

MoCha-Stereo: Motif Channel Attention Network for Stereo Matching

Ziyang Chen^{1*} Wei Long^{1*} He Yao^{1*} Yongjun Zhang^{1†}
Bingshu Wang² Yongbin Qin¹ Jia Wu¹

^{1†} College of Computer Science and Technology, The State Key Laboratory of Public Big Data,
Institute of Artificial Intelligence, Guizhou University

² College of Software, Northwest Polytechnical University

ziyangchen2000@gmail.com; zyj6667@126.com[†]



Figure 1. **Motivation. Addressing the issue of geometric structure loss in feature channels arising from deep learning.** (From left to right: Input image, visualization of a Normal Channel, visualization of a Motif [39] Channel.) Due to the fuzziness of geometric edges in certain channels, achieving accurate matching of stereo image edges is a challenging problem. MoCha-Stereo guides ordinary channels to focus on edge features through motif channels, achieving more accurate detail matching. Motif Channel refers to channel that composed of repeatedly occurring geometric contours. The regions delineated by the yellow border represent the magnified details.

Abstract

Learning-based stereo matching techniques have made significant progress. However, existing methods inevitably lose geometrical structure information during the feature channel generation process, resulting in edge detail mismatches. In this paper, the **Motif Channel Attention Stereo Matching Network (MoCha-Stereo)** is designed to address this problem. We provide the **Motif Channel Correlation Volume (MCCV)** to determine more accurate edge matching costs. MCCV is achieved by projecting motif channels, which capture common geometric structures in feature channels, onto feature maps and cost volumes. In addition, edge variations in the reconstruction error map also affect details matching, we propose the **Reconstruction Error Motif Penalty (REMP)** module to further refine the full-resolution disparity estimation. REMP integrates the frequency information of typical channel features from the reconstruction error. MoCha-Stereo ranks **1st** on the **KITTI-2015** and **KITTI-2012 Reflective** leaderboards. Our structure also shows excellent performance in **Multi-View Stereo**.

*Co-first author.

†Corresponding author.

Code is available at [MoCha-Stereo](#).

1. Introduction

Stereo matching remains a foundational challenge in computer vision, bearing significant relevance to autonomous driving, virtualization, rendering, and related sectors [40]. The primary goal of the assignment is to establish a pixel-wise displacement map, or disparity, which can be used to identify the depth of the pixels in the scene. Edge performance of disparity maps is particularly vital in techniques requiring pixel-level rendering, such as virtual reality and augmented reality, where precise fitting between the scene model and image mapping is essential [23]. This underscores the need for a close alignment between the edges of the disparity map and the original RGB image.

Traditional stereo matching relies on global [8], semi-global [13], or local [2] grayscale relationships between left and right view pixels. These methods struggle to fully leverage scene-specific prior knowledge. Achieving optimal results often involves human observation, this tuning process can be resource-intensive in scenes with complex images

[41]. With the advancement of deep learning, learning-based methods [15, 30, 40] have generally achieved better results. A case in point is RAFT-Stereo [17], which introduced a coarse-to-fine scheme by computing All-Pairs Correlation (APC). To comprehensively learn channel features, GwcNet [12] proposes a method of computing correlations by grouping left and right features along the channel dimension, called Group-Wise Correlation (GWC). IGEV-Stereo [36] introduces Combined Geometry Encoding Volume (CGEV), a cost calculation method that combines GWC [12] and APC [17], this cost calculation approach has achieved state-of-the-art results. There is also a body of research focused on obtaining more accurate matching results through post-processing of disparities. They [3, 16, 47] utilize CNN structures directly applied to the additional error maps in the hope of achieving better results.

Learning-based methods have achieved impressive results. However, numerous channels experience loss of geometric details during the generation of feature channels. This phenomenon leads to a mismatch in the representation of object edges. Loss of geometric edges in the channel is a challenging problem because each block of the neural network performs non-linear transformations [31], excessive nonlinearity can saturate activations in some channels, and insufficient nonlinearity leads to inadequate values. It is difficult for deep learning to directly recover geometric details. As shown in the middle picture of Fig. 1, certain normal channels suffer from severe blurring. Details loss in channels naturally complicates the matching of edges.

To address the above problems, we propose **Motif Channel Attention Stereo Matching Network (MoCha-Stereo)**. The core idea of MoCha-Stereo is to restore the lost detailed features in feature channels by utilizing the repeated geometric contours within normal channels. Channels that preserve the common features are referred to as **motif channels**. The following improvements are presented:

- 1) We introduce a novel stereo matching framework that incorporates repeated geometric contours. This architecture enables more accurate cost computation and disparity estimation through detail restoration of feature channels.
- 2) We propose Motif Channel Attention (MCA) to mitigate imbalanced nonlinear transformations in network training. MCA optimizes feature channels through motif channel projection instead of direct network optimization. Inspired by time-series motif mining, we capture motif channels using sliding windows.
- 3) To achieve more precise matching cost computation for edge matching, we construct the Channel Affinity Matrix Profile (CAMP)-guided correlation volume. This volume is derived from the correlation matrix between normal and motif channels, then mapped onto the base correlation volume to produce a more rational cost volume called Motif

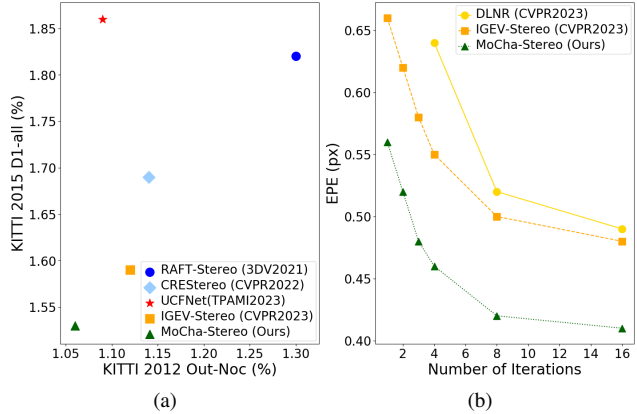


Figure 2. (a) Comparison with SOTA methods [15, 17, 28, 36] on KITTI 2012 [9] and 2015 leaderboards [22] (lower is better). (b) Performance evaluation of the Scene Flow test set [21] in comparison to IGEV-Stereo [36] and DLNR [47] as the number of iterations changes (lower EPE means better).

