

# Stereo Any Video: Temporally Consistent Stereo Matching

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

<https://tontomtommi.github.io/StereoAnyVideo/>




Figure 1. Prediction on a dynamic real-world stereo video. Our method produces temporally consistent and accurate disparities without flickering, enabling the reconstruction of dense and clean metric point clouds for dynamic scenes. See video for better visualization.

## Abstract

*This paper introduces **Stereo Any Video**, a powerful framework for video stereo matching. It can estimate spatially accurate and temporally consistent disparities without relying on auxiliary information such as camera poses or optical flow. The strong capability is driven by rich priors from monocular video depth models, which are integrated with convolutional features to produce stable representations. To further enhance performance, key architectural innovations are introduced: all-to-all-pairs correlation, which constructs smooth and robust matching cost volumes, and temporal convex upsampling, which improves temporal coherence. These components collectively enhance robustness, accuracy, and temporal consistency, establishing a new standard in video stereo matching. Extensive experiments demonstrate that our method achieves state-of-the-*

*art performance across multiple datasets both qualitatively and quantitatively in zero-shot settings, as well as strong generalization to real-world indoor and outdoor scenarios.*

## 1. Introduction

Video stereo matching is a task to reconstruct a 3D scene from a stereo camera. It estimates disparity by identifying the horizontal correspondences between rectified left and right images at each time step. Consistent and accurate disparity estimation is crucial for reconstructing and understanding real-world 3D scenes, benefiting various downstream tasks such as autonomous driving, robotic navigation, and virtual/augmented reality.

With the advances in deep learning, stereo matching has seen significant progress [27, 31, 55, 61], achieving promising performance on public benchmarks [13, 44, 45].

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

However, most existing approaches are designed for image-based estimation and do not explicitly leverage temporal information in video sequences. When applied directly to video data, these methods often suffer from temporal inconsistencies, leading to noticeable flickering and artifacts in disparity maps and reconstructed point clouds. To mitigate this, recent video-based methods incorporate additional temporal cues, such as camera poses or optical flow, through predefined models or learnable mechanisms. For example, CODD [29] employs a pose network to estimate relative camera movements and align adjacent disparities within a fusion framework. Similarly, TemporalStereo [73] and TC-Stereo [70] follow this paradigm with variations in architectural design and alignment operations. Another approach, BiDAStereo [22], integrates optical flow [51] for temporal feature alignment and introduces a temporal aggregation module to propagate information across time.

Despite these advances, existing video methods remain vulnerable to dynamic scenes, camera movement, and the accuracy of auxiliary information. In real-world scenarios, dynamic objects undergo complex motions and deformations that violate traditional multi-view constraints [18], disrupting translational invariance in disparity estimation. Moreover, models relying on auxiliary cues are constrained by the performance of camera pose or optical flow estimation modules. In environments with complex camera movements, these estimation modules degrade, becoming bottlenecks that limit the robustness of video stereo methods.

To address these limitations, we propose Stereo Any Video, a robust framework for video-based stereo matching without relying on camera pose or optical flow. Inspired by recent findings in video generation [33, 48], which emphasize the role of robust feature representations in ensuring temporal coherence, we take a new perspective to design the framework for consistency: feature robustness and stability. Specifically, during feature extraction, we integrate convolutional cues with a frozen monocular video depth foundation model [7] trained on million-scale data. This ensures robust feature maps with rich priors. For correspondence matching, we introduce an all-to-all-pair correlation, which enforces smooth and reliable matching by jointly considering bidirectional correspondences from both target and reference views. For cost aggregation, we propose temporal convex upsampling within a fully temporal-based recurrent unit, playing a key role in maintaining temporal consistency. As shown in Figure 1, our method demonstrates a strong ability to produce accurate and consistent disparities, benefiting the reconstruction for dynamic scenes.

The contributions of our work can be summarized as follows: (1) Stereo Any Video, a framework that integrates monocular video depth priors to achieve accurate and temporally consistent video stereo matching. (2) A novel all-to-all-pair correlation module and a temporal convex upsam-

pling mechanism, ensuring robust and stable cost volume representations and aggregation. (3) State-of-the-art performance on diverse benchmarks, setting a new standard for video stereo matching.

## 2. Related Work

In this section, we provide a brief literature review of stereo matching methods in the deep learning era. We first discuss image-based stereo matching methods, followed by video-based methods, and finally, we review recent advancements leveraging foundation priors.

**Image-based Stereo Matching.** Since the introduction of the first end-to-end stereo network [34], deep learning-based stereo matching has seen substantial advancements. Early methods typically construct 3D [16, 30, 38, 50, 65] or 4D [5, 9, 14, 25, 62, 67, 71] cost volumes from left and right image features, followed by 2D or 3D convolutions for cost aggregation. Inspired by RAFT [51], iterative refinement has become the mainstream approach. Representative advancements include multilevel recurrent units [31], adaptive correlation [27], geometry-encoding volumes [61, 64], and frequency information selection [55], all of which have significantly improved benchmark performance [13, 44, 45]. Transformer-based methods [15, 28, 49, 58, 66] have further refined stereo matching by modeling token correspondence and capturing long-range dependencies. Beyond accuracy, several methods focus on real-world challenges such as efficiency [17, 56, 60, 63], domain generalization [6, 12, 39, 42, 72], and robustness [21, 46, 47]. However, as all these methods are designed for static image pairs, they lack temporal awareness, leading to inconsistent disparity predictions in video sequences.

