

Learning Parallax Attention for Stereo Image Super-Resolution

Longguang Wang¹, Yingqian Wang¹, Zhengfa Liang², Zaiping Lin¹, Jungang Yang¹, Wei An¹, Yulan Guo^{1*}

¹College of Electronic Science and Technology, National University of Defense Technology, China

²National Key Laboratory of Science and Technology on Blind Signal Processing, China

{wanglongguang15, yulan.guo}@nudt.edu.cn

Abstract

Stereo image pairs can be used to improve the performance of super-resolution (SR) since additional information is provided from a second viewpoint. However, it is challenging to incorporate this information for SR since disparities between stereo images vary significantly. In this paper, we propose a parallax-attention stereo super-resolution network (PASSRnet) to integrate the information from a stereo image pair for SR. Specifically, we introduce a parallax-attention mechanism with a global receptive field along the epipolar line to handle different stereo images with large disparity variations. We also propose a new and the largest dataset for stereo image SR (namely, Flickr1024). Extensive experiments demonstrate that the parallax-attention mechanism can capture correspondence between stereo images to improve SR performance with a small computational and memory cost. Comparative results show that our PASSRnet achieves the state-of-the-art performance on the Middlebury, KITTI 2012 and KITTI 2015 datasets.

1. Introduction

Super-resolution (SR) aims to reconstruct high-resolution (HR) images from their low-resolution (LR) counterparts. Recovering an HR image from a single shot is a long-standing problem [1, 2, 3]. Recently, dual cameras are becoming increasingly popular in mobile phones and autonomous vehicles. It is already demonstrated that subpixel shifts contained in LR stereo images can be used to improve SR performance [4]. However, since disparities between stereo images can vary significantly for different baselines, focal lengths, depths and resolutions, it is highly challenging to incorporate stereo correspondence for SR.

Traditional multi-image SR methods [7, 8] use patch recurrence across images to obtain correspondence. However, these methods cannot exploit sub-pixel correspondence and their computational cost is high. Recent CNN-based frameworks [9, 10, 11] incorporate optical flow estimation and SR




Figure 1: Visual results achieved by bicubic interpolation, SRCNN [1], LapSRN [5], StereoSR [6] and our network for $2\times$ SR. These results are achieved on “test_image_002” of the KITTI 2015 dataset.

in unified networks to solve the video SR problem. However, these methods cannot be directly applied to stereo image SR since the disparity can be much larger than their receptive field.

Stereo matching has been investigated to obtain correspondence between a stereo image pair [12, 13, 14]. Recent CNN-based methods [15, 16, 17, 18] use 3D or 4D cost volumes in their networks to model long-range dependency between stereo image pairs. Intuitively, these CNN-based stereo matching methods can be integrated with SR to provide accurate correspondence. However, 4D cost volume based methods [15, 16] suffer from a high computational and memory burden, which is unbearable for stereo image SR. Although the efficiency of 3D cost volume based methods [17, 18] is improved, these methods cannot handle stereo images with large disparity variations since a fixed maximum disparity is used to construct a cost volume.

Recently, Jeon *et al.* proposed a stereo SR network (StereoSR) [6] to provide correspondence cues for SR using an image stack. Specifically, the image stack is obtained by concatenating the left image and the images generated by shifting the right image with different intervals. A direct mapping between parallax shifts and an HR image is then obtained. However, the flexibility of this method for different sensors and scenes is limited since the largest allowed disparity is fixed (*i.e.*, 64 in [6]) in this algorithm.

In this paper, we propose a parallax-attention stereo SR network (PASSRnet) to incorporate stereo correspondence for the SR task. Given a stereo image pair, a residual atrous spatial pyramid pooling (ASPP) module is first used to generate multi-scale features. Then, these features are fed to a parallax-attention module (PAM) to capture stereo correspondence. 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. Consequently, our PAM can capture global correspondence while maintaining high flexibility. Afterwards, attention-driven feature aggregation is performed to update the features of the left image. Finally, these features are used to generate the SR result. Ablation study is performed on the KITTI 2015 dataset to test our PASSRnet. Comparative experiments are further conducted on the Middlebury, KITTI 2012 and KITTI 2015 datasets to demonstrate the superior performance of our network (as shown in Fig. 1).

Our main contributions can be summarized as follows: 1) We propose a PASSRnet for SR by incorporating stereo correspondence; 2) We introduce a generic parallax-attention mechanism with a global receptive field along the epipolar line to handle different stereo images with large disparity variations. It is demonstrated that reliable correspondence can be efficiently generated by the parallax-attention mechanism for the improvement of SR performance; 3) We propose a new dataset, namely Flickr1024, for the training of stereo image SR networks. The Flickr1024 dataset consists of 1024 high-quality stereo image pairs and covers diverse scenes; 4) Our PASSRnet achieves the state-of-the-art performance as compared to recent single image SR and stereo image SR methods.

