This CVPR paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

MonSter: Marry Monodepth to Stereo Unleashes Power

Junda Cheng¹, Longliang Liu¹, Gangwei Xu¹, Xianqi Wang¹, Zhaoxing Zhang¹, Yong Deng², Jinliang Zang², Yurui Chen², Zhipeng Cai³, Xin Yang^{1†} ¹ Huazhong University of Science and Technology ² Autel Robotics ³ Intel Labs {jundacheng, longliangl, gwxu, xianqiw, zzx, xinyang2014}@hust.edu.cn, czptc2h@gmail.com



Figure 1. Zero-shot generalization comparison: all models are trained on Scene Flow and tested directly on KITTI. Compared to the baseline IGEV [37], our method MonSter shows significant improvement in challenging regions such as reflective surfaces, textureless areas, fine structures, and distant objects.

Abstract

Stereo matching recovers depth from image correspondences. Existing methods struggle to handle ill-posed regions with limited matching cues, such as occlusions and textureless areas. To address this, we propose Mon-Ster, a novel method that leverages the complementary strengths of monocular depth estimation and stereo matching. MonSter integrates monocular depth and stereo matching into a dual-branch architecture to iteratively improve each other. Confidence-based guidance adaptively selects reliable stereo cues for monodepth scale-shift recovery. The refined monodepth is in turn guides stereo effectively at illposed regions. Such iterative mutual enhancement enables MonSter to evolve monodepth priors from coarse objectlevel structures to pixel-level geometry, fully unlocking the potential of stereo matching. As shown in Fig. 2, Mon-*Ster ranks* 1^{*st*} *across five most commonly used leaderboards* - SceneFlow, KITTI 2012, KITTI 2015, Middlebury, and ETH3D. Achieving up to 49.5% improvements (Bad 1.0 on ETH3D) over the previous best method. Comprehensive analysis verifies the effectiveness of MonSter in ill-posed regions. In terms of zero-shot generalization, MonSter significantly and consistently outperforms state-of-theart across the board. The code is publicly available at: https://github.com/Junda24/MonSter.

1. Introduction

Stereo matching estimates a disparity map from rectified stereo images, which can be subsequently converted into metric depth. It is the core of many applications such as self-driving, robotic navigation, and 3D reconstruction.

Deep learning based methods [5, 13, 15, 18, 33, 34, 39, 47, 49] have demonstrated promising performance on standard benchmarks. These methods can be roughly categorized into *cost filtering-based* methods and *iterative optimization-based* methods. Cost filtering-based methods [5, 13, 14, 20, 36, 40, 42] construct 3D/4D cost volume using CNN features, followed by a series of 2D/3D con-

[†]Corresponding author.



Figure 2. Leaderboard performance. Our method ranks 1st across 5 leaderboards, advancing SOTA by a large margin.

volutions for regularization and filtering to minimize mismatches. Iterative optimization-based methods [15, 18, 37, 48] initially construct an all-pairs correlation volume, then index a local cost to extract motion features, which guide the recurrent units (ConvGRUs) [9] to iteratively refine the disparity map.

Both types of methods essentially derive disparity from similarity matching, assuming visible correspondences in both images. This poses challenges in *ill-posed* regions with limited matching cues, e.g., occlusions, textureless areas, repetitive/thin structures, and distant objects with low pixel representation. Existing methods [5, 29, 31, 36, 41, 43, 48] address this issue by enhancing matching with stronger feature representations during feature extraction or cost aggregation: [17, 41] employ transformer as feature extractors; [36, 37] incorporate geometric information into the cost volume using attention mechanisms to strengthen matching information; [33, 48] further improve performance by introducing high-frequency information during the iterative refinement stage. However, these methods still struggle to fundamentally resolve the issue of mismatching, limiting their practical performance.

Unlike stereo matching, monocular depth estimation directly recovers 3D from a single image, which does not encounter the challenge of mismatching. While monocular depth provides complementary structural information for stereo, pre-trained models often yield relative depth with scale and shift ambiguities. As shown in Fig.3, the prediction of monodepth models differ heavily from the groundtruth. Even after global scale and shift alignment, substantial errors still persist, complicating pixel-wise fusion of monocular depth and stereo disparity. Based on these insights, we propose *MonSter*, a novel approach that decomposes stereo matching into monocular depth estimation and per-pixel scale-shift recovery, which fully combines the strengths of monocular and stereo algorithms and overcomes the limitations from the lack of matching cues.

MonSter constructs separate branches for monocular



Figure 3. Distance between GT disparity and the disparity derived from monocular depth [45] on KITTI dataset. The red line indicates identical disparity maps. (a): Without any alignment. (b) Align depth with GT using global scale and global shift values (same for all pixels). (c) The monocular disparity produced by MonSter, with per-pixel shift refinement. Even globally aligned with GT, SOTA monocular depth models still exhibit severe noise. Our method MonSter effectively addresses this issue by refining monocular depth with per-pixel shift, which fully unlocks the power of monocular depth priors for stereo matching.

depth estimation and stereo matching, and adaptively fuses them through Stereo Guided Alignment (SGA) and Mono Guided Refinement (MGR) modules. SGA first rescales monodepth into a "monocular disparity" by aligning globally with stereo disparity. Then it uses condition-guided GRUs to adaptively select reliable stereo cues for updating the per-pixel monocular disparity shift. Symmetric to SGA, MGR uses the optimized monocular disparity as the condition to adaptively refine the stereo disparity in regions where matching fails. Through multiple iterations, the two branches effectively complement each other: 1) Though beneficial at coarse object-level, directly and unidirectionally fuse monodepth into stereo suffers from scale-shift ambiguities, which often introduces noise in complex regions such as slanted or curved surfaces. Refining monodepth with stereo resolves this issue effectively, ensuring the robustness of MonSter. 2) The refined monodepth in turn provides strong guidance to stereo in challenging regions. E.g., the depth perception ability of stereo matching degrades with distance due to smaller pixel proportions and the increased matching difficulty. Monodepth models pretrained on large-scale datasets are less affected by such issues, which can effectively improve stereo disparity in the corresponding region.