Channel Correlation Volume (MCCV).

4) To leverage the geometric information of the potential channels in the reconstruction error map, we develop Reconstruction Error Motif Penalty (REMP) to extract the motif channels from the error map, optimizing the disparity based on the high and low-frequency signals.

We validated the performance of MoCha-Stereo on several leaderboards. As shown in Fig. 2 (a), MoCha-Stereo ranks **1st** on KITTI 2015 [22] and KITTI 2012 Reflective [9], achieves SOTA result in Scene Flow [21], zero-shot performance [24, 25], and MVS domain [14]. Our designs also make iteration more efficient. As illustrated in Fig. 2 (b), MoCha-Stereo achieves superior results with fewer iterations, allowing users to choose between efficient or high-precision settings based on their preferences.

2. Related Work

2.1. Motif Mining for Time series analysis

The concept of motif originates from time series analysis. Motif mining has become one of the most commonly used primitives in time series data mining [6, 39]. In a time series T , there exists a subsequence $T_{i,L}$, which starts from the i -th position in the time series T and is a continuous subset of values with a length of L . Motif is the pair of subsequences $T_{a,L}$ and $T_{b,L}$ in a time series that are the most similar. In mathematical notation, For case $\forall i, j \in [1, 2, \dots, n - L + 1]$ and $a \neq b, i \neq j$, the motif [1] satisfies as Equ. 1.

$$\text{dist}(T_{a,L}, T_{b,L}) \leq \text{dist}(T_{i,L}, T_{j,L}) \quad (1)$$

where dist means a distance measure. The distances between all subsequences and their nearest neighbors are

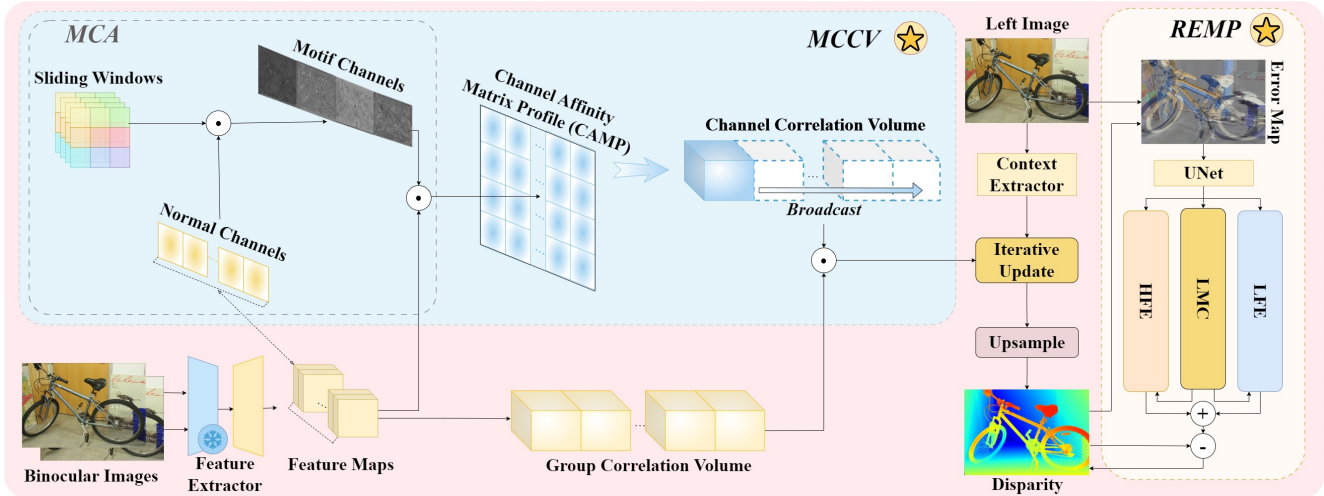


Figure 3. Overview of MoCha-Stereo. MoCha-Stereo initially constructs the Motif Channel Correlation Volume (MCCV) by projecting the relationship between motif channels and normal channels into the basic group correlation volume. Subsequently, based on this cost volume, we employ an iterative way to build the disparity map. Finally, the Reconstruction Error Motif Penalty (REMP) module is applied to penalize the generation of the full-resolution disparity map. In REMP, *LFE* refers to the Low-frequency Error branch, *LMC* denotes to the Latent Motif Channel branch, and *HFE* means the High-frequency Error branch. The **star symbol** means our primary innovations.

stored as **Matrix Profile (MP)** [39], representing the repetitive pattern in the time series. This numerical relational information is utilized to identify recurring patterns, detect anomalies, and perform pattern matching in time series.

For the channel features, the geometric structure of graphics theoretically repeats as well. Differing from time series, the repetitive patterns we aim to uncover represent the geometric structures in images, necessitating consideration of two-dimensional contextual information. Furthermore, the computation of MP requires multiple samplings of subsequences from the series. Selecting sub-patches from multi-channel image features for computing similarity is computationally expensive. Therefore, we develop a collection of adaptable sliding windows that are arranged into two-dimensional vectors. These windows are used for capturing repeating geometrical patterns. These repeated patterns are stored as motif channels for MoCha-Stereo.

