

Generalized Geometry Encoding Volume for Real-time Stereo Matching

Jiaxin Liu, Gangwei Xu, Xianqi Wang, Chengliang Zhang, Xin Yang*

Huazhong University of Science and Technology
{ljx123, gwxu, xianqi, z_chengliang, xinyang2014}@hust.edu.cn

Abstract

Real-time stereo matching methods primarily focus on enhancing in-domain performance but often overlook the critical importance of generalization in real-world applications. In contrast, recent stereo foundation models leverage monocular foundation models (MFMs) to improve generalization, but typically suffer from substantial inference latency. To address this trade-off, we propose Generalized Geometry Encoding Volume (GGEV), a novel real-time stereo matching network that achieves strong generalization. We first extract depth-aware features that encode domain-invariant structural priors as guidance for cost aggregation. Subsequently, we introduce a Depth-aware Dynamic Cost Aggregation (DDCA) module that adaptively incorporates these priors into each disparity hypothesis, effectively enhancing fragile matching relationships in unseen scenes. Both steps are lightweight and complementary, leading to the construction of a generalized geometry encoding volume with strong generalization capability. Experimental results demonstrate that our GGEV surpasses all existing real-time methods in zero-shot generalization capability, and achieves state-of-the-art performance on the KITTI 2012, KITTI 2015, and ETH3D benchmarks.

Introduction

Stereo matching aims to estimate dense, pixel-wise disparity maps from a pair of rectified stereo images. As a long-standing and challenging task in computer vision, it plays a fundamental role in a wide range of applications, including 3D reconstruction (Liang et al. 2025b), autonomous driving (Li et al. 2025; Liang et al. 2025c), and robotic navigation. These real-world scenarios impose strict demands on both generalization and inference latency.

Existing real-time stereo matching methods have adopted various strategies to achieve fast inference. These include using downsampled (Khamis et al. 2018; Xu et al. 2021) or sparse cost volume representations (Duggal et al. 2019; Xu et al. 2024a), lightweight aggregation networks (Bangunharcana et al. 2021), and replacing computationally expensive 3D convolutions with 2D convolutions (Xu and Zhang 2020; Li et al. 2024). However, most existing methods rely heavily on clear and unambiguous matching cues,




Figure 1: Zero-shot generalization comparison. All models are trained on Scene Flow and tested on KITTI, Middlebury, and ETH3D. GGEV achieves comparable speed to RT-IGEV while offering improved generalization on unseen scenes.

and struggle to effectively aggregate information in unseen domains—particularly in challenging regions such as occlusions, textureless areas, repetitive patterns, and thin structures.

Recent methods have introduced monocular foundation models (MFMs) into stereo matching, achieving remarkable zero-shot generalization performance. FoundationStereo (Wen et al. 2025) designs a higher-capacity aggregation network to better exploit monocular priors. MonSter (Cheng et al. 2025a) employs a dual-branch architecture that iteratively refines both monocular and stereo disparity estimates. These approaches typically rely on costly backbones to extract rich and detailed features for cost volume construction, and employ complex iterative mechanisms to address the scale-shift issue between monocular and stereo. Although these methods can improve generalization, they often overlook the critical importance of inference latency in real-world applications. To this end, a motivating question arises: *how to design a real-time stereo matching network that achieves strong generalization while maintaining high accuracy?*

To answer this question, we analyze the limitations of current geometry encoding volumes and identify two key limitations: 1) the critical regions vary significantly across

*Corresponding author.

disparity hypotheses; 2) the matching relationships within these regions are highly fragile due to unseen textures, occlusions, repetitive patterns, and thin structures (see Fig. 3). In this paper, we propose Generalized Geometry Encoding Volume (GGEV), a real-time stereo matching network that efficiently incorporates MFMs into the cost aggregation process to enhance cost volume representations. Specifically, the proposed GGEV first constructs depth-aware features by integrating texture and depth features extracted from Depth Anything V2 (Yang et al. 2024b) using a lightweight fusion network, thereby obtaining reliable structural priors that help stabilize fragile matching relationships. In contrast to conventional hourglass-based aggregation networks (Xu et al. 2022; Bangunharcana et al. 2021) that process all disparity hypotheses uniformly, our method adaptively integrates depth structural priors into the corresponding disparity hypotheses, thereby enhancing the structural representation and generalization of the cost volume. In particular, we first compute an affinity matrix between each disparity hypothesis and the depth feature map, where the disparity hypothesis provides positional cues and the depth features offer rich structural context. These affinity matrices are then used to generate dynamic convolutional kernels that adaptively filter the concatenated disparity hypotheses and depth features. Furthermore, we incorporate a combination of large and small convolution kernels to capture complementary low- and high-frequency information.

Our proposed GGEV outperforms all existing real-time stereo matching methods in both in-domain accuracy and zero-shot generalization capability. It achieves state-of-the-art results on the KITTI 2012, KITTI 2015, and ETH3D benchmarks. Remarkably, even when trained solely on the synthetic Scene Flow dataset, GGEV demonstrates strong cross-domain generalization to real-world scenarios, as illustrated in Fig. 1.

In summary, our main contributions are:

- We propose a novel generalized geometry encoding volume that efficiently integrates depth priors in a lightweight manner to enhance generalization.
- We propose a Depth-aware Dynamic Cost Aggregation (DDCA) module that adaptively generates dynamic convolution kernels based on the affinity between disparity hypotheses and depth features.
- Our method demonstrates strong generalization to real-world scenarios, even when trained solely on the synthetic datasets.
- Our method outperforms existing real-time approaches on public benchmarks such as KITTI 2012, KITTI 2015 and ETH3D.