**Video-based Stereo Matching.** Video stereo matching especially in dynamic scenes presents unique challenges. Early work, such as [75], employed LSTMs with unsupervised learning for stereo videos but did not explicitly incorporate temporal information. CODD [29] introduced a per-frame stereo network, a motion network, and a fusion network designed for dynamic scenes. Neighboring-frame disparities are aligned and fused using estimated camera poses. TemporalStereo [73] extended this approach with a coarse-to-fine network that leverages past context to enhance predictions in challenging scenarios. Cheng *et al.* [10] developed a framework tailored for XR systems, reducing computational costs through temporal cost aggregation. TC-Stereo [70] further introduced a temporal disparity completion mechanism and a temporal state fusion module to exploit temporal information. These methods rely on auxiliary camera motion or pre-defined scene geometry to incorporate temporal information. DynamicStereo [24] introduced a transformer-based architecture that integrates temporal, stereo, and spatial attention. BiDAStereo [22] proposed an optical-flow-based alignment mechanism to propagate

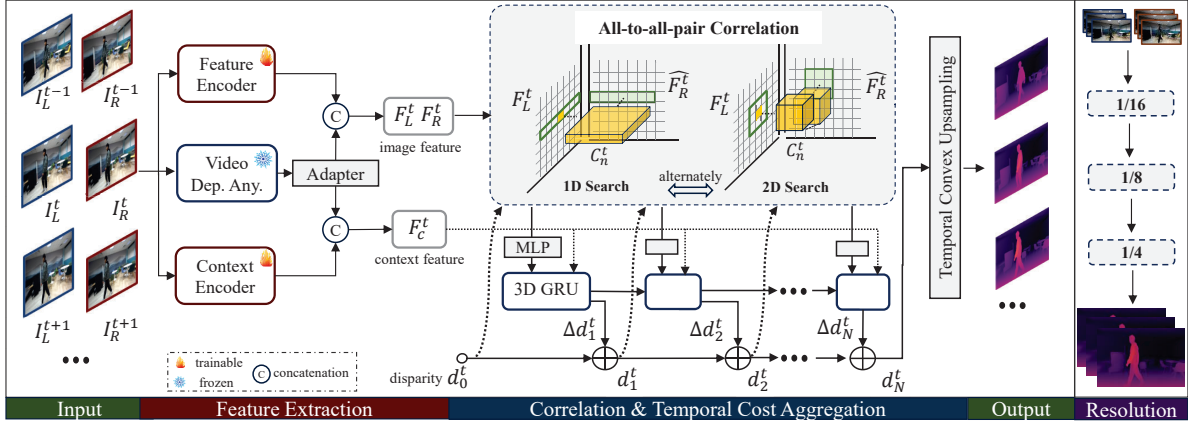


Figure 2. Overview of Stereo Any Video. Given a stereo video sequence as input, the model first extracts features and context information using trainable convolutional encoders and a frozen monocular video depth encoder (Video Depth Anything [7]). At each iteration, an all-to-all-pair correlation computes feature correlations, followed by a multilayer perceptron (MLP) encoder for compact representation. Disparities are iteratively refined using a 3D Gated Recurrent Unit (GRU), integrated with cost volume features, and upsampled via a temporal convex upsampling layer. This process is repeated within a cascaded pipeline to progressively recover full-resolution disparities.

temporal information across the entire sequence. Building on this alignment, BiDAStabilizer [23] introduced a lightweight plug-and-play stabilizer network to improve the temporal consistency of existing image-based stereo matching methods.

**Stereo Matching with Foundation Priors.** Vision foundation models, such as CLIP [41], the DINO series [4, 37], and SAM [26], have been widely applied across various vision tasks due to their rich priors. In stereo matching area, FormerStereo [74] proposed feature transformation and C4 space to adapt ViT-based foundation model to stereo pipeline, learning domain invariant representations. AIOStereo [76] introduced a selective knowledge transfer module to distill knowledge from multiple heterogeneous vision foundation models into a single stereo matching model. More recent approaches primarily focused on leveraging monocular depth priors from Depth Anything [68, 69] to enhance stereo performance. For instance, Stereo Anywhere [1] directly incorporated relative depth maps into the stereo network, demonstrating robustness in challenging conditions such as non-Lambertian surfaces. DEFOM-Stereo [20] employed these depth maps alongside a novel scale updater within a recurrent module for improved alignment. FoundationStereo [59] extracted feature maps from Depth Anything and introduced a side-tuning adapter to adapt monocular features for stereo backbone. MonSter [8] utilized both depth maps and feature maps, and proposed a dual-branched refinement module to fuse stereo and monocular cues effectively. These methods highlight the effectiveness of foundation priors in image-based stereo tasks. In our work, we build upon this direction by investigating video-based stereo matching and exploring how foundation priors can enhance consistency in dynamic scenes.

### 3. Method

In this section, we introduce Stereo Any Video, a feed-forward video-based framework. Given a rectified stereo sequence  $\{I_L^t, I_R^t\}_{t \in (1, T)} \in \mathbb{R}^{H \times W \times 3}$  where  $T$  denotes the total number of frames,  $H$  and  $W$  are the frame height and width, respectively, our goal is to estimate a temporally consistent sequence of disparity maps  $\{d^t\}_{t \in (1, T)} \in \mathbb{R}^{H \times W}$  aligned with the left view. As shown in Figure 2, our method consists of three key components: feature extraction (Section 3.1), correlation (Section 3.2), and temporal cost aggregation (Section 3.3).

#### 3.1. Feature Extraction with Foundation Priors

Unlike image-based tasks, maintaining temporal consistency is the primary challenge in video stereo matching. Recent studies in video generation [33, 48] indicate that substantial variations across different frames can lead to unstable feature representations, thereby degrading temporal coherence. Therefore, ensuring feature stability across adjacent frames is crucial for achieving temporally consistent disparity estimation. Towards this, we leverage Video Depth Anything (VDA) [7] for feature extraction, which offers two key advantages: (1) trained on internet-scale data, VDA provides rich priors that can enhance disparity estimation in regions with ambiguous stereo correspondence; (2) unlike image-based monocular depth models [68, 69], VDA generates temporally stable features through an explicit temporal alignment process before the final prediction head, thereby preserving coherence across frames.

Besides, we investigate whether the depth maps from monocular video depth models can further improve stereo performance. We compare DepthCrafter [19], a representative monocular video model, with RAFTStereo [31], an

Table 1. Quantitative comparison of DepthCrafter [19] and RAFTStereo [31] on Sintel dataset [2].