Our main contributions can be summarized as follows:

- We propose a novel stereo-matching method MonSter, fully leveraging the pixel-level monocular depth priors to significantly improve the depth perception performance of stereo matching in ill-posed regions and fine structures.
- As shown in Fig. 2, MonSter ranks 1st across 5 widelyused leaderboards: KITTI 2012 [10], KITTI 2015 [21], Scene Flow [20], Middlebury [25] and ETH3D [26]. Advancing SOTA by up to 49.5%.

• Compared to SOTA methods, MonSter achieves the best zero-shot generalization consistently across diverse datasets. MonSter trained solely on synthetic data demonstrates strong performance across diverse real-world datasets (see Fig.1).

2. Related Work

Matching based Methods. Mainstream stereo matching methods recover disparity from matching costs. These methods can generally be divided into two categories: cost filtering-based methods and iterative optimization-based methods. Cost filtering-based methods [6, 7, 11, 13, 20, 27, 42] typically construct a 3D/4D cost volume using feature maps and subsequently employing 2D/3D CNN to filter the volume and derive the final disparity map. Constructing a cost volume with strong representational capacity and accurately regressing disparity from a noisy cost volume are key challenges. [5] proposes a stacked hourglass 3D CNN for better cost regularization. [13, 28, 36] propose the group-wise correlation volume, the attention concatenation volume, and the pyramid warping volume respectively to improve the expressiveness of the matching cost. Recently, a novel class of methods based on iterative optimization [3, 15, 18, 33, 37, 48] has achieved SOTA performance in both accuracy and efficiency. These methods use ConvGRUs to iteratively update the disparity by leveraging local cost values retrieved from the all-pairs correlation volume. Similarly, iterative optimization-based methods also primarily focus on improving the matching cost construction and the iterative optimization stages. CREStereo [15] proposes a cascaded recurrent network to update the disparity field in a coarse-to-fine manner. IGEV [37] proposes a Geometry Encoding Volume that encodes geometry and context information for a more robust matching cost. Both types of methods essentially derive depth from matching costs, which are inherently limited by ill-posed regions.

Stereo Matching with Structural Priors. Matching in ill-posed regions is challenging, previous methods tried to leverage structural priors to address this issue. EdgeStereo [29] enhances performance in edge regions by incorporating edge detection cues into the disparity estimation pipeline. SegStereo [43] utilizes semantic cues from a segmentation network as guidance for stereo matching, improving performance in textureless regions. However, semantic and edge cues only provide object-level priors, which are insufficient for pixel-level depth perception. Consequently, they are not effective in challenging scenes with large curved or slanted surfaces. Therefore, some methods incorporate monocular depth priors to provide per-pixel guidance for stereo matching through dense relative depth. CLStereo [46] introduces a monocular branch that serves as a contextual constraint, transferring geometric priors from the monocular to the stereo branch. Los [16] uses monocular depth as local

structural prior to generating the slant plane, which can explicitly leverage structure information for updating disparities. However, monocular depth estimation suffers from severe scale and shift ambiguities as shown in Fig.3. Directly using it as a structural prior to constrain stereo matching can introduce heavy noise. Our MonSter adaptively selects reliable stereo disparity to correct the scale and shift of monodepth, which fully leverages the monocular depth priors while avoiding noise, thereby significantly enhancing stereo performance in ill-posed regions.

3. Method

As shown in Fig. 4, MonSter consists of 1) a *monocular depth branch*, 2) a *stereo matching branch* and 3) a *mutual refinement module*. The two branches estimate initial monocular depth and the stereo disparity, which are fed into mutual refinement to iteratively improve each other.

3.1. Monocular and Stereo Branches

The monocular depth branch can leverage most monocular depth models to achieve non-trivial performance improvements (see Sec.4.5 for analysis). Our best empirical configuration uses pretrained DepthAnythingV2 [45] as the monocular depth branch, which uses DINOv2 [22] as the encoder and DPT [24] as the decoder. The stereo matching branch follows IGEV [37] to obtain the initial stereo disp, with modifications only to the feature extraction component, as shown in Fig. 4. To efficiently and fully leverage the pretrained monocular model, the stereo branch shares the ViT encoder in DINOv2 with the monocular branch, with parameters frozen to prevent the stereo matching training from affecting its generalization ability. Moreover, the ViT architecture extracts feature at a single resolution, while recent stereo matching methods commonly utilize multiscale features at four scales (1/32, 1/16, 1/8, and 1/4 of the original image resolution). To fully align with IGEV, we employ a stack of 2D convolutional layers, denoted as the feature transfer network, to downsample and transform the ViT feature into a collection of pyramid features $\mathcal{F} = \{F_0, F_1, F_2, F_3\}$, where $F_k \in \mathbb{R}^{\frac{H}{25-k} \times \frac{W}{25-k} \times c_k}$. We follow IGEV to construct Geometry Encoding Volume and use the same ConvGRUs for iterative optimization. To balance accuracy and efficiency, we perform only N_1 iterations to obtain a initial stereo disparity with reasonable quality.

