

# Lite Any Stereo: Efficient Zero-Shot Stereo Matching

Junpeng Jing Weixun Luo Ye Mao<sup>†</sup> Krystian Mikolajczyk  
Imperial College London

<https://tomtomtommi.github.io/LiteAnyStereo/>



Figure 1. Zero-shot prediction on in-the-wild stereo images. The proposed method achieves accurate disparity estimation across diverse scenarios and maintains high efficiency, even on older-generation GPUs.

## Abstract

Recent advances in stereo matching have focused on accuracy, often at the cost of significantly increased model size. Traditionally, the community has regarded efficient models as incapable of zero-shot ability due to their limited capacity. In this paper, we introduce **Lite Any Stereo**, a stereo depth estimation framework that achieves strong zero-shot generalization while remaining highly efficient. To this end, we design a compact yet expressive backbone to ensure scalability, along with a carefully crafted hybrid cost aggregation module. We further propose a three-stage training strategy on million-scale data to effectively bridge the sim-to-real gap. Together, these components demonstrate that an ultra-light model can deliver strong generalization, ranking 1<sup>st</sup> across four widely used real-world benchmarks. Remarkably, our model attains accuracy comparable to or exceeding state-of-the-art non-prior-based accuracy methods while requiring less than 1% computational cost, setting a new standard for efficient stereo matching.

## 1. Introduction

From the foundational work of [37] through the classical advances culminating in works such as [51], stereo vision has

seen decades of steady progress built on a wide range of algorithmic ideas. In the past decade, deep learning methods have driven a dramatic leap in accuracy—leading many researchers to view stereo vision as largely solved—yet these gains often rely on large, computationally expensive models that are difficult to deploy on resource-limited hardware.

Learning-based methods [32, 35, 58, 67] achieve remarkable accuracy, continuously raising scores on standard benchmarks [14, 40, 46, 47]. These methods can be broadly categorized as accuracy-oriented. More recently, the emergence of foundation models trained on internet-scale data, such as the DepthAnything series [73, 74], has led to further advancement in the field. Stereo models augmented with depth priors [9, 23, 61] have demonstrated excellent zero-shot generalization: a single set of model weights can deliver strong results across diverse scenarios. However, despite these successes, such approaches prioritize accuracy over efficiency, which limits their practical applications.

In contrast, efficiency-oriented approaches [20, 48, 68] trade accuracy for faster inference and lower resource use, however, the accuracy gap to large models remains significant. The absence of light models with strong zero-shot generalization may lead to a conclusion that such models lack capacity for zero-shot applications. Consequently, most approaches still rely on domain-specific fine-tuning, thus falling short of being practical off-the-shelf solutions.

Training from large data is a key factor in enabling zero-

<sup>†</sup>Corresponding author: ye.mao21@imperial.ac.uk




Figure 2. Zero-shot performance. Our method achieves SOTA by a large margin, with even better or comparable non-prior-based accurate model, while requiring less than 1% of their MACs.

shot capability. StereoAnything [19] scaled training by using a monocular depth model [74] to generate pseudo-disparity maps, yielding 30M real-world samples. While large data improves zero-shot performance for efficient models [20], they remain far behind accuracy-oriented models, primarily due to the limited quality of monocular depth.

In this paper, we propose Lite Any Stereo, an ultra-lite stereo matching model designed for zero-shot generalization. We achieve this from two perspectives: architecture and training strategy. We design a backbone with a hybrid cost aggregation module that jointly leverages 2D and 3D representations, capturing complementary spatial and disparity cues. This carefully crafted backbone achieves very low computational cost and fast inference speed. To train for zero-shot capability, we scale the data to a million level with a three-stage strategy. Specifically, self-distillation on synthetic labeled data is performed after supervised training. We further make use of real-world unlabeled data, which has long been overlooked in stereo, for knowledge distillation. This combination effectively mitigates the sim-to-real gap and proves to be universally beneficial across varied model architectures. Together, these innovations significantly improve performance. As shown in Fig. 1, our method shows strong generalization, producing accurate disparities on in-the-wild images with high efficiency.

Our main contributions can be summarized as follows:

- We present Lite Any Stereo, an efficient stereo matching model that achieves zero-shot capability. As shown in Fig. 2, it delivers the highest zero-shot accuracy among all efficient methods by a large margin. This is the first efficient model to outperform or match accuracy-oriented models that do not use foundational priors, while requiring less than 1% of their computational cost.
- We propose a hybrid cost aggregation module that captures complementary disparity and spatial information,

enhancing representation at low cost.

- We develop a three-stage training strategy that integrates synthetic supervision, self-distillation, and real-world knowledge distillation, effectively improving zero-shot generalization across diverse architectures.

## 2. Related Work

In this section, we first review accurate stereo matching methods, followed by a discussion of approaches focused on computational efficiency.

### 2.1. Accurate Stereo Methods

Modern stereo matching methods based on deep learning typically construct either 3D or 4D cost volumes to capture pixel-level correspondences between the left and right views. These volumes serve as the foundation for disparity estimation, with convolutional neural networks used to aggregate contextual information. Broadly, CNN-based approaches can be categorized into two types: those that emphasize cost aggregation [7, 10, 15, 17, 18, 21, 34, 42, 52, 69, 71, 76], and those that adopt iterative refinement strategies [32, 35, 58, 66] to progressively enhance disparity predictions. Another line of research explores transformer-based stereo models [16, 33, 50, 60, 70], which leverage attention mechanisms to capture global context. In addition, some methods explicitly address temporal consistency in stereo videos [26–29], aiming to produce consistent disparities across time. While these models perform well on domain-specific benchmarks [14, 40, 46, 47], generalization to unseen domains in zero-shot settings remains a major challenge. To address this, some works have focused on learning domain-invariant representations [8, 11, 44, 77] and robust modules [19, 25, 49], aiming to improve performance under domain shifts. A recent new direction involves leveraging prior knowledge from monocular foundation models [73, 74]. By integrating monocular depth cues into stereo pipelines, several methods [3, 9, 23, 61, 78, 79] demonstrate enhanced zero-shot generalization. However, these improvements often come at the cost of increased computational complexity, limiting their deployment in real-time or resource-constrained scenarios.