Method	Sintel Clean		
	EPE ↓	TEPE ↓	MOS (human) ↓
DepthCrafter [19]	8.68	2.08	0.212
RAFTStereo [31]	1.42	0.84	0.933




Figure 3. Qualitative comparison of DepthCrafter [19] and RAFTStereo [31] on Sintel dataset [2].

image-based stereo model. Using open-source weights, we evaluate their performance on Sintel [2], as shown in Table 1 and Figure 3. Specifically, we first obtain relative depth maps from DepthCrafter (without scale and shift) and compute disparities using RAFTStereo on the same sequence. The relative depth maps are then aligned to RAFTStereo’s disparities, and objective metrics are computed against the ground truth. Additionally, we conduct a subjective user study (see supplementary material for details). As observed, while DepthCrafter produces visually plausible results—reflected in better human scores—its spatial and temporal errors are significantly larger than those of RAFTStereo. This suggests that its relative depth is “consistent but wrong”, making it unreliable for directly improving consistency within a stereo framework. Therefore, we only utilize its features as valuable priors.

As shown in Figure 2, our feature extraction module consists of two trainable encoders, each employing a series of residual blocks [31] to extract convolutional image features and context features. To integrate VDA priors, we adopt a shallow convolutional adapter [59] that downsamples depth features from the frozen VDA. The extracted depth features are then concatenated with convolutional features to form  $\{\mathbf{F}_L, \mathbf{F}_R\}$  and  $\mathbf{F}_C$ , providing more stable feature representations for subsequent processing.

### 3.2. All-to-all-pair Correlation

At the  $n$ -th iteration for each time step, the right feature  $\mathbf{F}_R$  is first warped toward the left feature  $\mathbf{F}_L$  using the previous disparity  $d_{n-1}$ , producing  $\hat{\mathbf{F}}_R$ . Previous methods [31, 61] typically construct cost volumes  $\mathbf{C}_n$  via all-pair correlation, employing a lookup mechanism to search for potential matches within a local radius. In [27], this approach was extended to an alternating local 1D and 2D search:

$$\mathbf{C}_n(x, y) = \left\langle \mathbf{F}_L(x, y), \hat{\mathbf{F}}_R(x + r_x, y + r_y) \right\rangle, \quad (1)$$

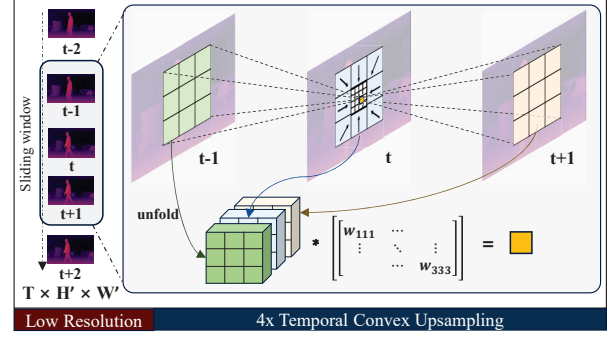


Figure 4. Workflow of the temporal convex upsampling. Each pixel of the high resolution disparity (yellow box) is calculated via its 27 low resolution neighbors using learnable weights.

where  $(x, y)$  denotes the position in  $\mathbf{F}_L$ ,  $\langle \cdot, \cdot \rangle$  represents the channel-wise product operation, and  $(r_x, r_y)$  defines the search range, typically set as  $(r_x, r_y) \in \{(\pm 4, 0), (\pm 1, \pm 1)\}$ . The fundamental principle of these methods is to construct one-directional correspondence, where a reference point is selected and directly matched with target points via multiplication. Inspired by recent advances in point tracking [11], we introduce bidirectional correspondence and develop a local all-to-all-pair correlation. Instead of computing correlation in a fixed direction, we calculate the similarity between every pair of potential matching points within two search windows, formulated as:

$$\mathbf{C}_n(x, y) = \left\langle \mathbf{F}_L(x + r_x, y + r_y), \hat{\mathbf{F}}_R(x + r_x, y + r_y) \right\rangle. \quad (2)$$

This approach enhances match verification and reduces ambiguity [32, 43]. While BiDAStereo [22] constructs temporal cost volumes by aligning adjacent frames for consistency, its performance is significantly constrained by the limitations of optical flow. In contrast, our method emphasizes robust spatial correlation, leveraging dense correspondences to enforce matching smoothness and enhance consistency and robustness across neighboring points [52, 53].

### 3.3. Temporal Cost Aggregation

At the  $n$ -th iteration, the all-to-all-pair correlation expands the potential matching points in  $\mathbf{C}_n$ , increasing its dimensionality from  $\mathbb{R}^{H' \times W' \times (2r_x+1)(2r_y+1)}$  to  $\mathbb{R}^{H' \times W' \times (2r_x+1)^2(2r_y+1)^2}$ , where  $H'$  and  $W'$  denote the height and width of frame feature, respectively. We first encode  $\mathbf{C}_n$  into a compact representation  $\mathbf{E}_n \in \mathbb{R}^{H' \times W' \times L}$  using a MLP:

$$\mathbf{E}_n = \text{MLP}(\mathbf{C}_n). \quad (3)$$

Next, we employ a 3D-GRU [24] to iteratively refine the disparity for each frame while preserving temporal consistency. The GRU aggregates information using separable 3D convolutions across both spatial and temporal dimensions. At iteration  $n$ , it takes as input the encoded cost volume  $\mathbf{E}_n$ ,

Table 2. Zero-shot generalization results. All methods were trained only on SceneFlow [34]. The representative image-based and video-based stereo matching models are compared on Sintel [2], Dynamic Replica [24], Infinigen SV [23], and Virtual KITTI2 [3]. The weights and parameters are fixed for evaluation. The **best** and second-best results are highlighted.