2. Related Work

In this section, we briefly review several major works for SR and long-range dependency learning.

2.1. Super-resolution

Single Image SR Since the seminal work of super-resolution convolutional neural network (SRCNN) [1], learning-based methods have dominated the research of single image SR. Kim *et al.* [19] proposed a very deep super-resolution network (VDSR) with 20 convolutional layers. Tai *et al.* [20] developed a deep recursive residual network (DRRN) to control model parameters. Recently, Zhang *et al.* [21] proposed a residual dense network (RDN) to facilitate effective feature learning through a contiguous memory mechanism.

Video SR Liao *et al.* [22] introduced the first CNN for video SR. They performed motion compensation to generate an ensemble of SR-drafts, and then employed a CNN to reconstruct HR frames from the ensemble. Caballero *et al.* [9] proposed an end-to-end video SR framework by incor-

porating a motion compensation module with an SR module. Tao *et al.* [10] integrated an encoder-decoder network with LSTM to fully use temporal correspondence. This architecture further facilitates the extraction of temporal context. Since correspondence between adjacent frames mainly exists within a local region, video SR methods focus on the exploitation of local dependency. Therefore, they cannot be directly applied to stereo image SR due to the non-local and long-range dependency in stereo images.

Light-field Image SR Light-field imaging can capture additional angular information of light at the cost of spatial resolution. To enhance spatial resolution, Yoon *et al.* [23] introduced the first light-field convolutional neural network (LFCNN). Yuan *et al.* [24] proposed a CNN framework with a single image SR module and an epipolar plane image enhancement module. To model the correspondence between images of adjacent sub-apertures, Wang *et al.* [25] developed a bidirectional recurrent CNN. Their network uses an implicit multi-scale feature fusion scheme to accumulate contextual information for SR. Note that, these methods are specifically proposed for light-field imaging with short baselines. Since stereo imaging usually has a much larger baseline than light-field imaging, these methods are unsuitable for stereo image SR.

Stereo Image SR Bhavsar *et al.* [26] argued that image SR and HR depth estimation are intertwined under stereo setting. Therefore, they proposed an integrated approach to jointly estimate the SR image and HR disparity from LR stereo images. Recently, Jeon *et al.* [6] proposed a StereoSR to employ parallax prior. 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 tensor is then fed to a plain CNN to generate the SR result by detecting similar patches within the disparity channel. However, StereoSR cannot handle different stereo images with large disparity variations since the number of shifted right images is fixed.

2.2. Long-range Dependency Learning

To handle different stereo images with varying disparities for SR, long-range dependency in stereo images should be captured. In this section, we review two types of methods for long-range dependency learning.

Cost Volume Cost volume is widely applied in stereo matching [15, 16, 17] and optical flow estimation [27, 28]. For stereo matching, several methods [15, 16] use naive concatenation to construct 4D cost volumes. These methods concatenate left feature maps with their corresponding right feature maps across all disparities to obtain a 4D cost volume (*i.e.*, $\text{height} \times \text{width} \times \text{disparity} \times \text{channel}$). Then, 3D CNNs are usually used for matching cost learning. However, learning matching costs from 4D cost volumes suffers from a high computational and memory burden. To

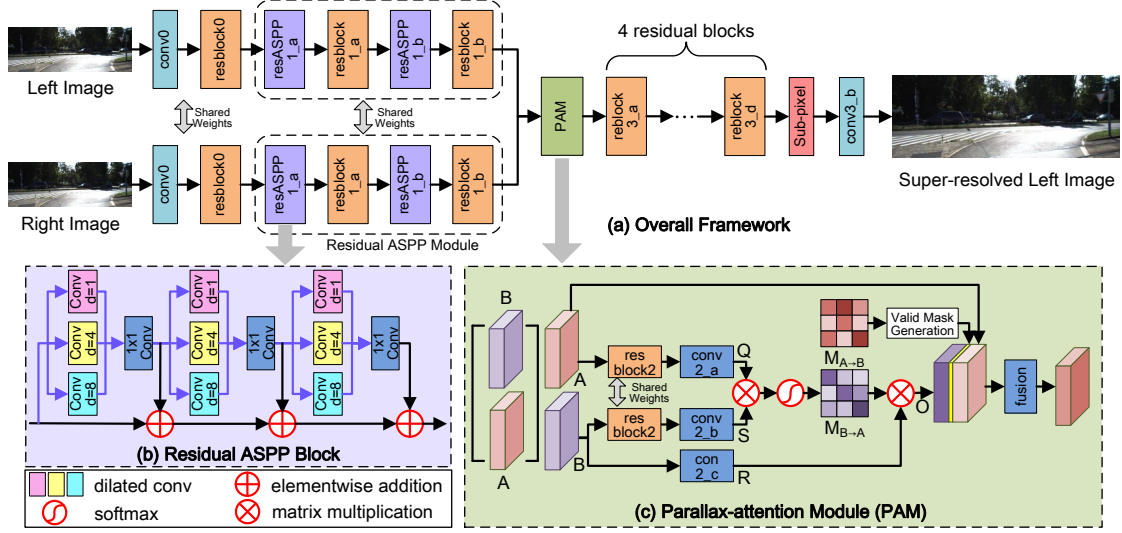


Figure 2: An overview of our PASSRnet.

achieve higher efficiency, dot product is used to reduce feature dimension [17, 18], resulting in 3D cost volumes (*i.e.*, height \times width \times disparity). However, due to the fixed maximum disparity in 3D cost volumes, these methods are unable to handle different stereo image pairs with large disparity variations.