Related Work

Real-time Stereo Matching

Recent works have focused on designing lightweight stereo matching networks while maintaining competitive accuracy. Some methods (Gu et al. 2020; Xu et al. 2024b; Chang, Chang, and Chen 2020; Wang et al. 2020) attempt to construct and aggregate cost volumes at lower resolutions to reduce computational overhead. However, this often leads to

a significant drop in accuracy. To preserve accuracy, some methods (Bangunharcana et al. 2021; Liang et al. 2025a) still construct high-resolution cost volumes, while using context-aware activations to enable lightweight yet effective aggregation. Alternatively, AANet (Xu and Zhang 2020) replaces computationally expensive 3D convolutions with deformable 2D convolutions. However, due to the lack of accurate structural guidance and the limited receptive field, these methods often suffer from insufficient accuracy. Besides aggregation-based approaches, RT-IGEV (Xu et al. 2025c) achieves a balance between accuracy and efficiency by leveraging the iterative framework with a reduced number of iterations. However, the issue of fragile matching relationships in the cost volume under unseen and challenging scenarios remains unresolved.

Zero-shot Generalized Stereo Matching

Zero-shot generalization in stereo matching has gained increasing attention. Some methods (Zhang et al. 2020, 2022) aim to learn domain-invariant representations by employing domain normalization or specialized losses. Others methods (Liu, Yu, and Qi 2022; Chuah et al. 2022) enhance generalization by avoiding shortcut learning. In contrast to these works, RAFT-Stereo (Lipson, Teed, and Deng 2021) enhances generalization by innovating the network framework. Recent methods (Zhou et al. 2025; Bartolomei et al. 2025) have further improved upon this foundation by incorporating monocular foundation models into stereo matching frameworks. FoundationStereo (Wen et al. 2025) constructs a large-scale training dataset to support the training of larger models. MonSter (Cheng et al. 2025a,b) employs a dual-branch architecture that iteratively refines monocular and stereo disparity estimates. DEFOM-Stereo (Jiang et al. 2025) initializes the disparity map with depth predicted by MFMs and introduces a scale update module. However, these methods typically rely on ViT-L backbones to extract fine-grained features for cost volume construction and require complex operations to address the scale-shift inherent in monocular affine depth. In contrast, our work emphasizes leveraging depth features to efficiently guide cost aggregation, avoiding the scale-shift issues.

Monocular Depth Foundation Model

MFMs (Xu et al. 2025a; Lin et al. 2025; Yang et al. 2024b) have demonstrated strong zero-shot generalization across diverse tasks without requiring target-domain fine-tuning. MiDaS (Ranftl et al. 2020) pioneers zero-shot relative depth estimation, while the DINO series (Caron et al. 2021; Oquab et al. 2023) leverages vision transformers for dense self-supervised representation learning. The Depth Anything series (Yang et al. 2024a,b) further pushes the boundary of generalization in depth estimation. In particular, Depth Anything V2 distills student models from a large-scale teacher model using unlabeled real-world data, yielding accurate and efficient depth predictors. This student model offers a lightweight and generalizable backbone, making it a suitable component for our framework.

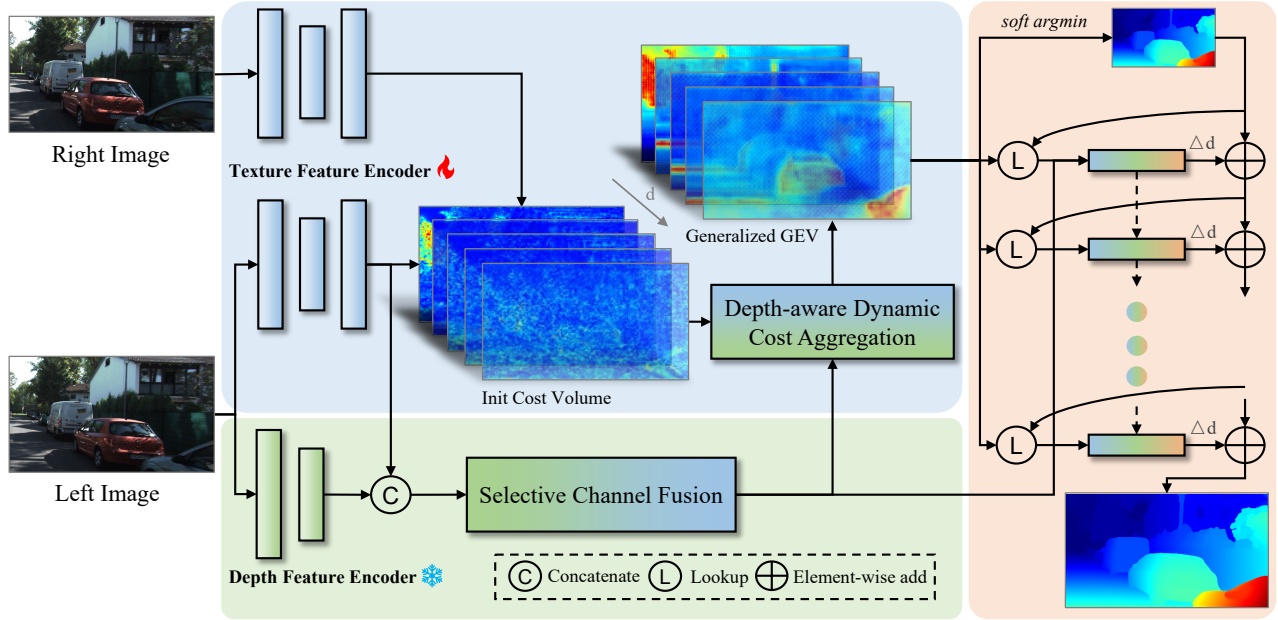


Figure 2: Overview of our proposed GGEV. The Selective Channel Fusion (SCF) module integrates texture features with depth features as a guidance for cost aggregation. Then, the Depth-aware Dynamic Cost Aggregation (DDCA) module adaptively incorporates depth structural priors to enhance the fragile matching relationships in the initial cost volume, resulting in a generalized geometry encoding volume.