3.2. Mutual Refinement

Once the initial (relative) monocular depth and stereo disparity are obtained, they are fed into the mutual refinement module to iteratively refine each other. We first perform *global scale-shift alignment*, which converts monocular depth into a disparity map and aligns it coarsely with stereo outputs. Then we iteratively perform a dual-branched refinement: *Stereo guided alignment (SGA)* leverages stereo



Figure 4. **Overview of MonSter.** MonSter consists of a monocular depth estimation branch, a stereo matching branch, and a mutual refinement module. It iteratively improves one branch with priors from the other, effectively resolving the ill-posedness in stereo matching.

cues to update the per-pixel shift of the monocular disparity; *Mono guided refinement (MGR)* leverages the aligned monocular prior to further refine stereo disparity.

Global Scale-Shift Alignment. Global Scale-Shift Alignment performs least squares optimization over a global scale s_G and a global shift t_G to coarsely align the inverse monocular depth with the stereo disparity:

$$s_G, t_G = \arg \min_{s_G, t_G} \sum_{i \in \Omega} \left(s_G \boldsymbol{D}_M(i) + t_G - \boldsymbol{D}_S^0(i) \right)^2$$
(1)
$$\boldsymbol{D}_M^0 = s_G \boldsymbol{D}_M + t_G,$$

where $D_M(i)$ and $D_S^0(i)$ are the inverse monocular depth and stereo disparity at the i-th pixel. Ω represents the region where stereo disparity values fall between 20% to 90% sorting in ascending order, which helps to filter unreliable regions such as the sky, extremely distant areas, and closerange outliers. Intuitively, this step converts the inverse monocular depth into a disparity map coarsely aligned with the stereo disparity, enabling effective mutual refinement in the same space. We call this aligned disparity map D_M^0 the *monocular disparity* in the remainder of the paper.

Stereo Guided Alignment (SGA). Though coarsely aligned, a unified shift t_G is not sufficient to recover accurate monocular disparity at different pixels. To fully release the potential of monocular depth prior, SGA leverages intermediate stereo cues to further estimate a per-pixel residual shift. To avoid introducing noisy stereo cues, SGA uses *confidence based guidance*. In each update step j, we compute the confidence using the flow residual map F_S^j , which is obtained by warping and subtracting features based on the stereo disparity D_S^j as:

$$\boldsymbol{F}_{S}^{j}(x,y) = \left\| \boldsymbol{F}_{3}^{L}(x,y) - \boldsymbol{F}_{3}^{R}(x - \boldsymbol{D}_{S}^{j},y) \right\|_{1}$$
(2)

where F_3^L, F_3^R represent the quarter-resolution features of left and right images in \mathcal{F} respectively. For each iteration, we also use the current stereo disparity D_S^j to index from the Geometry Encoding Volume to obtain geometry features of stereo branch G_S^j follow IGEV. Concatenated with F_S^j and D_S^j , we obtain the stereo condition as:

$$x_S^j = [\operatorname{En}_g([\boldsymbol{G}_S^j, \boldsymbol{F}_S^j, \boldsymbol{D}_S^j]), \operatorname{En}_d(\boldsymbol{D}_M^j), \boldsymbol{D}_M^j], \quad (3)$$

where En_g and En_d are two convolutional layers for feature encoding. We feed x_S^j into condition-guided ConvGRUs to update the hidden state h_m^{i-1} of monocular branch as:

$$z^{j} = \sigma(\text{Conv}([h_{M}^{j-1}, x_{S}^{j}], W_{z}) + c_{k}),$$

$$r^{j} = \sigma(\text{Conv}([h_{M}^{j-1}, x_{S}^{j}], W_{r}) + c_{r}),$$

$$\tilde{h}_{M}^{j} = \tanh(\text{Conv}([r^{j} \odot h_{M}^{j-1}, x_{S}^{j}], W_{h}) + c_{h}),$$

$$h_{M}^{j} = (1 - z^{j}) \odot h_{M}^{j-1} + z^{j} \odot \tilde{h}_{M}^{j},$$
(4)

where c_k , c_r , c_h are context features. Based on the hidden state h_M^j , we decode a *residual shift* Δt through two convolutional layers to update the monocular disparity:

$$\boldsymbol{D}_{M}^{j+1} = \boldsymbol{D}_{M}^{j} + \Delta \boldsymbol{t}.$$
 (5)

Mono Guided Refinement (MGR). Symmetric to SGA, MGR leverages the aligned monocular disparity to address stereo deficiencies in ill-posed regions, thin structures, and distant objects. Specifically, we employ the same conditionguided GRU architecture with independent parameters to refine stereo disparity. We simultaneously calculate the flow residual maps F_M^j , F_S^j and geometric features G_M^j , G_S^j for both the monocular and stereo branches, providing a comprehensive stereo refinement guidance:

$$\begin{aligned} \boldsymbol{F}_{M}^{j}(x,y) &= \left\| \boldsymbol{F}_{3}^{L}(x,y) - \boldsymbol{F}_{3}^{R}(x - \boldsymbol{D}_{M}^{j},y) \right\|_{1}, \\ \boldsymbol{x}_{M}^{j} &= [\mathrm{En}_{g}([\boldsymbol{G}_{M}^{j},\boldsymbol{F}_{M}^{j},\boldsymbol{D}_{M}^{j}]), \mathrm{En}_{d}(\boldsymbol{D}_{M}^{j}), \boldsymbol{D}_{M}^{j}, \\ \mathrm{En}_{g}([\boldsymbol{G}_{S}^{j},\boldsymbol{F}_{S}^{j},\boldsymbol{D}_{S}^{j}]), \mathrm{En}_{d}(\boldsymbol{D}_{S}^{j}), \boldsymbol{D}_{S}^{j}], \end{aligned}$$
(6)

where G_M^j is the geometry features of the monocular branch obtained by indexing from the Geometry Encoding



Figure 5. Qualitative results on ETH3D. MonSter outperforms IGEV in challenging areas with strong reflectance, fine structures etc.