Self-attention Mechanisms Attention mechanisms have been widely used to capture long-range dependency [29, 30]. For self-attention mechanisms [31, 32, 33], a weighted sum of all positions in spatial and/or temporal domain is calculated as the response at a position. Through matrix multiplication, self-attention mechanisms can capture the interaction between any two positions. Consequently, long-range dependency can be modeled with a small increase in computational and memory cost. Self-attention mechanisms have been successfully applied in image modeling [32] and semantic segmentation [33]. Recent non-local networks [34, 35] share a similar idea and can be considered as a generalization of self-attention mechanisms. Note that, since self-attention mechanisms model dependency across the whole image, directly applying these mechanisms to stereo image SR involves unnecessary calculations.

Inspired by self-attention mechanisms, we develop a parallax-attention mechanism to model global dependency in stereo images. Compared to cost volumes, our parallax-attention mechanism is more flexible and efficient. Compared to self-attention mechanisms, our parallax-attention mechanism takes full use of epipolar constraints to reduce search space and improve efficiency. Moreover, the parallax-attention mechanism enforces our network to focus on the most similar feature rather than collecting all similar features for correspondence generation. It is demonstrated that the parallax-attention mechanism can generate reliable

correspondence to improve SR performance (Section 4.3.1).

3. Method

Our PASSRnet takes a stereo image pair as input and super-resolves the left image. The architecture of our PASSRnet is shown in Fig. 2 and Table 1.

3.1. Residual Atrous Spatial Pyramid Pooling (ASPP) Module

Feature representation with rich context information is important for correspondence estimation [16]. Therefore, both large receptive field and multi-scale feature learning are required to obtain a discriminative representation. To this end, 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. 2 (a), our residual ASPP module is constructed by alternately cascading a residual ASPP block with a residual block. Input features are first fed to a residual ASPP block to generate multi-scale features. These 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. 2 (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. 4.3.1.

Table 1: The detailed architecture of our PASSRnet. LReLU represents leaky ReLU with a leakage factor of 0.1, dila stands for dilation rate, \otimes denotes batch-wise matrix multiplication, and s is the upscaling factor.

Name	Setting	Input	Output
input		$H \times W \times 3$	$H \times W \times 3$
conv0	3×3 LReLU	$H \times W \times 3$	$H \times W \times 64$
resblock0	$\begin{bmatrix} 3 \times 3 \\ 3 \times 3 \end{bmatrix}$	$H \times W \times 64$	$H \times W \times 64$
<i>Residual ASPP Module</i>			
resASPP 1.a	$\begin{bmatrix} 3 \times 3 & 3 \times 3 & 3 \times 3 \\ \text{LReLU, LReLU, LReLU} \\ \text{dila=1} & \text{dila=4} & \text{dila=8} \\ 1 \times 1 \end{bmatrix} \times 3$	$H \times W \times 64$	$H \times W \times 64$
resblock 1.a	$\begin{bmatrix} 3 \times 3 \\ 3 \times 3 \end{bmatrix}$	$H \times W \times 64$	$H \times W \times 64$
resASPP 1.b	$\begin{bmatrix} 3 \times 3 & 3 \times 3 & 3 \times 3 \\ \text{LReLU, LReLU, LReLU} \\ \text{dila=1} & \text{dila=4} & \text{dila=8} \\ 1 \times 1 \end{bmatrix} \times 3$	$H \times W \times 64$	$H \times W \times 64$
resblock 1.b	$\begin{bmatrix} 3 \times 3 \\ 3 \times 3 \end{bmatrix}$	$H \times W \times 64$	$H \times W \times 64$
<i>Parallax-Attention Module</i>			
resblock2	$\begin{bmatrix} 3 \times 3 \\ 3 \times 3 \end{bmatrix}$	$H \times W \times 64$	$H \times W \times 64$
conv2.a	1×1	$H \times W \times 64$	$H \times W \times 64$
conv2.b	1×1 , reshape	$H \times W \times 64$	$H \times 64 \times W$
conv2.c	1×1	$H \times W \times 64$	$H \times W \times 64$
att_map	conv2.a \otimes conv2.b	$H \times W \times 64$ $H \times 64 \times W$	$H \times W \times W$
mult	att_map \otimes conv2.c	$H \times W \times W$ $H \times W \times 64$	$H \times W \times 64$
fusion	1×1	$H \times W \times 129$	$H \times W \times 64$
<i>CNN</i>			
resblock3 $\times 4$	$\begin{bmatrix} 3 \times 3 \\ 3 \times 3 \end{bmatrix}$	$H \times W \times 64$	$H \times W \times 64$
sub-pixel	1×1 , pixel shuffle	$H \times W \times 64$	$sH \times sW \times 64$
conv3.b	3×3	$sH \times sW \times 64$	$sH \times sW \times 3$

3.2. Parallax-attention Module (PAM)

Inspired by self-attention mechanisms [32, 33], we develop PAM to capture global correspondence in stereo images. Our PAM efficiently integrates the information from a stereo image pair.

