

Spatial-Angular Interaction for Light Field Image Super-Resolution

Yingqian Wang¹, Longguang Wang¹, Jungang Yang¹,
Wei An¹, Jingyi Yu², Yulan Guo^{1,3}

¹National University of Defense Technology,
²ShanghaiTech University, ³Sun Yat-sen University.
{wangyingqian16, yangjungang, yulan.guo}@nudt.edu.cn

<https://github.com/YingqianWang/LF-InterNet>

Abstract. Light field (LF) cameras record both intensity and directions of light rays, and capture scenes from a number of viewpoints. Both information within each perspective (i.e., spatial information) and among different perspectives (i.e., angular information) is beneficial to image super-resolution (SR). In this paper, we propose a spatial-angular interactive network (namely, LF-InterNet) for LF image SR. Specifically, spatial and angular features are first separately extracted from input LFs, and then repetitively interacted to progressively incorporate spatial and angular information. Finally, the interacted features are fused to super-resolve each sub-aperture image. Experimental results demonstrate the superiority of LF-InterNet over the state-of-the-art methods, i.e., our method can achieve high PSNR and SSIM scores with low computational cost, and recover faithful details in the reconstructed images.

Keywords: Light Field Imaging, Super-Resolution, Feature Decoupling, Spatial-Angular Interaction

1 Introduction

Light field (LF) cameras provide multiple views of a scene, and thus enable many attractive applications such as post-capture refocusing [1], depth sensing [2], saliency detection [3,4], and de-occlusion [5]. However, LF cameras face a trade-off between spatial and angular resolution. That is, they either provide dense angular samplings with low image resolution (e.g., Lytro and RayTrix), or capture high-resolution (HR) sub-aperture images (SAIs) with sparse angular samplings (e.g., camera arrays [6,7]). Consequently, many efforts have been made to improve the angular resolution through LF reconstruction [8,9,10,11], or the spatial resolution through LF image super-resolution (SR) [12,13,14,15,16,17]. In this paper, we focus on the LF image SR problem, namely, to reconstruct HR SAIs from their corresponding low-resolution (LR) SAIs.

Image SR is a long-standing problem in computer vision. To achieve high reconstruction performance, SR methods need to incorporate as much useful information as possible from LR inputs. In the area of single image SR (SISR),

good performance can be achieved by fully exploiting the neighborhood context (i.e., spatial information) in an image. Using the spatial information, SISR methods [18,19,20,21,22,23,24] can successfully hallucinate missing details. In contrast, LFs record scenes from multiple views, and the complementary information among different views (i.e., angular information) can be used to further improve the performance of LF image SR.

However, due to the complicated 4D structures of LFs, many LF image SR methods fail to fully exploit both the angular information and the spatial information, resulting in inferior SR performance. Specifically, in [25,26,27], SAIs are first super-resolved separately using SISR methods [18,20], and then fine-tuned together to incorporate the angular information. The angular information is ignored by these two-stage methods [25,26,27] during their upsampling process. In [15,13], only part of SAIs are used to super-resolve one view, and the angular information in these discarded views is not incorporated. In contrast, Rossi et al. [14] proposed a graph-based method to consider all angular views in an optimization process. However, this method [14] cannot fully use the spatial information, and is inferior to recent deep learning-based SISR methods [20,21,22].

Since spatial and angular information are highly coupled in 4D LFs and contribute to LF image SR in different manners, it is difficult for networks to perform well using these coupled information directly. In this paper, we propose a spatial-angular interactive network (i.e., LF-InterNet) to efficiently use spatial and angular information for LF image SR. Specifically, we design two convolutions (i.e., spatial/angular feature extractor) to extract and decouple spatial and angular features from input LFs. Then, we develop LF-InterNet to progressively interact the extracted features. Thanks to the proposed spatial-angular interaction mechanism, information in an LF can be effectively used in an efficient manner, and the SR performance is significantly improved. We perform extensive ablation studies to demonstrate the effectiveness of our model, and compare our method with state-of-the-art SISR and LF image SR methods from different perspectives, which demonstrate the superiority of our LF-InterNet.

