

Parallax Attention for Unsupervised Stereo Correspondence Learning

Longguang Wang, Yulan Guo, Yingqian Wang, Zhengfa Liang, Zaiping Lin, Jungang Yang, Wei An

Abstract—Stereo image pairs encode 3D scene cues into stereo correspondences between the left and right images. To exploit 3D cues within stereo images, recent CNN based methods commonly use cost volume techniques to capture stereo correspondence over large disparities. However, since disparities can vary significantly for stereo cameras with different baselines, focal lengths and resolutions, the fixed maximum disparity used in cost volume techniques hinders them to handle different stereo image pairs with large disparity variations. In this paper, we propose a generic parallax-attention mechanism (PAM) to capture stereo correspondence regardless of disparity variations. Our PAM integrates epipolar constraints with attention mechanism to calculate feature similarities along the epipolar line to capture stereo correspondence. Based on our PAM, we propose a parallax-attention stereo matching network (PASMnet) and a parallax-attention stereo image super-resolution network (PASSRnet) for stereo matching and stereo image super-resolution tasks. Moreover, we introduce a new and large-scale dataset named Flickr1024 for stereo image super-resolution. Experimental results show that our PAM is generic and can effectively learn stereo correspondence under large disparity variations in an unsupervised manner. Comparative results show that our PASMnet and PASSRnet achieve the state-of-the-art performance.

Index Terms—Parallax Attention, Stereo Matching, Image Super-Resolution, Unsupervised Learning, Stereo Correspondence

1 INTRODUCTION

WITH the popularity of dual cameras in mobile phones, autonomous vehicles and robots, stereo vision has attracted increasingly attention in both academia and industry [1], [2]. Traditional studies [3], [4], [5] mainly focus on finding correspondences between stereo images to provide depth information, namely stereo matching. Recent works further use 3D cues within stereo images for various tasks including stereo image restoration [6], [7], stereo magnification [8], stereo video retargeting [9] and stereo neural style transfer [10]. In real-world applications such as mobile phones, the baselines and resolutions of stereo cameras in various devices are different. Moreover, the focal length of stereo cameras may change as the cameras adjust to a new scene. Since disparities between stereo images can vary significantly for stereo cameras with different baselines, focal lengths and resolutions, it is highly challenging to incorporate stereo correspondence for various real-world applications.

Motivated by the powerful feature representation capability of convolutional neural network (CNN), deep feature representations are widely used to calculate similarities for stereo correspondence. However, the local receptive field hinders plain CNNs to capture correspondence over large disparities. To overcome this limitation, cost volumes are

widely applied in CNN-based methods [4], [5], [11], [12]. Several methods [4], [5], [13] concatenate unary features from left and right images to generate a 4D cost volume (*i.e.*, height \times width \times disparity \times channel). Then, 3D CNNs are used to learn matching costs within the cost volume. However, learning matching costs from 4D cost volumes suffers from a high computational and memory burden. To achieve higher efficiency, 1D correlation is used to reduce feature dimension [11], [12], [14], [15], resulting in 3D cost volumes (*i.e.*, height \times width \times disparity). Although 4D/3D cost volumes enable CNNs to capture correspondence over large disparities, the fixed maximum disparity hinders them to handle different stereo image pairs with large disparity variations.

In this paper, we propose an unsupervised parallax-attention mechanism (PAM) to learn stereo correspondence under large disparity variations. Our PAM integrates epipolar constraints with attention mechanism to calculate feature similarities along the epipolar line. Specifically, for each pixel in the left image, its feature similarities with all possible disparities in the right image are computed to generate an attention map. By regularizing the attention map, our PAM can focus on the most similar feature to provide accurate stereo correspondence.

Compared to cost volume techniques, our PAM has three remarkable properties: **First**, the PAM gets access to different disparities through matrix multiplication instead of shift operation. Therefore, our PAM does not need to manually set a fixed maximum disparity and can handle large disparity variations. **Second**, cost volume based methods commonly regress disparities upon matching costs and then calculate losses upon these disparities. However, this may lead to unreasonable cost distribution due to the ambiguity of disparity regression (as shown in Fig. 1). In contrast, performing direct regularization on parallax-attention maps

- Longguang Wang, Yulan Guo, Yingqian Wang, Zaiping Lin, Jungang Yang, and Wei An are with the College of Electronic Science and Technology, Nation University of Defense Technology (NUDT), P. R. China. Yulan Guo is also with the School of Electronics and Communication Engineering, Sun Yat-sen University, Guangzhou, 510275, China. E-mail:{wanglongguang15, yulan.guo, wangyingqian16, linzaiping, yangjungang, anwei}@nudt.edu.cn. Zhengfa Liang is with the National Key Laboratory of Science and Technology on Blind Signal Processing, P. R. China. Email: liangzhengfa10@nudt.edu.cn.
- L. Wang and Y. Guo have equal contribution to this work and are co-first authors. Corresponding author: Yulan Guo.




Fig. 1: An illustration of the ambiguity for disparity regression in existing cost volume based methods. These three distributions have equal regression results d_0 but are quite different. A unimodal distribution peaked at the true disparity with high sharpness (*i.e.*, low uncertainty) is more reasonable, as shown in (a).

(*e.g.*, through smoothness loss and cycle loss, as described in Sec. 4.2) enables our PAM to achieve improved performance. **Third**, the PAM is compact and can be applied to various tasks such as stereo matching and stereo image super-resolution (SR). By using PAM, 3D cues from left and right images can be aggregated without explicit disparity calculation.

This paper is an extension of our previous conference version [6]. The contributions of this work can be summarized as follows:

- A generic parallax-attention mechanism is proposed to learn stereo correspondence in image pairs with large disparity variations in an unsupervised manner.
- The PAM is successfully applied to two specific tasks: stereo matching and stereo image SR. Our PAM based networks achieve the state-of-the-art performance in both stereo matching and stereo image SR.
- A new dataset, namely Flickr1024, is proposed for the training of stereo image SR networks. This dataset consists of 1024 high-quality stereo image pairs and covers various scenes.

The rest of this paper is organized as follows. In Section 2, we briefly review the related work. In Section 3, we present our PAM in details. In Sections 4 and 5, we apply the PAM to stereo matching and stereo image SR tasks, respectively. Finally, we conclude this paper in Section 6.

2 RELATED WORK

Stereo vision has been investigated for various tasks in recent decades [3], [5], [6], [7], [8], [10]. Here, we focus on two specific tasks, including stereo matching and stereo image SR. Moreover, we also review attention mechanisms that are highly related to our work.

2.1 Stereo Matching

2.1.1 Supervised Stereo Matching

Traditional stereo matching methods commonly follow a popular four-step pipeline, including matching cost calculation, cost aggregation, disparity calculation and disparity refinement [14].

Motivated by the success of CNN in many vision tasks such as object recognition [16], [17] and detection [18],

[19], some early learning-based methods [20], [21], [22] use CNNs to replace one or multiple steps in the stereo matching pipeline. Zbontar and Lecun [20] used a CNN to compute matching cost between two image patches. The resulting matching costs are then processed by several traditional techniques (including cross-based cost aggregation and semi-global matching) to predict the disparity map. Luo *et al.* [3] proposed a dot product operation to compute correlation between unary feature representations of left and right images. Their network has lower computational complexity and achieves higher efficiency than [20]. To improve feature representation capability, pyramid pooling [23] and highway networks [22] are used to enlarge the receptive field and deepen the network, respectively. These two methods produce more accurate disparities than [20].

To achieve better performance, several methods [4], [11], [14] are proposed to integrate all steps into an end-to-end architecture for joint optimization. Mayer [11] proposed an end-to-end CNN called DispNet to regress disparities from stereo images. Specifically, 1D correlation operation is used to generate 3D cost volumes for matching cost learning. Following DispNet, several works [14], [24] have been proposed to employ correlation operation and 3D cost volumes for disparity estimation. Liang *et al.* [14] used feature constancy to bridge the gap between disparity calculation and disparity refinement. Therefore, their network can seamlessly integrate the disparity calculation and disparity refinement steps into a compact network to improve efficiency.

Recent methods [4], [5], [13] commonly use concatenation to generate 4D cost volumes and then use 3D CNNs for matching cost aggregation. GC-Net [4] is the first work to use naive concatenation instead of correlation to build 4D cost volumes. Specifically, 3D convolutions are used to aggregate context information within the 4D cost volume for matching cost aggregation. Following GC-Net, Chang *et al.* [5] proposed a pyramid stereo matching network (PSMNet) with a spatial pyramid pooling module to enlarge receptive fields for more representative features. In PSMNet, stacked 3D hourglass networks are used to aggregate matching costs at multiple scales. Inspired by the semi-global matching (SGM) method [25], Zhang *et al.* [26] introduced semi-global and local guided aggregation layers to aggregate matching costs in different directions and local regions.

2.1.2 Unsupervised Stereo Matching

Due to the difficulty of collecting a large dataset with densely labeled groundtruth depth, several methods [24], [27], [28] have been developed to learn stereo matching in an unsupervised manner.

Zhou *et al.* [27] adopted left-right consistency check to produce a confidence map to guide the training of a network. Yang *et al.* [24] proposed a SegStereo network to use semantic cues for stereo matching. Specifically, they introduced a segmentation sub-network to provide semantic features and proposed a semantic loss regularization to improve the robustness of disparity estimation. Li *et al.* [28] proposed an occlusion aware stereo matching network (OASM) to exploit occlusion cues for stereo matching. Specifically, they introduced an occlusion inference module to provide occlusion cues and proposed a hybrid loss to




Fig. 2: Comparison between self-attention mechanism and parallax-attention mechanism.

use the interaction between disparity and occlusion as the supervision for network training.

Existing supervised and unsupervised stereo matching methods commonly used 3D/4D cost volumes to capture stereo correspondence over large disparities. However, the fixed maximum disparity hinders them to handle different stereo images with large disparity variations. Moreover, 4D cost volumes require high computational and memory costs.

2.2 Stereo Image Super-Resolution

Stereo image SR aims at recovering a high-resolution (HR) left image from a pair of low-resolution (LR) stereo images. Given additional information from a second viewpoint, stereo correspondence can be used to improve the SR performance. A recent work [29] further extend the stereo image SR task to a stereoscopic image SR task that reconstructs a pair of HR left/right images while preserving stereo consistency. Note that, we only focus on the stereo image SR task in this paper. Due to the unique characteristics of stereo images (such as epipolar constraints and non-local correspondence), video SR and light-field image SR methods are unsuitable for stereo image SR. Video SR methods [30], [31], [32] commonly focus on the exploitation of local correspondence since motion between adjacent frames are mainly limited within a local region. Therefore, these video SR methods are unsuitable for stereo image SR due to the non-local correspondence over large disparities in stereo images. Light-field image SR methods [33], [34], [35] are specifically proposed for light-field images with short baselines. Therefore, they cannot be directly used for stereo image SR since stereo images usually have much longer baselines than light-field images.

To enhance the resolution of stereo images, Bhavsar *et al.* [36] argued that image SR and HR depth estimation are intertwined under stereo setting. They proposed a joint framework to iteratively estimate SR image and HR disparity. Recently, Jeon *et al.* [37] proposed to employ a parallax prior in CNN for stereo image SR. Given a stereo image pair, the right image is shifted with different intervals and concatenated with the left image to generate a stereo tensor. The stereo tensor is then fed to a CNN called StereoSR to generate SR results. However, StereoSR cannot handle different stereo images with large disparity variations since the number of shifted right images is fixed.

2.3 Attention Mechanism

Attention mechanism was first introduced by Bahdanau *et al.* [38] and has been widely used in natural language processing tasks to capture long-range (long-term) dependencies [39], [40]. Recently, attention mechanism has also been

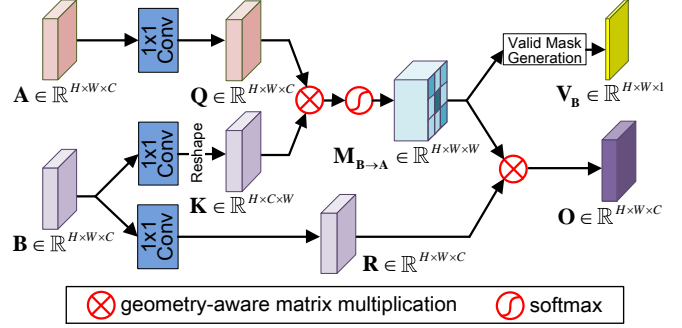


Fig. 3: An illustration of our PAM.

applied in many computer vision tasks including semantic segmentation [41], [42], image captioning [43], [44] and image generation [45]. Due to the limited receptive fields of CNNs, long-range dependencies cannot be captured. To overcome this limitation, self-attention mechanism is introduced to calculate the correlation of any two positions in an image. Wang *et al.* [46] proposed a non-local network to aggregate information from non-local regions, which can be viewed as a general format of the self-attention mechanism. Zhang *et al.* [45] introduced self-attention mechanism for image generation. The self-attention mechanism enables their network to use features in both local regions and distant regions to generate consistent scenarios. Fu *et al.* [42] proposed to use self-attention mechanism to aggregate long-range information for semantic segmentation.