Method

Overall Framework

As illustrated in the Fig. 2, our network architecture consists of four main stages: multi-cue feature extraction, cost volume construction, depth-aware dynamic cost aggregation, and depth-aware iterative refinement.

Multi-cue Feature Extraction

The feature extraction stage consists of two complementary cues: (a) texture features are extracted from both the left and right images to construct the cost volume; (b) depth features are extracted from the left image to guide cost aggregation and iterative refinement.

Texture Feature Encoder. Given rectified left and right images $I_l, I_r \in \mathbb{R}^{3 \times H \times W}$, we employ the MobileNetV2 pretrained on ImageNet (Deng et al. 2009) to extract multi-scale texture features $\mathbf{f}_{l,i}, \mathbf{f}_{r,i} \in \mathbb{R}^{C_i \times \frac{H}{i} \times \frac{W}{i}}, i \in \{4, 8, 16\}$.

Depth Feature Encoder. To incorporate the generalization ability of MFMs into our framework, we employ a frozen Depth Anything V2 Small to extract multi-scale depth features $\mathbf{f}_{d,i} \in \mathbb{R}^{C_i \times \frac{H}{i} \times \frac{W}{i}}, i \in \{2, 4, 8, 16\}$, using only the left image. These features serve as depth structural priors to guide cost aggregation. To enrich the feature representation, we additionally design a fusion module to integrate texture and depth features.

Selective Channel Fusion. We adopt a lightweight 1×1 convolution to enable selective feature integration while preserving structural details and avoiding spatial blurring. The




Figure 3: Effectiveness of our DDCA in generalization evaluation. The first row show the initial cost volume features across different disparity hypotheses, which are fragile in unseen scenes and contain many mismatches. In contrast, the second row shows the results after applying our DDCA, which effectively filters out incorrect matches and preserves accurate matching features at their corresponding disparity planes, leading to clearer and more reliable structures.

Selective Channel Fusion (SCF) module takes the concatenated \mathbf{f}_l and \mathbf{f}_d as input and generates depth-aware prior features $\mathbf{f}_{da,i} \in \mathbb{R}^{C_i \times \frac{H}{i} \times \frac{W}{i}}, i \in \{4, 8, 16\}$.

Cost Volume Construction

Given the texture features at $1/4$ resolution $\mathbf{f}_{l,4}$ and $\mathbf{f}_{r,4}$, we construct the group-wise correlation volume \mathbf{C} .

$$\mathbf{C}(g, d, x, y) = \frac{1}{N_c/N_g} \langle \mathbf{f}_{l,4}^g(x, y), \mathbf{f}_{r,4}^g(x-d, y) \rangle, \quad (1)$$

where $\langle \cdot, \cdot \rangle$ denotes the inner product, $d \in \mathcal{D} = \{0, 1, 2, \dots, D/4 - 1\}$ is the disparity index, N_c denotes the




Figure 4: The architecture of proposed DDCA.

number of feature channels, and $N_g = 8$ is the total number of groups.

Depth-aware Dynamic Cost Aggregation

Along the disparity dimension, the critical regions corresponding to different disparity hypotheses often vary, and tend to be highly fragile in areas with unseen textures, occlusions, repetitive patterns, and thin structures (see Fig. 3). Treating these regions uniformly often leads to mismatches, edge blurring, loss of fine details (see Fig. 6). To address this, inspired by OverLoCK (Lou and Yu 2025), we propose a novel Depth-aware Dynamic Cost Aggregation (DDCA) module that adaptively responds to cost volumes at different disparity hypotheses. Moreover, leveraging depth-aware features as structural guidance can effectively enhance fragile matching relationships. To dynamically model input-dependent behaviors, we propose to capture the response relationship between disparity hypotheses and depth features by computing their affinity. These affinity values are then used to define dynamic convolution kernels, effectively injecting depth-aware structural cues into every kernel weight across different disparity hypotheses. We employ dynamic convolution kernels to aggregate the cost maps using a sliding-window mechanism, similar to standard 2D convolutions, which makes our module lightweight and real-time.

Disparity-wise Depth Structural Representation. As shown in Fig. 4, given an input disparity hypothesis map $\mathbf{C}_d \in \mathbb{R}^{G \times H \times W}$ (d is the disparity index) and a depth-aware feature map $\mathbf{f}_{da} \in \mathbb{R}^{C \times H \times W}$, We first transform the input feature into two components, namely $\mathbf{Q} \in \mathbb{R}^{C \times HW}$ and $\mathbf{K} \in \mathbb{R}^{C \times S^2}$, \mathbf{K} represents the aggregation of \mathbf{f}_{da} into $S \times S$ regional centers via adaptive average pooling. Simple matrix multiplications between each pair of \mathbf{Q} and \mathbf{K} produce the affinity matrices $\mathbf{A} \in \mathbb{R}^{HW \times S^2}$. The process can be written as:

$$\mathbf{Q} = \text{Re}(W_q \mathbf{C}_d), \quad (2)$$

$$\mathbf{K} = \text{Re}(W_k \text{Pool}(\mathbf{f}_{da})), \quad (3)$$

$$\mathbf{A} = \mathbf{Q}^T \mathbf{K}, \quad (4)$$

where W_q and W_k denote 1×1 convolutional layers, and $\text{Re}(\cdot)$ refers to the reshape operation. Analogous to multi-head attention (Dosovitskiy et al. 2020), we divide \mathbf{Q} and \mathbf{K} along the channel dimension into G groups to obtain $\mathbf{Q}^g \in \mathbb{R}^{\frac{C}{G} \times HW}$ and $\mathbf{K}^g \in \mathbb{R}^{\frac{C}{G} \times S^2}$, yielding \mathbf{A}^g affinity matrices to capture the relationships within each channel group.