2 Related Works

Single Image SR. In the area of SISR, deep learning-based methods have been extensively explored. Readers can refer to recent surveys [28,29,30] for more details in SISR. Here, we only review several milestone works. Dong et al. [18] proposed the first CNN-based SR method (i.e., SRCNN) by cascading 3 convolutional layers. Although SRCNN is shallow and simple, it achieves significant improvements over traditional SR methods [31,32,33]. Afterwards, SR networks became increasingly deep and complex, and thus more powerful in spatial information exploitation. Kim et al. [19] proposed a very deep SR network (i.e., VDSR) with 20 convolutional layers. Global residual learning is applied to VDSR to avoid slow convergence. Lim et al. [20] proposed an enhanced deep SR network (i.e., EDSR) and achieved substantial performance improvements by applying both local and global residual learning. Subsequently, Zhang et al. [34]

proposed a residual dense network (i.e., RDN) by combining residual connection and dense connection. RDN can fully extract hierarchical features for image SR, and thus achieve further improvements over EDSR. More recently, Zhang et al. [21] and Dai et al. [22] further improved the performance of SISR by proposing residual channel attention network (i.e., RCAN) and second-order attention network (i.e., SAN). RCAN and SAN are the most powerful SISR methods to date and can achieve a very high reconstruction accuracy.

LF image SR. In the area of LF image SR, different paradigms have been proposed. Bishop et al. [35] first estimated the scene depth and then used a deconvolution approach to estimate HR SAIs. Wanner et al. [36] proposed a variational LF image SR framework using the estimated disparity map. Farrugia et al. [37] decomposed HR-LR patches into several subspaces, and achieved LF image SR via PCA analysis. Alain et al. [12] extended SR-BM3D [38] to LFs, and super-resolved SAIs using LFBM5D filtering. Rossi et al. [14] formulated LF image SR as a graph optimization problem. These traditional methods [35,36,37,12,14] use different approaches to exploit angular information, but perform inferior in spatial information exploitation as compared to recent deep learning-based methods. In the pioneering work of deep learning-based LF image SR (i.e., LFCNN [25]), SAIs are super-resolved separately using SRCNN and fine-tuned in pairs to incorporate angular information. Similarly, Yuan et al. [27] proposed LF-DCNN, in which they used EDSR [20] to super-resolve each SAI and then fine-tuned the results. LFCNN and LF-DCNN handle the LF image SR problem in two stages and do not use angular information in the first stage. Wang et al. [15] proposed LFNet by extending BRCN [39] to LF image SR. In their method, SAIs from the same row or column are fed to a recurrent network to incorporate angular information. Zhang et al. [13] stacked SAIs along different angular directions to generate input volumes, and then fed them to a multi-stream residual network named resLF. LFNet and resLF reduce 4D LF to 3D LF by using part of SAIs to super-resolve one view. Consequently, angular information in these discarded views cannot be incorporated. To consider all views for LF image SR, Yeung et al. [16] proposed LFSSR to alternately shuffle LF features between SAI pattern and MacPI pattern for convolution. Jin et al. [17] proposed an all-to-one LF image SR framework (i.e., LF-ATO) and performed structural consistency regularization to preserve the parallax structure among reconstructed views.

3 Method

3.1 Spatial-Angular Feature Decoupling

An LF has a 4D structure and can be denoted as $\mathcal{L} \in \mathbb{R}^{U \times V \times H \times W}$, where U and V represent the angular dimensions (e.g., $U = 3$, $V = 4$ for a 3×4 LF), H and W represent the height and width of each SAI. Intuitively, an LF can be considered as a 2D angular collection of SAIs, and the SAI at each angular coordinate (u, v) can be denoted as $\mathcal{L}(u, v, :, :) \in \mathbb{R}^{H \times W}$. Similarly, an LF can also be organized