Our work is motivated by the self-attention mechanism. We incorporate epipolar constraints with attention mechanism and develop a parallax-attention mechanism to capture correspondences in stereo images (as shown in Fig. 2). For each pixel in the left image, our parallax-attention mechanism attends to all disparities along the epipolar line and learns to focus on the most similar one.

3 PARALLAX-ATTENTION MECHANISM (PAM)

In this section, we present our PAM in details. We first introduce the formulation of our PAM, and then describe left-right consistency, cycle consistency and valid mask based on our PAM.

3.1 Formulation

3.1.1 Overview

In self-attention mechanism, a feature map of size $\mathbb{R}^{H \times W \times C}$ is first reshaped to the size of $\mathbb{R}^{HW \times C}$. Then, matrix multiplication ($\mathbb{R}^{HW \times C} \times \mathbb{R}^{C \times HW}$) is used to calculate the correlation of any two positions in an image. For stereo images, the corresponding pixel for a pixel in the left image only lies along its epipolar line in the right image. Taking this epipolar constraint into consideration, our PAM uses specifically designed reshape operation and geometry-aware matrix multiplication to calculate the correlation between a pixel in the left image and all positions along the epipolar line in the right image.

As illustrated in Fig. 3, given two feature maps $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{H \times W \times C}$ from a stereo image pair, they are first fed to 1×1 convolutions for feature adaption. Specifically, \mathbf{A} is




Fig. 4: An illustration of our PAM using a toy example. The parallax-attention maps (30×30) correspond to the regions (1×30) marked by a yellow stroke. The first row represents left/right stereo images and the second row stands for parallax-attention maps $\mathbf{M}_{right \rightarrow left}$ and $\mathbf{M}_{left \rightarrow right}$.

fed to a 1×1 convolution to produce a query feature map $\mathbf{Q} \in \mathbb{R}^{H \times W \times C}$. Meanwhile, \mathbf{B} is fed to another 1×1 convolution to produce a key feature map $\mathbf{K} \in \mathbb{R}^{H \times W \times C}$, which is then reshaped to $\mathbb{R}^{H \times C \times W}$. Then, matrix multiplication is performed between \mathbf{Q} and \mathbf{K} and a softmax layer is applied, resulting in a parallax-attention map $\mathbf{M}_{B \rightarrow A} \in \mathbb{R}^{H \times W \times W}$. Through matrix multiplication, our PAM can efficiently encode feature correlation between any two positions along the epipolar line into the parallax-attention map. The parallax-attention map has several remarkable properties, as described in Sec. 3.2 and 3.3. Next, \mathbf{B} is fed to a 1×1 convolution to generate response feature map $\mathbf{R} \in \mathbb{R}^{H \times W \times C}$, which is further multiplied by $\mathbf{M}_{B \rightarrow A}$ to produce output features $\mathbf{O} \in \mathbb{R}^{H \times W \times C}$. Meanwhile, a valid mask \mathbf{V}_B is also generated (Section 3.3).

It should be noted that all disparities are considered in the PAM. That is, our PAM does not need to manually set a fixed maximum disparity and can handle large disparity variations. Since PAM can learn to focus on the features at accurate disparities using feature similarities, correspondence can then be captured.

3.1.2 Toy Example

We further illustrate our PAM using a toy example, as shown in Fig. 4. Given a stereo image pair \mathbf{I}_{left} and \mathbf{I}_{right} of size $\mathbb{R}^{30 \times 30}$, parallax-attention maps $\mathbf{M}_{left \rightarrow right}$ and $\mathbf{M}_{right \rightarrow left}$ of size $\mathbb{R}^{30 \times 30 \times 30}$ can be generated by our PAM. Note that, each slice of the parallax-attention maps (e.g., $\mathbf{M}_{right \rightarrow left}(i, :, :)$) delivers the dependency between corresponding rows (i.e., $\mathbf{I}_{left}(i, :)$ and $\mathbf{I}_{right}(i, :)$). If the disparity between a pair of stereo images is zero, all parallax-attention maps are identity matrices, as shown in Fig. 4 (a). That is because, the j^{th} pixel in $\mathbf{I}_{left}(i, :)$ corresponds to the j^{th} pixel in $\mathbf{I}_{right}(i, :)$. Therefore, position (j, j) in $\mathbf{M}_{right \rightarrow left}$ is focused on. For regions with non-zero disparities, e.g., the red region in Fig. 4 (b) with a disparity of 5, the j^{th} pixel in $\mathbf{I}_{left}(i, :)$ corresponds to the $(j - 5)^{th}$ pixel in $\mathbf{I}_{right}(i, :)$. Therefore, position $(j, j - 5)$ in $\mathbf{M}_{right \rightarrow left}$ is focused on. In summary,

stereo correspondence can be depicted by the positions of focused pixels in the parallax-attention maps.

Furthermore, occlusion can be encoded by the parallax-attention maps. Specifically, it can be observed from $\mathbf{M}_{right \rightarrow left}$ in Fig. 4(b) that several vertical regions are “discarded” without any position being focused on. That is because, these regions in $\mathbf{I}_{right}(i, :)$ are occluded in $\mathbf{I}_{left}(i, :)$, thus no correspondence should be focused on. Similar “discarded” horizontal regions are also caused by occlusion.

It should be noted that, only integer disparities are considered in our toy example, which is not the real case. In practice, our PAM can focus on adjacent pixels to handle sub-pixel disparities. Due to the softmax layer used in PAM, several pixels in “discarded” horizontal regions may be incorrectly focused on. However, these occluded regions can be excluded using valid masks, as demonstrated in Sec. 3.3.

3.2 Left-Right Consistency and Cycle Consistency

To capture reliable and consistent correspondence, we introduce left-right consistency and cycle consistency to regularize our PAM.

Given feature representations extracted from a pair of stereo images \mathbf{I}_{left} and \mathbf{I}_{right} , two parallax-attention maps ($\mathbf{M}_{left \rightarrow right}$ and $\mathbf{M}_{right \rightarrow left}$) are generated by our PAM. Ideally, the following left-right consistency can be obtained if our PAM captures accurate correspondence:

$$\begin{cases} \mathbf{I}_{left} = \mathbf{M}_{right \rightarrow left} \otimes \mathbf{I}_{right} \\ \mathbf{I}_{right} = \mathbf{M}_{left \rightarrow right} \otimes \mathbf{I}_{left} \end{cases}, \quad (1)$$

where \otimes denotes geometry-aware matrix multiplication. Based on Eq. (1), we can further derive a cycle consistency:

$$\begin{cases} \mathbf{I}_{left} = \mathbf{M}_{left \rightarrow right \rightarrow left} \otimes \mathbf{I}_{left} \\ \mathbf{I}_{right} = \mathbf{M}_{right \rightarrow left \rightarrow right} \otimes \mathbf{I}_{right} \end{cases}, \quad (2)$$

where the cycle-attention maps $\mathbf{M}_{left \rightarrow right \rightarrow left}$ and $\mathbf{M}_{right \rightarrow left \rightarrow right}$ can be calculated as:

$$\begin{cases} \mathbf{M}_{left \rightarrow right \rightarrow left} = \mathbf{M}_{right \rightarrow left} \otimes \mathbf{M}_{left \rightarrow right} \\ \mathbf{M}_{right \rightarrow left \rightarrow right} = \mathbf{M}_{left \rightarrow right} \otimes \mathbf{M}_{right \rightarrow left} \end{cases}. \quad (3)$$

We further illustrate geometry-aware matrix multiplication \otimes^1 in Fig. 5. Take Eq. (1) for an example, the product of corresponding slices in $\mathbf{M}_{right \rightarrow left}$ and \mathbf{I}_{right} (e.g., $\mathbf{M}_{right \rightarrow left}(i, :, :)$ and $\mathbf{I}_{right}(i, :, :)$) determines the i^{th} slice of \mathbf{I}_{left} , i.e., $\mathbf{I}_{left}(i, :, :)$. All these slices are concatenated to obtain $\mathbf{I}_{left} \in \mathbb{R}^{H \times W \times C}$.

3.3 Valid Mask

Since left-right consistency and cycle consistency do not hold for occluded regions, we perform occlusion detection based on the parallax-attention map to generate valid masks and only enforce consistency regularization on those valid regions.

As illustrated in the Sec. 3.1.2, occlusion is encoded as the “discarded” regions in the parallax-attention maps. In practice, vertical “discarded” regions that correspond to occluded regions are usually assigned with small weights in

1. \otimes can be implemented using `tf.matmul()` or `torch.matmul()`.

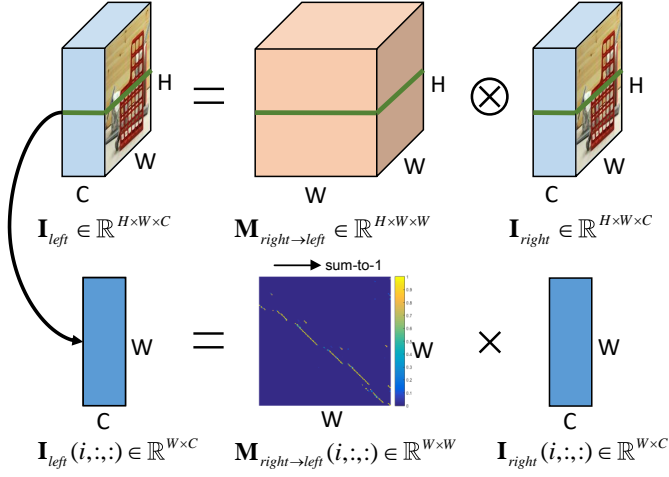


Fig. 5: An illustration of geometry-aware matrix multiplication \otimes .

the parallax-attention maps (e.g., $\mathbf{M}_{left \rightarrow right}$). That is because, occluded pixels in the left image cannot find their correspondences in the right image, their feature similarities with all disparities are low. Therefore, a valid mask $\mathbf{V}_{left} \in \mathbb{R}^{H \times W}$ can be obtained by:

$$\mathbf{V}_{left}(i, k) = \begin{cases} 1, & \text{if } \sum_{j \in [1, W]} \mathbf{M}_{left \rightarrow right}(i, j, k) > \tau \\ 0, & \text{otherwise} \end{cases}, \quad (4)$$

where τ is a threshold (empirically set to 0.1 in this paper). An examples of valid mask is shown in Fig. 6. In practice, we use several morphological operations to handle isolated pixels and holes in valid masks.

In summary, our PAM provides a flexible and effective approach for unsupervised stereo correspondence learning and can be applied to various stereo vision tasks. In this paper, we demonstrate the effectiveness of PAM on two typical stereo tasks: stereo matching and stereo image SR.

4 PAM FOR UNSUPERVISED STEREO MATCHING

Stereo matching aims at finding correspondence pixels between a stereo image pair. Since disparities can vary significantly in real-world applications, the fixed maximum disparity in cost volume techniques hinders them to handle large disparity variations. In contrast, our PAM can get rid of setting a fixed maximum disparity by efficiently calculating feature correlation between any two positions along the epipolar line. Therefore, the PAM can be used for unsupervised stereo matching to handle large disparity variations.

4.1 Network Architecture

4.1.1 Overview

Based on our PAM, we proposed a parallax-attention stereo matching network (PASMnet). The architecture of our PASMnet is shown in Fig. 7. Given a stereo image pair of size $H \times W$, they are first fed to an hourglass network for feature extraction. Then, the features extracted from left




Fig. 6: Visualization of valid masks. A stereo image pair and the occluded regions (i.e., yellow regions) are illustrated.

and right images are fed to a cascaded parallax-attention module to regress matching costs in a coarse-to-fine manner. Next, an initial disparity can be obtained from the matching cost using an output module. Finally, the initial disparity is further refined using an hourglass network to produce the output disparity.

4.1.2 Cascaded Parallax-Attention Module

After feature extraction, features from left and right images are fed to the cascaded parallax-attention module to regress matching costs in a coarse-to-fine manner. Specifically, our cascaded parallax-attention module consists of 3 stages (with 4 parallax-attention blocks in each stage), as shown in Fig. 7(b). First, features from \mathbf{F}^1 (i.e., $\mathbf{F}_{left}^1, \mathbf{F}_{right}^1 \in \mathbb{R}^{\frac{H}{16} \times \frac{W}{16} \times C}$) and initial matching costs $\mathbf{C}_{right \rightarrow left}^0, \mathbf{C}_{left \rightarrow right}^0 \in \mathbb{R}^{\frac{H}{16} \times \frac{W}{16} \times \frac{W}{16}}$ with all elements being set to 0 are passed to the first parallax-attention block. \mathbf{F}_{left}^1 and \mathbf{F}_{right}^1 are fed to two 3×3 convolutions to obtain \mathbf{F}_{left} and \mathbf{F}_{right} , respectively. Note that, the convolutions for left and right images share the same parameters. Then, query features \mathbf{Q} and key features \mathbf{K} are obtained from \mathbf{F}_{left} and \mathbf{F}_{right} through 1×1 convolutions. Next, \mathbf{K} is reshaped and multiplied with \mathbf{Q} , resulting in a matching cost $\mathbf{C}_{right \rightarrow left}^1 \in \mathbb{R}^{\frac{H}{16} \times \frac{W}{16} \times \frac{W}{16}}$. Once $\mathbf{C}_{right \rightarrow left}^1$ is ready, \mathbf{F}_{left} and \mathbf{F}_{right} are exchanged to obtain $\mathbf{C}_{left \rightarrow right}^1$. After that, $\mathbf{C}_{right \rightarrow left}^0$ and $\mathbf{C}_{left \rightarrow right}^0$ are added to $\mathbf{C}_{right \rightarrow left}^1$ and $\mathbf{C}_{left \rightarrow right}^1$, respectively. Meanwhile, \mathbf{F}_{left}^1 and \mathbf{F}_{right}^1 are added to \mathbf{F}_{left} and \mathbf{F}_{right} , respectively.