Parallax-attention Mechanism The architecture of our PAM is illustrated in Fig. 2 (c). Given two feature maps $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{H \times W \times C}$, they are fed to a transition residual block to generate \mathbf{A}_0 and \mathbf{B}_0 . Then, \mathbf{A}_0 is fed to a 1×1 convolution layer to produce a query feature map $\mathbf{Q} \in \mathbb{R}^{H \times W \times C}$. Meanwhile, \mathbf{B}_0 is fed to another 1×1 convolution layer to produce $\mathbf{S} \in \mathbb{R}^{H \times W \times C}$, which is then reshaped to $\mathbb{R}^{H \times C \times W}$. Batch-wise matrix multiplication is then performed between \mathbf{Q} and \mathbf{S} and a softmax layer is applied, resulting in a parallax attention map $\mathbf{M}_{\mathbf{B} \rightarrow \mathbf{A}} \in \mathbb{R}^{H \times W \times W}$. For more details, please refer to the supplemental material. Next, \mathbf{B} is fed to a 1×1 convolution to generate $\mathbf{R} \in \mathbb{R}^{H \times W \times C}$, which is further multiplied by $\mathbf{M}_{\mathbf{B} \rightarrow \mathbf{A}}$ to produce features $\mathbf{O} \in \mathbb{R}^{H \times W \times C}$. As a weighted sum of features at all possible disparities, \mathbf{O} is then integrated with




Figure 3: Visual comparison between parallax-attention maps $\mathbf{M}_{\text{right} \rightarrow \text{left}}$ generated by our PAM and the groundtruth. These attention maps (100×100) correspond to the regions (1×100) marked by blue and pink strokes in the left image.

its corresponding local features \mathbf{A} . Since PAM can gradually focus on the features at accurate disparities using feature similarities, correspondence can then be captured. Note that, once $\mathbf{M}_{\mathbf{B} \rightarrow \mathbf{A}}$ is ready, \mathbf{A} and \mathbf{B} are exchanged to produce $\mathbf{M}_{\mathbf{A} \rightarrow \mathbf{B}}$ for valid mask generation (as described below). Finally, stacked features and a valid mask are fed to a 1×1 convolution layer for feature fusion.

Different from self-attention mechanisms [32, 33], our parallax-attention mechanism enforces our network to focus on the most similar feature along the epipolar line rather than collecting all similar features, resulting in sparse attention maps. A comparison between parallax-attention maps generated by our PAM and the groundtruth is shown in Fig. 3. Note that, $\mathbf{M}_{\text{right} \rightarrow \text{left}}(i, j, k)$ represents the contribution of position (i, k) in the right image to position (i, j) in the left image. Consequently, the patterns in an attention map can reflect the correspondence between stereo pairs and also encode disparity information. For more details, please refer to the supplemental material. It can be observed that our PAM produces patterns similar to the groundtruth, which indicates that reliable stereo correspondence can be captured by our PAM. 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 [36]. Therefore, a transition block is introduced in our PAM to alleviate this problem. The effectiveness of the transition block is demonstrated in Sec. 4.3.1.

Left-right Consistency & Cycle Consistency Given deep features extracted from an LR stereo image pair ($\mathbf{I}_{\text{left}}^L$ and $\mathbf{I}_{\text{right}}^L$), two parallax-attention maps ($\mathbf{M}_{\text{left} \rightarrow \text{right}}$ and $\mathbf{M}_{\text{right} \rightarrow \text{left}}$) can be generated by PAM. Ideally, the following left-right consistency can be obtained if our PAM captures accurate correspondence:

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

where \otimes denotes batch-wise matrix multiplication. Based



Figure 4: Visualization of valid masks. Two left images and their occluded regions (i.e., yellow regions) are illustrated.

on Eq. (1), we can further derive a cycle consistency:

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

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

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

Here, we introduce left-right consistency and cycle consistency to regularize the training of our PAM for the generation of reliable and consistent correspondence.