### 2.2. Efficient Stereo Matching

Real-time performance is critical for stereo matching in practical applications. Early approaches [13, 30, 56] address this by estimating disparities at reduced spatial resolutions to lower computational cost, though at the expense of accuracy. Subsequent methods aim to improve this efficiency and accuracy trade-off based on 3D aggregation. Representative advancements include guided cost volume excitation in CoEx [1], edge-aware upsampling in BGNNet [64], and sparse attention for selective high-resolution matching in Fast-ACVNet [65, 67]. To

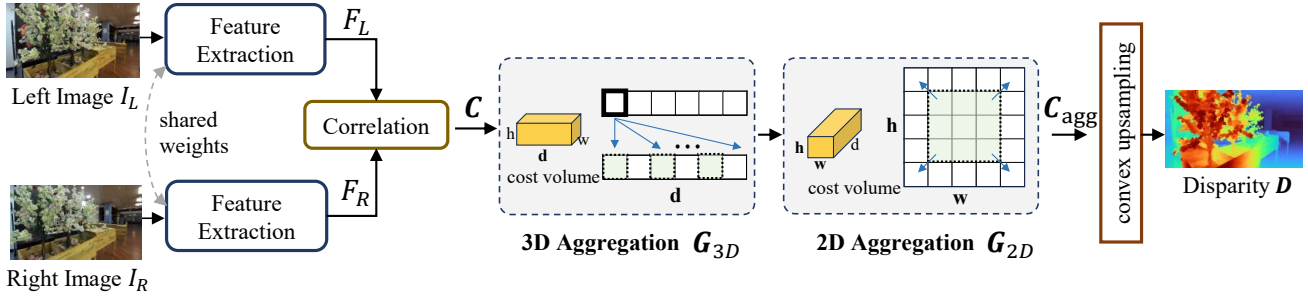


Figure 3. Overview of the proposed Lite Any Stereo. Given an input stereo image pair, features are first extracted using a shared-weight feature extraction module. A correlation module then constructs cost volume from extracted features, which is processed by a hybrid 3D-2D cost aggregation module to obtain aggregated cost volume along both disparity and spatial dimensions. Finally, low-resolution disparity map is estimated and a convex upsampling operation is applied to recover the full-resolution disparity map.

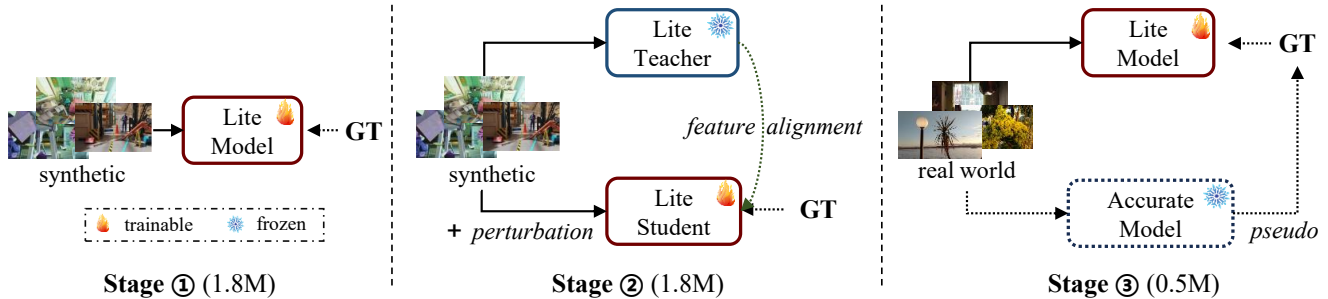


Figure 4. Overview of the proposed three-stage training strategy. **Stage ①**: The lite model is trained using a standard supervised setup on a mixed of synthetic datasets including 1.8M labeled stereo image pairs. **Stage ②**: Self-distillation is employed, where both teacher and student models are initialized from the Stage ① weights. The teacher receives clean data, while the student is fed perturbed inputs to encourage learning of domain-invariant representations via feature alignment. **Stage ③**: The lite model is fine-tuned on unlabeled real-world data using pseudo labels generated by a frozen accurate model.

further reduce complexity, another line of works have explored 2D-based alternatives to traditional 3D convolutions. Notable examples include AANet [69], which employs deformable 2D convolutions for adaptive cost aggregation; HITNet [52], which reconstructs disparities through iterative warping without explicit cost volumes; and MobileStereoNet-2D [48], which adopts lightweight MobileNet blocks. More recent efforts continue this trend with specialized strategies: uncertainty-aware adaptive warping in Lite-CREStereo++ [25], channel-wise enhancement in LightStereo [20], and frequency-disparity guided bilateral aggregation in BANet [68]. Despite these advances in efficiency, most of these models are optimized for a specific domain, particularly the KITTI online benchmarks [14, 40], and still struggle to generalize well across diverse, unseen scenarios. Achieving strong zero-shot ability while maintaining real-time efficiency remains an open challenge.

### 3. Method

In this section, we introduce Lite Any Stereo, an efficient feed-forward network for zero-shot stereo matching. We first demonstrate the overall framework (Sec. 3.1), and then go through the training strategy (Sec. 3.2).

#### 3.1. Framework

As shown in Fig. 3, the overall framework comprises four main stages: feature extraction, correlation, cost aggregation, and disparity estimation.

**Feature Extraction.** Recent methods [9, 23, 61] have demonstrated remarkable improvements by leveraging depth features from DepthAnything (DA) [73, 74]. However, we find that the additional computational overhead introduced by DA, even with the smallest DA-S variant, is prohibitive for an efficiency-oriented stereo model. Therefore, we adopt a conventional backbone [20, 68] pretrained on ImageNet [12] without external priors for feature extraction. Moreover, we find that compared with more recent CNN backbones [36, 62], the channel configuration of MobileNetV2 [45] is better aligned with the requirements of our task. Specifically, given a pair of rectified stereo images  $\{\mathbf{I}_L, \mathbf{I}_R\} \in \mathbb{R}^{H \times W \times 3}$ , we employ two weight-sharing networks to pyramidally generate multi-scale features  $\{\mathbf{F}_L^s\}$  and  $\{\mathbf{F}_R^s\}$ , where  $s \in \{\frac{1}{4}, \frac{1}{8}, \frac{1}{16}, \frac{1}{32}\}$  denotes the downsampling ratio. To unify the spatial resolution for subsequent processing, features at all scales are upsampled to  $\frac{1}{4}$  resolution using residual upsampling blocks, following [20].

**Correlation.** Given the extracted left and right feature maps

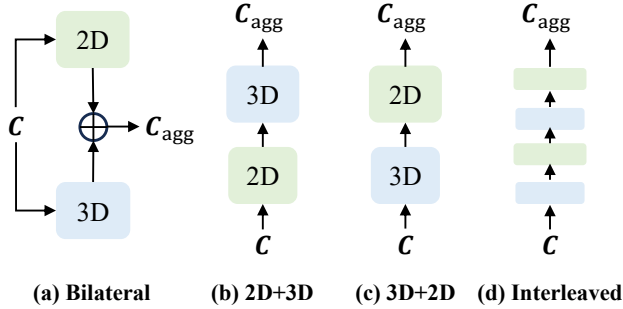


Figure 5. Design choices for hybrid cost aggregation module.