Fig. 1: SAI array (left) and MacPI (right) representations of LFs. Both the SAI array and the MacPI representations have the same size of $\mathbb{R}^{UH \times VW}$. Note that, to convert an SAI array representation into a MacPI representation, pixels at the same spatial coordinates of each SAI need to be extracted and organized according to their angular coordinates to generate a macro-pixel. Then, a MacPI can be generated by organizing these macro-pixels according to their spatial coordinates. More details are presented in the supplemental material.

into a 2D spatial collection of macro-pixels (namely, a MacPI). The macro-pixel at each spatial coordinate (h, w) can be denoted as $\mathcal{L}(:, :, h, w) \in \mathbb{R}^{U \times V}$. An illustration of these two LF representations is shown in Fig. 1.

Since most methods use SAIs distributed in a square array as their input, we follow [12,14,25,26,16,13,17] to set $U = V = A$ in our method, where A denotes the angular resolution. Given an LF of size $\mathbb{R}^{A \times A \times H \times W}$, both a MacPI and an SAI array can be generated by organizing pixels according to corresponding patterns. Note that, when an LF is organized as an SAI array, the angular information is implicitly contained among different SAIs and thus is hard to extract. Therefore, we use the MacPI representation in our method and design spatial/angular feature extractors (SFE/AFE) to extract and decouple spatial/angular information.

Here, we use a toy example in Fig. 2 to illustrate the angular and spatial feature extractors. Specifically, AFE is defined as a convolution with a kernel size of $A \times A$ and a stride of A . Padding is not performed so that features generated by AFE have a size of $\mathbb{R}^{H \times W \times C}$, where C represents the feature depth. In contrast, SFE is defined as a convolution with a kernel size of 3×3 , a stride of 1, and a dilation of A . We perform zero padding to ensure that the output features have the same spatial size $AH \times AW$ as the input MacPI. It is worth noting that, during angular feature extraction, each macro-pixel can be exactly convolved by AFE, while the information across different macro-pixels is not aliased. Similarly, during spatial feature extraction, pixels in each SAI can be convolved by SFE, while the angular information is not involved. In this way, the spatial and angular information in an LF is decoupled.




Fig. 2: An illustration of angular and spatial feature extractors. Here, an LF of size $\mathbb{R}^{3 \times 3 \times 3 \times 3}$ is used as a toy example. For better visualization, pixels from different SAIs are represented with different labels (e.g., red arrays or green squares), while different macro-pixels are paint with different background colors. Note that, AFE only extracts angular features and SFE only extracts spatial features, resulting in decoupling of spatial-angular information.

Due to the 3D property of real scenes, objects at different depths have different disparity values. Consequently, pixels of an object among different views cannot always locate at a single macro-pixel [40]. To address this problem, we enlarge the receptive field of our LF-InterNet by cascading multiple SFEs and AFEs in an interactive manner (see Fig. 4). Here, we use the Grad-CAM method [41] to visualize the receptive field of our LF-InterNet by highlighting contributive input regions. As shown in Fig. 3, the angular information indeed contributes to LF image SR, and the receptive field is enough to cover the disparities in LFs.

3.2 Network Design

Our LF-InterNet takes an LR MacPI of size $\mathbb{R}^{AH \times AW}$ as its input and produces an HR SAI array of size $\mathbb{R}^{\alpha AH \times \alpha AW}$, where α denotes the upscaling factor. Following [13,16,17], we convert images into YCbCr color space, and only super-resolve the Y channel of images. An overview of our network is shown in Fig. 4.