Valid Masks Since left-right consistency and cycle consistency do not hold for occluded regions, we use an occlusion detection method to generate valid masks. We only enforce consistency on valid regions. In the parallax-attention map generated by our PAM (e.g., $\mathbf{M}_{left \rightarrow right}$), it is observed that pixels in occluded regions are usually assigned with small weights. Therefore, a valid mask $\mathbf{V}_{left \rightarrow right} \in \mathbb{R}^{H \times W}$ can be obtained by:

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

where τ is a threshold (empirically set to 0.1) and W is the width of stereo images. Two examples of valid masks are shown in Fig. 4. According to the parallax-attention mechanism, $\mathbf{M}_{left \rightarrow right}(i, k, j)$ represents the contribution of position (i, j) in the left image to position (i, k) in the right image. Since occluded pixels in the left image cannot find their correspondences in the right image, their values (i.e., $\sum_{k \in [1, W]} \mathbf{M}_{left \rightarrow right}(i, k, j)$) are usually low. Thus, we consider these pixels as occluded ones. In practice, we use several morphological operations to handle isolated pixels and holes in valid masks. Note that, occluded regions in the left image cannot obtain additional information from the right image. Therefore, valid mask $\mathbf{V}_{left \rightarrow right}$ is further used to guide feature fusion, as shown in Fig. 2 (c).

3.3. Losses

We design four losses for the training of our PASSRnet. Other than an SR loss, we introduce three additional losses,

including photometric loss, smoothness loss and cycle loss, to help the network to fully exploit the correspondence between stereo images. The overall loss function is formulated as:

$$\mathcal{L} = \mathcal{L}_{SR} + \lambda(\mathcal{L}_{photometric} + \mathcal{L}_{smooth} + \mathcal{L}_{cycle}), \quad (5)$$

where λ is empirically set to 0.005. The performance of our network with different losses will be analyzed in Sec. 4.3.2.

SR Loss The mean square error (MSE) loss is used as the SR loss:

$$\mathcal{L}_{SR} = \|\mathbf{I}_{left}^{SR} - \mathbf{I}_{left}^H\|_2^2, \quad (6)$$

where \mathbf{I}_{left}^{SR} and \mathbf{I}_{left}^H represent the SR result and HR groundtruth of the left image, respectively.

Photometric Loss Since collecting a large stereo dataset with densely labeled groundtruth disparities is highly challenging, we train our PAM in an unsupervised manner. Note that, if the groundtruth disparities are available, we can generate the groundtruth attention maps accordingly (see the supplemental material for more details) and train our PAM in a supervised manner. Following [37], we introduce a photometric loss using the mean absolute error (MAE) loss. Note that, since the left-right consistency defined in Eq. (1) only holds in non-occluded regions, we introduce a photometric loss as:

$$\begin{aligned} \mathcal{L}_{photometric} = & \sum_{p \in \mathbf{V}_{left \rightarrow right}} \|\mathbf{I}_{left}^L(p) - (\mathbf{M}_{right \rightarrow left} \otimes \mathbf{I}_{right}^L)(p)\|_1 \\ & + \sum_{p \in \mathbf{V}_{right \rightarrow left}} \|\mathbf{I}_{right}^L(p) - (\mathbf{M}_{left \rightarrow right} \otimes \mathbf{I}_{left}^L)(p)\|_1, \end{aligned} \quad (7)$$

where p represents a pixel with a valid mask value.

Smoothness Loss To generate accurate and consistent attention in textureless regions, a smoothness loss is defined on the attention maps $\mathbf{M}_{left \rightarrow right}$ and $\mathbf{M}_{right \rightarrow left}$:

$$\begin{aligned} \mathcal{L}_{smooth} = & \sum_{\mathbf{M}} \sum_{i, j, k} (\|\mathbf{M}(i, j, k) - \mathbf{M}(i+1, j, k)\|_1 + \\ & \|\mathbf{M}(i, j, k) - \mathbf{M}(i, j+1, k+1)\|_1), \end{aligned} \quad (8)$$

where $\mathbf{M} \in \{\mathbf{M}_{left \rightarrow right}, \mathbf{M}_{right \rightarrow left}\}$. The first and second terms in Eq. (8) are used to achieve vertical and horizontal attention consistency, respectively.

Cycle Loss In addition to photometric loss and smoothness loss, we further introduce a cycle loss to achieve cycle consistency. Since $\mathbf{M}_{left \rightarrow right \rightarrow left}$ and $\mathbf{M}_{right \rightarrow left \rightarrow right}$ in Eq. (2) can be considered as identity matrices, we design a cycle loss as:

$$\begin{aligned} \mathcal{L}_{cycle} = & \sum_{p \in \mathbf{V}_{left \rightarrow right}} \|\mathbf{M}_{left \rightarrow right \rightarrow left}(p) - I(p)\|_1 + \\ & \sum_{p \in \mathbf{V}_{right \rightarrow left}} \|\mathbf{M}_{right \rightarrow left \rightarrow right}(p) - I(p)\|_1, \end{aligned} \quad (9)$$

where $I \in \mathbb{R}^{H \times W \times W}$ is a stack of H identity matrices.