As shown in Fig. 7(b), the features and matching cost produced by the preceding parallax-attention block are fed to the succeeding block. The features produced by stage 1 are bilinearly upsampled and concatenated with \mathbf{F}_{left}^2 and \mathbf{F}_{right}^2 (from \mathbf{F}^2 in Fig. 7(a)). Then, the concatenated features and the upsampled matching cost ($\mathbf{C}_{right \rightarrow left}^1 \uparrow, \mathbf{C}_{left \rightarrow right}^1 \uparrow$)

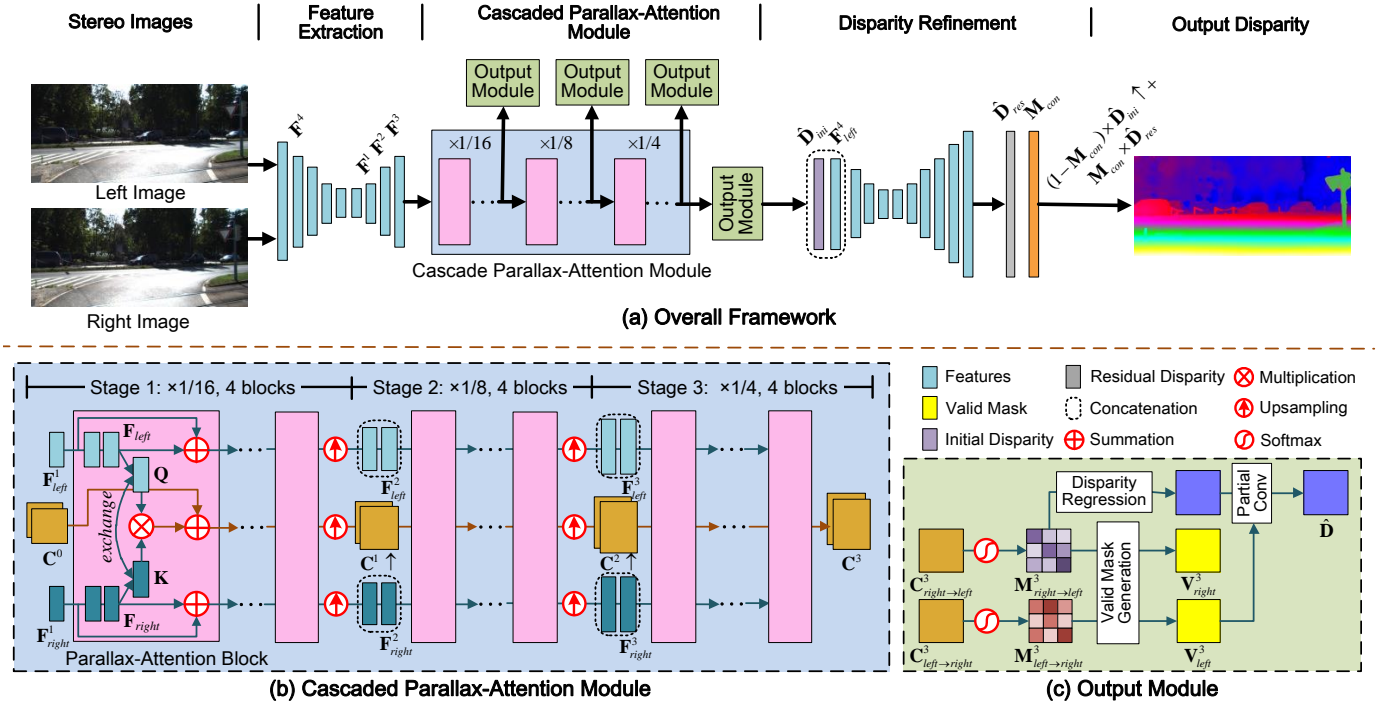


Fig. 7: An overview of our PASMnet.

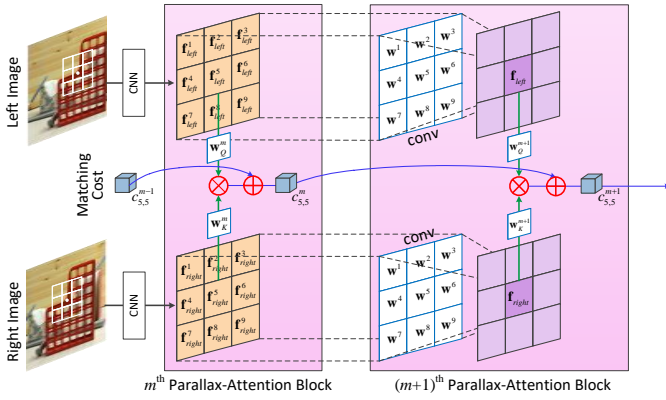


Fig. 8: An illustration of matching cost aggregation in our cascaded parallax-attention module.

are passed to stages 2 and 3 for further refinement, resulting in the final matching cost $\mathbf{C}_{right \rightarrow left}^3, \mathbf{C}_{left \rightarrow right}^3 \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times \frac{W}{4}}$.

Different from the cost volume based methods [4], [5] that explicitly aggregate matching costs using 3D convolutions, our network performs implicit matching cost aggregation by cascading several parallax-attention blocks, as shown in Fig. 8. Here, only one 3×3 convolutional layer (without batch normalization and ReLU layers) in the parallax-attention blocks is considered for simplicity. In the m^{th} parallax-attention block, the matching cost $c_{5,5}^m$ between the locations of \mathbf{f}_{left}^5 and \mathbf{f}_{right}^5 is computed as:

$$\begin{aligned} c_{5,5}^m &= c_{5,5}^{m-1} + (\mathbf{w}_Q^m \mathbf{f}_{left}^5)^T \mathbf{w}_K^m \mathbf{f}_{right}^5 \\ &= c_{5,5}^{m-1} + (\mathbf{f}_{left}^5)^T (\mathbf{w}_Q^m)^T \mathbf{w}_K^m \mathbf{f}_{right}^5, \\ &= c_{5,5}^{m-1} + \Psi^m(\mathbf{f}_{left}^5, \mathbf{f}_{right}^5) \end{aligned} \quad (5)$$

where $\mathbf{w}_Q^m, \mathbf{w}_K^m \in \mathbb{R}^{C \times C}$ are the kernels of two 1×1 convolutional layers, $\mathbf{f}_{left}^5, \mathbf{f}_{right}^5 \in \mathbb{R}^{C \times 1}$ are the corresponding features, and $\Psi^m(\cdot)$ represents the matching cost calculation between the pair of input features. In the $(m+1)^{\text{th}}$ block, features from the m^{th} block within a local neighborhood are aggregated to produce \mathbf{f}_{left} and \mathbf{f}_{right} :

$$\begin{cases} \mathbf{f}_{left} = \sum_{i=1}^9 \mathbf{w}^i \mathbf{f}_{left}^i \\ \mathbf{f}_{right} = \sum_{j=1}^9 \mathbf{w}^j \mathbf{f}_{right}^j \end{cases}, \quad (6)$$

where $\mathbf{w}^i, \mathbf{w}^j \in \mathbb{R}^{C \times C}$ are kernels at different sub-locations for two 3×3 convolutional layers with shared parameters, and $\mathbf{f}_{left}, \mathbf{f}_{right} \in \mathbb{R}^{C \times 1}$ represent the output features. Consequently, the matching cost $c_{5,5}^{m+1}$ is computed as:

$$\begin{aligned} c_{5,5}^{m+1} &= c_{5,5}^m + (\mathbf{w}_Q^{m+1} \mathbf{f}_{left})^T \mathbf{w}_K^{m+1} \mathbf{f}_{right} \\ &= c_{5,5}^m + \sum_i \sum_j (\mathbf{w}_Q^{m+1} \mathbf{w}^i \mathbf{f}_{left}^i)^T \mathbf{w}_K^{m+1} \mathbf{w}^j \mathbf{f}_{right}^j \\ &= c_{5,5}^m + \sum_i \sum_j (\mathbf{f}_{left}^i)^T (\mathbf{w}^i)^T (\mathbf{w}_Q^{m+1})^T \mathbf{w}_K^{m+1} \mathbf{w}^j \mathbf{f}_{right}^j, \\ &= c_{5,5}^m + \sum_i \sum_j \Psi_{i,j}^{m+1}(\mathbf{f}_{left}^i, \mathbf{f}_{right}^j) \end{aligned} \quad (7)$$

where $\mathbf{w}_Q^{m+1}, \mathbf{w}_K^{m+1} \in \mathbb{R}^{C \times C}$ are the kernels of two 1×1 convolutional layers, and $\Psi_{i,j}^{m+1}(\cdot)$ represents the matching cost calculation between the pair of input features. Note that, \mathbf{f}_{left}^i and \mathbf{f}_{right}^j are from the m^{th} parallax-attention block. That is, matching costs in a local neighborhood are implicitly aggregated as parallax-attention blocks are cascaded. Moreover, residual connection further enables the aggregation of matching costs from different network depths.

4.1.3 Output Module

As shown in Fig. 7(a), the matching cost of the last block at each stage in the cascaded parallax-attention module

is fed to an output module. Within the output module of stage 3, $\mathbf{C}_{right \rightarrow left}^3$ and $\mathbf{C}_{left \rightarrow right}^3$ are first fed to a softmax layer to produce parallax-attention maps $\mathbf{M}_{right \rightarrow left}^3$ and $\mathbf{M}_{left \rightarrow right}^3 \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times \frac{W}{4}}$, as shown in Fig. 7(c). Next, $\mathbf{M}_{right \rightarrow left}^3$ and $\mathbf{M}_{left \rightarrow right}^3$ are used to generate valid masks \mathbf{V}_{right}^3 and \mathbf{V}_{left}^3 . Once $\mathbf{M}_{right \rightarrow left}^3$ is obtained, the disparity can be regressed as:

$$\hat{\mathbf{D}}(i, j) = \sum_{k=0}^{W/4-1} (j-k) \times \mathbf{M}_{right \rightarrow left}^3(:, :, k). \quad (8)$$

Note that, the estimated disparity is a sum of all disparity candidates weighted by the parallax-attention map. Consequently, our PASMnet does not need to manually set a fixed maximum disparity and can handle large disparity variations. Since pixels in occluded regions cannot find their correspondences, disparity in occluded regions cannot be well handled. Therefore, we exclude these invalid disparities and fill occluded regions using partial convolution [47].

4.1.4 Disparity Refinement

After the cascaded parallax-attention module, a disparity refinement module is introduced using features from the left image as guidance to provide structural information like edges. As shown in Fig. 7(a), the initial disparity $\hat{\mathbf{D}}_{ini}$ is concatenated with \mathbf{F}_{left}^4 and fed to an hourglass network to produce a residual disparity map $\hat{\mathbf{D}}_{res}$ and a confidence map \mathbf{M}_{con} . Finally, the refined disparity is calculated as:

$$\hat{\mathbf{D}}_{refined} = (1 - \mathbf{M}_{con}) \times \hat{\mathbf{D}}_{ini} \uparrow + \mathbf{M}_{con} \times \hat{\mathbf{D}}_{res}, \quad (9)$$

where \uparrow is a bilinear upsampling operator.

4.2 Losses

4.2.1 Photometric Loss

Following [28], [48], [49], we introduce a photometric loss consisting of a mean absolute error (MAE) loss term and a structural similarity index (SSIM) loss term. Note that, since photometric consistency only holds in non-occluded regions, the photometric loss is formulated as:

$$\mathcal{L}_p = \frac{1}{N} \sum_{p \in \mathbf{V}_{left}^s} \alpha \frac{1 - \mathcal{S}(\mathbf{I}_{left}^s(p), \hat{\mathbf{I}}(p))}{2} + (1 - \alpha) \left\| \mathbf{I}_{left}^s(p) - \hat{\mathbf{I}}(p) \right\|_1, \quad (10)$$

where $\hat{\mathbf{I}} = \mathcal{W}(\mathbf{I}_{right}, \hat{\mathbf{D}}_{refined})$ and \mathcal{W} is a warping operator using the refined disparity. \mathcal{S} is an SSIM function, p represents a valid pixel covered by the valid mask, N is the number of valid pixels, and α is empirically set to 0.85 in this paper.