2.2. Learning-based Stereo Matching

The field of stereo matching has witnessed considerable advancements owing to learning-based techniques in recent years. Gwc-Net [12] introduced the novel concept of GWC volume, which served as a major inspiration for future landmark achievements in the field [18, 27]. RAFT-Stereo [17] advances the construction of the cost volume by deploying an APC pyramid Volume. Building upon the foundation set by Gwc-Net [12] and RAFT-Stereo [17], IGEV-Stereo [36] proposed a volume that combines APC [17] and GWC [12], utilizing an iterative indexing method for the update of the disparity map. Another portion of learning-based

methods are focus on obtaining more accurate matching results through post-processing of disparities. iResNet [16] and DLNR [47] utilize UNet and Hourglass, respectively, employing convolutional operations to directly output the convolutional results of the reconstruction errors. DRNet [3] opts not to attach additional reconstruction error maps but instead appends Geometric and Photometric Error.

Learning-based methods have made significant progress. However, state-of-the-art algorithms [36, 47] inevitably lose geometric details in certain channels. Few studies have considered using repeated geometric profiles in multiple feature maps to restore the edge texture of channel feature maps. Additionally, there is limited recognition regarding the significance of the frequency information in the potential motif channels from error maps for shaping the edges.

3. Method

To address the aforementioned issues, we respectively introduce MCCV and REMP as depicted in Fig. 3. These components utilise the projection of Motif Channels to re-instate the geometric structure of channel features.

3.1. Context Extractor & Feature Extractor

Following [17, 36], the context extractor consisting of a series of residual blocks and downsampling layers, generating context at 1/4, 1/8, and 1/16 scales. The feature extractor, which also follows [36], uses a backbone pretrained on ImageNet [7] as the frozen layer [29]. The upsampling blocks utilizes skip connections from outputs of downsampling blocks to obtain multi-scale outputs

$f_{l,i}(f_{r,i}) \in \mathbb{R}^{C_i \times \frac{H}{i} \times \frac{W}{i}}$ ($i = 4, 8, 16, 32$), C_i represents the feature channels, and $f_l(f_r)$ denotes the left (right) view features here.

3.2. Motif Channel Correlation Volume

Although multi-channel extractors contribute to the learning of intricate features, an excess of nonlinearity may saturate activation values in specific channels, while insufficient nonlinearity can yield suboptimal activation values. MCCV is proposed to address the imbalanced learning of geometric structures in feature channels.

Motif Channel Attention (MCA) for feature maps. Feature channels exhibit varying degrees of geometric structure loss, but the fundamental geometric structure of feature channels is theoretically invariant. Inspired by motif mining for time series [6], we use N_s sets of sliding windows SW ($N_s = 4$), each with a length of 9, and organized the windows into 3×3 , to mine repetitive patterns. Unlike the window used for mining temporal motifs, we designed a set of adaptive-weight windows, rather than extracting values directly from the feature map. The initial values of the sliding window are set as random parameters. Based on the gradient loss, we backpropagate to adjust the values of the window weights. Following Equ. 2, we obtain the s -th ($1 \leq s \leq N_s$) motif channel feature map f_{fre}^{mc} in the frequency domain.

$$f_{fre}^{mc}(s, h, w) = \sum_{c=1}^{N_c} \sum_{i=1}^3 \sum_{j=1}^3 (SW(s, h+i, w+j) \times f_{fre}(c, h, w)) \quad (2)$$

where (h, w) are the coordinates of the pixel, c denotes the c -th feature channel, N_c is the number of normal feature channels. The frequency domain feature $f_{fre} = F(f - G(f))$. F is the Fourier transform, G denotes Gaussian high-pass filter with 3×3 kernels. This approach aims to capture repeatedly occurring frequency-domain features. The rationale behind this is that edge textures often exhibit high-frequency expressions. We then transformed the results back to the spatial domain, accumulating and normalizing them to derive motif channels f^{mc} . This aggregation method takes into account the surrounding pixel information and strengthens attention to the geometric structure repeated across channels through accumulation, enhancing the reliability of matching for edge textures.

Channel Affinity Matrix Profile (CAMP) guided Correlation Volume. In order to enhance the accuracy of matching cost computation for edges with the assistance of motif channels, we propose the Correlation Volume guided by CAMP, as shown in Fig. 3. The foundational cost volume still employs the extraction of feature maps $f_{l(r),4}$ using GWC [12] according to Equ. 3.

$$\mathbb{C}_g(d, h, w, g) = \frac{1}{N_c/N_g} \langle f_{l,4}^g(h, w), f_{r,4}^g(h, w + d) \rangle \quad (3)$$

where d is the disparity level, $\langle \dots, \dots \rangle$ is the inner product, N_c is the number of channels, N_g is the group number ($N_g = 8$). Calculating Correlation Volume solely from feature maps makes it challenging to accurately match details. This difficulty arises because certain channels lose geometric structure. MoCha-Stereo exploits the motif channels f^{mc} obtained by MCA and the normal channels $f_{l,4}$ obtained by feature extractor. Using their affinity, MoCha-Stereo constructs CAMP, which is a matrix that stores the relationship between motif channels and normal channels, as shown in Equ. 4.

$$CAMP(s, c, h, w) = f^{mc}(s, h, w) \times f_{l,4}(c, h, w), \quad \text{where } 1 \leq s \leq N_s, 1 \leq c \leq N_c \quad (4)$$

where s denotes the s -th motif channel, and c denotes the c -th normal channel. CAMP allows the projection of motif channels onto normal channels, serving as a coefficient to modulate the spatial domain information of normal channels. This enhances attention to the geometric structure of channels. We can construct a Channel Correlation (\mathbb{C}_c) based on CAMP. As shown in Equ. 5, \mathbb{C}_c performs spatial interaction across channels, theoretically reinforcing frequent patterns in space, which are the desired geometric structure.

$$\mathbb{C}_c(d, h, w) = \sum_{s=1}^{N_s} \sum_{c=1}^{N_c} (3DCConv(CAMP(s, c, h, w)), 3DCConv(CAMP(s, c, h, w + d))) \quad (5)$$

where $3DCConv$ means 3D convolution operator.

