirschmüller, York Kitajima, Greg Krathwohl, Nera Nešić, Xi Wang, and Porter Westling. High-resolution stereo datasets with subpixel-accurate ground truth. In *German conference on pattern recognition*, pages 31–42. Springer, 2014. 1, 7, 8
- [42] Thomas Schops, Johannes L Schonberger, Silvano Galliani, Torsten Sattler, Konrad Schindler, Marc Pollefeys, and Andreas Geiger. 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*, pages 3260–3269, 2017. 1, 7, 8
- [43] Zhelun Shen, Yuchao Dai, and Zhibo Rao. Cfnets: Cascade and fused cost volume for robust stereo matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13906–13915, 2021. 8
- [44] Zhelun Shen, Yuchao Dai, Xibin Song, Zhibo Rao, Dingfu Zhou, and Liangjun Zhang. Pcw-net: Pyramid combination and warping cost volume for stereo matching. In *European conference on computer vision*, pages 280–297. Springer, 2022. 7
- [45] Zhelun Shen, Xibin Song, Yuchao Dai, Dingfu Zhou, Zhibo Rao, and Liangjun Zhang. Digging into uncertainty-based pseudo-label for robust stereo matching. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. 8
- [46] Vladimir Tankovich, Christian Hane, Yinda Zhang, Adarsh Kowdle, Sean Fanello, and Sofien Bouaziz. Hitnet: Hierarchical iterative tile refinement network for real-time stereo matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14362–14372, 2021. 7
- [47] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, pages 402–419. Springer, 2020. 3, 5, 6
- [48] Fabio Tosi, Luca Bartolomei, and Matteo Poggi. A survey on deep stereo matching in the twenties. *arXiv preprint arXiv:2407.07816*, 2024. 1
- [49] Jonathan Tremblay, Thang To, and Stan Birchfield. Falling things: A synthetic dataset for 3d object detection and pose estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 2038–2041, 2018. 7, 8

- [50] Qiang Wang, Shizhen Zheng, Qingsong Yan, Fei Deng, Kaiyong Zhao, and Xiaowen Chu. Irs: A large naturalistic indoor robotics stereo dataset to train deep models for disparity and surface normal estimation. In *2021 IEEE International Conference on Multimedia and Expo (ICME)*, pages 1–6. IEEE, 2021. 8
- [51] Wenshan Wang, DeLong Zhu, Xiangwei Wang, Yaoyu Hu, Yuheng Qiu, Chen Wang, Yafei Hu, Ashish Kapoor, and Sebastian Scherer. Tartanair: A dataset to push the limits of visual slam. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4909–4916. IEEE, 2020. 7, 8
- [52] Xianqi Wang, Gangwei Xu, Hao Jia, and Xin Yang. Selective-stereo: Adaptive frequency information selection for stereo matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19701–19710, 2024. 1, 2, 3, 5, 6, 7, 8
- [53] Philippe Weinzaepfel, Thomas Lucas, Vincent Leroy, Yohann Cabon, Vaibhav Arora, Romain Brégier, Gabriela Csurka, Leonid Antsfeld, Boris Chidlovskii, and Jérôme Revaud. Croco v2: Improved cross-view completion pre-training for stereo matching and optical flow. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 17969–17980, 2023. 3, 7
- [54] Bowen Wen, Matthew Trepte, Joseph Aribido, Jan Kautz, Orazio Gallo, and Stan Birchfield. Foundationstereo: Zero-shot stereo matching. *arXiv preprint arXiv:2501.09898*, 2025. 3
- [55] Gangwei Xu, Junda Cheng, Peng Guo, and Xin Yang. Attention concatenation volume for accurate and efficient stereo matching. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12981–12990, 2022. 7
- [56] Gangwei Xu, Xianqi Wang, Xiaohuan Ding, and Xin Yang. Iterative geometry encoding volume for stereo matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21919–21928, 2023. 1, 2, 3, 5, 6, 7
- [57] Gangwei Xu, Xianqi Wang, Zhaoxing Zhang, Junda Cheng, Chunyuan Liao, and Xin Yang. Igev++: Iterative multi-range geometry encoding volumes for stereo matching. *arXiv preprint arXiv:2409.00638*, 2024. 7
- [58] Haofei Xu and Juyong Zhang. Aanet: Adaptive aggregation network for efficient stereo matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1959–1968, 2020. 1, 2
- [59] Haofei Xu, Jing Zhang, Jianfei Cai, Hamid Reza Tofighi, Fisher Yu, Dacheng Tao, and Andreas Geiger. Unifying flow, stereo and depth estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. 3, 7, 8
- [60] Peng Xu, Zhiyu Xiang, Chengyu Qiao, Jingyun Fu, and Tianyu Pu. Adaptive multi-modal cross-entropy loss for stereo matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5135–5144, 2024. 7
- [61] Gengshan Yang, Joshua Manela, Michael Happold, and Deva Ramanan. Hierarchical deep stereo matching on high-resolution images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5515–5524, 2019. 7, 8
- [62] Lihe Yang, Bingyi Kang, Zilong Huang, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything: Unleashing the power of large-scale unlabeled data. In *CVPR*, 2024. 1, 3
- [63] Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything v2. *arXiv:2406.09414*, 2024. 1, 2, 3, 4
- [64] Pierluigi Zama Ramirez, Fabio Tosi, Matteo Poggi, Samuele Salti, Luigi Di Stefano, and Stefano Mattoccia. Open challenges in deep stereo: the booster dataset. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2022. CVPR. 8
- [65] Feihu Zhang, Victor Prisacariu, Ruigang Yang, and Philip HS Torr. Ga-net: Guided aggregation net for end-to-end stereo matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 185–194, 2019. 2, 3
- [66] Feihu Zhang, Xiaojuan Qi, Ruigang Yang, Victor Prisacariu, Benjamin Wah, and Philip Torr. Domain-invariant stereo matching networks. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, pages 420–439. Springer, 2020. 1, 3, 6
- [67] Jiawei Zhang, Xiang Wang, Xiao Bai, Chen Wang, Lei Huang, Yimin Chen, Lin Gu, Jun Zhou, Tatsuya Harada, and Edwin R Hancock. Revisiting domain generalized stereo matching networks from a feature consistency perspective. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13001–13011, 2022. 3
- [68] Jiawei Zhang, Lei Huang, Xiao Bai, Jin Zheng, Lin Gu, and Edwin Hancock. Exploring the usage of pre-trained features for stereo matching. *International Journal of Computer Vision*, pages 1–22, 2024. 2, 7
- [69] Haoliang Zhao, Huizhou Zhou, Yongjun Zhang, Jie Chen, Yitong Yang, and Yong Zhao. High-frequency stereo matching network. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1327–1336, 2023. 1, 2, 3, 6, 7