4.2.2 Smoothness Loss

Following [28], [48], we use an edge-aware smoothness loss to encourage local smoothness of the disparity, which is defined as:

$$\mathcal{L}_s = \frac{1}{N} \sum_p \left(\left\| \nabla_x \hat{\mathbf{D}}_{refined}(p) \right\|_1 e^{-\|\nabla_x \mathbf{I}_{left}^s(p)\|_1} + \left\| \nabla_y \hat{\mathbf{D}}_{refined}(p) \right\|_1 e^{-\|\nabla_y \mathbf{I}_{left}^s(p)\|_1} \right), \quad (11)$$

where ∇_x and ∇_y are gradients in the x and y axes, respectively.

4.2.3 PAM Loss

We introduce three additional losses to regularize our PAM at multiple scales to capture accurate and consistent stereo correspondence. The PAM loss term for scale s ($s = 1, 2, 3$) is defined as:

$$\mathcal{L}_{PAM}^s = \mathcal{L}_{PAM-p}^s + \lambda_{PAM-s} \mathcal{L}_{PAM-s}^s + \lambda_{PAM-c} \mathcal{L}_{PAM-c}^s. \quad (12)$$

• \mathcal{L}_{PAM-p}^s

Different from the aforementioned photometric loss, here, we introduce a photometric loss based on the parallax-attention maps as:

$$\mathcal{L}_{PAM-p}^s = \frac{1}{N_{left}^s} \sum_{p \in \mathbf{V}_{left}^s} \left\| \mathbf{I}_{left}^s(p) - (\mathbf{M}_{right \rightarrow left}^s \otimes \mathbf{I}_{right}^s)(p) \right\|_1 + \frac{1}{N_{right}^s} \sum_{p \in \mathbf{V}_{right}^s} \left\| \mathbf{I}_{right}^s(p) - (\mathbf{M}_{left \rightarrow right}^s \otimes \mathbf{I}_{left}^s)(p) \right\|_1. \quad (13)$$

where N_{left}^s and N_{right}^s are the numbers of valid pixels in \mathbf{V}_{left}^s and \mathbf{V}_{right}^s , respectively. \mathbf{I}_{left}^s and \mathbf{I}_{right}^s are bilinearly downsampled images at corresponding scale level.

• \mathcal{L}_{PAM-s}^s

Different from the aforementioned smoothness loss, here, another smoothness loss is introduced to directly regularize parallax-attention maps:

$$\mathcal{L}_{PAM-s}^s = \frac{1}{N^s} \sum_{\mathbf{M}} \sum_{i,j,k} \left(\left\| \mathbf{M}^s(i, j, k) - \mathbf{M}^s(i+1, j, k) \right\|_1 + \left\| \mathbf{M}^s(i, j, k) - \mathbf{M}^s(i, j+1, k+1) \right\|_1 \right), \quad (14)$$

where $\mathbf{M}^s \in \left\{ \mathbf{M}_{left \rightarrow right}^s, \mathbf{M}_{right \rightarrow left}^s \right\}$ and N^s is the number of pixels in $\mathbf{M}_{left \rightarrow right}^s$. The first and the second terms are used to achieve vertical and horizontal attention consistency, respectively. Take $\mathbf{M}_{right \rightarrow left}^s$ as an example, $\mathbf{M}_{right \rightarrow left}^s(i, j, k)$ measures the contribution of (i, k) in \mathbf{I}_{right}^s to (i, j) in \mathbf{I}_{left}^s using their feature similarity $\text{Sim}(\mathbf{F}_{left}^s(i, j), \mathbf{F}_{right}^s(i, k))$. Our smoothness loss enforces $\text{Sim}(\mathbf{F}_{left}^s(i+1, j), \mathbf{F}_{right}^s(i+1, k))$ and $\text{Sim}(\mathbf{F}_{left}^s(i, j+1), \mathbf{F}_{right}^s(i, k+1))$ to be close to $\text{Sim}(\mathbf{F}_{left}^s(i, j), \mathbf{F}_{right}^s(i, k))$. Consequently, smoothness in correspondence (disparity) space can be encouraged.

• \mathcal{L}_{PAM-c}^s

In addition to \mathcal{L}_{PAM-p}^s and \mathcal{L}_{PAM-s}^s , we further introduce a cycle loss to achieve cycle consistency. Since $\mathbf{M}_{left \rightarrow right}^s \rightarrow left$ and $\mathbf{M}_{right \rightarrow left}^s \rightarrow right$ in Eq. (2) can be considered as identity matrices, we design a cycle loss as:

$$\mathcal{L}_{PAM-c}^s = \frac{1}{N_{left}^s} \sum_{p \in \mathbf{V}_{left}^s} \left\| \mathbf{M}_{left \rightarrow right \rightarrow left}^s(p) - I^s(p) \right\|_1 + \frac{1}{N_{right}^s} \sum_{p \in \mathbf{V}_{right}^s} \left\| \mathbf{M}_{right \rightarrow left \rightarrow right}^s(p) - I^s(p) \right\|_1, \quad (15)$$

where I^s is a stack of identity matrices.

In summary, the overall loss is defined as:

$$\mathcal{L}_{unsup} = \mathcal{L}_p + \lambda_s \mathcal{L}_s + \lambda_{PAM} (0.2 \mathcal{L}_{PAM}^1 + 0.3 \mathcal{L}_{PAM}^2 + 0.5 \mathcal{L}_{PAM}^3). \quad (16)$$

Note that, groundtruth disparities are not used in the overall loss. That is, our network is trained in an unsupervised manner.

TABLE 1: Comparative results achieved on SceneFlow by our PASMnet trained with different settings.

Model	Hourglass	Coarse-to-Fine	# Blocks	EPE	> 1px	> 3px
PASMnet		✓	1	4.86	19.58	16.31
PASMnet	✓		1	4.78	19.50	16.21
PASMnet	✓	✓	1	4.68	19.32	16.09
PASMnet	✓	✓	2	4.63	19.19	16.05
PASMnet	✓	✓	3	4.58	19.05	15.97
PASMnet	✓	✓	4	4.54	18.99	15.91

4.3 Experimental Results

4.3.1 Datasets and Metrics

We used the SceneFlow [11] and KITTI 2015 [50] datasets for training and test.

SceneFlow: The SceneFlow dataset is a synthetic dataset consisting of 35454 training image pairs and 4370 test image pairs of size 540×960 .

KITTI 2015: The KITTI 2015 dataset is a real-world dataset with street views collected from a driving car. This dataset contains 200 training image pairs and 200 test image pairs of size 375×1242 .

For evaluation, we used end-point-error (EPE) and t-pixel error rate ($> tpx$) as metrics. Metrics for both non-occluded regions (Noc) and all pixels (All) are evaluated.

4.3.2 Implementation Details

We first trained our network on the SceneFlow dataset and performed ablation study on this dataset. During training phase, patches of size 256×512 were randomly cropped as inputs. $\lambda_{\text{PAM-s}}$, $\lambda_{\text{PAM-c}}$, λ_s and λ_{PAM} are set to 1, 1, 0.1, and 1, respectively. All models were optimized using the Adam method [51] with $\beta_1 = 0.9$, $\beta_2 = 0.999$ and a batch size of 14. The initial learning rate was set to 1×10^{-3} for 5 epochs and decreased to 1×10^{-4} for another 5 epochs. Note that, pixels with disparities over 192 are excluded for training to demonstrate the generalization of our network to large disparity variations, since our test set contains image pairs with disparities over 192. Following [28], pixels with disparities over 192 are also excluded for evaluation unless specified otherwise. Next, we fine-tuned our network on the KITTI 2015 dataset to obtain the final model for submission. $\lambda_{\text{PAM-s}}$, $\lambda_{\text{PAM-c}}$, λ_s , and λ_{PAM} are set to 5, 5, 0.5, and 1 for this dataset, respectively. The initial learning rate was set to 1×10^{-4} for 60 epochs and decreased to 1×10^{-5} for another 20 epochs. All experiments were conducted on a PC with two Nvidia RTX 2080Ti GPUs.

4.3.3 Ablation Study

Hourglass Feature Extraction. Hourglass network is used to aggregate features from multiple scales in our PASMnet. To demonstrate its effectiveness, we replaced the hourglass network with a pyramid network for feature extraction. It can be observed from Table 1 that the 1-pixel/3-pixel error rates are increased from 19.32/16.09 to 19.58/16.31 if hourglass network is replaced with a pyramid network. As compared to pyramid network, hourglass network can aggregate features from multiple scales, resulting in features rich of multi-scale geometric information. Therefore, better performance can be achieved.

TABLE 2: Ablation results for the disparity refinement module in our PASMnet on SceneFlow.

Model	EPE	> 1px	> 3px
PASMnet_wo_Refinement	4.56	19.15	16.11
PASMnet	4.54	18.99	15.91

Coarse-to-fine Manner. Our cascaded parallax-attention module regresses matching costs in a coarse-to-fine manner. To demonstrate its effectiveness, we introduced a variant that regresses matching costs on a single scale (*i.e.*, $\times \frac{1}{4}$). It can be observed from Table 1 that if matching costs are regressed on a single scale, the 1-pixel/3-pixel error rates are increased from 19.32/16.09 to 19.50/16.21. By using a coarse-to-fine manner, disparities can be progressively refined and better performance can be achieved.

Number of Cascaded Blocks. We tested the performance of our PASMnet with different number of parallax-attention blocks. From Table 1 we can see that, the performance of our network improves as the number of parallax-attention blocks increases. Specifically, our PASMnet with 4 blocks outperforms the variant with 1 block, with 1-pixel/3-pixel error rates being improved from 19.32/16.09 to 18.99/15.91. That is because, by cascading more blocks, our network can progressively refine the matching cost for better performance. We further fed the matching costs produced by the last two parallax-attention blocks in stage 3 (*i.e.*, blocks 3 and 4) to the output module for disparity regression. The disparities, matching cost distributions and parallax-attention distributions are visualized in Fig. 9. As more parallax-attention blocks are cascaded, matching cost is aggregated to produce higher peak values at the groundtruth disparity in the parallax-attention distributions. Therefore, more accurate and smoother disparities are produced.

Disparity Refinement. After the cascaded parallax-attention module, disparity refinement is performed by using features from the left image to provide structural information. To demonstrate the effectiveness of disparity refinement, we introduced a variant by removing this disparity refinement module and upsampling \hat{D}_{mi} to obtain the final disparity. It can be observed from Table 2 that the 1-pixel/3-pixel error rates are increased from 18.99/15.91 to 19.15/16.11 without disparity refinement. That is because, structural information helps to effectively refine the disparity. We further visualize the confidence map learned within the disparity refinement module in Fig. 10. We can see that occluded regions and edge regions benefit a lot from disparity refinement while other regions in the initial disparity \hat{D}_{mi} are "good enough".

PAM vs. Cost Volume. To demonstrate the efficiency of the proposed PAM, we compare our parallax-attention block with 3D convolution in terms of model size, FLOPs and memory consumption in Table 3. With the same number of channels, our parallax-attention block has fewer parameters than 3D convolution (20C vs. 27C). Assume that $D = 192$, $H = 540$, $W = 960$ (*i.e.*, the size of images in SceneFlow), when performed at the $\frac{1}{4}$ resolution level (*i.e.*, $D = 48$, $H = 135$, $W = 240$), our block achieves a $16.2 \times$ reduction in FLOPs and a comparable memory cost as compared to 3D convolution when $C = 12$. As C increases,

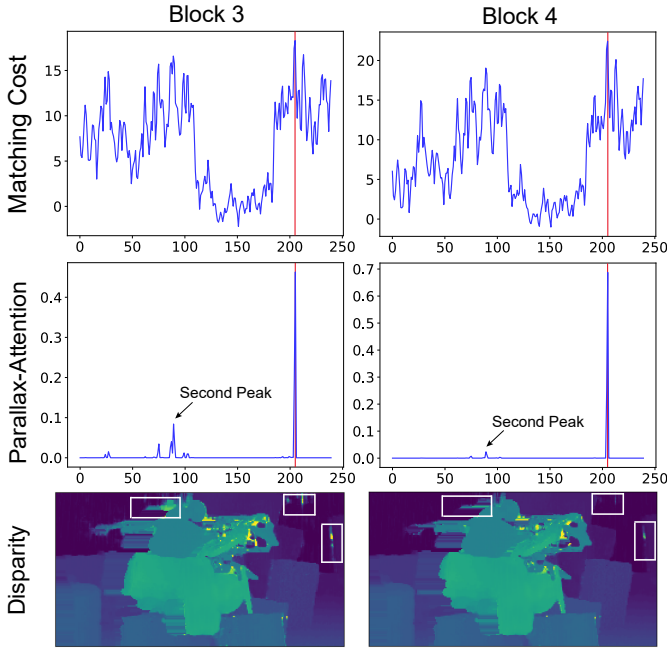


Fig. 9: Matching cost distributions, parallax-attention distributions and disparities produced by blocks 3 and 4 in stage 3 of our cascaded parallax-attention module. The red lines in the first and second rows represent the groundtruth disparity. As compared to block 3, block 4 produces more accurate and smoother disparities by cascading more parallax-attention blocks, especially in regions denoted with white boxes.