To obtain the final cost volume \mathbb{C} , MoCha-Stereo uses \mathbb{C}_c as a weight-adjusted basis for the basic Correlation Volume. Since geometric structures are theoretically invariant, there is no need to learn new \mathbb{C}_c by adding extra groups. Broadcasting is sufficient to achieve the interaction between \mathbb{C}_c and GWC \mathbb{C}_g , as illustrated in Equ. 6, enabling different groups to learn the same set of geometric structure features.

$$\mathbb{C}(d, h, w) = \sum_{g=1}^{N_g} (\mathbb{C}_g(d, h, w, g) \times \mathbb{C}_c(d, h, w)) \quad (6)$$

3.3. Iterative Update Operator

Following [36, 46, 47], MoCha-Stereo utilizes the iterative update operator [47] to obtain the disparity map $d_k = d_{k-1} + \Delta d_k$ at 1/4 resolution for the k -th iteration.

3.4. Reconstruction Error Motif Penalty

The disparity map output by the iteration is at a resolution of 1/4 of the original image. There is still room for optimization after upsampling the disparity map. Several works [16, 46, 47] have been dedicated to refinement networks based on reconstruction error. However, none of these works have addressed the separation of the low-frequency

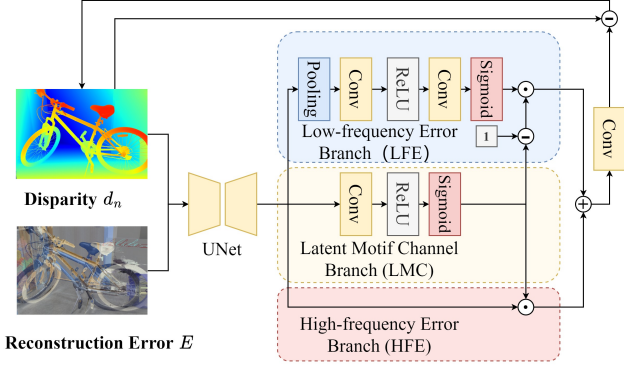


Figure 4. REMP module for Full-Resolution Refine. The upper branch (LFE) obtains low-frequency information through pooling, the lower branch (HFE) retains the original high-resolution image as high-frequency detailed information, and the middle branch (LMC) learns motif features through CNN.

and high-frequency components of the reconstruction error, making the refinement process challenging to achieve more effective results. We propose the Reconstruction Error Motif Penalty (REMP) module for Full-Resolution Refine. The disparity map is a single-channel image. To obtain multi-channel information, we draw inspiration from [3], using the error as the input to the network and outputting multi-channel features. In contrast to [3], we use reconstruction error E obtained by Equ. 7 and the disparity map d_n from the last iteration as inputs.

$$E = K_l(R - \frac{TN^T}{D})K_r^{-1}I_r - I_l \quad (7)$$

$K_{l(r)}$ represents the intrinsic matrix of the left (right) camera in the stereo system, R is the rotation matrix from the right view coordinate system to the left view coordinate system, T is the translation matrix from the right view coordinate system to the left view coordinate system, N is the normal vector of the object plane in the right view coordinate system, D is the perpendicular distance between the object plane and the camera light source (this distance is obtained from the computed disparity), $I_l(r)$ is the left (right) image.

As illustrated in Fig. 4, the UNet in REMP is solely designed to obtain multi-channel features related to the reconstruction error E and disparity map d_n . The core idea of REMP is to optimize both high-frequency and low-frequency errors in the disparity map using representative motif information. REMP divides the features output by the UNet into three branches. The pooling operation in the upper branch (LFE) effectively acts as a low-pass filter, preserving low-frequency information in the image while attenuating high-frequency information to some extent. The middle branch (LMC) guides the network in learning typical motif information, and the lower branch (HFE) undergoes no transformation, aiming to preserve the orig-

inal high-frequency details of the high-resolution image. Through the mappings of these three branches, we learn the feature errors as penalties to refine the disparity map d_n , as shown in Equ. 8.

$$\begin{aligned} o &= UNet(Concat(d'_n, E)) \\ HFE(o) &= o \odot LMC(o) \\ d_n &= d'_n - Conv(LFE(o) \odot (1 - LMC(o)) + HFE(o)) \end{aligned} \quad (8)$$

where \odot means Hadamard product, d'_n represents the disparity map before refinement, lfe denotes the computation process of low-frequency error, and lmc refers to the computation in the LMC branch.

3.5. Loss Function

The computation of the loss function requires the disparity maps outputted at each iteration as well as the initial disparity map d_0 . The initial disparity d_0 is obtained from the Volume $V_g(g \in \{1, 2, \dots, N_g\})$ projected by C_c and GWC C_g . For each group g ($g \leq N_g$), the cost calculation method C_g is uniquely associated with a corresponding volume V_g . We generate d_0 through N_g groups of correlation volumes, expressed by Equ. 9.

$$d_0 = SoftMax(3DCConv(V_1 \oplus V_2 \oplus \dots \oplus V_{N_g})) \quad (9)$$

where \oplus refers to the concatenation operation performed along the group dimension. The initial disparity d_0 serves as the starting point for the iteration and is input to the update module. Following [36], the total loss is defined as Equ. 10.

$$L = Smooth_{L1}(d_0 - d_{gt}) + \sum_{i=1}^n \gamma^{n-i} \|d_i - d_{gt}\|_1 \quad (10)$$

where $Smooth_{L1}$ is defined by [10], $\gamma = 0.9$, d_{gt} is ground truth disparity.