$\mathbf{F}_L^{\frac{1}{4}}$  and  $\mathbf{F}_R^{\frac{1}{4}}$ , the cost volume  $\mathbf{C}$  for each disparity  $d$  within the range  $[0, D_{\max}/4]$  is constructed as:

$$\mathbf{C}(d, h, w) = \frac{1}{N_c} \left\langle \mathbf{F}_L^{\frac{1}{4}}(h, w), \mathbf{F}_R^{\frac{1}{4}}(h, w - d) \right\rangle, \quad (1)$$

where  $D_{\max}$  denotes the predefined maximum disparity value,  $\langle \cdot, \cdot \rangle$  denotes the inner product,  $N_c$  is the number of channels, and  $(h, w)$  represents the pixel location.

**Cost Aggregation.** Recent efficient models adopt 2D cost aggregation [20, 48, 68], which reduces computation by collapsing the disparity into channels and performing spatial-only convolutions. However, this strategy does not explicitly model structured continuity along the disparity dimension, potentially limiting generalization. To address this, we introduce a hybrid aggregation module that combines complementary 2D and 3D reasoning. The 3D block  $\mathbf{G}_{3D}(\cdot)$  jointly aggregates over disparity and spatial dimensions to capture cross-disparity structure, while the 2D block  $\mathbf{G}_{2D}(\cdot)$  provides efficient spatial refinement. This design preserves geometric awareness with minimal overhead.

Specifically, multi-scale 3D convolutions are employed in the 3D block, and ConvNeXt layers [36] are used in the 2D block. A naive half-half hybrid design is ineffective, as 3D convolutions dominate the compute budget along the disparity dimension, even though many disparity levels contribute little to the final results. Instead, we retain only a small proportion of 3D component to preserve disparity perception while maintaining efficiency. We also explore multiple design choices for integrating 2D and 3D branches while keeping the overall computational cost fixed, as illustrated in Fig. 5. In particular: (a) follows [68], where bilateral aggregation is adopted and the outputs of 2D and 3D branches are summed; (b) and (c) apply serial connections, but with different ordering of 2D and 3D blocks; (d) interleaves 2D and 3D blocks. Through ablation studies (Section 4.3), we find that design (c) achieves the best performance. Therefore, we adopt (c) as the default configuration of our cost aggregation module, formulated as follows,

$$\mathbf{C}_{\text{agg}} = \mathbf{G}_{2D}(\mathbf{G}_{3D}(\mathbf{C})). \quad (2)$$

**Disparity Estimation.** Similar to other efficient methods

Table 1. Overview of the real-world stereo datasets (0.5M samples in total) used for training our model.

Dataset	Indoor	Outdoor	MPix	Images
Flickr1024 [59]	✓	✓	0.73	1K
InStereo2k [2]	✓		0.93	2K
Holopix50K [22]	✓	✓	0.74	49K
Driving Stereo [72]		✓	0.40	174K
SouthKenSV [27]	✓	✓	0.92	113K
UASOL [4]		✓	2.74	156K




Figure 6. Effects of the proposed three stages training strategy.

[20, 68], we apply the soft-argmax operation to regress the disparity map  $\mathbf{d}$  at  $\frac{1}{4}$  scale:

$$\mathbf{d} = \sum_{d=0}^{D_{\max}/4-1} d \times \sigma(\mathbf{C}_{\text{agg}}(d)), \quad (3)$$

where  $\sigma(\cdot)$  is a softmax layer. Convex upsampling is then used to upsample  $\mathbf{d}$  to the full-resolution  $\mathbf{D} \in \mathbb{R}^{H \times W}$ .

### 3.2. Training Strategy

To achieve zero-shot capability, the proposed model is trained on a large amount of high-quality data by following carefully designed training stages, which we illustrate in Fig. 4. We compile datasets into two categories: synthetic annotated data and realistic unannotated data. Since synthetic datasets provide accurate ground-truth annotations, we first train our model on them from scratch to enhance model’s matching ability. Our synthetic datasets combine SceneFlow [38] (35K), FallingThings [53] (30K), FSD [61] (1.1M), CREStereo [32] (0.2M), VKITTI2 [6] (21K), TartanAir [57] (0.31M), and Dynamic Replica [29] (0.14M), resulting in a total of 1.8M annotated samples. While other synthetic datasets are available, such as IRS [55], Sintel [5], Spring [39], and InfinigenSV [27], they are excluded due to annotation quality issues (e.g., inaccurate disparities for transparent objects in IRS) or domain mismatch (e.g., cinematic or natural scenes in Spring and InfinigenSV).

In **Stage 1**, the model is trained in a supervised, end-to-end manner without data augmentation, using the commonly used disparity loss  $\mathcal{L}_{\text{disp}}$ :

$$\mathcal{L}_{\text{disp}} = \text{smooth}_{L_1}(\mathbf{D} - \mathbf{D}_{gt}), \quad (4)$$

where  $\mathbf{D}_{gt}$  denotes the ground-truth disparity.

Table 2. Ablation study on KITTI 2012 [14] (K.12), KITTI 2015 [40] (K.15), ETH3D [47] (E.), and Middlebury [46] (M.). We report D1 for KITTI, Bad 1 for ETH3D, and Bad 2 for Middlebury. We also report MACs (Multiply-Accumulate Operations) for backbone design choices. Unless otherwise specified, models are trained for 200K iterations without data augmentation on a 1.4M-image subset from synthetic datasets [6, 32, 38, 53, 61] using the default operations in [68]. Default settings in final model are marked in gray.