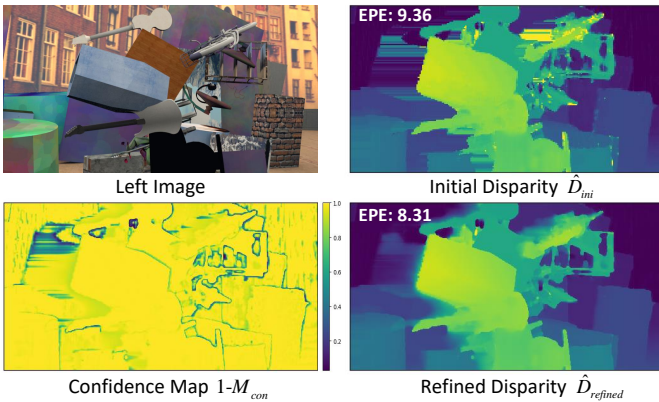


Fig. 10: An example of initial disparity \hat{D}_{ini} , confidence map $1 - M_{con}$ and refined disparity $\hat{D}_{refined}$ produced by our PASMnet.

better efficiency is achieved by our parallax-attention block, with $23.5 \times / 2.1 \times$ reductions in FLOPs/memory cost being achieved for $C = 32$.

We further introduced a network variant by replacing PAM with a 4D cost volume technique and then compared the stereo matching performance of this variant to our PASMnet. Specifically, we replaced the four cascaded parallax-attention blocks at each stage of our cascaded parallax-attention module with four 3D convolutions. Note that, the number of channels are adjusted to ensure that this variant has comparable parameters to our PASMnet.

TABLE 3: Comparison between our parallax-attention block and 3D convolution. Here, \otimes refers to geometry-aware matrix multiplication.

Module	Params.	FLOPs	Memory
Parallax-Attention Block	3×3 Conv	$C^2 \times 3^2 \times 2$	$HW \times C \times 2 \times 2$
	1×1 Conv	$C^2 \times 2$	$HW \times C \times 2 \times 2$
	\otimes	-	$HW \times W \times 2$
	Total	$C^2 \times 20$	$HW \times C^2 \times (40 + \frac{2W}{C})$
3D Conv	$C^2 \times 3^3$	$HW \times C^2 \times 3^3 \times D$	$HW \times C \times D$

TABLE 4: Comparison between our PAM and cost volume formation on SceneFlow.

Model	Params.	Time	Memory	EPE	$> 1px$	$> 3px$
Cost Volume	7.35M	1 \times	1 \times	6.02	19.79	17.42
PAM	7.61M	1.1 \times	0.5 \times	4.54	18.99	15.91

TABLE 5: Comparative results achieved on SceneFlow by our PASMnet trained with different losses.

Model	\mathcal{L}_p	\mathcal{L}_s	\mathcal{L}_{PAM-p}	\mathcal{L}_{PAM-s}	\mathcal{L}_{PAM-c}	EPE	$> 1px$	$> 3px$
PASMnet	✓					6.07	20.78	17.11
PASMnet	✓	✓				4.87	19.69	16.54
PASMnet	✓	✓	✓			4.63	19.46	16.12
PASMnet	✓	✓	✓		✓	4.58	19.13	15.99
PASMnet	✓	✓	✓	✓	✓	4.54	18.99	15.91

The maximum disparity is set to 192 for the cost volume technique. Comparative results achieved on SceneFlow are shown in Table 4. Compared to the 4D cost volume technique, our PAM achieves much better performance with comparable model size and inference time. With the proposed PAM loss, we can directly regularize our PAM to capture accurate and consistent stereo correspondence. Therefore, better performance can be achieved. Moreover, our PAM also achieves higher efficiency, with its memory consumption being half of the cost volume technique.

Losses. We retrained our PASMnet using different losses to test the effectiveness of our loss function. It can be observed from Table 5 that the EPE/1-pixel error rate/3-pixel error rate achieved by our network are increased from 4.54/18.99/15.91 to 6.07/20.78/17.11 if only photometric loss is used for training. That is because, photometric loss cannot handle texture-less regions well. If smoothness loss is included for training, the EPE value is decreased to 4.87. Further, the performance is gradually improved if PAM loss is added. That is because, PAM loss regularizes our PAM to capture accurate and consistent stereo correspondence.

An example of matching cost and parallax-attention distributions produced by our network trained with different losses is shown in Fig. 11. Using \mathcal{L}_{PAM} for regularization, our network produces more reasonable matching cost and parallax-attention distributions. Both of them have peak values at the groundtruth disparity with high sharpness. This clearly demonstrates that performing direct regularization on parallax-attention maps enables our PAM to achieve better performance.




Fig. 11: An example of matching cost and parallax-attention distributions. (a) Results produced by our PASMnet trained without \mathcal{L}_{PAM} . (b) Results produced by our PASMnet. The red lines represent the groundtruth disparity. Using \mathcal{L}_{PAM} for regularization, our PASMnet produces more reasonable cost and parallax-attention distributions, which have peak values at the groundtruth disparity with higher sharpness.

TABLE 6: Comparison between PAM and cost volume formation on SceneFlow with different resolutions.

Resolution	Cost Volume			PAM		
	EPE	$> 1px$	$> 3px$	EPE	$> 1px$	$> 3px$
960×540	44.04	48.53	47.04	22.02	33.17	31.43
480×270	14.60	35.75	31.77	10.57	35.56	30.85
240×135	7.87	39.00	30.31	6.04	39.93	29.64
Average	22.17	41.09	36.37	12.88	36.22	30.64

4.3.4 Flexibility to Disparity Variations

We tested the flexibility of our PAM with respect to large disparity variations. We chose 20 image pairs from the test set of SceneFlow with more than 20% disparity values larger than 200 for evaluation. Note that, all pixels were used for evaluation.

Resolutions. We resized chosen image pairs to different resolutions to test the flexibility of our PAM. Results achieved on images with different resolutions are shown in Table 6. It can be observed that our PAM outperforms the cost volume technique on all metrics except 1-pixel error rate at the resolution of 240×135 . Moreover, the performance improvement achieved by our PAM is enhanced as the resolution increases. That is because, the fixed maximum disparity hinders the cost volume technique to capture longer-range correspondence. In contrast, our PAM is more flexible and robust to large disparity variations and achieves better performance.

Depths. We compared the performance on regions with different depths (disparities) to test the flexibility of our PAM. Results achieved on images with different depths are shown in Table 7. It can be observed that the performance of our PAM is significantly better than the cost volume technique on regions with disparities larger than 100. That is because, our PAM can get rid of setting a fixed maximum disparity to handle large disparities. In contrast, the cost

TABLE 7: Comparison between PAM and cost volume formation on SceneFlow with different depths.

Disparity	Cost Volume			PAM		
	EPE	$> 1px$	$> 3px$	EPE	$> 1px$	$> 3px$
> 200	150.21	99.99	99.99	57.90	41.62	41.60
100-200	36.24	45.80	45.79	25.16	33.15	33.15
< 100	13.71	33.92	31.92	16.31	34.57	31.81
Average	44.04	48.53	47.04	22.02	33.17	31.43

TABLE 8: Comparison to existing supervised and unsupervised stereo matching methods on KITTI 2015.

		Noc			All		
		D1-bg	D1-fg	D1-all	D1-bg	D1-fg	D1-all
Supervised	DispNet [11]	4.11	3.72	4.05	4.32	4.41	4.34
	GC-Net [4]	2.02	5.58	2.61	2.21	6.16	2.87
	CRL [52]	2.32	3.12	2.45	2.48	3.59	2.67
	iResNet [14]	2.07	2.76	2.19	2.25	3.40	2.44
	PSMNet [5]	1.71	4.31	2.14	1.86	4.62	2.32
	PASMnet_192 (ours)	1.88	3.91	2.22	2.04	4.33	2.41
Unsupervised	USCNet [53]	-	-	11.71	-	-	16.55
	Yu et al. [54]	-	-	8.35	-	-	19.14
	Zhou et al. [27]	-	-	8.61	-	-	9.91
	SegStereo [24]	-	-	7.70	-	-	8.79
	OASM [28]	5.44	17.30	7.39	6.89	19.42	8.98
	PASMnet (ours)	5.35	15.24	6.99	5.89	16.74	7.70
	PASMnet_192 (ours)	5.02	15.16	6.69	5.41	16.36	7.23

volume technique cannot capture the correspondences with disparities larger than the fixed maximum disparity.

4.3.5 Comparison to State-of-the-arts

We compare our PASMnet with existing unsupervised stereo matching methods [24], [27], [28], [53], [54] on the KITTI 2015 dataset. Note that, [24] and [28] use a maximum disparity of 192 as a prior knowledge for both training and test. For fair comparison, we also trained a variant with a maximum disparity of 192. Specifically, we excluded positions with disparities larger than the maximum disparity in the attention maps. Moreover, we include several supervised stereo matching methods [4], [5], [11], [14], [52] and trained our PASMnet_192 in a supervised manner for comparison. Specifically, we first finetuned our PASMnet_192 (which was trained on SceneFlow in an unsupervised manner) on SceneFlow using groundtruth disparities as its supervision. Then, this model was further finetuned on the KITTI 2015 dataset. As we can see from Table 8, our PASMnet outperforms other unsupervised methods by notable margins and significantly reduces the performance gap between supervised and unsupervised methods. With a prior maximum disparity, the performance of our PASMnet is further improved. Specifically, our PASMnet_192 achieves a much lower 3-pixel error rate (6.69/7.23) than OASM (7.39/8.98). If supervision is provided, our PASMnet_192 produces competitive results to iResNet and PSMNet. Figure 12 shows several visual examples. From the white boxes in Fig. 12, it can be observed that our PASMnet produces much more accurate and smoother results.

5 PAM FOR STEREO IMAGE SUPER-RESOLUTION

Stereo image SR aims at recovering an HR left image from a pair of LR stereo images. The key challenge for stereo




Fig. 12: Results achieved on the KITTI 2015 dataset.

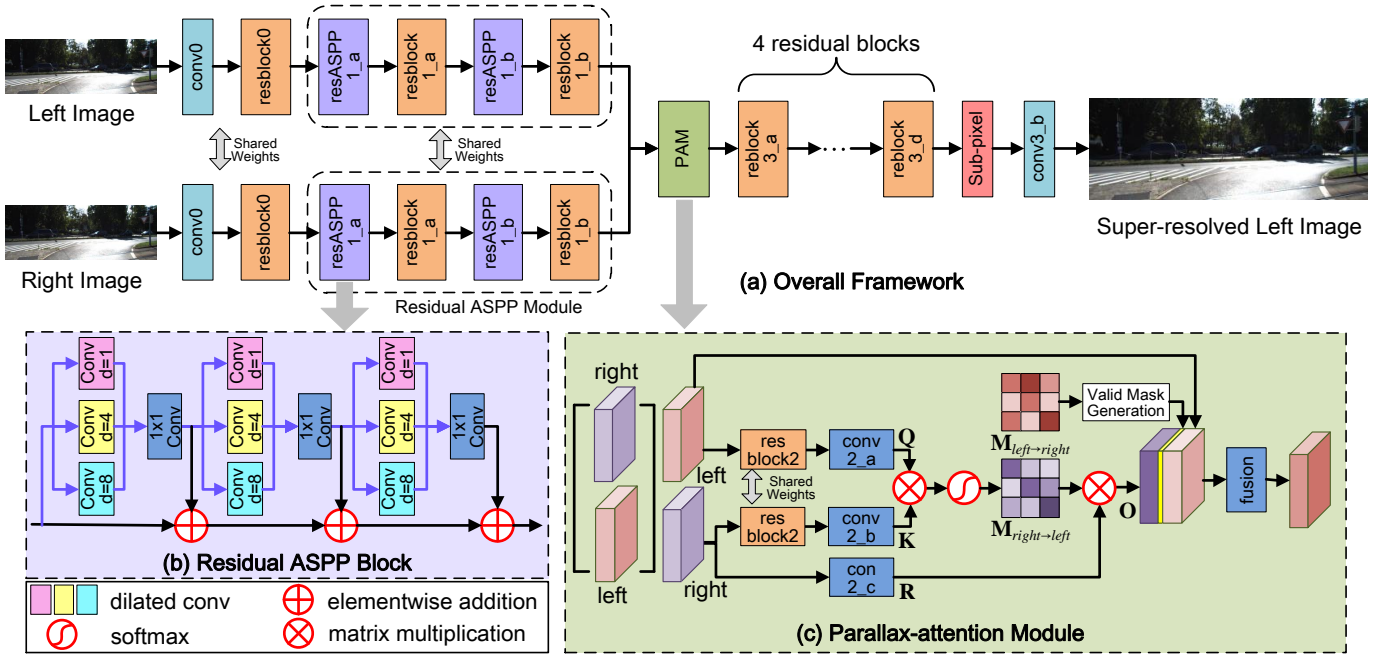


Fig. 13: An overview of our PASSRnet.

image SR lies in exploiting stereo correspondence to aggregate information from the pair of stereo images. Our PAM can effectively aggregate information from a pair of stereo images using the parallax-attention map as its guidance. Therefore, the PAM can be applied for the stereo image SR task.