Table 2: 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 PAM	Left-Right	25.28	0.771	1.32M	135ms
PASSRnet without transition residual block	Left-Right	25.36	0.773	1.34M	160ms
PASSRnet	Left-Right	25.43	0.776	1.42M	176ms

4. Experimental Results

In this section, we first introduce the datasets and implementation details, and then conduct ablation experiments to test our network. We further compare our network to recent single image SR and stereo image SR methods.

4.1. Datasets

For training, we followed [6] and downsampled 60 Middlebury [38] images by a factor of 2 to generate HR images. We further collected 1024 stereo images from Flickr to construct a new Flickr1024 dataset. This dataset was used as the augmented training data for our PASSRnet. Please see the supplemental material for more details about the Flickr1024 dataset. For test, we used 5 images from the Middlebury dataset, 20 images from the KITTI 2012 dataset [39] and 20 images from the KITTI 2015 dataset [40] as benchmark datasets. We further collected 10 close-shot stereo images (with disparities larger than 200) from Flickr to test the flexibility of our network to large disparity variations. For validation, we selected another 20 images from the KITTI 2012 dataset.

4.2. Implementation Details

During the training phase, 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 also cropped. The horizontal patch size was increased 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 to maintain epipolar constraints. We used peak signal-to-noise ratio (PSNR) and structural similarity index (SSIM) to test SR performance. Similar to [6], we cropped borders to achieve fair comparison.

Our PASSRnet was implemented in Pytorch on a PC with an Nvidia GTX 1080Ti GPU. All models were optimized using the Adam method [41] 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 consistent improvement.

4.3. Ablation Study

In this section, we present ablation experiments to justify our design choices, including the network architecture and the losses.

4.3.1 Network Architecture

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 performance improvement, we removed PAM 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. Results achieved on the KITTI 2015 dataset are listed in Table 2.

Compared to the original PASSRnet, the network trained with single images suffers a decrease of 0.16 dB (25.43 to 25.27) in PSNR. Further, 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 images can be used to improve the performance of PASSRnet.

Residual ASPP Module Residual ASPP module is used in our network to extract multi-scale features. To demonstrate the effectiveness of residual ASPP, 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 2, we can see that SR performance benefits from both residual connections and atrous convolutions. If residual connections are removed, the PSNR value is decreased from 25.43 dB to 25.40 dB. That is because, residual connections enable our residual ASPP module to extract features at more scales,

resulting in more robust feature representations. Furthermore, if atrous convolutions are replaced by ordinary ones, the PSNR value is decreased from 25.43 dB to 25.38 dB. That is because, 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 PAM is introduced to integrate the information from stereo images. To demonstrate its effectiveness, we introduced a variant by removing PAM and directly stacking the output features of the residual ASPP module. It can be observed from Table 2 that the PSNR value is decreased from 25.43 dB to 25.28 dB if PAM is removed. That is because, long spatial distance between local features in the left image and their dependency in the right image hinders plain CNNs to integrate these features effectively.

Transition Block in PAM Transition block in PAM is introduced to alleviate the training conflict in shared layers. To demonstrate the effectiveness of transition block, we removed it from our PAM and retrained the network. It can be observed from Table 2 that the PSNR value is decreased from 25.43 dB to 25.36 dB if the transition block is removed. That is because, the transition block enhances task-specific feature learning in our PAM and alleviates training conflict in shared layers. Therefore, more representative features can be learned in shared layers.

PAM vs. Cost Volume Cost volume and 3D convolutions are commonly used to obtain stereo correspondence [15, 16]. To demonstrate the efficiency of our PAM in stereo correspondence generation, we replaced PAM with a 4D cost volume and two 3D convolutional layers ($3 \times 3 \times 3$). It can be observed from Table 3 that our PAM has less than half of the parameters in the cost volume formation. Moreover, our PAM achieves superior computational efficiency, with FLOPs being reduced by over 150 times. With PAM, our PASSRnet achieves better SR performance (*i.e.*, PSNR value is increased from 25.23 dB to 25.43 dB) and 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.

4.3.2 Losses

To test the effectiveness of our losses, we retrained PASSRnet using different losses.

It can be observed from Table 4 that the PSNR value of our PASSRnet is decreased from 25.43 to 25.35 if PASSRnet is trained with only SR loss. That is because, with only this loss, our PAM learns to collect all similar features along the epipolar line and cannot focus on the most

Table 3: 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×	1×	25.43	0.776
Cost Volume	221K	151×	1.5×	25.23	0.768

Table 4: 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

similar feature to provide accurate correspondence. Further, the performance is gradually improved if photometric loss, smoothness loss and cycle loss are added. That is because, these losses encourage our PAM to generate reliable and consistent correspondence. Overall, our PASSRnet achieves the best performance (*i.e.*, PSNR=25.43 dB and SSIM=0.776) when it is trained with all these losses.