(a) Design choices for cost aggregation module, as shown in Fig. 5. Introducing 3D block before 2D block in serial connection is effective.

case	K.12	K.15	E.	M.	MACs (G)
2D	5.02	5.01	6.48	11.29	<b>32.9</b>
bilateral	5.10	5.10	8.55	12.00	35.8
2D-3D	4.73	<b>4.44</b>	28.85	11.52	33.9
3D-2D	4.78	4.64	<b>5.39</b>	<b>10.89</b>	33.9
interleaved	<b>4.61</b>	4.73	6.20	11.34	35.2

(c) Layer choices for 2D block in cost aggregation module. ConvNext is effective with a lowest computation cost.

case	K.12	K.15	E.	M.	MACs (G)
MobileNet v2	4.78	<b>4.64</b>	5.39	10.89	33.9
ConvNeXt	<b>4.38</b>	4.84	5.75	<b>9.50</b>	<b>32.8</b>
ConvNeXt v2	4.79	4.81	<b>5.03</b>	10.52	34.0

(e) Training strategy. knowledge distillation is more effective than direct data augmentation.

case	K.12	K.15	E.	M.
none	4.38	4.84	5.75	9.50
data aug.	4.31	4.79	<b>5.73</b>	11.14
know. dis.	<b>3.64</b>	<b>4.63</b>	6.82	<b>8.86</b>

(f) Knowledge distillation choices in stage ②. Fixed teacher’s weights is more effective.

case	K.12	K.15	E.	M.
EMA	3.97	4.82	6.91	9.78
hard copy	4.22	4.71	<b>6.35</b>	10.07
fixed	<b>3.64</b>	<b>4.63</b>	6.82	<b>8.86</b>

(g) Effects of the three stages training strategy on full training set. See Fig. 6 for visualization.

case	K.12	K.15	E.	M.
stage ①	4.05	4.55	4.43	8.49
stage ②	3.66	4.53	4.69	<b>7.03</b>
stage ③	<b>3.04</b>	<b>3.87</b>	<b>3.53</b>	7.51

(b) 3D kernel size for cost aggregation module. Enlarging the kernel size along disparity dimension is not effective. We keep the general (3, 3, 3).

3D kernel	K.12	K.15	E.	M.	MACs (G)
(3, 3, 3)	<b>4.38</b>	<b>4.84</b>	<b>5.75</b>	<b>9.50</b>	32.8
(5, 3, 3)	4.65	4.88	6.29	10.36	32.5
(7, 3, 3)	4.84	4.85	6.37	9.64	32.6
(11, 3, 3)	4.86	4.87	6.87	10.39	32.8
(3, 1, 1)	4.70	4.85	6.63	10.06	<b>32.1</b>

(d) 3D proportion for cost aggregation module. 4.8% is effective, while a larger one can be detrimental under the same computational budget.

3D proportion	K.12	K.15	E.	M.	MACs (G)
4.8%	<b>4.38</b>	4.84	5.75	<b>9.50</b>	32.8
9.5%	4.71	<b>4.51</b>	5.81	10.06	<b>32.6</b>
15.6%	4.49	4.68	<b>5.54</b>	10.34	32.7

In **Stage ②**, we introduce a self-distillation strategy to improve feature robustness. Both teacher and student models have the same architectures initialized from the first stage. The teacher model receives clean inputs, while the student model is exposed to strongly perturbed inputs (see supplementary materials), encouraging domain-invariant representation learning. We employ a feature alignment loss  $\mathcal{L}_{feat}$  in addition to  $\mathcal{L}_{disp}$ :

$$\mathcal{L}_{feat} = 1 - \frac{1}{HW} \sum_{i=1}^{HW} \cos(F_i, F'_i), \quad (5)$$

where  $F_i$  and  $F'_i$  are feature vectors from the teacher and student models, respectively. Here, we evaluate several distillation schemes: (a) training only the student model with fixed teacher model weights, (b) updating the teacher model via Exponential Moving Average (EMA) [43], and (c) directly copying student model weights to the teacher model at each iteration. Our ablation study (Section 4.3) shows that the simplest strategy (a) yields the best performance, and is therefore adopted in the second stage.

High-quality real-world stereo data with annotations remain scarce and sparse (e.g., Lidar-based ground truth), which constrains training scalability. In contrast, the vast amount of unlabeled real-world stereo data has not been used by the community. With the increasing availability and diversity of such data, in **Stage ③**, we leverage unannotated stereo pairs to improve model generalization, collecting a total of 0.5M pairs, as summarized in Tab. 1. Pseudo labels

are generated using FoundationStereo [61], a strong accurate model. Importantly, even for datasets with ground-truth annotations (e.g., DrivingStereo [72]), we use pseudo but dense labels for stereo pairs (excluding the weather subset), as they provide stronger supervision than sparse data.

We further observe that data quality is more critical than scale: incorporating low-quality or domain-specific data in the third stage can degrade zero-shot performance. For example, Stereo4D [24] contains 18M stereo pairs mined from internet videos but only at a limited resolution ( $512 \times 512$ ), HRWSI [63] suffers from poor rectification quality, and datasets such as SCOD [41] are restricted to narrow domains. Finally, we do not apply self-distillation here, as it provides no observable performance gains when training with pseudo labels. As shown in Fig. 6, the proposed strategy progressively improves the disparity estimation, making the leaves and background regions noticeably clearer.

## 4. Experiments

### 4.1. Benchmarks, Metrics, and Baselines

**Benchmarks.** We evaluate our method on five widely-used real-world stereo datasets. Middlebury [46] is an indoor dataset containing 15 stereo pairs with high-quality ground truth captured using structured light. In the main paper, we report results under half-resolution and non-occluded settings; full-resolution and quarter-resolution results are provided in the supplementary materials. ETH3D [47] consists

Table 3. Zero-shot generalization results on four public benchmarks: KITTI 2012 [14], KITTI 2015 [40], ETH3D [47], and Middlebury (H) [46]. The most commonly used metrics are adopted. In the first block, all efficient methods are trained only on Scene Flow [38]. In the second and third blocks, methods are allowed to train on any existing datasets excluding the four target domains. Accuracy-based methods are shown as reference. The weights and parameters are fixed for evaluation. MACs (G) are measured at the KITTI resolution of  $1242 \times 375$ . \* indicates models retrained with original code on the same synthetic data as ours. † denotes results trained on 30M pseudo-labeled samples using the strategy in [19]. ‡ marks results reported by [61]. The **best** and **second best** are marked with colors.