Disparity-wise Adaptive Cost Aggregation. We leverage G groups affinity matrices \mathbf{A}^g to generate G distinct $K \times K$ dynamic convolution kernels, enabling the network to adaptively aggregate features with diverse structural patterns. First, we use another learnable linear layer W_m to map \mathbf{A}^g into corresponding convolution kernel weights $\mathbf{M}^g \in \mathbb{R}^{HW \times K^2}$, which are then normalized via a softmax operation. Then, each row of \mathbf{M}^g is reshaped into a $K \times K$ kernel, resulting in spatially adaptive convolution filters tailored to each pixel based on the input. During the convolution operation, the channels of the fused features \mathbf{C}_d and \mathbf{f}_{da} are also divided into G groups, where channels within the same group share the same dynamic kernel to reduce computational overhead. The process can be written as:

$$\mathbf{M}^g = \text{softmax}(\mathbf{A}^g W_m), \quad (5)$$

$$\mathbf{C}'_d = \mathbf{C}_d * \mathbf{M}^g_{\text{dynamic}}(\mathbf{C}_d, \mathbf{f}_{da}), \quad (6)$$

where $*$ denotes the group-wise convolution operation. Meanwhile, we adopt a combination of large and small convolutional kernels to facilitate the fusion of low- and high-frequency information. We rearrange the independently aggregated disparity hypotheses to reconstruct the generalized geometry encoding volume \mathbf{C}' .

Initial Disparity Prediction. We apply the *soft-argmin* to \mathbf{C}' to regress the initial disparity estimation.

$$\mathbf{d}_0 = \sum_{d \in \mathcal{D}} d \times \text{Softmax}(\mathbf{C}'(d)), \quad (7)$$

where \mathbf{d}_0 is at 1/4 scale of the original image resolution.

Depth-aware Iterative Refinement

Given \mathbf{d}_0 , we employ an iterative GRU to progressively refine the disparity map. The hidden state h_0 is initialized using the depth feature $\mathbf{f}_{da,4}$, injecting structural priors into the recurrent refinement process. The single-layer GRU updates its hidden state based on the latest disparity \mathbf{d}_k and the geometry features \mathbf{f}_G indexed from \mathbf{C}' . The update operator can be formulated as:

$$z_k = \sigma(\text{Conv}([h_{k-1}, x_k], W_z)), \quad (8)$$

$$r_k = \sigma(\text{Conv}([h_{k-1}, x_k], W_r)), \quad (9)$$

$$\tilde{h}_k = \tanh(\text{Conv}([r_k \odot h_{k-1}, x_k], W_h)), \quad (10)$$

$$h_k = (1 - z_k) \odot h_{k-1} + z_k \odot \tilde{h}_k, \quad (11)$$

where x_k denotes the concatenation of \mathbf{d}_k and \mathbf{f}_G ; \odot denotes element-wise product; σ denotes sigmoid; W_z , W_r and W_h are the parameters of the network. Based on the h_k , we decode a residual disparity $\Delta \mathbf{d}_k$ using two convolutional layers, and update the current disparity accordingly.

$$\mathbf{d}_{k+1} = \mathbf{d}_k + \Delta \mathbf{d}_k. \quad (12)$$

Spatial Upsampling. We leverage depth features to assist in recovering full-resolution disparity map. We first convolve the h_k to generate features and upsample them to half resolution. The upsampled features are then concatenated with the depth features \mathbf{f}_d from left image to produce a weight map $W \in \mathbb{R}^{H \times W \times 9}$. Finally, we obtain the full-resolution disparity map by performing a weighted combination over the local neighboring points of the low-resolution disparity map \mathbf{d}_k .

Training Set		Scene Flow			
Target	Method	KITTI 2012 2015		Middlebury quarter	ETH3D
Accuracy	RAFT-Stereo (2021)	4.5	5.7	9.3	3.2
	FC-GANet (2022)	4.6	5.3	7.8	5.8
	DEFOM-Stereo (2025)	3.7	4.9	5.6	2.3
	DEFOM-Stereo (ViT-S)	4.2	5.3	6.3	2.6
	FoundationStereo (2025)	3.2	4.9	-	1.8
Speed	DeepPrunerFast (2019)	16.8	15.9	18.3	11.0
	CoEx (2021)	13.5	10.6	14.5	9.0
	BGNet+ (2021)	5.3	6.6	11.2	10.3
	Fast-ACVNet (2024a)	12.4	10.6	13.5	7.9
	IINet (2024)	11.6	8.5	-	-
	RT-IGEV (2025c)	5.8	6.6	7.8	5.8
	GGEV (Ours)	4.1	5.5	6.5	2.8
Training Set	Scene Flow + CREStereo + Tartan Air				
	RT-IGEV (2025c)	4.0	5.4	8.6	3.4
	GGEV (Ours)	3.6	4.7	5.7	2.2

Table 1: Zero-shot generalization from synthetic to real. Standard threshold-based error metrics are employed for evaluation: 3-pixel for KITTI, 2-pixel for Middlebury, and 1-pixel for ETH3D. **Bold**: Best.

Loss Function

We use Smooth L1 loss to supervise the initial disparity \mathbf{d}_0 regression from GGEV, and employ L1 loss to supervise the disparity $\{\mathbf{d}_i\}_{i=1}^N$ estimated through iterative refinement.

$$\mathcal{L} = |\mathbf{d}_0 - \mathbf{d}_{gt}|_{smooth} + \sum_{i=1}^N \gamma^{N-i} \|\mathbf{d}_i - \mathbf{d}_{gt}\|_1, \quad (13)$$

where \mathbf{d}_{gt} denotes the ground truth disparity, $\gamma = 0.9$ is a decay factor, and N is the number of iterations used during training.