4. Experiment

4.1. Implementation Details

MoCha-Stereo is implemented using the PyTorch framework, with the AdamW [19] optimizer employed during training. For training and ablation experiments, our model was trained on the Scene Flow [21] for 200k epochs, with a batch size of 8, which is equipped with 2 NVIDIA A6000 GPUs. To evaluate the performance of our model, we conducted assessments using the KITTI-2012 [9], KITTI-2015 [22], Scene Flow [21], ETH3D [25], and Middlebury [24]. The training and testing settings are consistent with [17, 36].

Method	Lac-GwcNet[18]	UPFNet [5]	ACVNet [35]	DLNR [47]	IGEV-Stereo [36]	MoCha-Stereo (Ours)
EPE (px)↓	0.75	0.71	0.48	0.48	<u>0.47</u>	0.41 (-12.77%)
Time (s)↓	0.65	0.27	0.48	<u>0.30</u>	0.37	0.34

Table 1. Quantitative evaluation on Scene Flow test set. The **best** result is bolded, and the second-best result is underscored. The variations in the performance of our method compared to the optimal results of other methods are indicated in red font.

Method	All				Reflective			
	Out-Noc (%)↓	Out-All (%)↓	Avg-Noc (px)↓	Avg-All (px)↓	Out-Noc (%)↓	Out-All (%)↓	Avg-Noc (px)↓	Avg-All (px)↓
GwcNet[12]	1.32	1.70	0.5	0.5	7.80	9.28	1.3	1.4
AcfNet[45]	1.17	1.54	0.5	0.5	6.93	8.52	1.8	1.9
RAFT-Stereo [17]	1.30	1.66	0.4	0.5	5.40	6.48	1.3	1.3
HITNet[30]	1.41	1.89	0.4	0.5	5.91	7.54	<u>1.0</u>	1.2
CREStereo[15]	1.14	1.46	0.4	0.5	6.27	7.27	1.4	1.4
Lac-GwcNet [18]	1.13	1.49	0.5	0.5	6.26	8.02	1.5	1.7
IGEV-Stereo [36]	<u>1.12</u>	<u>1.44</u>	0.4	0.4	<u>4.35</u>	<u>5.00</u>	<u>1.0</u>	<u>1.1</u>
MoCha-Stereo(Ours)	1.06 (-5.36%)	1.36	0.4	0.4	3.83 (-11.95%)	4.50	0.8	0.9

Table 2. Results on the KITTI-2012 leaderboard. Out-Noc represents the percentage of erroneous pixels in non-occluded areas, Out-All denotes the percentage of erroneous pixels in the entire image. Avg-Noc refers to the end-point error in non-occluded areas, Avg-All indicates the average disparity error across the entire image. Error threshold is 3 px.

Method	All pixels (%)↓			Noc pixels (%)↓		
	bg	fg	all	bg	fg	all
GwcNet[12]	1.74	3.93	2.11	1.61	3.49	1.92
RAFT-Stereo[17]	1.58	3.05	1.82	1.45	2.94	1.69
CREStereo[15]	1.45	2.86	1.69	1.33	2.60	1.54
Lac-GwcNet[18]	1.43	3.44	1.77	1.30	3.29	1.63
CFNet[26]	1.54	3.56	1.81	1.43	3.25	1.73
UPFNet[5]	<u>1.38</u>	2.85	1.62	1.26	2.70	1.50
CroCo-Stereo[34]	<u>1.38</u>	<u>2.65</u>	<u>1.59</u>	1.30	2.56	1.51
IGEV-Stereo[36]	<u>1.38</u>	2.67	<u>1.59</u>	1.27	2.62	<u>1.49</u>
DLNR[47]	1.60	2.59	1.76	1.45	2.39	1.61
MoCha-Stereo (Ours)	1.36	2.43	1.53 (-3.77%)	1.24	<u>2.42</u>	1.44 (-3.36%)

Table 3. Results on the KITTI-2015 leaderboard. Error threshold is 3 px. Background error is indicated by bg, and front-ground error by fg.

4.2. Comparisons with State-of-the-art

We contrast the SOTA techniques on KITTI-2015 [22], KITTI-2012 [9], and Scene Flow [21]. MoCha-Stereo achieves excellent performance on each of the aforementioned datasets. On the Scene Flow [21] dataset, MoCha-Stereo achieves a new SOTA EPE of 0.41, which surpasses IGEV-Stereo [36] by a margin of **12.77%**. The quantitative comparisons are presented in Tab. 1.

In order to validate the performance of MoCha-Stereo in real-world scenarios, we conducted experiments on the KITTI-2012 [9] and KITTI-2015 [22] benchmarks. MoCha-Stereo ranks *1st* among all the methods submitted to these online benchmarks. Evaluation details are shown in Tab. 2 and Tab. 3. We also provide visualizations of

MoCha-Stereo and compare it with existing SOTA algorithms [17, 18, 36] in Fig. 5. Moreover, MoCha-Stereo achieves *1st* result in reflective regions, where determining geometric edges is often more challenging, as shown in Tab. 2. MoCha-Stereo is the **first** algorithm to control Avg-Noc to within 0.8 px under 5 px error threshold and to control Out-Noc to less than 4% for an error threshold of 3 px among all published methods.

4.3. Zero-shot Generalization

Due to the difficulty in obtaining a large amount of ground truth for real-world scenes, generalization ability is also crucial. We evaluate the generalization ability of MoCha-Stereo by testing it on the Middlebury [24] and ETH3D [25] datasets without fine-tune. As illustrated in Fig. 6 and Tab. 4, our method exhibits SOTA performance in the zero-shot scenarios.