Method	KITTI 2012		KITTI 2015		ETH3D		Middlebury		MACs (G)
	D1	EPE	D1	EPE	Bad 1.0	EPE	Bad 2.0	EPE	
<i>Efficient methods: SceneFlow</i>									
CoEX [1]	22.30	3.37	17.33	3.00	31.97	14.05	26.42	4.90	54
MobileStereoNet-2D [48]	19.30	2.51	21.88	2.85	17.10	1.83	37.98	7.54	127
MobileStereoNet-3D [48]	19.59	2.67	17.89	3.01	18.42	3.02	25.26	4.61	564
FastACV [65]	13.90	2.05	11.83	2.21	7.84	1.05	19.61	4.66	72
FastACV+ [67]	16.69	2.29	15.28	3.27	11.50	2.43	27.34	7.16	85
Lite-CREStereo++ [25]	5.93	1.29	7.37	1.36	<b>8.95</b>	1.18	14.91	3.32	101
LightStereo-M [20]	6.76	1.26	6.79	1.42	13.93	0.73	16.99	2.06	<b>33</b>
LightStereo-L [20]	6.80	1.27	6.62	<b>1.29</b>	9.66	<b>0.50</b>	17.23	2.88	84
BANet-2D [68]	14.75	2.23	16.98	3.91	44.89	35.95	26.79	6.96	36
BANet-3D [68]	17.39	2.41	17.28	3.36	29.27	14.06	28.79	8.05	78
<b>Lite Any Stereo</b>	<b>5.45</b>	<b>1.18</b>	<b>6.45</b>	1.32	15.38	0.77	<b>13.13</b>	<b>1.60</b>	<b>33</b>
<i>Efficient methods: Million-scale</i>									
LightStereo-M* [20]	4.10	0.99	4.97	1.13	5.33	0.41	10.85	1.51	<b>33</b>
BANet-2D* [68]	3.90	0.93	4.71	1.07	5.92	0.38	10.05	1.34	36
StereoAnything-L† [19]	4.00	0.92	4.81	1.10	3.81	<b>0.31</b>	9.82	1.21	84
<b>Lite Any Stereo</b>	<b>3.04</b>	<b>0.79</b>	<b>3.87</b>	<b>0.99</b>	<b>3.53</b>	0.32	<b>7.51</b>	<b>0.94</b>	<b>33</b>
<i>Accurate methods: Million-scale</i>									
Selective-IGEV‡ [58]	3.20	–	4.50	–	3.40	–	7.50	–	3619
FoundationStereo [61]	2.51	0.67	2.83	0.86	0.49	0.14	1.12	0.37	12824

Table 4. Comparison of model performance on DrivingStereo weather [72]. Lower value is better for both metrics.

Method	Rainy		Sunny		Foggy		Cloudy		Overall		MACs (G)
	D1	EPE	D1	EPE	D1	EPE	D1	EPE	D1	EPE	
FoundationStereo [61]	27.01	3.96	4.31	1.57	7.67	1.82	3.85	1.53	10.71	2.22	12824
<b>Lite Any Stereo</b>	<b>20.69</b>	<b>2.61</b>	<b>3.84</b>	<b>1.47</b>	<b>6.78</b>	<b>1.64</b>	<b>3.65</b>	<b>1.47</b>	<b>8.74</b>	<b>1.80</b>	<b>33</b>

of 27 grayscale stereo pairs with laser-scanned ground truth, covering both indoor and outdoor scenes. KITTI 2012 [14] and KITTI 2015 [40] include 194 and 200 stereo pairs, respectively, captured in outdoor driving environments, with ground truth obtained via LiDAR. DrivingStereo weather [72] contains driving scene images under four different weather conditions, where each class of weather contains 500 frames. We report results at full resolution. Note that none of these datasets are used in the training process.

**Evaluation Metrics.** For all datasets, we report the average End-Point Error (EPE), which measures the mean per-pixel disparity error. For Middlebury and ETH3D, we also report the percentage of pixels with disparity errors greater than a threshold  $X$  (Bad- $X$ ). For KITTI datasets, we report the D1 error, defined as the percentage of pixels with disparity error exceeding 3 pixels and 5% of the ground truth disparity.

**Baselines.** To ensure fair comparison, we re-evaluate all baseline methods under consistent settings on our local machine. This avoids discrepancies caused by differing bench-

mark configurations, such as occlusion masking and metric definitions (e.g., D1 vs. Bad-3.0). Unless otherwise noted, we use the official checkpoints provided in each method.

## 4.2. Implementation Details

Lite Any Stereo is implemented using PyTorch. The model is trained for 150K, 50K, and 100K steps in stage 1, stage 2, and stage 3 with a total batch size of 176 on NVIDIA A100 GPUs. We adopt the AdamW optimizer [31] with a one-cycle learning rate schedule, where peak learning rate is set to  $2 \times 10^{-4}$ . During training, input images are randomly cropped to  $256 \times 512$ , and then finetuned on  $320 \times 736$ . The maximum disparity value  $D_{\max}$ , is set to 192, following the same configuration in previous works [20, 68].

## 4.3. Ablation Study

We investigate various design choices and training strategies for our model. Unless otherwise specified, all models are trained for 100K iterations without data augmentation




Figure 7. Qualitative comparison of zero-shot inference on in-the-wild images. Each column shows disparity predictions from different methods using the same checkpoint (Tab. 3, second block), trained on million-scale data. The selected scenarios contain diverse real-world challenges, including reflections, complex lighting conditions, non-textured and repetitive textures. Comparison methods produce blurry boundaries, texture-copy artifacts, or fail to recover geometric details, while our method consistently preserves structural integrity, produces smooth yet accurate disparity maps, and demonstrates superior zero-shot generalization across diverse scenarios.

on a 1.4M-image subset of synthetic datasets [6, 32, 38, 53, 61]. We report D1 errors on KITTI 2012 [14] (K.12) and KITTI 2015 [40] (K.15), Bad 1.0 error on ETH3D [47] (E.), and Bad 2.0 error on Middlebury [46] (M.), as in Tab. 2.

**Cost Aggregation Design.** In Tab. 2(a), we first compare different aggregation choices (Fig. 5) and observe that the 3D–2D hybrid design is effective. We then evaluate alternative 2D layer choices in Tab. 2(c), and find that ConvNeXt layers [36] provide the best trade-off between accuracy and efficiency. Building upon the ConvNeXt-based 2D layers, we further evaluate the influence of 3D convolution kernel sizes (Tab. 2(b)). Unlike the findings in [61], where larger kernels improved performance, our results show that standard (3, 3, 3) configuration achieves the highest accuracy in our baseline. Finally, as shown in Tab. 2(d), increasing the proportion of 3D blocks leads to performance degradation, particularly on the Middlebury dataset under limited MACs. Therefore, we retain only a small 3D component to preserve disparity perception while maintaining efficiency.

**Training Strategy Choices.** As shown in Tab. 2(e) and (f), we compare different training strategies in stage ②. Our self-distillation approach yields more robust feature representations than data augmentation, leading to stronger performance. Furthermore, keeping the teacher model’s weights fixed proves more effective than using EMA or hard copy updates. Tab. 2(g) highlights the benefits of our three-stage training scheme, which overall improves performance. The slight drop on Middlebury in stage 3 may be

attributed to the limited scale of the indoor real-world data.