Experiments

Implementation Details

We implement our GGEV with PyTorch and perform our experiments using NVIDIA RTX 3090 GPUs. For all training, we use the AdamW optimizer and clip gradients to the range $[-1, 1]$, and adopt a one-cycle learning rate schedule. We perform 11 update iterations during training and use 8 iterations during inference. Following standard (Chang and Chen 2018; Min et al. 2025; Yao et al. 2025), we pretrain GGEV on the Scene Flow (Mayer et al. 2016) for most experiments. For finetuning on KITTI (Geiger, Lenz, and Urtasun 2012; Menze and Geiger 2015), we use the mixed dataset of KITTI 2012 and KITTI 2015. For finetuning on ETH3D (Schops et al. 2017), we follow the CREStereo (Li et al. 2022) and GMStereo (Xu et al. 2023b), using a collection of public stereo datasets. Detailed information is provided in the Supplementary Material.

Zero-Shot Generalization

We follow MonSter by training our model on the Scene Flow and directly evaluating it on four real-world datasets. The re-




Figure 5: Qualitative comparison on ETH3D.

sults are shown in Tab. 1, our GGEV achieves the best performance among all real-time models and demonstrates generalization capability comparable to that of some accuracy-oriented methods. This improvement is attributed to the incorporation of monocular depth priors and the adaptive handling of different disparity hypotheses. Specifically, compared to the SOTA real-time method RT-IGEV, our GGEV reduces the error rates by 29% on KITTI 2012, 16% on KITTI 2015, and 16% on Middlebury-quarter. More notably, it achieves a 51% error reduction on the real-world indoor-outdoor ETH3D dataset. Our GGEV also outperforms domain generalization methods, such as FC-GANet. Using the same ViT-S backbone, our GGEV achieves comparable performance to DEFOM-Stereo (ViT-S, 255ms), while reducing inference time by 81%, highlighting the efficiency of our approach.

By adding more synthetic datasets, the generalization performance of our GGEV is further improved, still surpassing that of RT-IGEV under the same training conditions. This demonstrates that high-quality training data can better guide our network in learning effective cost aggregation.

Benchmark Comparisons

To demonstrate the outstanding performance of our method, we conduct comprehensive comparisons with prior methods on three widely-used stereo benchmarks: KITTI 2012, KITTI 2015, and ETH3D. The fine-tuning settings are in the Supplementary Material.

KITTI. As shown in Tab. 2 and Fig. 8, among all real-time models, our method achieves 1st performance. On the KITTI 2012, our proposed GGEV outperforms both RT-IGEV and BANet-3D by 13% on the 2-noc and 3-noc metrics. On the KITTI 2015, our proposed GGEV achieves the top performance on both D1-bg and D1-all metrics.

ETH3D. As shown in Tab. 3 and Fig. 5, our GGEV significantly outperforms all existing real-time stereo matching methods across all evaluation metrics, with nearly 50% error reduction on each metric. Notably, our method surpasses both GMStereo and Selective-IGEV on the Bad 1.0 metric, while requiring less than one-fourth of their inference time.

Target	Method	KITTI 2012						KITTI 2015			Time (ms)
		2-pixels		3-pixels		EPE		D1-bg	D1-fg	D1-all	
		noc	all	noc	all	noc	all	all			
Accuracy	RAFT-Stereo (2021)	1.92	2.42	1.30	1.66	0.4	0.5	1.58	3.05	1.82	380
	IGEVS (2023a)	1.71	2.17	1.12	1.44	0.4	0.4	1.38	2.67	1.59	180
	Selective-IGEVS (2024)	1.59	2.05	1.07	1.38	0.4	0.4	1.33	2.61	1.55	240
	Moster (2025a)	1.36	1.75	0.84	1.09	0.4	0.4	1.13	2.81	1.41	450
Speed	DeepPrunerFast (2019)	-	-	-	-	-	-	2.32	3.91	2.59	50
	AANet (2020)	2.30	2.96	1.55	2.05	0.4	0.5	1.65	3.96	2.03	60
	HITNet (2021)	2.00	2.65	1.41	1.89	0.4	0.5	1.74	3.20	1.98	20
	BGNet+ (2021)	2.78	3.35	1.62	2.03	0.5	0.6	1.81	4.09	2.19	35
	DecNet (2021)	-	-	-	-	-	-	2.07	3.87	2.37	50
	CoEx (2021)	2.54	3.09	1.55	1.93	0.5	0.5	1.74	3.41	2.02	27
	TemporalStereo (2023)	-	-	-	-	-	-	1.61	2.78	1.81	45
	Fast-ACVNet+ (2024a)	2.39	2.97	1.45	1.85	0.5	0.5	1.70	3.53	2.01	45
	IINet (2024)	2.76	3.34	1.81	2.21	0.5	0.5	2.02	3.39	2.25	26
	RT-IGEVS (2025c)	1.93	2.51	1.29	1.68	0.4	0.5	1.48	3.37	1.79	40
	BANet-3D (2025b)	2.08	2.71	1.27	1.72	0.5	0.5	1.52	3.02	1.77	30
	GGEVS (Ours)	1.66	2.17	1.10	1.44	0.4	0.4	1.38	3.28	1.70	47

Table 2: Quantitative evaluation on KITTI 2012 and KITTI 2015. **Bold**: Best.

Tgt	Method	ETH3D			
		Bad 0.5	Bad 1.0	Bad 2.0	AvgErr
Accuracy	RAFT-Stereo (2021)	7.04	2.44	0.44	0.18
	GMStereo (2023b)	5.94	1.83	0.25	0.19
	Selective-IGEVS (2024)	3.06	1.23	0.22	0.12
	FoundationStereo (2025)	1.26	0.26	0.08	0.09
Speed	Fast-ACVNet (2024a)	14.25	5.62	1.41	0.31
	HITNet (2021)	7.83	2.79	0.80	0.20
	GGEVS (Ours)	3.70	1.19	0.34	0.14

Table 3: Quantitative evaluation on ETH3D. **Bold**: Best.