4.4. Extension to MVS

MoCha-Stereo has been extended as MoCha-MVS for application in the field of MVS. Compared to recent learning-based MVS methods, MoCha-MVS achieves excellent performance by balancing accuracy and completeness. As shown in Tab. 5, our method outperforms SOTA methods specifically designed for MVS, indicating the excellent scalability of our approach.

4.5. Ablations

To validate and comprehensively understand the architecture of our model, we conducted certain ablation exper-




Figure 5. Visualisation on the KITTI dataset. We conducted comparisons with existing SOTA methods [17, 18, 36]. Our method accurately captures the edges of the left car and right roof in the first scene. In the second scene, it avoids confusion in the positions of the left two cars, and achieving a complete match of the edges of the right car doors. It is evident that our method excels in matching edge details.




Figure 6. Visualisation on the Middlebury dataset. All results presented in this section demonstrate zero-shot generalization on the Scene Flow dataset. The odd-numbered columns show the original images, the even-numbered columns present zoomed-in details.

Method	Middlebury			ETH↓
	Full↓	Half↓	Quarter↓	
PSMNet[4]	39.5	15.8	9.8	10.2
GANet[42]	32.2	13.5	8.5	6.5
DSMNet[43]	21.8	13.8	8.1	6.2
CFNet[26]	28.2	15.3	9.8	5.8
DLNR [47]	14.5	9.5	7.6	23.1
IGEVStereo [36]	15.2	7.1	6.2	3.6
MoCha-Stereo (Ours)	12.4	6.2	4.9	3.2
	-14.5%	-12.7%	-21.0%	-11.1%

Table 4. Zero-shot evaluation on [24, 25]. Every model undergoes scene flow training without fine-tuning on Middlebury and ETH3D dataset. The 2-pixel error rate is employed for Middlebury, and 1-pixel error rate for ETH3D.

iments. Following [17, 36], all hyperparameter settings remained consistent with the pretraining phase for the Scene Flow dataset.

Motif Channel Correlation Attention (MCCV). For all models in the ablation studies, we perform 16 iterations of updating at inference. As shown in Tab. 6, MCCV contributes to improved prediction accuracy. The decomposition of MCCV from coarse to fine stages into the actions on

Method	Ove.↓	Acc.↓	Comp.↓
MVSNet[38]	0.462	0.396	0.527
CasMVSNet[11]	0.355	0.325	0.385
PatchmatchNet[32]	0.352	0.427	0.277
IterMVS[33]	0.363	0.373	0.354
CER-MVS[20]	0.332	0.359	0.305
Vis-MVSNet[44]	0.365	0.369	0.361
Miper-MVS[48]	0.345	0.364	0.327
DispMVS[37]	0.339	0.354	0.324
MoCha-MVS	0.319 (-5.90%)	0.314	0.325

Table 5. Quantitative evaluation on DTU [14] dataset expanded in MVS domain. Acc. means an indicator of accuracy, Comp. means an indicator of completeness, and Ove. means an indicator of the overall consideration of Acc. and Comp. (lower means better).

the feature map by MCA and on the Correlation Volume by CAMP results in a synergistic effect, leading to a reduction of 7.6% in EPE (from 0.458 to 0.423). The effective improvement achieved by MCCV is attributed to its ability to address the bottleneck of existing feature extractors, which severely lose geometric edge information in some channels. This results in a more reasonable computation of the matching cost for geometric edges. As shown in Fig. 7, directing




Figure 7. An example of one of the feature channels in visual form. The first picture shows the initial normal channel, and the last picture shows the visualization after paying attention to Motif Channels. The middle picture visualize a motif channel. It can be observed that the edge texture details are emphasized in the new feature channels.

Model	MCA for Feature Maps	Correlation Volume guided by CAMP	REMP	EPE (px)	D1-Error >3px(%)	Time (s)	Params.(M)
Baseline				0.458	2.536	0.33	18.31
+MCA	✓			0.449	2.492	0.33	18.32
+REMP			✓	0.445	2.469	0.33	20.70
+MCA+REMP	✓		✓	0.438	2.434	0.33	20.71
+MCCV	✓	✓		0.423	2.358	0.34	18.35
Full model	✓	✓	✓	0.412	2.302	0.34	20.74

Table 6. Ablation study for MoCha-Stereo. The baseline employed in these experiments utilized EfficientNet [29] as the backbone for IGEV-Stereo [36] with 16 iterations. The Time denotes the inference time on single NVIDIA A6000.

	Number of Iterations					
	1	2	3	4	8	16
EPE (px)	0.56	0.52	0.48	0.46	0.42	0.41
Time (s)	0.19	0.20	0.21	0.22	0.26	0.34

Table 7. Ablation study for number of iterations.

the attention of normal channels to motif channels summarizing repeated geometric textures in channel features enhances the clarity of edge textures in normal channels.

Reconstruction Error Motif Penalty (REMP). As shown in Tab. 6, REMP, by incorporating reconstruction error, learns motif information on the channels to understand high and low-frequency errors. The learned errors are then utilized as a penalty term to adjust the disparity map, resulting in a reduction of EPE by 2.8% (from 0.458 to 0.445). This experimental result validates the effectiveness of REMP.

Number of Iterations. MoCha-Stereo enhances the efficiency of iterations. As shown in Tab. 7, MoCha-Stereo, with the information recovered by the Motif Channel, achieves SOTA results without the need for a large number of iterations. For instance, in a comparable inference time, we achieve a **40.8%** reduction in EPE compared to UPFNet [5] (EPE 0.71 px, time 0.27 s) with **8** iterations. With only **4** iterations, MoCha-Stereo outperforms IGEV-Stereo [36] (EPE 0.47 px, time 0.37 s) by over 2.1% in accuracy and saves **40.5%** of the inference time. Information about [5, 36] can be obtained in Tab. 1. Overall, MoCha-

Stereo achieves SOTA performance even with a small number of iterations, allowing users to balance time efficiency and performance based on their specific needs.