Method	Sintel Final			Dynamic Replica			Infinigen SV			Virtual KITTI2		
	TEPE↓	$\delta_{1px}^t$ ↓	EPE↓	TEPE↓	$\delta_{1px}^t$ ↓	EPE↓	TEPE↓	$\delta_{1px}^t$ ↓	EPE↓	TEPE↓	$\delta_{1px}^t$ ↓	EPE↓
<i>Image-based</i>												
RAFTStereo [31]	2.09	13.70	4.89	0.145	2.02	0.38	2.33	19.80	3.77	1.25	18.39	1.45
IGEVStereo [61]	1.74	12.65	2.44	0.081	0.80	<u>0.20</u>	2.35	19.33	3.81	1.42	18.88	1.39
Selective-IGEV [55]	1.96	12.95	2.82	0.133	1.29	0.29	2.85	21.83	5.69	1.41	19.10	1.38
DEFOM-Stereo [20]	2.33	12.43	2.38	<u>0.080</u>	0.75	<b>0.13</b>	2.01	18.91	3.24	1.07	16.63	1.34
MonSter [8]	1.70	<u>11.06</u>	3.30	0.157	0.86	0.45	<b>1.65</b>	15.73	<b>2.27</b>	<b>0.73</b>	<b>10.93</b>	<b>0.67</b>
<i>Video-based</i>												
DynamicStereo [24]	1.43	11.80	4.26	0.083	<u>0.73</u>	0.24	<u>1.98</u>	<u>15.68</u>	4.67	1.06	15.63	1.47
BiDAStereo [22]	<u>1.26</u>	11.65	<u>1.99</u>	0.091	1.04	0.65	1.99	16.50	6.52	1.02	14.82	1.43
Stereo Any Video (ours)	<b>1.07</b>	<b>9.98</b>	<b>1.75</b>	<b>0.057</b>	<b>0.67</b>	0.25	<b>1.65</b>	<b>14.27</b>	<u>3.16</u>	<u>0.74</u>	<u>12.45</u>	<u>1.01</u>

context feature  $\mathbf{F}_c$ , and the disparities from the previous iteration  $d_{n-1}$ . To enhance representation, we integrate super kernels [22], temporal and spatial attention [24] in this process. The disparity update in 3D-GRU is as follows:

$$x_n = [\text{Conv2d}(\mathbf{E}_n), \text{Conv2d}(d_{n-1}), d_{n-1}, \mathbf{F}_c], \quad (4)$$

$$\tilde{x}_n = \text{Attention}(x_n), \quad (5)$$

$$z_n = \sigma(\text{Conv3d}_z([h_{n-1}, \tilde{x}_n])), \quad (6)$$

$$r_n = \sigma(\text{Conv3d}_r([h_{n-1}, \tilde{x}_n])), \quad (7)$$

$$q_n = \tanh(\text{Conv3d}_q([h_{n-1} \odot r_n, \tilde{x}_n])), \quad (8)$$

$$h_n = (1 - z_n) \odot h_{n-1} + z_n \odot q_n, \quad (9)$$

$$d_n = d_{n-1} + \text{Conv3d}(h_n), \quad (10)$$

where  $[\cdot]$  is concatenation,  $\odot$  denotes element-wise multiplication, and  $\sigma$  is the sigmoid activation function.

While prior methods emphasize temporal consistency in correlation and cost aggregation, the upsampling process remains overlooked, despite its crucial role in achieving accurate and temporally consistent disparities. Towards this, we adapt the commonly used convex upsampling approach [51] to a temporal variant. In Figure 4, each high-resolution disparity pixel (yellow box) is computed as a convex combination of its 27 low resolution neighbors across the 3 frames (green, blue, and faint yellow boxes). The network predicts the corresponding  $3 \times 3 \times 3$  weights using 3D convolutions, enabling a temporally consistent upsampling:

$$\mathbf{w} = \text{softmax}(\text{Conv3d}(h_n)), \quad (11)$$

$$D_n^t = \alpha \cdot \sum_{ijk} \mathbf{w}_{ijk} \odot \text{unfold}([d_n^{t-1}, d_n^t, d_n^{t+1}]), \quad (12)$$

where  $\alpha$  is upsampling rate,  $D_n^t$  is the upsampled disparity.

### 3.4. Loss Function

Recent works have introduced OPW loss [57] and TGM loss [7] in monocular video depth estimation to enhance

temporal consistency. However, we found out that these losses impose a trade-off between accuracy and consistency without yielding substantial improvement in the video stereo task. Thus, we opt for a purely image-based  $L_1$  loss while still achieving strong temporal consistency as:

$$\mathcal{L} = \sum_{t=1}^T \sum_{n=1}^N \gamma^{N-n} \|D_{\text{gt}}^t - D_n^t\|, \quad (13)$$

where  $T$  is the number of frames,  $N$  is the number of iterations, and  $\gamma$  is a weighting coefficient set to 0.9. Here,  $D_{\text{gt}}^t$  denotes the ground truth disparity for the  $t$ -th frame. The final output is obtained from the last iteration prediction.

## 4. Experiments

In this section, we first introduce the datasets and metrics used in the evaluation, as well as implementation details. We then present the zero-shot generalization capabilities of our method, followed by ablation studies to validate the efficacy of each component in our framework. Please refer to the supplementary material for more details.

### 4.1. Benchmark Datasets and Metric

**Dataset.** Our work focuses on videos captured with moving cameras, making commonly used image benchmarks unsuitable, such as Middlebury [44], ETH3D [45], and KITTI 2012/2015 [13] (two frames per scene). To evaluate our method, we use five synthetic and two real-world stereo video datasets, all featuring dynamic scenes. Sintel [2] is a synthetic dataset comprising movie sequences available in both clean and final passes. Spring [35] is a high resolution synthetic dataset. Dynamic Replica [24] is a synthetic indoor dataset featuring non-rigid objects, such as animals and people. Infinigen SV [23] is a synthetic dataset focusing on outdoor natural environments. Virtual KITTI2 [3] is a synthetic dataset simulating outdoor driving scenarios. KITTI Depth [54] is a real-world outdoor dataset captured for autonomous driving applications, with sparse

Table 3. Quantitative comparisons on Spring [35], Sintel [2], and KITTI Depth [54] where methods are allowed to train on any existing datasets. Model weights remain fixed during evaluation. Training datasets sharing the same domain as the test sets are highlighted. “—” represents the test set has been used in training. Abbreviations: SF – SceneFlow [34], DR – Dynamic Replica [24], ISV – Infinigen SV [23], M – Middlebury [44], KD – KITTI Depth [54], VK – Virtual KITTI2 [3]. The **best** and second-best results are highlighted.