Ablation Study

We conduct a series of ablation studies to verify the effectiveness and rationality of the proposed modules. We use a simplified RT-IGEVS as our baseline by reducing the aggregation depth to two downsampling stages and the number of inference iterations is set to 8. All model variants are trained for 200k steps on the Scene Flow and evaluated under both in-domain and zero-shot generalization settings. Additionally, we report the inference time and trainable parameters for each variant. The results are shown in Tab. 4.

Effectiveness of Depth Feature Encoder. The incorporation of depth features enhances generalization performance, owing to the strong generalization of the pretrained MFMs. However, treating all disparity hypotheses uniformly fails to fully leverage the potential of the depth features.

Effectiveness of Selective Channel Fusion. The fusion of texture and depth features enhances the model’s in-domain fitting, but its effect on generalization is mixed.




Figure 6: Qualitative comparison of disparity regression results between the GEVS from RT-IGEVS and our generalized GEVS on Scene Flow, KITTI, Middlebury and ETH3D. GGEVS shows improved performance in edges, occlusions, and textureless areas, even without iterative operations.

Effectiveness of Dynamic Cost Aggregation. We deploy Dynamic Cost Aggregation module guided by texture features, which significantly improves in-domain performance. However, texture features are sensitive to textureless regions and appearance variations, resulting in limited improvements in generalization.

Effectiveness of All Components. By integrating all the proposed modules, we observe comprehensive improve-

Models	Proposed Modules			Scene Flow	KITTI 2012	KITTI 2015	Middlebury	ETH3D	Params. (M)	Time (ms)
	DFE	SCF	DCA	EPE	D1	D1	Bad 2.0	Bad 1.0		
Baseline				0.54	6.63	8.01	7.84	7.54	3.60	30
+DFE (ViT-S)	✓			0.52	4.40	6.32	6.47	5.02	3.57	37
+DFE+SCF	✓	✓		0.49	5.14	7.58	5.57	4.85	3.65	38
+DCA			✓	0.47	6.80	6.75	7.73	5.61	3.63	39
Full Model (ViT-L)	✓	✓	✓	0.45	4.20	5.65	4.41	1.69	3.75	110
Full Model (ViT-S)	✓	✓	✓	0.46	4.11	5.56	6.53	2.84	3.68	47

Table 4: Ablation study of proposed networks on the Scene Flow test set and zero-shot generation. The parameters counted here are the trainable ones. The run-time is measured at the KITTI resolution of 1248×384 .

Method	KITTI 2012 (Reflective)			
	2-noc	2-all	3-noc	3-all
RAFT-Stereo (2021)	8.41	9.87	5.40	6.48
CREStereo (2022)	9.71	11.26	6.27	7.27
HITNet (2021)	9.75	11.85	5.91	7.54
RT-IGEV (2025c)	9.56	11.54	5.76	7.26
BANet-3D (2025b)	9.64	11.97	5.37	7.07
GGEV (Ours)	7.33	9.27	4.04	5.34

Table 5: Evaluation on reflective (Ill-Posed) regions of the KITTI 2012 benchmark. **Bold**: Best.

ments in both accuracy and generalization. The SCF introduces the generalization capability of MFMs into the cost aggregation, while the DDCA adaptively incorporates depth features to enhance fragile matching relationships. Qualitative results of the initial disparity from RT-IGEV and GGEV are presented in Fig. 6, where our full model demonstrates both high accuracy and strong generalization capability, even without iterative refinement. Additionally, employing a ViT-L backbone further boosts model performance, while introducing higher inference latency.

Trainable Parameters. Our full model (ViT-S) introduces only a marginal increase of 0.08M (+2%) trainable parameters compared to the baseline. Specifically, the proposed DFE and SCF add 0.05M parameters, while the DCA contributes an additional 0.03M.

Inference Times. Our full model (ViT-S) satisfies HITNet’s real-time constraint of 100ms, with a 17ms increase over the baseline, including 8ms from the DFE and SCF and 9ms from the DCA, while achieving comparable speed to RT-IGEV and delivering significantly better performance.

Number of Iterations. Our GGEV achieves better in-domain and zero-shot performance with fewer iterations. As shown in Fig. 7, our method consistently outperforms RT-IGEV under the same number of iterations.

Performance in Ill-posed Regions

We evaluate GGEV on reflective areas in KITTI 2012. As shown in Tab. 5, our GGEV achieves the best performance among all real-time stereo matching methods, even sur-




Figure 7: Performance comparison with different numbers of inference iterations.




Figure 8: Qualitative comparison on KITTI. Our GGEV shows significant improvement in challenging regions such as fine structures and reflective surfaces.

passing several accuracy-oriented approaches. Compared to RAFT-Stereo, it achieves improvements of 12% and 25% on the 2-noc and 3-noc metrics, respectively.

Conclusion

We presented GGEV, a real-time stereo matching framework that achieves impressive zero-shot generalization. The proposed Selective Channel Fusion and Depth-aware Dynamic Cost Aggregation modules extract depth-aware features as generalized structural guidance and adaptively aggregate disparity-specific critical regions, collaboratively generating a generalized geometry encoding volume. Experimental results demonstrate that our method consistently outperforms all existing real-time approaches. Future work could explore leveraging metric depth foundation models to provide more accurate depth guidance or extending our method to real-time video stereo matching.

Acknowledgments

This research is supported by the National Key R&D Program of China (2024YFE0217700), the National Natural Science Foundation of China (62472184,623B2036), the Fundamental Research Funds for the Central Universities, and the Innovation Project of Optics Valley Laboratory (Grant No. OVL2025YZ005).