Volume using the monocular disparity D_M^j . x_M^j represents the monocular condition for ConvGRUs, and we use the Eq. (4) to update the hidden state h_S^i similarly, with only the condition input changed from x_S^i to x_M^i . We use the same two convolutional layers to decode the residual disparity Δd and update the current stereo disparity following Eq. (5). After N_2 rounds of dual-branched refinement, the disparity of the stereo branch is the final output of MonSter.

3.3. Loss Function

We use the L1 loss to supervise the output from two branches. We denote the set of disparities from the first N_1 iterations of the stereo branch as $\{d_i\}_{i=0}^{N_1-1}$ and follow [18] to exponentially increase the weights as the number of iterations increases. The total loss is defined as the sum of the monocular branch loss \mathcal{L}_{Mono} and the stereo branch loss \mathcal{L}_{Stereo} as follows:

$$\mathcal{L}_{Stereo} = \sum_{i=0}^{N_1-1} \gamma^{N_1+N_2-i} || \boldsymbol{d}_i - \boldsymbol{d}_{gt} ||_1 + \sum_{i=N_1}^{N_1+N_2-1} \gamma^{N_1+N_2-i} || \boldsymbol{D}_S^{i-N_1} - \boldsymbol{d}_{gt} ||_1, \quad (7)$$
$$\mathcal{L}_{Mono} = \sum_{i=N_1}^{N_1+N_2-1} \gamma^{N_1+N_2-i} || \boldsymbol{D}_m^{i-N_1} - \boldsymbol{d}_{gt} ||_1$$

where $\gamma = 0.9$, and d_{gt} is the ground truth.

4. Experiment

4.1. Implementation Details

We implement MonSter with PyTorch and perform experiments using NVIDIA RTX 3090 GPUs. We use the AdamW [19] optimizer and clip gradients to the range [-1, 1] following baseline [37]. We use the one-cycle learning rate schedule with a learning rate of 2e-4 and train Mon-Ster with a batch size of 8 for 200k steps as the pretrained model. For the monocular branch, we use the ViT-large version of DepthAnythingV2 [45] and freeze its parameters to prevent training of stereo-matching tasks from affecting its generalization and accuracy.

Following the standard [15, 16, 33], we pretrain Mon-Ster on Scene Flow [20] for most experiments. For finetuning on ETH3D and Middlebury, we follow SOTA methods [15, 16, 33] to create the Basic Training Set (BTS) from various public datasets for pretraining, including Scene Flow [20], CREStereo [15], Tartan Air [32], Sintel Stereo [4], FallingThings [30] and InStereo2k [2].

4.2. Benchmark Performance

To demonstrate the outstanding performance of our method, we evaluate MonSter on the five most commonly used benchmarks: KITTI 2012 [10], KITTI 2015 [21], ETH3D [26], Middlebury [25] and Scene Flow [20].

Scene Flow [20]. As shown in Tab. 1, we achieve a new state-of-the-art performance with an EPE metric of 0.37 on Scene Flow, surpassing our baseline [37] by 21.28% and outperforming the SOTA method [33] by 15.91%.

ETH3D[26]. Following the SOTA methods [15, 16, 33, 38, 41], the full training set is composed of the BTS and ETH3D training set. Our MonSter ranks 1st on the ETH3D leaderboard. As shown in Tab. 2, MonSter achieves the best performance among all published methods for all metrics. Compared with baseline IGEV, we achieve improvements of 58.93%, 52.32%, and 41.18% in the three reported metrics, respectively. Qualitative comparisons shown in Fig.5

Method	GwcNet[13]	LEAStereo[8]	ACVNet[36]	IGEV[37]	Selective-IGEV[33]	MonSter (Ours)
EPE $(px)\downarrow$	0.76	0.78	0.48	0.47	<u>0.44</u>	0.37 (-15.91%)

Table 1. Quantitative evaluation on Scene Flow test set. The best result is bolded, and the second-best result is underscored.