Method	Training Data	Scale	Spring			Sintel Final			KITTI Depth		
			TEPE↓	$\delta_{1px}^t$ ↓	EPE↓	TEPE↓	$\delta_{1px}^t$ ↓	EPE↓	TEPE↓	$\delta_{1px}^t$ ↓	EPE↓
<i>Image-based</i>											
RAFTStereo [31]	10 mixed datasets	0.6M	1.49	9.79	5.37	—	—	—	0.41	7.96	<u>0.51</u>
IGEVStereo [61]	SF+M	0.04M	1.53	9.72	5.37	1.53	13.49	2.23	0.46	9.12	0.60
Selective-IGEV [55]	SF+M	0.04M	1.46	9.33	5.46	1.64	13.78	2.14	0.48	9.66	0.60
DEFOM-Stereo [20]	15 mixed datasets	0.7M	1.29	9.12	4.25	—	—	—	<u>0.39</u>	<u>7.08</u>	<b>0.48</b>
MonSter [8]	10 mixed datasets	0.6M	1.25	8.65	5.19	—	—	—	0.41	7.73	0.60
FoundationStereo [59]	7 mixed datasets	1.3M	1.78	9.82	5.40	—	—	—	0.40	7.20	0.57
<i>Video-based</i>											
DynamicStereo [24]	SF+DR+ISV+KD	0.1M	1.07	9.64	9.36	1.38	11.84	3.38	0.42	8.21	0.57
BiDAStereo [22]	SF+DR+ISV+KD	0.1M	<u>0.90</u>	<u>7.90</u>	<u>3.92</u>	<u>1.33</u>	<u>11.34</u>	<u>1.97</u>	0.42	8.33	0.58
Stereo Any Video (ours)	SF+DR+ISV+VK	0.1M	<b>0.77</b>	<b>6.52</b>	<b>3.17</b>	<b>0.99</b>	<b>9.09</b>	<b>1.48</b>	<b>0.35</b>	<b>6.05</b>	0.61

depth maps obtained from a LiDAR sensor. South Kensington SV [23] is a real-world dataset capturing daily scenarios from a stereo camera for qualitative evaluation.

**Metric.** For accuracy evaluation, we use End-Point Error (EPE), which measures the average per-pixel disparity error. To assess temporal consistency, we employ Temporal EPE (TEPE) [24], which quantifies the variation in EPE across the temporal dimension. We also compute  $\delta_{n\text{-px}}^t$ , which represents the proportion of pixels with TEPE exceeding  $n$ . For all metrics, lower values indicate better performance.

## 4.2. Implementation Details

We implement Stereo Any Video in PyTorch and train on NVIDIA A100 GPUs. The model is trained with a batch size of 8 on SceneFlow [34] for 120k iterations and fine-tuned on a mixed dataset for 80k iterations consisting of Dynamic Replica [24] training set, Infinigen SV [23] training set, and Virtual KITTI2 [3]. Before being fed into the network, sequences are randomly cropped to  $256 \times 512$ . We use the AdamW optimizer with a learning rate of  $2e^{-4}$  and apply a one-cycle learning rate schedule. The full training process takes approximately 6 days. Following [24], we apply various data augmentation techniques, including random cropping, rescaling, and saturation shifts. The sequence length is set to  $T = 5$  during training, and  $T = 20$  for evaluation. The number of iterations in GRU is set to  $N = 10$  in training, and  $N = 20$  in evaluation.

## 4.3. Evaluation

**Quantitative Results.** Table 2 presents the results when models are trained solely on SceneFlow [34] and evaluated on four dynamic synthetic datasets including diverse indoor and outdoor environments. The results clearly demonstrate that our model achieves state-of-the-art zero-shot generalization across nearly all metrics, excelling in both accuracy and temporal consistency. Table 3 reports the results

when models are allowed to train on any existing datasets. Notably, despite being trained solely on synthetic datasets, our model surpasses comparison methods that are trained on real-world datasets or share the same domain as the test sets by at least 15% on diverse benchmarks. Specifically, compared to the robust version of RAFT-Stereo, which is trained on 10 mixed datasets, including in-domain sets such as Sintel [2] and KITTI [13], our approach achieves superior performance across most metrics. These results highlight the strong generalization capability of our method when applied to unseen domains.

**Qualitative Results.** Figure 5 and Figure 6 present qualitative results using the checkpoints from Table 3 on dynamic real-world stereo videos from the South Kensington SV dataset [23]. For the indoor scenario shown in Figure 5, the upper rows display three input left frames along with their predicted disparities, while the subsequent rows visualize the globally aligned reconstructed point cloud. The red arrows highlight that our method produces temporally consistent and accurate disparities, whereas other methods exhibit noticeable artifacts. Specifically, Selective-IGEV and DynamicStereo suffer from flickering artifacts, with cabinets and walls displaying low-frequency oscillations. BiDAStereo fails to maintain an accurate shape, resulting in a distorted table. For the outdoor scenario in Figure 6, the driving car presents a challenging case due to its transparent surfaces and rapid motion. As shown, all other methods struggle with this case, failing to produce reliable disparity estimations. In contrast, our method demonstrates superior performance, maintaining accurate and stable predictions.