#### 4.4. Evaluation

**Zero-Shot Generalization.** Tab. 3 presents a quantitative comparison of zero-shot generalization performance across four public datasets. In the first block of the table, we restrict our model training to stage 2 on SceneFlow to ensure a fair comparison. Under this setting, our method achieves superior performance on most metrics while requiring the fewest MACs, highlighting the effectiveness of both model design and training strategy. We then evaluate under a more realistic setting as in [61], where methods are allowed to train on any available dataset except the target domains. In this setup, our approach consistently outperforms the other methods by a large margin. Our model achieves comparable or better results than non-prior-based accurate method [58] with less than 1% of its MACs. Fig. 7 presents qualitative comparison on in-the-wild, high-resolution (2K) images. Our method generates smooth yet accurate disparity maps across different real-world scenes, whereas the other methods struggle to preserve geometric details.

**In-Domain Performance.** Tab. 4 reports the results on the DrivingStereo weather subset. Interestingly, although our model is only trained via knowledge distillation on general datasets, it surpasses the teacher model (FoundationStereo [61]) by a significant margin. This highlights the strong effectiveness of our proposed approach. Although in-domain performance is not the primary focus of this work, we also

Table 5. Results on KITTI 2012 [14] and KITTI 2015 [40] leaderboard. The MACs are measured for an input size of  $1242 \times 375$ . Our finetuned version ranks 1st on the leaderboards among all efficient methods at the time of submission.

Method	KITTI 2012					KITTI 2015			MACs (G)
	3-noc	3-all	4-noc	4-all	EPE-noc / all	D1-bg	D1-fg	D1-all	
DispNetC [38]	4.11	4.65	2.77	3.20	0.9 / 1.0	4.32	4.41	4.34	-
DeepPruner-Fast [13]	-	-	-	-	-	2.32	3.91	2.59	194
AANet+ [69]	1.55	2.04	1.20	1.58	<b>0.4 / 0.5</b>	1.65	3.96	2.03	-
DecNet [75]	-	-	-	-	-	2.07	3.87	2.37	-
BGNet+ [64]	1.62	2.03	1.16	1.48	0.5 / 0.6	1.81	4.09	2.19	77
HITNet [52]	1.41	1.89	1.14	1.53	<b>0.4 / 0.5</b>	1.74	3.20	1.98	47
CoEx [1]	1.55	1.93	1.15	1.42	0.5 / 0.5	1.79	3.82	2.13	54
MobileStereoNet-2D [48]	-	-	-	-	-	2.49	4.53	2.83	127
MobileStereoNet-3D [48]	-	-	-	-	-	2.75	3.87	2.10	564
Fast-ACVNet [65]	1.68	2.13	1.23	1.56	0.5 / 0.6	1.82	3.93	2.17	72
Fast-ACVNet+ [67]	1.45	1.85	1.06	1.36	0.5 / 0.5	1.70	3.53	2.01	85
Lite-CREStereo++ [25]	1.43	1.82	1.12	1.44	0.5 / 0.5	1.79	3.53	2.08	93
LightStereo-M [20]	1.56	1.91	1.10	1.36	0.5 / 0.5	1.81	3.22	2.04	<b>33</b>
LightStereo-L [20]	1.55	1.87	1.10	1.33	0.5 / 0.5	1.78	<b>2.64</b>	1.93	84
BANet-2D [68]	1.38	1.79	1.01	1.32	0.5 / 0.5	1.59	3.03	1.83	36
BANet-3D [68]	1.27	1.72	0.95	1.27	0.5 / 0.5	1.52	3.02	1.77	78
<b>Lite Any Stereo</b>	<b>1.09</b>	<b>1.49</b>	<b>0.76</b>	<b>1.04</b>	<b>0.4 / 0.5</b>	<b>1.36</b>	3.45	<b>1.71</b>	<b>33</b>

Table 6. Performance of the proposed training strategy.

Method	K.12	K.15	E.	M.
<i>LightStereo-M</i> [20]				
Stage ①	4.34	5.27	6.68	10.29
Stage ②	3.80	4.62	5.44	<b>8.96</b>
Stage ③	<b>3.35</b>	<b>4.14</b>	<b>4.22</b>	9.85
<i>BANet-2D</i> [68]				
Stage ①	4.34	4.78	7.71	10.54
Stage ②	3.87	4.80	4.86	<b>9.54</b>
Stage ③	<b>3.28</b>	<b>4.08</b>	<b>4.05</b>	10.30

evaluate our model on the KITTI online leaderboard (test set) trained on [54] without using its original annotations. As can be seen in Tab. 5, our model achieves the highest accuracy among all published efficient methods across almost all metrics at the time of submission.

**Effectiveness of the Training Strategy.** To verify the generality of our training strategy, we apply it to two representative efficient methods, LightStereo-M [20] and BANet-2D [68]. In Tab. 6, our strategy consistently improves both methods across multiple benchmarks, which shows its broad applicability to different architectures.

**Inference Time Analysis.** Prior works report inference times measured on different hardware, making direct comparisons unreliable. To ensure fairness, we benchmark all methods under the same hardware setup. In Tab. 7, our method consistently achieves the fastest speed across different GPUs, highlighting its superior efficiency. Notably, it requires only 2.5 GB of GPU memory for 2K inputs, indicating strong potential for deployment on embedded devices, drones, and other resource-constrained platforms.

Table 7. Inference time (ms) evaluated locally on different GPU hardware under the same experimental setup.

Method	1080	4090	A5000	A100
Lite-CREStereo++ [25]	63	38	55	40
LightStereo-M [20]	24	22	27	21
LightStereo-L [20]	34	30	35	27
BANet-2D [68]	58	30	63	27
BANet-3D [68]	40	25	44	22
<b>Lite Any Stereo</b>	<b>21</b>	<b>19</b>	<b>23</b>	<b>17</b>

## 4.5. Limitations

While our method improves generalization, it still trails depth-prior-based approaches. A major bottleneck is the limited availability of high-quality real-world stereo data, which restricts performance gains. For instance, accuracy on Middlebury drops. Robustness under challenging cases such as transparency and reflection also can be improved.

## 5. Conclusion

We presented Lite Any Stereo, an efficient zero-shot stereo matching model. By integrating a compact backbone with a three-stage training strategy, our model generalizes well to in-the-wild scenarios while maintaining high efficiency. Extensive evaluations demonstrate state-of-the-art performance across multiple benchmarks, underscoring the potential of lightweight models for broad real-world deployment. Future work includes scaling real-world data and developing a model zoo with different model sizes for flexible deployment across diverse computational budgets.