	E'	TH3D[2	5]	Mid	ldlebury[25]	1	KITTI 2	015[<mark>10</mark>]		ŀ	KITTI 2	2012[10]
Method	Bad1.0	Bad1.0	RMSE	Bad2.0	Bad2.0	RMSE	D1-bg	D1-all	D1-bg	D1-all	Out-2	Out-2	Out-3	Out-3
	Noc	All	Noc	Noc	All	Noc	Noc	Noc	All	All	Noc	All	Noc	All
GwcNet[13]	6.42	6.95	0.69	-	-	-	1.61	1.92	1.74	2.11	2.16	2.71	1.32	1.70
GANet[47]	6.22	6.86	0.75	-	-	-	1.40	1.73	1.55	1.93	1.89	2.50	1.19	1.60
LEAStereo[8]	-	-	-	7.15	12.10	8.11	1.29	1.51	1.40	1.65	1.90	2.39	1.13	1.45
ACVNet[36]	2.58	2.86	0.45	13.70	19.50	32.2	1.26	1.52	1.37	1.65	1.83	2.35	1.13	1.47
RAFT-Stereo[18]	2.44	2.60	0.36	4.74	9.37	8.41	1.44	1.69	1.58	1.82	1.92	2.42	1.30	1.66
CREStereo[15]	0.98	1.09	0.28	3.71	8.13	7.70	1.33	1.54	1.45	1.69	1.72	2.18	1.14	1.46
IGEV[37]	1.12	1.51	0.34	4.83	8.16	12.80	1.27	1.49	1.38	1.59	1.71	2.17	1.12	1.44
CroCo-Stereo[35]	0.99	1.14	0.30	7.29	11.11	8.91	1.30	1.51	1.38	1.59	-	-	-	-
DLNR[48]	-	-	-	3.20	6.98	7.78	1.42	1.61	1.60	1.76	-	-	-	-
Selective-IGEV[33]	1.23	1.56	0.29	2.51	6.04	7.26	1.22	1.44	1.33	1.55	1.59	2.05	1.07	1.38
LoS[16]	0.91	1.03	0.31	4.20	8.03	6.99	1.29	1.52	1.42	1.65	1.69	2.12	1.10	1.38
NMRF-Stereo[12]	-	-	-	-	-	-	1.18	1.46	1.28	1.57	1.59	2.07	1.01	1.35
MonSter(Ours)	0.46	0.72	0.20	2.64	6.14	6.71	1.05	1.33	1.13	1.41	1.36	1.75	0.84	1.09

Table 2. **Results on four popular benchmarks.** All results are derived from official leaderboard publications or corresponding papers. All metrics are presented in percentages, except for RMSE, which is reported in pixels. For testing masks, "All" denotes being tested with all pixels while "Noc" denotes being tested with a non-occlusion mask. The **best** and second best are marked with colors.

exhibit a similar trend. Notably, even compared with the previous best method LoS, we improved the Bad 1.0 (Noc) metric from 0.91 to 0.46, achieving a 49.45% improvement.

Middlebury. Also following [15, 16, 33, 41], the training set of Middlebury is the combination of the BTS and Middlebury training set. As shown in Tab. 2, MonSter outperforms all existing methods in terms of RMSE metric. Compared with baseline [37], we achieved an improvement of 45.34% on the Bad 2.0 (Noc) metric.

KITTI. Following the training of SOTA methods [12, 15, 33], we finetuned Scene Flow pretrained model on the mixed dataset of KITTI 2012 and KITTI 2015 for 50k steps. At the time of writing, MonSter ranks 1st simultaneously on both the KITTI 2012 and KITTI 2015 leaderboards. As shown in Tab. 2, we achieve the best performance for all metrics. On KITTI 2015, MonSter surpasses CREStereo [15] and Selective-IGEV [33] by 16.57% and 9.03% on D1-all metric of all regions, respectively. As for KITTI 2012, we significantly outperform the existing SOTA by a large margin. Compared to the current highest-precision method, NMRF-Stereo [12] and Selective-IGEV [33], we achieve a 19.26% improvement and a 21.01% improvement respectively in the Out-3 metric across all regions.

4.3. Performance in Ill-posed Regions

MonSter fully leverages the advantages of monocular depth priors, effectively overcoming challenges in ill-posed regions. To validate this, we conduct comparisons on several representative ill-posed regions, such as reflective ar-

	KITTI 2012 Reflective Region							
Method	Out-2	Out-2	Out-3	Out-3	Out-4	Out-4		
	Noc	All	Noc	All	Noc	All		
ACVNet[36]	11.42	13.53	7.03	8.67	5.18	6.48		
CREStereo[15]	9.71	11.26	6.27	7.27	4.93	5.55		
IGEV[37]	7.57	8.80	4.35	5.00	3.16	3.57		
Selective-IGEV[33]	6.73	7.84	3.79	4.38	2.66	3.05		
LoS[16]	6.31	7.84	3.47	4.45	2.41	3.01		
NMRF-Stereo[12]	10.02	12.34	6.35	8.11	4.80	6.09		
MonSter(Ours)	5.66	6.81	2.75	3.38	1.73	2.13		

Table 3. Results of the reflective regions on KITTI 2012 leader-board. Thebestandsecond bestare marked with colors.

Method	Ec	lges	Non-Edges		
Wiethou	EPE	>1px	EPE	>1px	
IGEV [37]	2.23	20.42	0.41	4.58	
Selective-IGEV [33]	2.18	20.01	0.38	4.35	
MonSter(Ours)	1.91	18.59	0.31	3.57	

Table 4. Comparison on SceneFlow test set in different regions.

eas, edge and non-edge regions, and distant backgrounds:

• **Reflective Regions**: We conducted comparisons on the reflective regions of KITTI 2012 benchmark. MonSter ranks 1st on the KITTI 2012 leaderboard for all metrics of reflective regions. As shown in Tab.3, we have elevated the existing SOTA to a new level. MonSter surpasses Selective-IGEV and LoS by 30.16% and 29.24% on Out-4(All) metric, respectively. Notably, compared to the SOTA method NMRF-Stereo, we achieved significant