## 4.4. Ablation Study

**Foundation Priors.** As shown in Table 4, incorporating foundation priors from a monocular depth model improves performance. Notably, VDA [7] outperforms DepthAnythingV2 [68] due to its stronger temporal consistency. In-

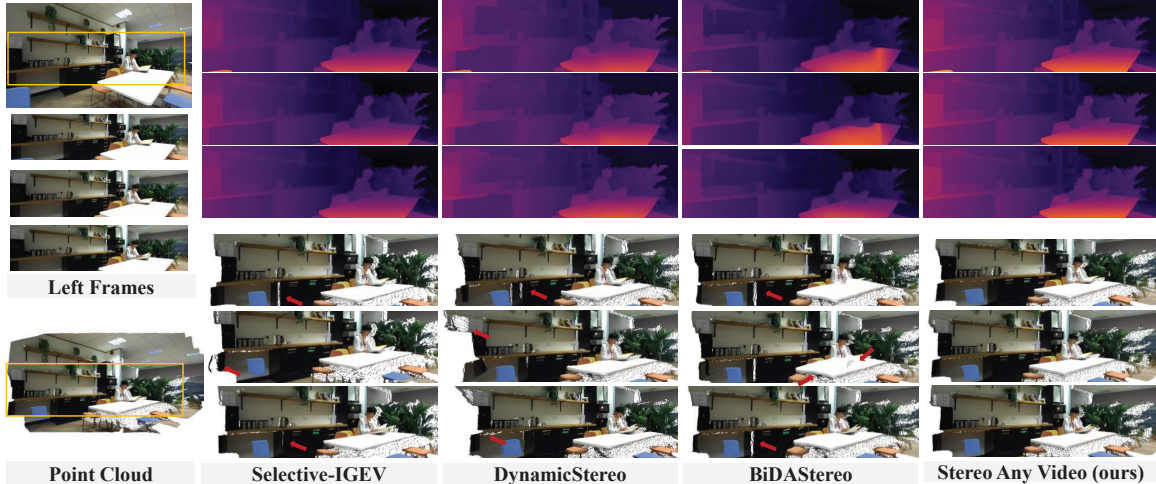


Figure 5. Qualitative comparison on a dynamic indoor scenario from the South Kensington SV dataset [23]. **Top:** Input left frames and predicted disparities for the same region across three frames using different methods. **Bottom:** Disparities converted into globally aligned point clouds and rendered from a camera view for comparison. Failure regions are highlighted with red arrows. Our method produces consistent and accurate disparities without flickering. See video in the supplementary material for better visualization.




Figure 6. Qualitative comparison on a dynamic outdoor scenario from the South Kensington SV dataset [23].

terestingly, increasing the backbone size does not yield additional performance gains. To investigate this, we conducted further experiments (Table 5) by varying depth and CNN feature channels based on VDA-L. However, the results remained similar. According to [7], VDA-L has  $13\times$  more parameters than VDA-S (381.8M vs. 28.4M), yet the performance gain is marginal ( $\delta_1(\uparrow)$ : 0.944 vs. 0.942 on KITTI; 0.971 vs. 0.959 on NYUv2). This suggests that the larger backbone does not provide richer priors. Moreover, large backbone requires additional CNN channels and parameters for adaptation, increasing training complexity. Therefore, we adopt VDA-S in the final model.

**Correlation.** To evaluate the effectiveness of our proposed all-to-all-pairs correlation, we replace it with local correlation [27] and experiment with using only 1D or 2D correlation instead of alternating between them. As shown in Table 4, all these modifications degrade performance. Notably, constructing 3D correlations—where the search window spans height, width, and time—further worsens results. This suggests that frames cannot be effectively utilized without proper alignment [23].

**Cost Volume Encoder.** Instead of directly encoding cost volumes with convolutions, we find that introducing an MLP layer to extract a compact representation before ap-

Table 4. Ablation study on the foundation priors, correlation, cost volume encoder, convex upsampling, attention, and inference settings of our model on Sintel [2] and Dynamic Replica [24]. The approach used in our final model is underlined.

Setting Variations	Sintel Final		Dynamic Replica	
	TEPE↓	$\delta_{1px}^t$ ↓	TEPE↓	$\delta_{1px}^t$ ↓
<i>Foundation Priors</i>				
w/o Depth	1.47	12.35	0.092	1.05
DepthAnythingV2-S [68]	1.33	11.27	0.088	0.99
VDA-S [7]	1.27	<b>10.98</b>	<b>0.083</b>	<b>0.78</b>
VDA-L [7]	<b>1.20</b>	12.47	<b>0.083</b>	0.82
<i>Correlation</i>				
1D + 2D local [27]	1.27	<b>10.98</b>	0.083	0.78
1D all-to-all-pairs	1.23	11.36	0.079	0.78
2D all-to-all-pairs	1.29	11.95	0.085	0.77
3D all-to-all-pairs	4.51	22.90	1.313	2.01
1D + 2D all-to-all-pairs	<b>1.21</b>	11.33	<b>0.076</b>	<b>0.75</b>
<i>Cost Volume Encoder</i>				
Convolution	1.26	11.54	0.081	0.77
<u>MLP + Convolution</u>	<b>1.21</b>	<b>11.33</b>	<b>0.076</b>	<b>0.75</b>
<i>Convex Upsampling</i>				
Bilinear	1.26	12.74	0.085	1.05
Convex [51]	1.21	11.33	0.076	<b>0.75</b>
Temporal Convex	<b>1.04</b>	<b>10.21</b>	<b>0.067</b>	0.79
<i>Attention in GRU</i>				
w/o Attention	<b>1.04</b>	10.21	0.067	0.79
Temporal	1.07	10.23	0.064	0.79
Temporal + Spatial	1.07	<b>9.98</b>	<b>0.057</b>	<b>0.67</b>
<i>Inference Stage</i>				
$\frac{1}{16} \rightarrow \frac{1}{8} \rightarrow \frac{1}{4}$	<b>1.07</b>	<b>9.98</b>	<b>0.057</b>	<b>0.67</b>
$\frac{1}{32} \rightarrow \frac{1}{16} \rightarrow \frac{1}{8} \rightarrow \frac{1}{4}$	1.08	10.82	0.058	0.70

plying convolutions enhances cost aggregation, leading to improved final results.

**Convex Upsampling.** We compare bilinear upsampling and general convex upsampling [51] to our proposed temporal convex upsampling module. As observed, our approach achieves better temporal consistency by incorporating learnable weights with adjacent frames.

**Attention in GRU.** We further investigate different attention mechanisms within the GRU update block. As shown in Table 4, combining temporal and spatial attention together leads to the best performance.

**Inference Stage.** We also explore inference starting at different resolutions. As can be seen in Table 4, a three-stage inference setup already achieves strong results.

## 4.5. Application

Our method enhances downstream applications that rely on accurate metric depth, such as 3D point tracking. As shown in Figure 7, we visualize results using the state-of-the-art 3D point tracking method DELTA [36]. Our predicted metric depth yields more points with precise trajectories than UniDepth [40], which gives less missing points as highlighted by the red arrows. Additional application, such as visual effects, are available in the supplementary material.