4.4. Comparison to State-of-the-arts

We compared our PASSRnet to a number of CNN-based SR methods on three benchmark datasets. Recent single image SR methods under comparison include SRCNN [1], VDSR [19], DRCN [42], LapSRN [5] and DRRN [20]. We also compared our PASSRnet to the latest stereo image SR method StereoSR [6]. The codes provided by the authors of these methods were used to conduct experiments. Note that, similar to [6, 43], EDSR [44], RDN [21] and D-DBPN [45] 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 5. It can be observed that our PASSRnet achieves the best performance on the Middlebury, KITTI 2012 and KITTI 2015 datasets. Specifically, compared to single image SR methods, our PASSRnet outperforms the second best approach (*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 reliable correspondence can be captured by our parallax-attention mechanism.

Qualitative Results Figure 5 illustrates the qualitative results achieved on two scenarios. It can be observed from zoom-in regions that single image SR methods cannot re-

Table 5: 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 $2\times$ SR results of StereoSR are presented on the KITTI 2012 and KITTI 2015 datasets since a $4\times$ SR model is unavailable.

Dataset	Scale	Single Image SR					Stereo Image SR	
		SRCNN [1]	VDSR [19]	DRCN [42]	LapSRN [5]	DRRN [20]	StereoSR [6]	Ours
Middlebury (5 images)	$\times 2$	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
	$\times 4$	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)	$\times 2$	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
	$\times 4$	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)	$\times 2$	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
	$\times 4$	24.68/0.744	25.01/0.760	25.04/0.759	25.03/0.760	25.05/0.756	-	25.43/0.776




Figure 5: Visual comparison for $2\times$ SR. These results are achieved on “test_image_013” of the KITTI 2012 dataset and “test_image_019” of the KITTI 2015 dataset.

Table 6: Comparison between our PASSRnet and StereoSR [6] on stereo images with different resolutions for $2\times$ SR.

Resolution	StereoSR [6]		Ours	
	PSNR	FLOPs	PSNR	FLOPs
High (500×500)	39.27	$1\times$	41.45 ($\uparrow 2.18$)	$0.57\times$
Middle (100×100)	34.21	$1\times$	35.04 ($\uparrow 0.83$)	$0.58\times$
Low (20×20)	29.48	$1\times$	29.88 ($\uparrow 0.40$)	$0.36\times$

cover reliable details. In contrast, our PASSRnet uses stereo correspondence to produce finer details with fewer artifacts, such as the railings and stripe in Fig. 5. Compared to StereoSR, our PASSRnet explicitly captures stereo correspondence for SR. Consequently, superior visual performance is achieved.

Flexibility We further tested the flexibility of our PASSRnet and StereoSR [6] with respect to large disparity variations. Results achieved on images with different resolutions are shown in Table 6. More results under different baselines and depths are available in the supplemental material. It can be observed that our PASSRnet is significantly bet-

ter 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.

5. Conclusion

In this paper, we propose a parallax-attention stereo super-resolution network (PASSRnet) to incorporate stereo correspondence for the SR task. Our PASSRnet introduces a parallax-attention mechanism with global receptive field to handle different stereo images with large disparity variations. We also introduce a new and the largest dataset for stereo image SR. It is demonstrated that our PASSRnet can effectively capture stereo correspondence for the improvement of SR performance. Comparison to recent single image SR and stereo image SR methods has shown that our network achieves the state-of-the-art performance.