Madha d	KITTI-12	KITTI-15	Middlebury	ETH3D				
Method	(>3px)	(>3px)	(>2px)	(>1px)				
Training Set	Scene Flow							
CFNet[27]	4.72	5.81	9.80	5.80				
RAFT-Stereo[18]	5.12	5.74	9.36	3.28				
CREStereo[15]	5.03	5.79	12.88	8.98				
Selective-IGEV[33]	5.64	6.05	12.04	5.40				
NMRF-Stereo[12]	4.23	5.10	7.54	3.82				
IGEV[37]	4.84	5.51	6.23	3.62				
MonSter(Ours)	3.62	3.97	5.17	2.03				
Training Set	Scene Fle	ow + CRE	Stereo + Tar	tan Air				
IGEV[37]	3.95	4.30	5.56	2.38				
MonSter(Ours)	2.95	3.23	2.94	1.21				

Table 5. **Zero-shot generalization benchmark.** The top half is generalization comparisons only trained on Scene Flow. The bottom half compares generalizations for models trained on a combination of Scene Flow, CREStereo, and TartanAir datasets.

improvements of 58.32% and 65.02% in the Out-3(All) and Out-4(All) metrics respectively.

- Edge & Non-edge Regions: Stereo matching faces challenges in edge and low-texture regions, which our method addresses by leveraging the strengths of monocular depth. To evaluate MonSter in these areas, we divide the Scene Flow test set into edge and non-edge regions using the Canny operator following [33]. As shown in Tab. 4, Mon-Ster outperforms the baseline [37] by 14.35% and 24.39% in edge regions and non-edge regions. Even when compared to [33], which is specifically designed to address edge and textureless regions, our method achieves improvements of 12.39% and 18.42% in two types of regions on EPE metric, respectively.
- Distant Backgrounds: Stereo matching struggles with depth perception for distant objects, we improve it by incorporating our SGA and MGR modules. As shown in Tab.2, MonSter improves the D1-bg metric by 18.12% compared to the baseline on KITTI 2015 benchmark. The D1-bg metric reflects the percentage of stereo disparity outliers averaged specifically over background regions.

4.4. Zero-shot Generalization

Our method not only improves the accuracy of in-domain datasets but also enhances generalization ability. We evaluate the generalization performance of MonSter from synthetic datasets to unseen real-world scenes. All models were trained on the same synthetic datasets and then tested directly on real-world datasets: KITTI[10, 21], Middlebury[25], and ETH3D[26] training sets. We conducted a more thorough generalization comparison by training with different datasets:

Scene Flow: We compared with SOTA methods which are only trained on the Scene Flow training set. As shown in the top half of Tab. 5, MonSter achieves the best gener-



Figure 6. Visual comparisons on our captured real-world zeroshot data provide a more comprehensive assessment of generalization capability. All models are trained solely on the synthetic SceneFlow dataset. MonSter significantly outperforms IGEV in textureless regions, reflective areas and fine structures, etc.

alization performance across four datasets. MonSter even surpasses CFNet and CREStereo, both of which are specifically designed for cross-domain generalization.

Mix 3 datasets: Interestingly, only adding CREStereo and TartanAir to the training set, the generalization of Mon-Ster improves significantly. This is because the SceneFlow dataset differs substantially from real-world distributions, whereas the virtual data in TartanAir more closely resemble real-world scenes. This similarity helps our SGA and MGR modules better in learning how to effectively integrate monocular and stereo cues. As shown in the bottom half of Tab. 5, we surpass the baseline[37] by 47.12% and 49.16% on Middlebury and ETH3D respectively.

The strong sim-to-real generalization of MonSter is highly encouraging. In the future, we plan to scale up the size and diversity of simulation training data to further boost the performance of MonSter.

4.5. Ablation Study

To demonstrate the effectiveness of each component of MonSter, we ablation on Scene Flow [20] and the following strategies are discussed:

• **Disparity Fusion**: Monocular depth often carries substantial noise, which complicates its fusion with the stereo branch's disparity. To demonstrate the efficiency of our Mono Guided Refinement (MGR) fusion method, we replaced the MGR module with a convolution-based hourglass network of equal parameter numbers, directly concatenating the monocular and stereo disparities for fusion, denoted as Mono+Conv. As shown in Tab.6, compared with 'Mono+Conv', 'Mono+MGR' improves the

Madal	Monocular	Disparity	Scale & Shift	Feature	EPE	>1px	Run-time
Model	Depth	Fusion	Refinement	Sharing	(px)	(%)	(S)
Baseline (IGEV)					0.47	5.21	0.37
Mono+Conv	\checkmark	Conv			0.46	5.12	0.64
Mono+MGR	\checkmark	MGR			0.43	4.96	0.65
Mono+MGR+Conv	\checkmark	MGR	Conv		0.42	4.82	0.65
Mono+MGR+SGA	\checkmark	MGR	SGA		0.39	4.43	0.66
Full model (MonSter)	\checkmark	MGR	SGA	\checkmark	0.37	4.25	0.64

Table 6. Ablation study of the effectiveness of proposed modules on the Scene Flow test set.

Model	Mono Depth Model	Iteration Number	Sceneflow (EPE)	Run-time (S)
IGEV[37]		32	0.47	0.37
Full model	DepthAnythingV2[45]	32	0.37	0.64
Full model-4iter	DepthAnythingV2[45]	4	0.42	0.34
Full model-V1	DepthAnythingV1[44]	32	0.39	0.64
Full model-midas	MiDaS[23]	32	0.41	0.51

Table 7. Efficiency and universality of MonSter.

EPE metric by 6.52%.