Table 5. Ablation study on the effect of different channel configurations for depth features and CNN features, based on VDA-L [7], on Sintel [2]. Trainable parameters are also reported.

Depth & CNN Channel Number	Sintel Final		Para. (M)
	TEPE↓	$\delta_{1px}^t$ ↓	
<i>VDA-L [7]</i>			
(32, 96)	1.20	12.47	8.4
(128, 64)	1.58	14.52	13.3
(128, 128)	1.31	13.21	19.3

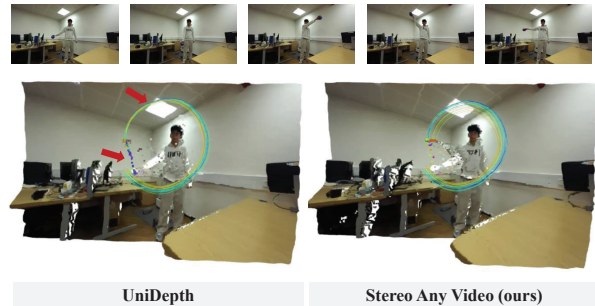


Figure 7. Visualization of 3D point tracking based on metric depth generated from UniDepth and our method.

Table 6. Comparison of different methods on model trainable [non-trainable] parameters, inference GPU memory, and multiply-accumulate (MAC) on 20 frames with resolution  $720 \times 1280$ .

Methods	Para. (M)	Mem. (G)	MACs (T)
DynamicStereo [24]	20.5	35.1	182.3
BiDAStereo [22]	12.2	41.1	186.6
Stereo Any Video (ours)	9.4 [28.4]	41.1	303.4

## 4.6. Limitation and Future Work

We report the inference efficiency results on a single NVIDIA A6000 GPU in Table 6. As shown, due to the use of foundation priors, our method incurs relatively high computational and memory costs. Looking ahead, we aim to develop a model zoo for this method, including both a large and a light version.

## 5. Conclusion

In this paper, we introduce Stereo Any Video, a robust framework for video stereo matching. By leveraging rich monocular priors, an all-to-all-pairs correlation mechanism, and a temporal convex upsampling layer, our approach produces temporally consistent and accurate disparities without relying on additional inputs such as camera pose or optical flow. Extensive evaluations show that our model achieves state-of-the-art performance on multiple standard benchmarks under zero-shot settings. Furthermore, our method generalizes well to real-world stereo videos, handling diverse content, motion, and camera movements effectively.

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

## References

- [1] 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
- [2] 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, 5, 6, 8
- [3] Johann Cabon, Naila Murray, and Martin Humenberger. Virtual kitti 2, 2020. 5, 6
- [4] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9650–9660, 2021. 3
- [5] Jia-Ren Chang and Yong-Sheng Chen. Pyramid stereo matching network. In *CVPR*, pages 5410–5418, 2018. 2
- [6] 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
- [7] Sili Chen, Hengkai Guo, Shengnan Zhu, Feihu Zhang, Zilong Huang, Jiashi Feng, and Bingyi Kang. Video depth anything: Consistent depth estimation for super-long videos, 2025. 2, 3, 5, 6, 7, 8
- [8] 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. 3, 5, 6
- [9] 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
- [10] Ziang Cheng, Jiayu Yang, and Hongdong Li. Stereo matching in time: 100+ fps video stereo matching for extended reality. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 8719–8728, 2024. 2
- [11] Seokju Cho, Jiahui Huang, Jisu Nam, Honggyu An, Seungryoung Kim, and Joon-Young Lee. Local all-pair correspondence for point tracking. *arXiv preprint arXiv:2407.15420*, 2024. 4
- [12] 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
- [13] 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, 5, 6
- [14] 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
- [15] 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
- [16] Xiaoyang Guo, Kai Yang, Wukui Yang, Xiaogang Wang, and Hongsheng Li. Group-wise correlation stereo network. In *CVPR*, pages 3273–3282, 2019. 2
- [17] Xianda Guo, Chenming Zhang, Youmin Zhang, Wenzhao Zheng, Dujun Nie, Matteo Poggi, and Long Chen. Light-stereo: Channel boost is all your need for efficient 2d cost aggregation, 2024. 2
- [18] Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003. 2
- [19] Wenbo Hu, Xiangjun Gao, Xiaoyu Li, Sijie Zhao, Xiaodong Cun, Yong Zhang, Long Quan, and Ying Shan. Depthcrafter: Generating consistent long depth sequences for open-world videos. *arXiv preprint arXiv:2409.02095*, 2024. 3, 4
- [20] Hualie Jiang, Zhiqiang Lou, Laiyan Ding, Rui Xu, Minglang Tan, Wenjie Jiang, and Rui Huang. Defom-stereo: Depth foundation model based stereo matching, 2025. 3, 5, 6
- [21] Junpeng Jing, Jiankun Li, Pengfei Xiong, Jiangyu Liu, Shuaicheng Liu, Yichen Guo, Xin Deng, Mai Xu, Lai Jiang, and Leonid Sigal. Uncertainty guided adaptive warping for robust and efficient stereo matching. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3318–3327, 2023. 2
- [22] Junpeng Jing, Ye Mao, and Krystian Mikolajczyk. Match-stereo-videos: Bidirectional alignment for consistent dynamic stereo matching. In *European Conference on Computer Vision*, pages 415–432. Springer, 2024. 2, 4, 5, 6, 8
- [23] Junpeng Jing, Ye Mao, Anlan Qiu, and Krystian Mikolajczyk. Match stereo videos via bidirectional alignment, 2024. 3, 5, 6, 7
- [24] Nikita Karaev, Ignacio Rocco, Benjamin Graham, Natalia Neverova, Andrea Vedaldi, and Christian Rupprecht. 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, 5, 6, 8
- [25] 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 *CVPR*, pages 66–75, 2017. 2
- [26] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4015–4026, 2023. 3
- [27] 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, 7, 8