References

- [1] 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.
- [2] Wenzhe Shi, Jose Caballero, Ferenc Huszar, Johannes Totz, Andrew P. Aitken, Rob Bishop, Daniel Rueckert, and Zehan Wang. Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In *CVPR*, pages 1874–1883, 2016.
- [3] Zheng Hui, Xiumei Wang, and Xinbo Gao. Fast and accurate single image super-resolution via information distillation network. In *CVPR*, 2018.
- [4] Sung Cheol Park, Min Kyu Park, and Moon Gi Kang. Super-resolution image reconstruction: a technical overview. *IEEE signal processing magazine*, 20(3):21–36, 2003.
- [5] 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.
- [6] 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*, 2018.
- [7] Matan Protter, Michael Elad, Hiroyuki Takeda, and Peyman Milanfar. Generalizing the nonlocal-means to super-resolution reconstruction. *IEEE Trans. Image Processing*, 18(1):36–51, 2009.
- [8] Hiroyuki Takeda, Peyman Milanfar, Matan Protter, and Michael Elad. Super-resolution without explicit sub-pixel motion estimation. *IEEE Trans. Image Processing*, 18(9):1958–1975, 2009.
- [9] Jose Caballero, Christian Ledig, Andrew P. Aitken, Alejandro Acosta, Johannes Totz, Zehan Wang, and Wenzhe Shi. Real-time video super-resolution with spatio-temporal networks and motion compensation. In *CVPR*, pages 2848–2857, 2017.
- [10] Xin Tao, Hongyun Gao, Renjie Liao, Jue Wang, and Jiaya Jia. Detail-revealing deep video super-resolution. In *ICCV*, pages 4482–4490, 2017.
- [11] Longguang Wang, Yulan Guo, Zaiping Lin, Xinpu Deng, and Wei An. Learning for video super-resolution through HR optical flow estimation. In *ACCV*, 2018.
- [12] Stephen T. Barnard and Martin A. Fischler. Computational stereo. *ACM Comput. Surv.*, 14(4):553–572, 1982.
- [13] Dan. Scharstein and R. Szeliski. A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *International Journal of Computer Vision*, 47(1-3), 2002.
- [14] Wenjie Luo, Alexander G. Schwing, and Raquel Urtasun. Efficient deep learning for stereo matching. In *CVPR*, pages 5695–5703, 2016.
- [15] 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.
- [16] Jia-Ren Chang and Yong-Sheng Chen. Pyramid stereo matching network. In *CVPR*, 2018.
- [17] 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*, 2018.
- [18] 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*, 2018.
- [19] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Accurate image super-resolution using very deep convolutional networks. In *CVPR*, pages 1646–1654, 2016.
- [20] Ying Tai, Jian Yang, and Xiaoming Liu. Image super-resolution via deep recursive residual network. In *CVPR*, pages 2790–2798, 2017.
- [21] Yulun Zhang, Yapeng Tian, Yu Kong, Bineng Zhong, and Yun Fu. Residual dense network for image super-resolution. In *CVPR*, 2018.
- [22] Renjie Liao, Xin Tao, Ruiyu Li, Ziyang Ma, and Jiaya Jia. Video super-resolution via deep draft-ensemble learning. In *ICCV*, pages 531–539, 2015.
- [23] Youngjin Yoon, Hae-Gon Jeon, Donggeun Yoo, Joon-Young Lee, and In So Kweon. Learning a deep convolutional network for light-field image super-resolution. In *ICCV Workshops*, pages 57–65, 2015.
- [24] Yan Yuan, Ziqi Cao, and Lijuan Su. Light-field image super-resolution using a combined deep CNN based on EPI. *IEEE Signal Process. Lett.*, 25(9):1359–1363, 2018.
- [25] 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 Processing*, 27(9):4274–4286, 2018.
- [26] Arnav V. Bhavsar and A. N. Rajagopalan. Resolution enhancement in multi-image stereo. *IEEE Trans. Pattern Anal. Mach. Intell.*, 32(9):1721–1728, 2010.
- [27] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. PWC-Net: CNNs for optical flow using pyramid, warping, and cost volume. In *CVPR*, 2017.
- [28] Jia Xu, Rene Ranftl, and Vladlen Koltun. Accurate optical flow via direct cost volume processing. In *CVPR*, pages 5807–5815, 2017.
- [29] Karol Gregor, Ivo Danihelka, Alex Graves, Danilo Jimenez Rezende, and Daan Wierstra. DRAW: A recurrent neural network for image generation. In *ICML*, volume 37, pages 1462–1471, 2015.
- [30] 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*, volume 37, pages 2048–2057, 2015.
- [31] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NIPS*, pages 6000–6010, 2017.

- [32] Han Zhang, Ian J. Goodfellow, Dimitris N. Metaxas, and Augustus Odena. Self-attention generative adversarial networks. In *NIPS*, 2018.
- [33] Jun Fu, Jing Liu, Haijie Tian, Zhiwei Fang, and Hanqing Lu. Dual attention network for scene segmentation. *arXiv preprint arXiv:1809.02983*, 2018.
- [34] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *CVPR*, 2018.
- [35] Ding Liu, Bihan Wen, Yuchen Fan, Chen Change Loy, and Thomas S. Huang. Non-local recurrent network for image restoration. In *NIPS*, 2018.
- [36] Ozan Sener and Vladlen Koltun. Multi-task learning as multi-objective optimization. In *NIPS*, 2018.
- [37] Clément Godard, Oisín Mac Aodha, and Gabriel J. Brostow. Unsupervised monocular depth estimation with left-right consistency. In *CVPR*, pages 6602–6611, 2017.
- [38] 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*, volume 8753, pages 31–42, 2014.
- [39] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the KITTI vision benchmark suite. In *CVPR*, pages 3354–3361, 2012.
- [40] Moritz Menze and Andreas Geiger. Object scene flow for autonomous vehicles. In *CVPR*, pages 3061–3070, 2015.
- [41] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- [42] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Deeply-recursive convolutional network for image super-resolution. In *CVPR*, pages 1637–1645, 2016.
- [43] Namhyuk Ahn, Byungkoo Kang, and Kyung-Ah Sohn. Fast, accurate, and lightweight super-resolution with cascading residual network. In *ECCV*, 2018.
- [44] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single image super-resolution. In *CVPR Workshops*, 2017.
- [45] Muhammad Haris, Greg Shakhnarovich, and Norimichi Ukita. Deep back-projection networks for super-resolution. In *CVPR*, 2018.