References

- Bangunharcana, A.; Cho, J. W.; Lee, S.; Kweon, I. S.; Kim, K.-S.; and Kim, S. 2021. Correlate-and-Excite: Real-Time Stereo Matching via Guided Cost Volume Excitation. In *IROS*, 3542–3548. IEEE.
- Bartolomei, L.; Tosi, F.; Poggi, M.; and Mattocchia, S. 2025. Stereo anywhere: Robust zero-shot deep stereo matching even where either stereo or mono fail. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 1013–1027.
- Caron, M.; Touvron, H.; Misra, I.; Jégou, H.; Mairal, J.; Bojanowski, P.; and Joulin, A. 2021. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, 9650–9660.
- Chang, J.-R.; Chang, P.-C.; and Chen, Y.-S. 2020. Attention-aware feature aggregation for real-time stereo matching on edge devices. In *ACCV*.
- Chang, J.-R.; and Chen, Y.-S. 2018. Pyramid stereo matching network. In *CVPR*, 5410–5418.
- Cheng, J.; Liu, L.; Xu, G.; Wang, X.; Zhang, Z.; Deng, Y.; Zang, J.; Chen, Y.; Cai, Z.; and Yang, X. 2025a. Monster: Marry monodepth to stereo unleashes power. In *CVPR*, 6273–6282.
- Cheng, J.; Liu, L.; Xu, G.; Wang, X.; Zhang, Z.; Deng, Y.; Zang, J.; Chen, Y.; Cai, Z.; and Yang, X. 2025b. Monster: Marry Monodepth to Stereo Unleashes Power. *arXiv preprint arXiv:2501.08643*.
- Chuah, W.; Tennakoon, R.; Hoseinnezhad, R.; Bab-Hadiashar, A.; and Suter, D. 2022. Itsa: An information-theoretic approach to automatic shortcut avoidance and domain generalization in stereo matching networks. In *CVPR*, 13022–13032.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *CVPR*, 248–255. Ieee.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Duggal, S.; Wang, S.; Ma, W.-C.; Hu, R.; and Urtasun, R. 2019. DeepPruner: Learning Efficient Stereo Matching via Differentiable PatchMatch. In *ICCV*, 4384–4393.
- Geiger, A.; Lenz, P.; and Urtasun, R. 2012. Are we ready for autonomous driving? the kitti vision benchmark suite. In *CVPR*, 3354–3361. IEEE.
- Gu, X.; Fan, Z.; Zhu, S.; Dai, Z.; Tan, F.; and Tan, P. 2020. Cascade cost volume for high-resolution multi-view stereo and stereo matching. In *CVPR*, 2495–2504.
- Jiang, H.; Lou, Z.; Ding, L.; Xu, R.; Tan, M.; Jiang, W.; and Huang, R. 2025. Defom-stereo: Depth foundation model based stereo matching. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 21857–21867.
- Khamis, S.; Fanello, S.; Rhemann, C.; Kowdle, A.; Valentin, J.; and Izadi, S. 2018. Stereonet: Guided hierarchical refinement for edge-aware depth prediction. In *ECCV*.
- Li, J.; Wang, P.; Xiong, P.; Cai, T.; Yan, Z.; Yang, L.; Liu, J.; Fan, H.; and Liu, S. 2022. Practical stereo matching via cascaded recurrent network with adaptive correlation. In *CVPR*, 16263–16272.
- Li, X.; Zhang, C.; Su, W.; and Tao, W. 2024. Iinet: Implicit intra-inter information fusion for real-time stereo matching. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 3225–3233.
- Li, Y.; Xiong, K.; Guo, X.; Li, F.; Yan, S.; Xu, G.; Zhou, L.; Chen, L.; Sun, H.; Wang, B.; et al. 2025. Recogdrive: A reinforced cognitive framework for end-to-end autonomous driving. *arXiv preprint arXiv:2506.08052*.
- Liang, B.; Wang, Y.; Huang, Z.; Hu, Z.; Hu, H.; Xu, J.; and Chen, D. 2025a. ARUNet: Advancing Real-Time Stereo Matching for Robotic Perception on Edge Devices. *IEEE Robotics and Automation Letters*.
- Liang, D.; Feng, T.; Zhou, X.; Zhang, Y.; Zou, Z.; and Bai, X. 2025b. Parameter-efficient fine-tuning in spectral domain for point cloud learning. *IEEE TPAMI*.
- Liang, D.; Hua, W.; Shi, C.; Zou, Z.; Ye, X.; and Bai, X. 2025c. Sood++: Leveraging unlabeled data to boost oriented object detection. *IEEE TPAMI*.
- Lin, H.; Peng, S.; Chen, J.; Peng, S.; Sun, J.; Liu, M.; Bao, H.; Feng, J.; Zhou, X.; and Kang, B. 2025. Prompting depth anything for 4k resolution accurate metric depth estimation. In *CVPR*, 17070–17080.
- Lipson, L.; Teed, Z.; and Deng, J. 2021. Raft-stereo: Multi-level recurrent field transforms for stereo matching. In *3DV*, 218–227. IEEE.
- Liu, B.; Yu, H.; and Qi, G. 2022. GraftNet: Towards Domain Generalized Stereo Matching with a Broad-Spectrum and Task-Oriented Feature. In *CVPR*, 13012–13021.
- Lou, M.; and Yu, Y. 2025. OverLoCK: An Overview-first-Look-Closely-next ConvNet with Context-Mixing Dynamic Kernels. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 128–138.
- Mayer, N.; Ilg, E.; Hausser, P.; Fischer, P.; Cremers, D.; Dosovitskiy, A.; and Brox, T. 2016. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In *CVPR*, 4040–4048.
- Menze, M.; and Geiger, A. 2015. Object scene flow for autonomous vehicles. In *CVPR*, 3061–3070.
- Min, J.; Jeon, Y.; Kim, J.; and Choi, M. 2025. S^2M^2 : Scalable Stereo Matching Model for Reliable Depth Estimation. *arXiv preprint arXiv:2507.13229*.