5. Conclusion and Future Work

We propose MoCha-Stereo, a novel stereo matching framework. MoCha-Stereo aims to alleviate edge mismatch caused by the geometric structure blurring of channel features. MCCV utilizes the geometric structure of repeated patterns in channel features to restore missing edge details and reconstructs the cost volume based on this novel channel feature structure. REMP penalizes the generation of the full-resolution disparity map based on the high and low-frequency information of the potential motif channel in the reconstruction error. MoCha-Stereo showcases robust cross-dataset generalization capabilities. It ranks **1st** on the KITTI-2015 and KITTI-2012 Reflective online benchmarks and demonstrates SOTA performance on ETH3D, Middlebury, Scene Flow datasets and MVS domain. In the future, we plan to extend the motif channel attention mechanism to more processes in stereo matching, further enhancing the capability of algorithm for edge matching.

Acknowledgement. This research is supported by Science and Technology Planning Project of Guizhou Province, Department of Science and Technology of Guizhou Province, China (Project No. [2023]159). Natural Science Research Project of Guizhou Provincial Department of Education, China (QianJiaoJi[2022]029, QianJiaoHeKY[2021]022).

References

- [1] Sara Alaei, Kaveh Kamgar, and Eamonn Keogh. Matrix profile xxii: exact discovery of time series motifs under dtw. In *Int. Conf. Data Mining*, pages 900–905. IEEE, 2020. **2**
- [2] Michael Bleyer and Margrit Gelautz. Simple but effective tree structures for dynamic programming-based stereo matching. In *VISAPP (2)*, pages 415–422, 2008. **1**
- [3] Rohan Chabra, Julian Straub, Christopher Sweeney, Richard Newcombe, and Henry Fuchs. Stereodnet: Dilated residual stereonet. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 11786–11795, 2019. **2, 3, 5**
- [4] Jia-Ren Chang and Yong-Sheng Chen. Pyramid stereo matching network. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5410–5418, 2018. **7**
- [5] Qibo Chen, Baozhen Ge, and Jianing Quan. Unambiguous pyramid cost volumes fusion for stereo matching. *IEEE Trans. Circuit Syst. Video Technol.*, 2023. **6, 8**
- [6] Hoang Anh Dau and Eamonn Keogh. Matrix profile v: A generic technique to incorporate domain knowledge into motif discovery. In *ACM Special Interest Group Knowl. Discovery Data Mining*, pages 125–134, 2017. **2, 4**
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 248–255. Ieee, 2009. **3**
- [8] Pedro F Felzenszwalb and Daniel P Huttenlocher. Efficient belief propagation for early vision. *Int. J. Comput. Vis.*, 70: 41–54, 2006. **1**
- [9] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2012. **2, 5, 6**
- [10] Ross Girshick. Fast r-cnn. In *Int. Conf. Comput. Vis.*, pages 1440–1448, 2015. **5**
- [11] Xiaodong Gu, Zhiwen Fan, Siyu Zhu, Zuozhuo Dai, Feitong Tan, and Ping Tan. Cascade cost volume for high-resolution multi-view stereo and stereo matching. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2495–2504, 2020. **7**
- [12] Xiaoyang Guo, Kai Yang, Wukui Yang, Xiaogang Wang, and Hongsheng Li. Group-wise correlation stereo network. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3273–3282, 2019. **2, 3, 4, 6**
- [13] Heiko Hirschmuller. Stereo processing by semiglobal matching and mutual information. *IEEE Trans. Pattern Anal. Mach. Intell.*, 30(2):328–341, 2007. **1**
- [14] Rasmus Jensen, Anders Dahl, George Vogiatzis, Engin Tola, and Henrik Aanæs. Large scale multi-view stereopsis evaluation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 406–413, 2014. **2, 7**
- [15] Jiankun Li, Peisen Wang, Pengfei Xiong, Tao Cai, Ziwei Yan, Lei Yang, Jiangyu Liu, Haoqiang Fan, and Shuaicheng Liu. Practical stereo matching via cascaded recurrent network with adaptive correlation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 16263–16272, 2022. **2, 6**
- [16] 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 *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2811–2820, 2018. **2, 3, 4**
- [17] Lahav Lipson, Zachary Teed, and Jia Deng. Raft-stereo: Multilevel recurrent field transforms for stereo matching. In *Int. Conf. 3D Vision*, pages 218–227. IEEE, 2021. **2, 3, 5, 6, 7**
- [18] Biyang Liu, Huimin Yu, and Yangqi Long. Local similarity pattern and cost self-reassembling for deep stereo matching networks. In *Assoc. Advancement Artif. Intell.*, pages 1647–1655, 2022. **3, 6, 7**
- [19] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *Int. Conf. Learn. Represent.*, 2018. **5**
- [20] Zeyu Ma, Zachary Teed, and Jia Deng. Multiview stereo with cascaded epipolar raft. In *Eur. Conf. Comput. Vis.*, pages 734–750. Springer, 2022. **7**
- [21] 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 *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4040–4048, 2016. **2, 5, 6**
- [22] Moritz Menze, Christian Heipke, and Andreas Geiger. Joint 3d estimation of vehicles and scene flow. *ISPRS annals of the photogrammetry, remote sensing and spatial information sciences*, 2:427, 2015. **2, 5, 6**
- [23] Daniel Scharstein and Richard Szeliski. A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *Int. J. Comput. Vis.*, 47:7–42, 2002. **1**
- [24] Daniel Scharstein, Heiko Hirschmüller, York Kitajima, Greg Krathwohl, Nera Nešić, Xi Wang, and Porter Westling. High-resolution stereo datasets with subpixel-accurate ground truth. In *Pattern Recog. German Conf.*, pages 31–42. Springer, 2014. **2, 5, 6, 7**
- [25] 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 *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3260–3269, 2017. **2, 5, 6, 7**
- [26] Zhelun Shen, Yuchao Dai, and Zhibo Rao. Cfnets: Cascade and fused cost volume for robust stereo matching. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 13906–13915, 2021. **6, 7**
- [27] Zhelun Shen, Yuchao Dai, Xibin Song, Zhibo Rao, Dingfu Zhou, and Liangjun Zhang. Pcw-net: Pyramid combination and warping cost volume for stereo matching. In *Eur. Conf. Comput. Vis.*, pages 280–297. Springer, 2022. **3**
- [28] Zhelun Shen, Xibin Song, Yuchao Dai, Dingfu Zhou, Zhibo Rao, and Liangjun Zhang. Digging into uncertainty-based pseudo-label for robust stereo matching. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2023. **2**
- [29] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *Int. Conf. Mach. Learn.*, pages 6105–6114. PMLR, 2019. **3, 8**