5.1 Network Architecture

5.1.1 Overview

Based on our PAM, we propose a parallax-attention stereo super-resolution network (PASSRnet) for stereo image SR. Specifically, the stereo images are first fed to a convolutional layer and a residual block for feature extraction. Then, the resulting features are passed to a residual atrous spatial pyramid pooling (ASPP) module to extract deep feature

representations. Next, a parallax-attention module is used to aggregate features from a stereo pair. Finally, 4 residual blocks and a sup-pixel convolutional layer are used to generate the SR results. The architecture of our PASSRnet is shown in Fig. 13.

5.1.2 Residual Atrous Spatial Pyramid Pooling (ASPP) Module

Since feature representation with rich context information is beneficial to correspondence estimation [5], we propose a residual ASPP module to enlarge the receptive field and extract hierarchical features with dense pixel sampling rate and scales. As shown in Fig. 13 (a), input features are first fed to a residual ASPP block to generate multi-scale features. The resulting features are then sent to a residual block for feature fusion. This structure is repeated twice to produce

final features. Within each residual ASPP block (as shown in Fig. 13 (b)), we first combine three dilated convolutions (with dilation rates of 1, 4, 8) to form an ASPP group, and then cascade three ASPP groups in a residual manner. Our residual ASPP block not only enlarges the receptive field, but also enriches the diversity of convolutions, resulting in an ensemble of convolutions with different receptive regions and dilation rates. The highly discriminative feature learned by our residual ASPP module is beneficial to the overall SR performance, as demonstrated in Sec. 5.4.

5.1.3 Parallax-Attention Module

Based on our PAM, we introduced a parallax-attention module to exploit stereo correspondence for the aggregation of features from a pair of stereo images. As shown in Fig. 13(c), features from left/right images are first fed to a transition residual block for feature adaption. Then, the resulting features are used to generate the output feature \mathbf{O} and the parallax-attention map $\mathbf{M}_{right \rightarrow left}$ based on PAM. Next, features from left/right images are exchanged to produce $\mathbf{M}_{left \rightarrow right}$ for valid mask generation. Finally, the concatenation of the output feature, the feature from the left image and the valid mask is fed to a 1×1 convolution for feature fusion.

It should be noted that our PASSRnet can be considered as a multi-task network to learn both stereo correspondence and SR. However, using shared features for different tasks usually suffers from training conflict [55]. Therefore, a transition block is used in our parallax-attention module to alleviate this problem. The effectiveness of the transition block is demonstrated in Sec. 5.4.

5.2 Losses

We used the mean square error (MSE) loss as the SR loss:

$$\mathcal{L}_{SR} = \left\| \mathbf{I}_{left}^{SR} - \mathbf{I}_{left}^H \right\|_2^2, \quad (17)$$

where \mathbf{I}_{left}^{SR} and \mathbf{I}_{left}^H represent the SR result and HR groundtruth of the left image, respectively. Other than the SR loss, we also used the PAM loss (as defined in Sec. 4.2.3) to help our network exploit the correspondence between stereo images. The overall loss is formulated as:

$$\mathcal{L} = \mathcal{L}_{SR} + \lambda \mathcal{L}_{PAM}, \quad (18)$$

where λ is empirically set to 0.005. The performance of our network with different losses will be analyzed in Sec. 5.4. Note that, stereo correspondence is learned in an unsupervised manner using \mathcal{L}_{PAM} while SR is learned in a supervised manner using \mathcal{L}_{SR} .

5.3 The Flickr1024 Dataset

Although several stereo datasets such as Middlebury [56] and KITTI [50], [57] are already available, these datasets are mainly proposed for stereo matching. Further, the Middlebury dataset only consists of close shots of man-made objects, while the KITTI 2012 and KITTI 2015 datasets only consist of road scenes. For stereo image SR task, a large dataset which covers diverse scenes and consists of images with high quality and rich details is required. Therefore, we introduce a new Flickr1024 dataset [58] for stereo image

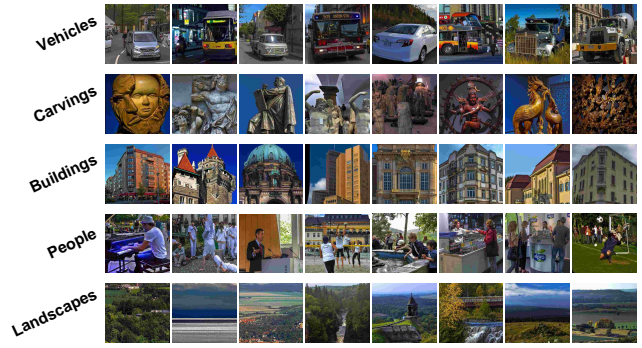


Fig. 14: Samples of different scenes covered in the Flickr1024 dataset.



Fig. 15: An example stereogram collected from Flickr.

SR, which covers a large diversity of scenes, including landscapes, urban scenes, people and man-made objects, as shown in Fig. 14. Our Flickr1024 dataset is available at: <https://yingqianwang.github.io/Flickr1024>.

5.3.1 Data Collection and Processing

We manually collected 1024 pairs of RGB stereo images from Flickr using tags such as stereophotography, stereoscopic and cross-eye 3D. All images are stereograms (as shown in Fig. 15) taken by amateur photographers using dual lens or dual cameras.

First, we cut each stereogram into a pair of images and crop their black margins. The resulting left and right images from a stereogram are exchanged since the stereograms are provided in a cross-eye mode. Note that, these stereo image pairs are originally shifted to a common focus plane by the amateurs to produce a perception of 3D for viewers. In other words, both positive and negative disparities exist in these image pairs. Therefore, we roughly shift these images back to ensure that zero disparity corresponds to infinite depth. For close shots, since regions with infinite depth are unavailable, we just shift these images to make the minimum disparity larger than a threshold (empirically set to 40 pixels in our dataset). Finally, we crop each resulting image to multiple of 12 pixels on both axes following [59].

5.3.2 Comparison to Existing Datasets

The Flickr1024 dataset is compared to three widely used stereo datasets including Middlebury, KITTI 2012 and KITTI 2015. CNNIQA [60] and SR-metric [61] are used to evaluate the image quality. It can be observed from Table 9 that our Flickr1024 dataset is larger than other datasets by at least 5 times. Besides, the pixel per image (ppi) value of our

TABLE 9: Comparison between the Middlebury, KITTI 2012, KITTI 2015 and Flickr1024 datasets. Only the training sets of the KITTI 2012 and KITTI 2015 datasets are considered.

Dataset	image pairs	ppi	CNNIQA (\downarrow)	SR-metric (\uparrow)
Middlebury	65	3511605	20.18	6.01
KITTI 2012	194	462564	20.32	7.15
KITTI 2015	200	465573	22.86	7.06
Flickr1024	1024	800486	19.75	7.12

Flickr1024 dataset is nearly 2 times those of the KITTI 2012 and KITTI 2015 datasets. Although the Middlebury dataset has the highest ppi value, the number of image pairs in this dataset is very limited. Further, our Flickr1024 dataset achieves comparable or even better CNNIQA and SR-metric values than other datasets, which demonstrates the good image quality of our Flickr1024 dataset.

5.4 Experimental Results

5.4.1 Datasets

For training, we followed [37] and downsampled 60 Middlebury [56] images by a factor of 2 to generate HR images. We further used Flickr1024 as the additional training data for our PASSRnet. For test, we used 5 images from the Middlebury dataset², 20 images from the KITTI 2012 dataset [57]³ and 20 images from the KITTI 2015 dataset [50]⁴. For validation, we selected another 20 images from the KITTI 2012 dataset.

5.4.2 Implementation Details

Training Details. We first downsampled HR images using bicubic interpolation to generate LR images, and then cropped 30×90 patches with a stride of 20 from these LR images. Meanwhile, their corresponding patches in HR images were cropped. The horizontal patch size was set to 90 to cover most disparities ($\sim 96\%$) in our training dataset. These patches were randomly flipped horizontally and vertically for data augmentation. Note that, rotation was not performed for augmentation, since it would destroy the epipolar constraints.

Our PASSRnet was implemented in Pytorch on a PC with an Nvidia GTX 1080Ti GPU. All models were optimized using the Adam method [51] with $\beta_1 = 0.9$, $\beta_2 = 0.999$ and a batch size of 32. The initial learning rate was set to 2×10^{-4} and reduced to half after every 30 epochs. The training was stopped after 80 epochs since more epochs do not provide further improvement.

Metrics. We used peak signal-to-noise ratio (PSNR) and SSIM to test SR performance. Following [37], we cropped borders to achieve fair comparison.

5.4.3 Ablation Study

Single Input vs. Stereo Input. Compared to single images, stereo image pairs provide additional information observed from a different viewpoint. To demonstrate the effectiveness of stereo information for SR, we removed the

parallax-attention module from our PASSRnet and retrained the network with single images (*i.e.*, the left images). For comparison, we also used pairs of replicated left images as the input to the original PASSRnet. It can be observed from Table 10 that the network trained with single images suffers a decrease of 0.16 dB (from 25.43 to 25.27) in terms of PSNR. If pairs of replicated left images are fed to the original PASSRnet, the PSNR value is decreased to 25.29 dB. Without extra information introduced by stereo images, our PASSRnet with replicated images achieves comparable performance to the network trained with single images. This clearly demonstrates that stereo information contributes to the performance improvement of our PASSRnet.

Residual ASPP Module. Residual ASPP module is used in our network to extract multi-scale features. To demonstrate its effectiveness, two variants were introduced. First, to test the effectiveness of residual connections, we removed them to obtain a cascading ASPP module. Then, to test the effectiveness of atrous convolutions, we replaced them with ordinary convolutions. From the comparative results shown in Table 10, we can see that SR performance benefits from both residual connections (25.43 vs. 25.40) and atrous convolutions (25.43 vs. 25.38). That is because, residual connections enable our residual ASPP module to extract features at more scales, resulting in more representative features. Furthermore, large receptive field of atrous convolutions facilitates our PASSRnet to employ context information in a large area. Therefore, more accurate correspondence can be obtained to improve SR performance.

Parallax-attention Module. Parallax-attention module is used to exploit stereo correspondence for feature aggregation. To demonstrate its effectiveness, we introduced a variant by removing the parallax-attention module and directly stacking the features produced by the residual ASPP module. It can be observed from Table 10 that the PSNR value is decreased from 25.43 dB to 25.28 dB if parallax-attention module is removed. Without parallax-attention module, local receptive field hinders our network to capture stereo correspondence over large disparities. By using the parallax-attention module, our network can effectively aggregate features from left and right images with global receptive field along the epipolar line to achieve better performance.

Transition Block. Transition block is used in the parallax-attention module to alleviate the training conflict in shared layers. To demonstrate the effectiveness of transition block, we removed it from our parallax-attention module and retrained the network. It can be observed from Table 10 that the PSNR value is decreased from 25.43 dB to 25.36 dB without transition block. That is because, the transition block enhances task-specific feature learning and alleviates training conflict in shared layers.

PAM vs. Cost Volume Cost volume and 3D convolutions are commonly used to obtain stereo correspondence [4], [5]. To demonstrate the efficiency of our PAM in stereo correspondence generation, we replaced parallax-attention module with a 4D cost volume followed two 3D convolutional layers ($3 \times 3 \times 3$). Comparative results achieved on the KITTI 2015 dataset are listed in Table 11. As compared to the cost volume formation, the number of parameters and FLOPs of our parallax-attention module are reduced by over 2 and 150

2. Middlebury: cloth2, motorcycle, piano, pipes and sword2.

3. KITTI2012: 000000_10 to 000019_10.

4. KITTI2015: 000000_10 to 000019_10.

TABLE 10: Comparative results achieved on the KITTI 2015 dataset by PASSRnet with different settings for $4\times$ SR.

Model	Input	PSNR	SSIM	Params.	Time
PASSRnet with single input	Left	25.27	0.770	1.32M	114ms
PASSRnet with replicated inputs	Left-Left	25.29	0.771	1.42M	176ms
PASSRnet without residual manner	Left-Right	25.40	0.774	1.42M	176ms
PASSRnet without atrous convolution	Left-Right	25.38	0.773	1.42M	176ms
PASSRnet without parallax-attention module	Left-Right	25.28	0.771	1.32M	135ms
PASSRnet without transition block	Left-Right	25.36	0.773	1.34M	160ms
PASSRnet	Left-Right	25.43	0.776	1.42M	176ms

TABLE 11: Comparison between our PAM and the cost volume formation for $4\times$ SR. FLOPs are calculated on $128\times 128\times 64$ input features, while Time/PSNR/SSIM values are achieved on the KITTI 2015 dataset.