- Scale & Shift Refinement: As shown in Fig.3, even after global scale-shift alignment, the monocular depth still suffers from significant scale and shift ambiguities. Therefore, leveraging high-confidence regions from the stereo branch to correct the scale and shift of the monocular disparity is essential. We incorporated our Stereo Guided Refinement module (SGA) to optimize the scale and shift based on 'Mono+MGR'. As shown in Tab.6, compared with 'Mono+MGR', our SGA module achieves improvements of 9.30% in EPE and 10.69% in 1-pixel error, demonstrating that monocular depth with accurate scale and shift can provide further enhancements to stereo matching. Similarly, we validated that our SGA module is more effective than direct convolution-based fusion. Compared to 'Mono+MGR+Conv', our SGA still achieves an 8.09% improvement in 1-pixel error.
- Feature Sharing: The feature extraction component of monocular depth estimation models, pre-trained on large datasets, contains rich contextual information that can be shared with the stereo branch. We employ a feature transfer network to map these features, transforming them into a representation compatible with stereo matching. As shown in the last 2 rows of Tab.6, feature sharing improved the EPE metric by 5.13%.

Compatibility with Other Monocular Models. To validate the generalizability of MonSter to different depth models, we replace the monocular branch with the ViT-large version of [44] and the dpt_beit_large version of Mi-DaS [23]. As shown in Tab.7, all MonSter variants consistently outperform the baseline [37]. This demonstrates the versatility of MonSter and its potential to benefit from future advancements in monocular depth models.

Efficiency. We use the ViT-large version of DepthAnythingV2 [45] as the monocular branch, which introduces

time and memory overheads. Compared to the baseline [37], our inference time increases from 0.37s to 0.64s. With the monocular branch containing 335.3M parameters, the stereo branch having 12.6M parameters, and the SGA and MGR modules having 8.2M parameters, our full model has a total of 356.1M parameters. Existing methods [1, 35] also use ViT-based encoders with a large number of parameters. For example, CroCo-Stereo [35] has 437.4M parameters, larger than MonSter, yet its performance is inferior, as shown in Tab.2, this demonstrates that the number of parameters does not determine the final effectiveness. Our research focuses on the accuracy and generalization capabilities, given the significant improvement in accuracy and generalization (as shown in Tab.2 and Tab.5), the additional inference cost is acceptable. Notably, our approach simplifies stereo matching and enables us to achieve SOTA performance with only 4 iterations $(N_1=N_2=2)$. As shown in Tab.7, while the baseline requires 32 iterations, our method achieves a 10.64% improvement in EPE with just 4 iterations, achieving a better accuracy-speed trade-off. Further memory reduction can be achieved through encoder quantization or distillation, which we leave as future work.

5. Conclusion

We propose MonSter, a novel pipeline that decouples the stereo matching task into a simpler paradigm of recovering scale and shift from relative depth. This approach fully leverages the contextual and geometric priors provided by monocular methods while avoiding issues such as noise and scale ambiguity. As a result, our method significantly enhances the accuracy and robustness of stereo matching in ill-posed regions. MonSter achieves substantial improvements over the state-of-the-art on five widely used benchmarks. Additionally, MonSter also demonstrates the best generalization performance, and experimental results show that there is still considerable room for improvement. Therefore, our future work will focus on scaling up Mon-Ster to serve as a stereo foundation model, enabling a wide range of downstream applications that require metric depth.

Acknowledgement. This work is supported by National Natural Science Foundation of China(62472184) and the Fundamental Research Funds for the Central Universities.

References

- Jihye Ahn, Hyesong Choi, Soomin Kim, and Dongbo Min. Madis-stereo: Enhanced stereo matching via distilled masked image modeling. *arXiv preprint arXiv:2409.02846*, 2024. 8
- [2] 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. 5
- [3] 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
- [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. 5
- [5] Jia-Ren Chang and Yong-Sheng Chen. Pyramid stereo matching network. In CVPR, pages 5410–5418, 2018. 1, 2, 3
- [6] Junda Cheng, Xin Yang, Yuechuan Pu, and Peng Guo. Region separable stereo matching. *IEEE Transactions on Multimedia*, 25:4880–4893, 2022. 3
- [7] Junda Cheng, Gangwei Xu, Peng Guo, and Xin Yang. Coatrsnet: Fully exploiting convolution and attention for stereo matching by region separation. *International Journal of Computer Vision*, 132(1):56–73, 2024. 3
- [8] Xuelian Cheng, Yiran Zhong, Mehrtash Harandi, Yuchao Dai, Xiaojun Chang, Hongdong Li, Tom Drummond, and Zongyuan Ge. Hierarchical neural architecture search for deep stereo matching. *NeurIPS*, 33:22158–22169, 2020. 6
- [9] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. arXiv preprint arXiv:1406.1078, 2014. 2
- [10] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *CVPR*, pages 3354–3361. IEEE, 2012. 2, 5, 6, 7
- [11] Xiaodong Gu, Zhiwen Fan, Siyu Zhu, Zuozhuo Dai, Feitong Tan, and Ping Tan. Cascade cost volume for high-resolution multi-view stereo and stereo matching. In *Proceedings of* the IEEE/CVF conference on computer vision and pattern recognition, pages 2495–2504, 2020. 3
- [12] 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. 6, 7
- [13] Xiaoyang Guo, Kai Yang, Wukui Yang, Xiaogang Wang, and Hongsheng Li. Group-wise correlation stereo network. In *CVPR*, pages 3273–3282, 2019. 1, 3, 6
- [14] 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 *ICCV*, pages 66–75, 2017. 1