- [30] 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 *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 14362–14372, 2021. [2](#), [6](#)
- [31] Jan L van Hemmen and Reimer Kühn. Nonlinear neural networks. *Physical Rev. Lett.*, 57(7):913, 1986. [2](#)
- [32] Fangjinhua Wang, Silvano Galliani, Christoph Vogel, Pablo Speciale, and Marc Pollefeys. Patchmatchnet: Learned multi-view patchmatch stereo. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 14194–14203, 2021. [7](#)
- [33] Fangjinhua Wang, Silvano Galliani, Christoph Vogel, and Marc Pollefeys. Itermvs: iterative probability estimation for efficient multi-view stereo. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 8606–8615, 2022. [7](#)
- [34] Philippe Weinzaepfel, Thomas Lucas, Vincent Leroy, Yohann Cabon, Vaibhav Arora, Romain Brégier, Gabriela Csurka, Leonid Antsfeld, Boris Chidlovskii, and Jerome Revaud. Croco v2: Improved cross-view completion pre-training for stereo matching and optical flow. In *Int. Conf. Comput. Vis.*, pages 17969–17980, 2023. [6](#)
- [35] Gangwei Xu, Junda Cheng, Peng Guo, and Xin Yang. Attention concatenation volume for accurate and efficient stereo matching. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 12981–12990, 2022. [6](#)
- [36] Gangwei Xu, Xianqi Wang, Xiaohuan Ding, and Xin Yang. Iterative geometry encoding volume for stereo matching. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 21919–21928, 2023. [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [37] Qingsong Yan, Qiang Wang, Kaiyong Zhao, Bo Li, Xiaowen Chu, and Fei Deng. Rethinking disparity: a depth range free multi-view stereo based on disparity. In *Assoc. Advancement Artif. Intell.*, pages 3091–3099, 2023. [7](#)
- [38] Yao Yao, Zixin Luo, Shiwei Li, Tian Fang, and Long Quan. Mvsnet: Depth inference for unstructured multi-view stereo. In *Eur. Conf. Comput. Vis.*, pages 767–783, 2018. [7](#)
- [39] Chin-Chia Michael Yeh, Yan Zhu, Liudmila Ulanova, Nurjahan Begum, Yifei Ding, Hoang Anh Dau, Diego Furtado Silva, Abdullah Mueen, and Eamonn Keogh. Matrix profile i: all pairs similarity joins for time series: a unifying view that includes motifs, discords and shapelets. In *Int. Conf. Data Mining*, pages 1317–1322. Ieee, 2016. [1](#), [2](#), [3](#)
- [40] Jure Zbontar and Yann LeCun. Computing the stereo matching cost with a convolutional neural network. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1592–1599, 2015. [1](#), [2](#)
- [41] Jure Zbontar, Yann LeCun, et al. Stereo matching by training a convolutional neural network to compare image patches. *J. Mach. Learn. Res.*, 17(1):2287–2318, 2016. [2](#)
- [42] Feihu Zhang, Victor Prisacariu, Ruigang Yang, and Philip HS Torr. Ga-net: Guided aggregation net for end-to-end stereo matching. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 185–194, 2019. [7](#)
- [43] Feihu Zhang, Xiaojuan Qi, Ruigang Yang, Victor Prisacariu, Benjamin Wah, and Philip Torr. Domain-invariant stereo matching networks. In *Eur. Conf. Comput. Vis.*, pages 420–439. Springer, 2020. [7](#)
- [44] Jingyang Zhang, Shiwei Li, Zixin Luo, Tian Fang, and Yao Yao. Vis-mvsnet: Visibility-aware multi-view stereo network. *Int. J. Comput. Vis.*, 131(1):199–214, 2023. [7](#)
- [45] Youmin Zhang, Yimin Chen, Xiao Bai, Suihanjin Yu, Kun Yu, Zhiwei Li, and Kuiyuan Yang. Adaptive unimodal cost volume filtering for deep stereo matching. In *Assoc. Advancement Artif. Intell.*, pages 12926–12934, 2020. [6](#)
- [46] Haoliang Zhao, Huizhou Zhou, Yongjun Zhang, Yong Zhao, Yitong Yang, and Ting Ouyang. Eai-stereo: Error aware iterative network for stereo matching. In *Asian Conf. Comput. Vision*, pages 315–332, 2022. [4](#)
- [47] Haoliang Zhao, Huizhou Zhou, Yongjun Zhang, Jie Chen, Yitong Yang, and Yong Zhao. High-frequency stereo matching network. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1327–1336, 2023. [2](#), [3](#), [4](#), [6](#), [7](#)
- [48] Huizhou Zhou, Haoliang Zhao, Qi Wang, Gefei Hao, and Liang Lei. Miper-mvs: Multi-scale iterative probability estimation with refinement for efficient multi-view stereo. *Neural Netw.*, 162:502–515, 2023. [7](#)