**Acknowledgments.** This work was funded by the IC-CSC.

## References

- [1] Antyanta Bangunharcana, Jae Won Cho, Seokju Lee, In So Kweon, Kyung-Soo Kim, and Soohyun Kim. Correlate-and-excite: Real-time stereo matching via guided cost volume excitation. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3542–3548. IEEE, 2021. 2, 6, 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(11):1–11, 2020. 4
- [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. 2
- [4] Zuria Bauer, Francisco Gomez-Donoso, Edmanuel Cruz, Sergio Orts-Escolano, and Miguel Cazorla. Uasol, a large-scale high-resolution outdoor stereo dataset. *Scientific data*, 6(1):162, 2019. 4
- [5] D. J. Butler, J. Wulff, G. B. Stanley, and M. J. Black. A naturalistic open source movie for optical flow evaluation. In *ECCV*, pages 611–625, 2012. 4
- [6] Yohann Cabon, Naila Murray, and Martin Humenberger. Virtual kitti 2, 2020. 4, 5, 7
- [7] Jia-Ren Chang and Yong-Sheng Chen. Pyramid stereo matching network. In *CVPR*, pages 5410–5418, 2018. 2
- [8] 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. 2
- [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, 2025. 1, 2, 3
- [10] Xuelian Cheng, Yiran Zhong, Mehrtash Harandi, Yuchao Dai, Xiaojun Chang, Tom Drummond, Hongdong Li, and Zongyuan Ge. Hierarchical neural architecture search for deep stereo matching. *arXiv preprint arXiv:2010.13501*, 2020. 2
- [11] 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. 2
- [12] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 3
- [13] Shivam Duggal, Shenlong Wang, Wei-Chiu Ma, Rui Hu, and Raquel Urtasun. Deeppruner: Learning efficient stereo matching via differentiable patchmatch. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4384–4393, 2019. 2, 8
- [14] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *CVPR*, pages 3354–3361, 2012. 1, 2, 3, 5, 6, 7, 8
- [15] 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. 2
- [16] Weiyu Guo, Zhaoshuo Li, Yongkui Yang, Zheng Wang, Russell H Taylor, Mathias Unberath, Alan Yuille, and Yingwei Li. Context-enhanced stereo transformer. In *European Conference on Computer Vision*, pages 263–279. Springer, 2022. 2
- [17] Xiaoyang Guo, Kai Yang, Wukui Yang, Xiaogang Wang, and Hongsheng Li. Group-wise correlation stereo network. In *CVPR*, pages 3273–3282, 2019. 2
- [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. 2
- [19] Xianda Guo, Chenming Zhang, Youmin Zhang, Dujun Nie, Ruilin Wang, Wenzhao Zheng, Matteo Poggi, and Long Chen. Stereo anything: Unifying stereo matching with large-scale mixed data. *arXiv preprint arXiv:2411.14053*, 2024. 2, 6
- [20] Xianda Guo, Chenming Zhang, Youmin Zhang, Wenzhao Zheng, Dujun Nie, Matteo Poggi, and Long Chen. Lightstereo: Channel boost is all your need for efficient 2d cost aggregation, 2024. 1, 2, 3, 4, 6, 8
- [21] Xianda Guo, Chenming Zhang, Ruilin Wang, Zhang Youmin, Wenzhao Zheng, Matteo Poggi, Hao Zhao, Qin Zou, and Long Chen. Stereocarla: A high-fidelity driving dataset for generalizable stereo. *arXiv preprint arXiv:2509.12683*, 2025. 2
- [22] Yiwen Hua, Puneet Kohli, Prithvi Uplavikar, Anand Ravi, Saravana Gunaseelan, Jason Orozco, and Edward Li. Holopix50k: A large-scale in-the-wild stereo image dataset. In *CVPR Workshop on Computer Vision for Augmented and Virtual Reality, Seattle, WA, 2020.*, 2020. 4
- [23] Hualie Jiang, Zhiqiang Lou, Laiyan Ding, Rui Xu, Minglang Tan, Wenjie Jiang, and Rui Huang. Defom-stereo: Depth foundation model based stereo matching, 2025. 1, 2, 3
- [24] Linyi Jin, Richard Tucker, Zhengqi Li, David Fouhey, Noah Snavely, and Aleksander Holynski. Stereo4d: Learning how things move in 3d from internet stereo videos. *arXiv preprint*, 2024. 5
- [25] 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 (ICCV)*, pages 3318–3327, 2023. 2, 3, 6, 8
- [26] Junpeng Jing, Ye Mao, and Krystian Mikolajczyk. Matchstereo-videos: Bidirectional alignment for consistent dynamic stereo matching. In *European Conference on Computer Vision*, pages 415–432. Springer, 2024. 2