- [15] 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 *CVPR*, pages 16263–16272, 2022. 1, 2, 3, 5, 6, 7
- [16] 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. 3, 5, 6
- [17] Zhaoshuo Li, Xingtong Liu, Nathan Drenkow, Andy Ding, Francis X Creighton, Russell H Taylor, and Mathias Unberath. Revisiting stereo depth estimation from a sequenceto-sequence perspective with transformers. In *Proceedings* of the IEEE/CVF international conference on computer vision, pages 6197–6206, 2021. 2
- [18] Lahav Lipson, Zachary Teed, and Jia Deng. Raft-stereo: Multilevel recurrent field transforms for stereo matching. In *3DV*, pages 218–227. IEEE, 2021. 1, 2, 3, 5, 6, 7
- [19] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101, 2017. 5
- [20] Nikolaus Mayer, Eddy Ilg, Philip Hausser, Philipp Fischer, Daniel Cremers, Alexey Dosovitskiy, and Thomas Brox. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In *CVPR*, pages 4040–4048, 2016. 1, 2, 3, 5, 7
- [21] Moritz Menze and Andreas Geiger. Object scene flow for autonomous vehicles. In *CVPR*, pages 3061–3070, 2015. 2, 5, 7
- [22] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. arXiv preprint arXiv:2304.07193, 2023. 3
- [23] 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. 8
- [24] 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. 3
- [25] 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 *GCPR*, pages 31–42. Springer, 2014. 2, 5, 6, 7
- [26] Thomas Schops, Johannes L Schonberger, Silvano Galliani, Torsten Sattler, Konrad Schindler, Marc Pollefeys, and Andreas Geiger. A multi-view stereo benchmark with highresolution images and multi-camera videos. In *CVPR*, pages 3260–3269, 2017. 2, 5, 6, 7
- [27] Zhelun Shen, Yuchao Dai, and Zhibo Rao. Cfnet: Cascade and fused cost volume for robust stereo matching. In *CVPR*, pages 13906–13915, 2021. 3, 7
- [28] 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. 3

- [29] Xiao Song, Xu Zhao, Liangji Fang, Hanwen Hu, and Yizhou Yu. Edgestereo: An effective multi-task learning network for stereo matching and edge detection. *IJCV*, 128(4):910–930, 2020. 2, 3
- [30] 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. 5
- [31] A Vaswani. Attention is all you need. Advances in Neural Information Processing Systems, 2017. 2
- [32] 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. 5
- [33] 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
- [34] Xianqi Wang, Hao Yang, Gangwei Xu, Junda Cheng, Min Lin, Yong Deng, Jinliang Zang, Yurui Chen, and Xin Yang. Stereogen: High-quality stereo image generation from a single image. arXiv preprint arXiv:2501.08654, 2025. 1
- [35] 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 pretraining for stereo matching and optical flow. In *Proceedings* of the IEEE/CVF International Conference on Computer Vision, pages 17969–17980, 2023. 6, 8
- [36] Gangwei Xu, Junda Cheng, Peng Guo, and Xin Yang. Attention concatenation volume for accurate and efficient stereo matching. In *CVPR*, pages 12981–12990, 2022. 1, 2, 3, 6
- [37] 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, 8
- [38] 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. 5
- [39] Gangwei Xu, Jiaxin Liu, Xianqi Wang, Junda Cheng, Yong Deng, Jinliang Zang, Yurui Chen, and Xin Yang. Banet: Bilateral aggregation network for mobile stereo matching. *arXiv preprint arXiv:2503.03259*, 2025. 1
- [40] Haofei Xu and Juyong Zhang. Aanet: Adaptive aggregation network for efficient stereo matching. In CVPR, pages 1959– 1968, 2020. 1
- [41] Haofei Xu, Jing Zhang, Jianfei Cai, Hamid Rezatofighi, Fisher Yu, Dacheng Tao, and Andreas Geiger. Unifying flow,

stereo and depth estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. 2, 5, 6

- [42] Haofei Xu, Jing Zhang, Jianfei Cai, Hamid Rezatofighi, Fisher Yu, Dacheng Tao, and Andreas Geiger. Unifying flow, stereo and depth estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. 1, 3
- [43] Guorun Yang, Hengshuang Zhao, Jianping Shi, Zhidong Deng, and Jiaya Jia. Segstereo: Exploiting semantic information for disparity estimation. In *ECCV*, pages 636–651, 2018. 2, 3
- [44] Lihe Yang, Bingyi Kang, Zilong Huang, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything: Unleashing the power of large-scale unlabeled data. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 10371–10381, 2024. 8
- [45] Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything v2. arXiv preprint arXiv:2406.09414, 2024. 2, 3, 5, 8
- [46] Chenghao Zhang, Gaofeng Meng, Bing Su, Shiming Xiang, and Chunhong Pan. Monocular contextual constraint for stereo matching with adaptive weights assignment. *Image* and Vision Computing, 121:104424, 2022. 3
- [47] Feihu Zhang, Victor Prisacariu, Ruigang Yang, and Philip HS Torr. Ga-net: Guided aggregation net for endto-end stereo matching. In *CVPR*, pages 185–194, 2019. 1, 6
- [48] 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. 2, 3, 6
- [49] Xuan Zhu, Jijun Xiang, Xianqi Wang, Longliang Liu, Yu Wang, Hong Zhang, Fei Guo, and Xin Yang. Svdc: Consistent direct time-of-flight video depth completion with frequency selective fusion. arXiv preprint arXiv:2503.01257, 2025. 1