Model	Params.	FLOPs	Time	PSNR	SSIM
PAM	94K	$1\times$	$1\times$	25.43	0.776
Cost Volume	221K	$151\times$	$1.5\times$	25.23	0.768

TABLE 12: Comparative results achieved on KITTI 2015 by our PASSRnet trained with different losses for $4\times$ SR.

Model	\mathcal{L}_{SR}	$\mathcal{L}_{photometric}$	\mathcal{L}_{smooth}	\mathcal{L}_{cycle}	PSNR	SSIM
PASSRnet	✓				25.35	0.771
PASSRnet	✓	✓			25.38	0.773
PASSRnet	✓	✓	✓		25.40	0.774
PASSRnet	✓	✓	✓	✓	25.43	0.776

times, respectively. Using the parallax-attention module, our PASSRnet achieves better SR performance (*i.e.*, PSNR value is increased from 25.23 dB to 25.43 dB) and higher efficiency (*i.e.*, running time is decreased by 1.5 times). That is because, two 3D convolutional layers are insufficient to capture long-range correspondence within the cost volume. However, adding more layers will lead to a significant increase of computational cost.

Losses To test the effectiveness of our losses, we retrained PASSRnet using different losses. It can be observed from Table 12 that the PSNR value of our PASSRnet is decreased from 25.43 dB to 25.35 dB if PASSRnet is trained with only SR loss. That is because, with only this loss, our PAM cannot capture accurate stereo correspondence for feature aggregation. If photometric loss, smoothness loss and cycle loss are added, the performance is gradually improved from 25.35/0.771 to 25.43/0.776. That is because, these losses encourage our PAM to generate reliable and consistent correspondence.

5.4.4 Comparison to State-of-the-arts

We compared our PASSRnet to recent CNN-based single image SR and stereo image SR methods, including SRCNN [62], VDSR [63], DRCN [64], LapSRN [65], DRRN [66] and StereoSR [37]. The codes provided by the authors of these methods were used to conduct experiments. Note that, similar to [37], [67], EDSR [68], RDN [69] and D-DBPN [70] are not included in our comparison since their model sizes are larger than our PASSRnet by at least 8 times.

Quantitative Results. The quantitative results are shown in Table 13. It can be observed that our PASSRnet achieves the best performance on the Middlebury, KITTI 2012 and

KITTI 2015 datasets. Specifically, our PASSRnet outperforms the best single image SR method (*i.e.*, DRRN) by 1.04 dB in terms of PSNR on the Middlebury dataset for $2\times$ SR. Moreover, the PSNR value achieved by our network is higher than that of StereoSR by 1.00 dB. That is because, more accurate correspondence can be captured by our PAM. **Qualitative Results.** Figure 16 illustrates the qualitative results achieved on two scenarios. It can be observed from zoom-in regions that single image SR methods cannot reliably recover missing details and suffer from obvious artifacts. In contrast, our PASSRnet uses stereo correspondence to produce finer details with fewer artifacts, such as the railings and stripe in Fig. 16. Compared to StereoSR [37], our PASSRnet captures more accurate stereo correspondence for feature aggregation. Consequently, visually more promising results are produced.

5.4.5 Flexibility to Disparity Variations

We further tested the flexibility of our PASSRnet with respect to large disparity variations.

Resolutions. We collected 10 close-shot stereo images (with disparities larger than 200) from Flickr and resized them to different resolutions to test the flexibility of our network. Results achieved on images with different resolutions are shown in Table 14.

It can be observed that our PASSRnet is significantly better than StereoSR in terms of efficiency (*i.e.*, FLOPs) on low resolution images. Meanwhile, our PASSRnet outperforms StereoSR by a large margin in terms of PSNR on high resolution images. That is because, StereoSR needs to perform padding for images with horizontal resolution lower than 64 pixels, which involves unnecessary calculations. For high resolution images, the fixed maximum disparity hinders StereoSR to capture longer-range correspondence. Therefore, the SR performance of StereoSR is limited.

Baselines. We used a view-by-view scanning scheme as in [71] to obtain stereo image pairs with different baselines. The device used for image acquisition is shown in Fig. 17. Specifically, we first installed a camera on the gantry and acquired an image of a static scene. This image was referred to as the left image. Then, the camera was controlled to move rightward along the horizontal rail with different distances. The images acquired at different locations were referred to as right images with different baselines. Five different scenes were used for image acquisition. For each scene, we obtained three stereo image pairs with short baselines (with disparities around 20 pixels), medium baselines (with disparities around 60 pixels) and large baselines (with disparities around 180 pixels). Results achieved by StereoSR

TABLE 13: Comparative PSNR/SSIM values achieved on the Middlebury, KITTI 2012 and KITTI 2015 datasets. Results marked with * are directly copied from the corresponding paper. Note that, only 2× SR results of StereoSR are presented on the KITTI 2012 and KITTI 2015 datasets since a 4× SR model is unavailable.

Dataset	Scale	Single Image SR					Stereo Image SR	
		SRCNN [62]	VDSR [63]	DRCN [64]	LapSRN [65]	DRRN [66]	StereoSR [37]	Ours
Middlebury (5 images)	×2	32.05/0.935	32.66/0.941	32.82/0.941	32.75/0.940	32.91/0.945	33.05/0.955*	34.05/0.960
	×4	27.46/0.843	27.89/0.853	27.93/0.856	27.98/0.861	27.93/0.855	26.80/0.850*	28.63/0.871
KITTI 2012 (20 images)	×2	29.75/0.901	30.17/0.906	30.19/0.906	30.10/0.905	30.16/0.908	30.13/0.908	30.65/0.916
	×4	25.53/0.764	25.93/0.778	25.92/0.777	25.96/0.779	25.94/0.773	-	26.26/0.790
KITTI 2015 (20 images)	×2	28.77/0.901	28.99/0.904	29.04/0.904	28.97/0.903	29.00/0.906	29.09/0.909	29.78/0.919
	×4	24.68/0.744	25.01/0.760	25.04/0.759	25.03/0.760	25.05/0.756	-	25.43/0.776




Fig. 16: Visual comparison for 2× SR. These results are achieved on “test_image_013” of the KITTI 2012 dataset and “test_image_019” of the KITTI 2015 dataset.

TABLE 14: Comparison between our PASSRnet and StereoSR [37] on stereo images with different resolutions for 2× SR.


Resolution	StereoSR [37]		Ours	
	PSNR	FLOPs	PSNR	FLOPs
High (500 × 500)	39.27	1×	41.45 (↑ 2.18)	0.57×
Middle (100 × 100)	34.21	1×	35.04 (↑ 0.83)	0.58×
Low (20 × 20)	29.48	1×	29.88 (↑ 0.40)	0.36×
Average	34.32	1×	35.46 (↑ 1.14)	0.50×

TABLE 15: Comparison between our PASSRnet and StereoSR [37] on stereo images with different baselines for 2× SR.

Baseline	StereoSR [37]		Ours	
	PSNR	SSIM	PSNR	SSIM
Large	37.13	0.9605	38.43 (↑1.30)	0.9690 (↑0.085)
Medium	37.35	0.9628	38.50 (↑1.15)	0.9692 (↑0.064)
Short	37.36	0.9628	38.50 (↑1.14)	0.9693 (↑0.065)
Average	37.28	0.9620	38.48 (↑1.20)	0.9692 (↑0.072)

[37] and our PASSRnet on these stereo image pairs with different baselines are shown in Table 15.

It can be observed that our PASSRnet outperforms StereoSR by 1.14 dB in terms of PSNR for stereo image pairs with short baselines. For stereo images pairs with large baselines, the improvement is increased to 1.30 dB. When the baseline is increased, StereoSR suffers from a more significant performance drop in terms of PSNR as compared to our PASSRnet



(a) Structure chart of the camera gantry



(b) Real gantry with a camera

(c) Controller

Fig. 17: An illustration of the device used for stereo image acquisition.

(0.23 vs. 0.07). This clearly demonstrates that our PASSRnet is more flexible and robust to large disparity variations than StereoSR.

Depths. We collected 30 stereo image pairs from Flickr. These image pairs cover different scenes with different depths. We divided these image pairs into three groups with small depths (with disparities around 150 pixels), medium

TABLE 16: Comparison between our PASSRnet and StereoSR [37] on stereo images with different depths for $2\times$ SR.

Depth	StereoSR [37]		Ours	
	PSNR	SSIM	PSNR	SSIM
Small	37.60	0.9652	39.03 (\uparrow 1.43)	0.9749 (\uparrow 0.0097)
Medium	31.08	0.9145	32.37 (\uparrow 1.29)	0.9219 (\uparrow 0.0074)
Large	36.36	0.9596	37.55 (\uparrow 1.19)	0.9646 (\uparrow 0.0050)
Large	35.01	0.9464	36.32 (\uparrow 1.31)	0.9538 (\uparrow 0.0074)

depths (with disparities around 80 pixels) and large depths (with disparities around 20 pixels). Results achieved by StereoSR [37] and our PASSRnet on these stereo image pairs with different depths are shown in Table 16.

It can be observed that our PASSRnet achieves high improvement on stereo image pairs with small depths. That is because, the fixed maximum disparity used in StereoSR hinders longer-range correspondence to be employed. In contrast, by using PAM, our PASSRnet can get rid of setting a fixed maximum disparity and incorporate stereo correspondence under large disparity variations.

6 CONCLUSION

In this paper, we propose a generic parallax-attention mechanism (PAM) to capture stereo correspondence under large disparity variations. Our PAM integrates epipolar constraints with attention mechanism to calculate feature similarities along the epipolar line. Our PAM is compact and can be applied to variant stereo tasks. We apply our PAM to stereo matching and stereo image SR tasks and therefore propose two networks (namely, PASMnet and PASSRnet). We also introduce a new and large-scale dataset named Flickr1024 for stereo image SR. Extensive experiments have demonstrated the effectiveness of our PAM and the state-of-the-art performance of our PASMnet and PASSRnet.

ACKNOWLEDGMENTS

This work was partially supported by the National Natural Science Foundation of China (No. 61972435). The authors would like to thank Sascha Becher and Tom Bentz for the approval of using their cross-eye stereo photographs on Flickr.

REFERENCES

- [1] Yulan Guo, Mohammed Bennamoun, Ferdous Sohel, Min Lu, and Jianwei Wan. 3d object recognition in cluttered scenes with local surface features: a survey. *IEEE Trans. Pattern Anal. Mach. Intell.*, 36(11):2270–2287, 2014.
- [2] Yulan Guo, Hanyun Wang, Qingyong Hu, Hao Liu, Li Liu, and Mohammed Bennamoun. Deep learning for 3d point clouds: A survey. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2020.
- [3] Wenjie Luo, Alexander G. Schwing, and Raquel Urtasun. Efficient deep learning for stereo matching. In *CVPR*, pages 5695–5703, 2016.
- [4] Alex Kendall, Hayk Martirosyan, Saumitro Dasgupta, and Peter Henry. End-to-end learning of geometry and context for deep stereo regression. In *ICCV*, pages 66–75, 2017.
- [5] Jia-Ren Chang and Yong-Sheng Chen. Pyramid stereo matching network. In *CVPR*, pages 5410–5418, 2018.
- [6] Longguang Wang, Yingqian Wang, Zhengfa Liang, Zaiping Lin, Jungang Yang, Wei An, and Yulan Guo. Learning parallax attention for stereo image super-resolution. In *CVPR*, pages 12250–12259, 2019.
- [7] Shangchen Zhou, Jiawei Zhang, Wangmeng Zuo, Haozhe Xie, Jinshan Pan, and Jimmy Ren. DAVANet: Stereo deblurring with view aggregation. In *CVPR*, 2019.
- [8] Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. Stereo magnification: learning view synthesis using multiplane images. *ACM Transactions on Graphics (TOG)*, 37(4):65, 2018.
- [9] Bing Li, Chia-Wen Lin, Boxin Shi, Tiejun Huang, Wen Gao, and C-C Jay Kuo. Depth-aware stereo video retargeting. In *CVPR*, pages 6517–6525, 2018.
- [10] Dongdong Chen, Lu Yuan, Jing Liao, Nenghai Yu, and Gang Hua. Stereoscopic neural style transfer. In *CVPR*, pages 6654–6663, 2018.
- [11] Nikolaus Mayer, Eddy Ilg, Philip Häusser, 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.
- [12] Zequn Jie, Pengfei Wang, Yonggen Ling, Bo Zhao, Yunchao Wei, Jiashi Feng, and Wei Liu. Left-right comparative recurrent model for stereo matching. In *CVPR*, pages 3838–3846, 2018.
- [13] Xiaoyang Guo, Kai Yang, Wukui Yang, Xiaogang Wang, and Hongsheng Li. Group-wise correlation stereo network. In *CVPR*, 2019.
- [14] Zhengfa Liang, Yiliu Feng, Yulan Guo, Hengzhu Liu, Linbo Qiao, Wei Chen, Li Zhou, and Jianfeng Zhang. Learning for disparity estimation through feature constancy. In *CVPR*, pages 2811–2820, 2018.
- [15] Zhengfa Liang, Yulan Guo, Yiliu Feng, Wei Chen, Linbo Qiao, Li Zhou, Jianfeng Zhang, and Hengzhu Liu. Stereo matching using multi-level cost volume and multi-scale feature constancy. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2019.
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Spatial pyramid pooling in deep convolutional networks for visual recognition. *IEEE Trans. Pattern Anal. Mach. Intell.*, 37(9):1904–1916, 2015.
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
- [18] Joseph Redmon, Santosh Kumar Divvala, Ross B. Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *CVPR*, pages 779–788, 2016.
- [19] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *ICCV*, pages 2961–2969, 2017.
- [20] J. Žbontar and Y. LeCun. Computing the stereo matching cost with a convolutional neural network. In *CVPR*, pages 1592–1599, 2015.
- [21] Sergey Zagoruyko and Nikos Komodakis. Learning to compare image patches via convolutional neural networks. In *CVPR*, pages 4353–4361, 2015.
- [22] Amit Shaked and Lior Wolf. Improved stereo matching with constant highway networks and reflective confidence learning. In *CVPR*, pages 6901–6910, 2017.
- [23] Haesol Park and Kyoung Mu Lee. Look wider to match image patches with convolutional neural networks. *IEEE Signal Process. Lett.*, 24(12):1788–1792, 2017.
- [24] Guorun Yang, Hengshuang Zhao, Jianping Shi, Zhidong Deng, and Jiaya Jia. Segstereo: Exploiting semantic information for disparity estimation. In *ECCV*, pages 636–651, 2018.
- [25] Heiko Hirschmuller. Stereo processing by semiglobal matching and mutual information. *IEEE Trans. Pattern Anal. Mach. Intell.*, 30(2):328–341, 2007.
- [26] 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.
- [27] Chao Zhou, Hong Zhang, Xiaoyong Shen, and Jiaya Jia. Unsupervised learning of stereo matching. In *ICCV*, pages 1576–1584, 2017.
- [28] Ang Li and Zejian Yuan. Occlusion aware stereo matching via cooperative unsupervised learning. In *ACCV*, pages 197–213, 2018.
- [29] Wonil Song, Sungil Choi, Somi Jeong, and Kwanghoon Sohn. Stereoscopic image super-resolution with stereo consistent feature. In *AAAI*, pages 12031–12038, 2020.
- [30] Xin Tao, Hongyun Gao, Renjie Liao, Jue Wang, and Jiaya Jia. Detail-revealing deep video super-resolution. In *ICCV*, pages 4482–4490, 2017.
- [31] Xintao Wang, Kelvin C.K. Chan, Ke Yu, Chao Dong, and Chen Change Loy. EDVR: Video restoration with enhanced deformable convolutional networks. In *CVPRW*, 2019.
- [32] Longguang Wang, Yulan Guo, Li Liu, Zaiping Lin, Xinpu Deng, and Wei An. Deep video super-resolution using HR optical flow estimation. *IEEE Trans. Image Process.*, 29:4323–4336, 2020.