- Oquab, M.; Darcet, T.; Moutakanni, T.; Vo, H.; Szafraniec, M.; Khalidov, V.; Fernandez, P.; Haziza, D.; Massa, F.; El-Nouby, A.; et al. 2023. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*.
- Ranftl, R.; Lasinger, K.; Hafner, D.; Schindler, K.; and Koltun, V. 2020. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE transactions on pattern analysis and machine intelligence*, 44(3): 1623–1637.
- Schops, T.; Schonberger, J. L.; Galliani, S.; Sattler, T.; Schindler, K.; Pollefeys, M.; and Geiger, A. 2017. A multi-view stereo benchmark with high-resolution images and multi-camera videos. In *CVPR*, 3260–3269.
- Tankovich, V.; Hane, C.; Zhang, Y.; Kowdle, A.; Fanello, S.; and Bouaziz, S. 2021. Hitnet: Hierarchical iterative tile refinement network for real-time stereo matching. In *CVPR*, 14362–14372.
- Wang, Q.; Shi, S.; Zheng, S.; Zhao, K.; and Chu, X. 2020. Fadnet: A fast and accurate network for disparity estimation. In *ICRA*, 101–107. IEEE.
- Wang, X.; Xu, G.; Jia, H.; and Yang, X. 2024. Selective-stereo: Adaptive frequency information selection for stereo matching. In *CVPR*, 19701–19710.
- Wen, B.; Trepte, M.; Aribido, J.; Kautz, J.; Gallo, O.; and Birchfield, S. 2025. Foundationstereo: Zero-shot stereo matching. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 5249–5260.
- Xu, B.; Xu, Y.; Yang, X.; Jia, W.; and Guo, Y. 2021. Bilateral Grid Learning for Stereo Matching Networks. In *CVPR*, 12497–12506.
- Xu, G.; Cheng, J.; Guo, P.; and Yang, X. 2022. Attention Concatenation Volume for Accurate and Efficient Stereo Matching. In *CVPR*, 12981–12990.
- Xu, G.; Lin, H.; Luo, H.; Wang, X.; Yao, J.; Zhu, L.; Pu, Y.; Chi, C.; Sun, H.; Wang, B.; et al. 2025a. Pixel-perfect depth with semantics-prompted diffusion transformers. *arXiv preprint arXiv:2510.07316*.
- Xu, G.; Liu, J.; Wang, X.; Cheng, J.; Deng, Y.; Zang, J.; Chen, Y.; and Yang, X. 2025b. BANet: Bilateral Aggregation Network for Mobile Stereo Matching. *arXiv preprint arXiv:2503.03259*.
- Xu, G.; Wang, X.; Ding, X.; and Yang, X. 2023a. Iterative geometry encoding volume for stereo matching. In *CVPR*, 21919–21928.
- Xu, G.; Wang, X.; Zhang, Z.; Cheng, J.; Liao, C.; and Yang, X. 2025c. Igev++: Iterative multi-range geometry encoding volumes for stereo matching. *IEEE TPAMI*.
- Xu, G.; Wang, Y.; Cheng, J.; Tang, J.; and Yang, X. 2024a. Accurate and Efficient Stereo Matching via Attention Concatenation Volume. *IEEE TPAMI*, 46(4): 2461–2474.
- Xu, G.; Wang, Y.; Gu, J.; Xue, T.; and Yang, X. 2024b. Hdr-flow: Real-time hdr video reconstruction with large motions. In *CVPR*, 24851–24860.
- Xu, H.; and Zhang, J. 2020. Aanet: Adaptive aggregation network for efficient stereo matching. In *CVPR*, 1959–1968.
- Xu, H.; Zhang, J.; Cai, J.; Rezatofighi, H.; Yu, F.; Tao, D.; and Geiger, A. 2023b. Unifying flow, stereo and depth estimation. *IEEE TPAMI*.
- Yang, L.; Kang, B.; Huang, Z.; Xu, X.; Feng, J.; and Zhao, H. 2024a. Depth anything: Unleashing the power of large-scale unlabeled data. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10371–10381.
- Yang, L.; Kang, B.; Huang, Z.; Zhao, Z.; Xu, X.; Feng, J.; and Zhao, H. 2024b. Depth anything v2. *Advances in Neural Information Processing Systems*, 37: 21875–21911.
- Yao, C.; Jia, Y.; Di, H.; Li, P.; and Wu, Y. 2021. A Decomposition Model for Stereo Matching. In *CVPR*, 6091–6100.
- Yao, C.; Yu, L.; Liu, Z.; Zeng, J.; Wu, Y.; and Jia, Y. 2025. Diving into the Fusion of Monocular Priors for Generalized Stereo Matching. *arXiv preprint arXiv:2505.14414*.
- Zhang, F.; Qi, X.; Yang, R.; Prisacariu, V.; Wah, B.; and Torr, P. 2020. Domain-invariant stereo matching networks. In *ECCV*, 420–439. Springer.
- Zhang, J.; Wang, X.; Bai, X.; Wang, C.; Huang, L.; Chen, Y.; Gu, L.; Zhou, J.; Harada, T.; and Hancock, E. R. 2022. Revisiting Domain Generalized Stereo Matching Networks from a Feature Consistency Perspective. In *CVPR*, 13001–13011.
- Zhang, Y.; Poggi, M.; and Mattocchia, S. 2023. Temporal-stereo: Efficient spatial-temporal stereo matching network. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 9528–9535. IEEE.
- Zhou, J.; Zhang, H.; Yuan, J.; Ye, P.; Chen, T.; Jiang, H.; Chen, M.; and Zhang, Y. 2025. All-in-One: Transferring Vision Foundation Models into Stereo Matching. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 10797–10805.