- [28] 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
- [29] Zhaoshuo Li, Wei Ye, Dilin Wang, Francis X Creighton, Russell H Taylor, Ganesh Venkatesh, and Mathias Unberath. Temporally consistent online depth estimation in dynamic scenes. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 3018–3027, 2023. 2
- [30] 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
- [31] Lahav Lipson, Zachary Teed, and Jia Deng. Raft-stereo: Multilevel recurrent field transforms for stereo matching. *arXiv preprint arXiv:2109.07547*, 2021. 1, 2, 3, 4, 5, 6
- [32] David G Lowe. Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60:91–110, 2004. 4
- [33] Yang Luo, Xuanlei Zhao, Mengzhao Chen, Kaipeng Zhang, Wenqi Shao, Kai Wang, Zhangyang Wang, and Yang You. Enhance-a-video: Better generated video for free, 2024. 2, 3
- [34] 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. 2, 5, 6
- [35] Lukas Mehl, Jenny Schmalfluss, Azin Jahedi, Yaroslava Naliwayko, 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. 5, 6
- [36] Tuan Duc Ngo, Peiye Zhuang, Chuang Gan, Evangelos Kalogerakis, Sergey Tulyakov, Hsin-Ying Lee, and Chaoyang Wang. Delta: Dense efficient long-range 3d tracking for any video, 2025. 8
- [37] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafranec, 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
- [38] 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
- [39] Jiahao Pang, Wenxiu Sun, Chengxi Yang, Jimmy Ren, Ruichao Xiao, Jin Zeng, and Liang Lin. Zoom and learn: Generalizing deep stereo matching to novel domains. In *CVPR*, pages 2070–2079, 2018. 2
- [40] Luigi Piccinelli, Yung-Hsu Yang, Christos Sakaridis, Mattia Segu, Siyuan Li, Luc Van Gool, and Fisher Yu. Unidepth: Universal monocular metric depth estimation, 2024. 8
- [41] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Aspell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 3
- [42] 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
- [43] Ignacio Rocco, Mircea Cimpoi, Relja Arandjelović, Akihiko Torii, Tomas Pajdla, and Josef Sivic. Neighbourhood consensus networks. *Advances in neural information processing systems*, 31, 2018. 4
- [44] 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
- [45] 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
- [46] Zhelun Shen, Yuchao Dai, and Zhibo Rao. Cfnet: Cascade and fused cost volume for robust stereo matching. In *CVPR*, pages 13906–13915, 2021. 2
- [47] Zhelun Shen, Xibin Song, Yuchao Dai, Dingfu Zhou, Zhibo Rao, and Liangjun Zhang. Digging into uncertainty-based pseudo-label for robust stereo matching. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–18, 2023. 2
- [48] Chenyang Si, Weichen Fan, Zhengyao Lv, Ziqi Huang, Yu Qiao, and Ziwei Liu. Repvideo: Rethinking cross-layer representation for video generation, 2025. 2, 3
- [49] 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
- [50] 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
- [51] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *ECCV*, pages 402–419, 2020. 2, 5, 8
- [52] Prune Truong, Martin Danelljan, Luc V Gool, and Radu Timofte. Gocor: Bringing globally optimized correspondence volumes into your neural network. *Advances in Neural Information Processing Systems*, 33:14278–14290, 2020. 4
- [53] Prune Truong, Martin Danelljan, and Radu Timofte. Glunet: Global-local universal network for dense flow and correspondences. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6258–6268, 2020. 4
- [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. 5, 6
- [55] 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, 5, 6
- [56] Yan Wang, Zihang Lai, Gao Huang, Brian H Wang, Laurens Van Der Maaten, Mark Campbell, and Kilian Q Weinberger. Anytime stereo image depth estimation on mobile devices. In *2019 international conference on robotics and automation (ICRA)*, pages 5893–5900. IEEE, 2019. 2
- [57] Yiran Wang, Min Shi, Jiaqi Li, Zihao Huang, Zhiguo Cao, Jianming Zhang, Ke Xian, and Guosheng Lin. Neural video depth stabilizer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9466–9476, 2023. 5
- [58] Philippe Weinzaepfel, Thomas Lucas, Vincent Leroy, Yann 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
- [59] Bowen Wen, Matthew Trepte, Joseph Aribido, Jan Kautz, Orazio Gallo, and Stan Birchfield. Foundationstereo: Zero-shot stereo matching, 2025. 3, 4, 6
- [60] 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
- [61] 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, 4, 5, 6
- [62] 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. 2
- [63] Gangwei Xu, Huan Zhou, and Xin Yang. Cgi-stereo: Accurate and real-time stereo matching via context and geometry interaction. *arXiv preprint arXiv:2301.02789*, 2023. 2
- [64] Gangwei Xu, Xianqi Wang, Zhaoxing Zhang, Junda Cheng, Chunyuan Liao, and Xin Yang. Igev++: Iterative multi-range geometry encoding volumes for stereo matching. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025. 2
- [65] Haofei Xu and Juyong Zhang. Aanet: Adaptive aggregation network for efficient stereo matching. In *CVPR*, pages 1959–1968, 2020. 2
- [66] 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
- [67] Gengshan Yang, Joshua Manela, Michael Happold, and Deva Ramanan. Hierarchical deep stereo matching on high-resolution images. In *CVPR*, pages 5515–5524, 2019. 2
- [68] 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. 3, 6, 8
- [69] Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything v2. *arXiv preprint arXiv:2406.09414*, 2024. 3
- [70] Jiayi Zeng, Chengtang Yao, Yuwei Wu, and Yunde Jia. Temporally consistent stereo matching. In *European Conference on Computer Vision*, pages 341–359. Springer, 2024. 2
- [71] 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
- [72] 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
- [73] Youmin Zhang, Matteo Poggi, and Stefano Mattoccia. Temporalstereo: Efficient spatial-temporal stereo matching network. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 9528–9535. IEEE, 2023. 2
- [74] 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. 3
- [75] Yiran Zhong, Hongdong Li, and Yuchao Dai. Open-world stereo video matching with deep rnn. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 101–116, 2018. 2
- [76] 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. 3