- [33] Yunlong Wang, Fei Liu, Kunbo Zhang, Guangqi Hou, Zhenan Sun, and Tieniu Tan. LFNet: A novel bidirectional recurrent convolutional neural network for light-field image super-resolution. *IEEE Trans. Image Process.*, 27(9):4274–4286, 2018.
- [34] Shuo Zhang, Youfang Lin, and Hao Sheng. Residual networks for light field image super-resolution. In *CVPR*, pages 11046–11055, 2019.
- [35] Yingqian Wang, Longguang Wang, Jungang Yang, Wei An, Jingyi Yu, and Yulan Guo. Spatial-angular interaction for light field image super-resolution. In *ECCV*, 2020.
- [36] Arnab V. Bhavsar and A. N. Rajagopalan. Resolution enhancement in multi-image stereo. *IEEE Trans. Pattern Anal. Mach. Intell.*, 32(9):1721–1728, 2010.
- [37] Daniel S. Jeon, Seung-Hwan Baek, Inchang Choi, and Min H. Kim. Enhancing the spatial resolution of stereo images using a parallax prior. In *CVPR*, pages 1721–1730, 2018.
- [38] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. In *ICLR*, 2014.
- [39] Ankur Parikh, Oscar Täckström, Dipanjan Das, and Jakob Uszkoreit. A decomposable attention model for natural language inference. In *EMNLP*, pages 2249–2255, 2016.
- [40] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, pages 6000–6010, 2017.
- [41] Hengshuang Zhao, Yi Zhang, Shu Liu, Jianping Shi, Chen Change Loy, Dahua Lin, and Jiaya Jia. PSANet: Point-wise spatial attention network for scene parsing. In *ECCV*, page 2, 2018.
- [42] Jun Fu, Jing Liu, Haijie Tian, Yong Li, Yongjun Bao, Zhiwei Fang, and Hanqing Lu. Dual attention network for scene segmentation. In *CVPR*, pages 3146–3154, 2019.
- [43] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron C. Courville, Ruslan Salakhutdinov, Richard S. Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *ICML*, pages 2048–2057, 2015.
- [44] Long Chen, Hanwang Zhang, Jun Xiao, Liqiang Nie, Jian Shao, Wei Liu, and Tat-Seng Chua. SCA-CNN: spatial and channel-wise attention in convolutional networks for image captioning. In *CVPR*, pages 6298–6306, 2017.
- [45] Han Zhang, Ian J. Goodfellow, Dimitris N. Metaxas, and Augustus Odena. Self-attention generative adversarial networks. In *NeurIPS*, 2018.
- [46] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *CVPR*, 2018.
- [47] Guilin Liu, Fitsum A Reda, Kevin J Shih, Ting-Chun Wang, Andrew Tao, and Bryan Catanzaro. Image inpainting for irregular holes using partial convolutions. In *ECCV*, pages 85–100, 2018.
- [48] Zhichao Yin and Jianping Shi. Geonet: Unsupervised learning of dense depth, optical flow and camera pose. In *CVPR*, pages 1983–1992, 2018.
- [49] Clément Godard, Oisín Mac Aodha, and Gabriel J. Brostow. Unsupervised monocular depth estimation with left-right consistency. In *CVPR*, pages 6602–6611, 2017.
- [50] Moritz Menze and Andreas Geiger. Object scene flow for autonomous vehicles. In *CVPR*, pages 3061–3070, 2015.
- [51] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- [52] J. Pang, W. Sun, J. S. Ren, C. Yang, and Q. Yan. Cascade residual learning: A two-stage convolutional neural network for stereo matching. In *ICCVW*, pages 878–886, 2017.
- [53] A. Ahmadi and I. Patras. Unsupervised convolutional neural networks for motion estimation. In *ICIP*, pages 1629–1633, 2016.
- [54] Jason J. Yu, Adam W. Harley, and Konstantinos G. Derpanis. Back to basics: Unsupervised learning of optical flow via brightness constancy and motion smoothness. In *ECCVW*, pages 3–10, 2016.
- [55] Ozan Sener and Vladlen Koltun. Multi-task learning as multi-objective optimization. In *NeurIPS*, page 5, 2018.
- [56] Daniel Scharstein, Heiko Hirschmüller, York Kitajima, Greg Krathwohl, Nera Nesić, Xi Wang, and Porter Westling. High-resolution stereo datasets with subpixel-accurate ground truth. In *GCPR*, pages 31–42, 2014.
- [57] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the KITTI vision benchmark suite. In *CVPR*, pages 3354–3361, 2012.
- [58] Yingqian Wang, Longguang Wang, Jungang Yang, Wei An, and Yulan Guo. Flickr1024: A large-scale dataset for stereo image super-resolution. In *ICCVW*, 2019.
- [59] Eirikur Agustsson and Radu Timofte. NTIRE 2017 challenge on single image super-resolution: Dataset and study. In *CVPR*, pages 1122–1131, 2017.
- [60] Le Kang, Peng Ye, Yi Li, and David S. Doermann. Convolutional neural networks for no-reference image quality assessment. In *CVPR*, pages 1733–1740, 2014.
- [61] Chao Ma, Chih-Yuan Yang, Xiaokang Yang, and Ming-Hsuan Yang. Learning a no-reference quality metric for single-image super-resolution. *Computer Vision and Image Understanding*, 158:1–16, 2017.
- [62] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Learning a deep convolutional network for image super-resolution. In *ECCV*, pages 184–199, 2014.
- [63] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Accurate image super-resolution using very deep convolutional networks. In *CVPR*, pages 1646–1654, 2016.
- [64] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Deeply-recursive convolutional network for image super-resolution. In *CVPR*, pages 1637–1645, 2016.
- [65] Wei-Sheng Lai, Jia-Bin Huang, Narendra Ahuja, and Ming-Hsuan Yang. Deep laplacian pyramid networks for fast and accurate super-resolution. In *CVPR*, pages 5835–5843, 2017.
- [66] Ying Tai, Jian Yang, and Xiaoming Liu. Image super-resolution via deep recursive residual network. In *CVPR*, pages 2790–2798, 2017.
- [67] Namhyuk Ahn, Byungkon Kang, and Kyung-Ah Sohn. Fast, accurate, and lightweight super-resolution with cascading residual network. In *ECCV*, pages 252–268, 2018.
- [68] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single image super-resolution. In *CVPR*, 2017.
- [69] Yulun Zhang, Yapeng Tian, Yu Kong, Bineng Zhong, and Yun Fu. Residual dense network for image super-resolution. In *CVPR*, pages 2472–2481, 2018.
- [70] Muhammad Haris, Gregory Shakhnarovich, and Norimichi Ukita. Deep back-projection networks for super-resolution. In *CVPR*, pages 1664–1673, 2018.
- [71] Yingqian Wang, Jungang Yang, Yulan Guo, Chao Xiao, and Wei An. Selective light field refocusing for camera arrays using bokeh rendering and superresolution. *IEEE Signal Processing Letters*, 26(1):204–208, 2019.




Longguang Wang received the B.E. degree in Electrical Engineering from Shandong University (SDU), Jinan, China, in 2015, and the M.E. degree in Information and Communication Engineering from National University of Defense Technology (NUDT), Changsha, China, in 2017. He is currently pursuing the Ph.D. degree with the College of Electronic Science and Technology, NUDT. His current research interests include low-level vision and deep learning.



Yulan Guo is currently an associate professor. He received the B.Eng. and Ph.D. degrees from National University of Defense Technology (NUDT) in 2008 and 2015, respectively. He was a visiting Ph.D. student with the University of Western Australia from 2011 to 2014. He worked as a postdoctoral research fellow with the Institute of Computing Technology, Chinese Academy of Sciences from 2016 to 2018. He has authored over 90 articles in journals and conferences, such as the IEEE TPAMI and IJCV. His current research interests focus on 3D vision, particularly on 3D feature learning, 3D modeling, 3D object recognition, and scene understanding. Dr. Guo received the ACM China SIGAI Rising Star Award in 2019, Wu-Wenjun Outstanding AI Youth Award in 2019, and the CAAI Outstanding Doctoral Dissertation Award in 2016. He served as an associate editor for IET Computer Vision and IET Image Processing, a guest editor for IEEE TPAMI, and an area chair for CVPR 2021 and ICPR 2020.



Yingqian Wang received the B.E. degree in Electrical Engineering from Shandong University (SDU), Jinan, China, in 2016, and the M.E. degree in Information and Communication Engineering from National University of Defense Technology (NUDT), Changsha, China, in 2018. He is currently pursuing the Ph.D. degree with the College of Electronic Science and Technology, NUDT. His current research interests focus on low-level vision, particularly on light field imaging and image super-resolution.



Zhengfa Liang received the B.Eng. degree from Tsinghua University, Beijing, China, in 2010, and the Ph.D. degree from National University of Defense Technology (NUDT), Changsha, China, in 2017. He is currently an Assistant Professor with National Key Laboratory of Science and Technology on Blind Signal Processing. His current research interests focus on 3D vision, particularly on stereo matching, depth estimation, face reconstruction, and face swapping.




Zaiping Lin received the B.Eng. and Ph.D. degrees from National University of Defense Technology (NUDT) in 2007 and 2012, respectively. He is currently an Assistant Professor with the College of Electronic Science and Technology, NUDT. His current research interests include infrared image processing and signal processing.



Jungang Yang received the B.Eng. and Ph.D. degrees from National University of Defense Technology (NUDT), in 2007 and 2013 respectively. He was a visiting Ph.D. student with the University of Edinburgh, Edinburgh from 2011 to 2012. He is currently an associate professor with the College of Electronic Science and Technology, NUDT. His research interests include computational imaging, image processing, compressive sensing and sparse representation.

Dr. Yang received the New Scholar Award of Chinese Ministry of Education in 2012, the Youth Innovation Award and the Youth Outstanding Talent of NUDT in 2016.



Wei An received the Ph.D. degree from the National University of Defense Technology (NUDT), Changsha, China, in 1999. She was a Senior Visiting Scholar with the University of Southampton, Southampton, U.K., in 2016. She is currently a Professor with the College of Electronic Science and Technology, NUDT. She has authored or co-authored over 100 journal and conference publications. Her current research interests include signal processing and image processing.