- [27] Junpeng Jing, Ye Mao, Anlan Qiu, and Krystian Mikolajczyk. Match stereo videos via bidirectional alignment. *arXiv preprint arXiv:2409.20283*, 2024. 4
- [28] Junpeng Jing, Weixun Luo, Ye Mao, and Krystian Mikolajczyk. Stereo any video: Temporally consistent stereo matching. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 20836–20846, 2025.
- [29] Nikita Karaev, Ignacio Rocco, Benjamin Graham, Natalia Neverova, Andrea Vedaldi, and Christian Ruppert. Dynamicstereo: Consistent dynamic depth from stereo videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13229–13239, 2023. 2, 4
- [30] Sameh Khamis, Sean Fanello, Christoph Rhemann, Adarsh Kowdle, Julien Valentin, and Shahram Izadi. Stereonet: Guided hierarchical refinement for real-time edge-aware depth prediction. In *ECCV*, pages 573–590, 2018. 2
- [31] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [32] 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, 4, 5, 7
- [33] 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. 2
- [34] 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 *CVPR*, pages 2811–2820, 2018. 2
- [35] Lahav Lipson, Zachary Teed, and Jia Deng. Raft-stereo: Multilevel recurrent field transforms for stereo matching. *arXiv preprint arXiv:2109.07547*, 2021. 1, 2
- [36] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11976–11986, 2022. 3, 4, 7
- [37] D Marr and T Poggio. Cooperative computation of stereo disparity. In *Neurocomputing: foundations of research*, pages 259–267. 1988. 1
- [38] 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. 4, 5, 6, 7, 8, 2
- [39] Lukas Mehl, Jenny Schmalfluss, Azin Jahedi, Yaroslava Nalivayko, and Andrés Bruhn. Spring: A high-resolution high-detail dataset and benchmark for scene flow, optical flow and stereo. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 4
- [40] Moritz Menze and Andreas Geiger. Object scene flow for autonomous vehicles. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 1, 2, 3, 5, 6, 7, 8
- [41] Junhong Min and Youngpil Jeon. Confidence aware stereo matching for realistic cluttered scenario. In *2024 IEEE International Conference on Image Processing (ICIP)*, pages 3491–3497. IEEE, 2024. 5
- [42] Jiahao Pang, Wenxiu Sun, Jimmy SJ Ren, Chengxi Yang, and Qiong Yan. Cascade residual learning: A two-stage convolutional neural network for stereo matching. In *CVPRW*, pages 887–895, 2017. 2
- [43] Boris T Polyak and Anatoli B Juditsky. Acceleration of stochastic approximation by averaging. *SIAM journal on control and optimization*, 30(4):838–855, 1992. 5
- [44] Zhibo Rao, Bangshu Xiong, Mingyi He, Yuchao Dai, Renjie He, Zhelun Shen, and Xing Li. Masked representation learning for domain generalized stereo matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5435–5444, 2023. 2
- [45] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4510–4520, 2018. 3
- [46] Daniel Scharstein and Richard Szeliski. A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *IJCV*, 47(1):7–42, 2002. 1, 2, 5, 6, 7
- [47] 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 *CVPR*, pages 3260–3269, 2017. 1, 2, 5, 6, 7
- [48] Faranak Shamsafar, Samuel Woerz, Rafia Rahim, and Andreas Zell. Mobilestereonet: Towards lightweight deep networks for stereo matching. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2417–2426, 2022. 1, 3, 4, 6, 8, 2
- [49] Zhelun Shen, Yuchao Dai, and Zhibo Rao. Cfnet: Cascade and fused cost volume for robust stereo matching. In *CVPR*, pages 13906–13915, 2021. 2
- [50] Qing Su and Shihao Ji. Chitransformer: Towards reliable stereo from cues. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1939–1949, 2022. 2
- [51] Tatsunori Tanai, Yasuyuki Matsushita, Yoichi Sato, and Takeshi Naemura. Continuous 3d label stereo matching using local expansion moves. *IEEE TPAMI*, 40(11):2725–2739, 2017. 1
- [52] 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 *CVPR*, pages 14362–14372, 2021. 2, 3, 8
- [53] Jonathan Tremblay, Thang To, and Stan Birchfield. Falling things: A synthetic dataset for 3d object detection and pose estimation. In *CVPRW*, pages 2038–2041, 2018. 4, 5, 7
- [54] Jonas Uhrig, Nick Schneider, Lukas Schneider, Uwe Franke, Thomas Brox, and Andreas Geiger. Sparsity invariant cnns.

- In *2017 international conference on 3D Vision (3DV)*, pages 11–20. IEEE, 2017. 8
- [55] 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. *arXiv preprint arXiv:1912.09678*, 2019. 4
- [56] Qiang Wang, Shaohuai Shi, Shizhen Zheng, Kaiyong Zhao, and Xiaowen Chu. FADNet: A fast and accurate network for disparity estimation. In *2020 IEEE International Conference on Robotics and Automation (ICRA 2020)*, pages 101–107, 2020. 2
- [57] 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. 2020. 4
- [58] 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, 6, 7
- [59] Yingqian Wang, Longguang Wang, Jungang Yang, Wei An, and Yulan Guo. Flickr1024: A large-scale dataset for stereo image super-resolution. In *International Conference on Computer Vision Workshops*, pages 3852–3857, 2019. 4
- [60] Philippe Weinzaepfel, Thomas Lucas, Vincent Leroy, Johann 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. 2
- [61] Bowen Wen, Matthew Trepte, Joseph Aribido, Jan Kautz, Orazio Gallo, and Stan Birchfield. Foundationstereo: Zero-shot stereo matching, 2025. 1, 2, 3, 4, 5, 6, 7
- [62] Sanghyun Woo, Shoubhik Debnath, Ronghang Hu, Xinlei Chen, Zhuang Liu, In So Kweon, and Saining Xie. Convnext v2: Co-designing and scaling convnets with masked autoencoders. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16133–16142, 2023. 3
- [63] Ke Xian, Jianming Zhang, Oliver Wang, Long Mai, Zhe Lin, and Zhiguo Cao. Structure-guided ranking loss for single image depth prediction. In *The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 5
- [64] Bin Xu, Yuhua Xu, Xiaoli Yang, Wei Jia, and Yulan Guo. Bilateral grid learning for stereo matching networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12497–12506, 2021. 2, 8
- [65] 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. 2, 6, 8
- [66] 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. 2
- [67] Gangwei Xu, Yun Wang, Junda Cheng, Jinhui Tang, and Xin Yang. Accurate and efficient stereo matching via attention concatenation volume. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. 1, 2, 6, 8
- [68] 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, 3, 4, 5, 6, 8, 2
- [69] Haofei Xu and Juyong Zhang. Aanet: Adaptive aggregation network for efficient stereo matching. In *CVPR*, pages 1959–1968, 2020. 2, 3, 8
- [70] Haofei Xu, Jing Zhang, Jianfei Cai, Hamid Rezaatofghi, Fisher Yu, Dacheng Tao, and Andreas Geiger. Unifying flow, stereo and depth estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. 2
- [71] Gengshan Yang, Joshua Manela, Michael Happold, and Deva Ramanan. Hierarchical deep stereo matching on high-resolution images. In *CVPR*, pages 5515–5524, 2019. 2
- [72] Guorun Yang, Xiao Song, Chaoqin Huang, Zhidong Deng, Jianping Shi, and Bolei Zhou. Drivingstereo: A large-scale dataset for stereo matching in autonomous driving scenarios. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 899–908, 2019. 4, 5, 6
- [73] 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. 1, 2, 3
- [74] Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything v2. *arXiv preprint arXiv:2406.09414*, 2024. 1, 2, 3
- [75] Chengtang Yao, Yunde Jia, Huijun Di, Pengxiang Li, and Yuwei Wu. A decomposition model for stereo matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6091–6100, 2021. 8
- [76] Feihu Zhang, Victor Prisacariu, Ruigang Yang, and Philip HS Torr. Ga-net: Guided aggregation net for end-to-end stereo matching. In *CVPR*, pages 185–194, 2019. 2
- [77] 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. 2
- [78] Yongjian Zhang, Longguang Wang, Kunhong Li, Yun Wang, and Yulan Guo. Learning representations from foundation models for domain generalized stereo matching. In *European Conference on Computer Vision*, pages 146–162. Springer, 2024. 2
- [79] Jingyi Zhou, Haoyu Zhang, Jiakang Yuan, Peng Ye, Tao Chen, Hao Jiang, Meiya Chen, and Yangyang Zhang. All-in-one: Transferring vision foundation models into stereo matching. *arXiv preprint arXiv:2412.09912*, 2024. 2