Overall Architecture Given an LR MacPI $\mathcal{I}_{LR} \in \mathbb{R}^{AH \times AW}$, the angular and spatial features are first extracted by AFE and SFE, respectively.

$$\mathcal{F}_{A,0} = H_A(\mathcal{I}_{LR}), \quad \mathcal{F}_{S,0} = H_S(\mathcal{I}_{LR}), \quad (1)$$

where $\mathcal{F}_{A,0} \in \mathbb{R}^{H \times W \times C}$ and $\mathcal{F}_{S,0} \in \mathbb{R}^{AH \times AW \times C}$ represent the extracted angular and spatial features, respectively. H_A and H_S represent the angular and spatial feature extractors (as described in Section 3.1), respectively. Once initial features are extracted, features $\mathcal{F}_{A,0}$ and $\mathcal{F}_{S,0}$ are further processed by a set of interaction groups (i.e., Inter-Groups) to achieve spatial-angular feature interaction:

$$(\mathcal{F}_{A,n}, \mathcal{F}_{S,n}) = H_{IG,n}(\mathcal{F}_{A,n-1}, \mathcal{F}_{S,n-1}), \quad (n = 1, 2, \dots, N), \quad (2)$$

where $H_{IG,n}$ denotes the n^{th} Inter-Group and N denotes the total number of Inter-Groups. In our LF-InterNet, we cascade all these Inter-Groups to fully use

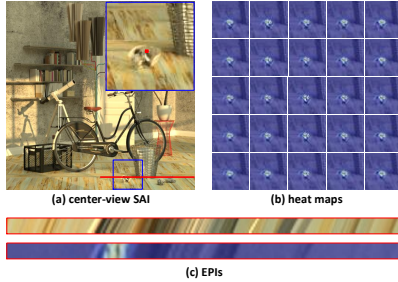


Fig. 3: A visualization of the receptive field of our LF-InterNet. We performed $2\times$ SR on the 5×5 center views of scene *HCInew_bicycle* [42]. (a) Center-view HR SAI. We select a target pixel (marked in red in the zoom-in region) at a shallow depth. (b) Highlighted input SAIs generated by the Grad-CAM method [41]. A cluster of pixels in each SAI are highlighted as contributive pixels, which demonstrates the contribution of angular information. (c) Epipolar-plane images (EPIs) of the input LF (top) and the highlighted SAIs (bottom). It can be observed that the highlighted pixels in the bottom EPI have an enough receptive field to cover the slopes in the top EPI, which demonstrates that our LF-InterNet can well handle the disparity problem in LF image SR.

the information interacted at different stages. Specifically, features generated by each Inter-Group are concatenated and fed to a bottleneck block to fuse the interacted information. The feature generated by the bottleneck block is further added with the initial feature $\mathcal{F}_{S,0}$ to achieve global residual learning. The fused feature $\mathcal{F}_{S,t}$ can be obtained by

$$\mathcal{F}_{S,t} = H_B([\mathcal{F}_{A,1}, \dots, \mathcal{F}_{A,N}], [\mathcal{F}_{S,1}, \dots, \mathcal{F}_{S,N}]) + \mathcal{F}_{S,0}, \quad (3)$$

where H_B denotes the bottleneck block, $[\cdot]$ denotes the concatenation operation. Finally, the fused feature $\mathcal{F}_{S,t}$ is fed to the reconstruction module, and an HR SAI array $\mathcal{I}_{SR} \in \mathbb{R}^{\alpha AH \times \alpha AW}$ can be obtained by

$$\mathcal{I}_{SR} = H_{1\times 1}(S_{pix}(R_{lf}(H_S(\mathcal{F}_{S,t})))), \quad (4)$$

where R_{lf} , S_{pix} , and $H_{1\times 1}$ represent LF reshape, pixel shuffling, and 1×1 convolution, respectively.

Spatial-Angular Feature Interaction The basic module for spatial-angular interaction is the interaction block (i.e., Inter-Block). As shown in Fig. 4(b), the Inter-Block takes a pair of angular and spatial features as its inputs to achieve feature interaction. Specifically, the input angular feature is first upsampled by a factor of A . Since pixels in a MacPI can be unevenly distributed due to edges and occlusions in real scenes [43], we learn this discontinuity using a 1×1 convolution and a pixel shuffling layer for angular-to-spatial upsampling. The upsampled angular feature is concatenated with the input spatial feature, and further fed

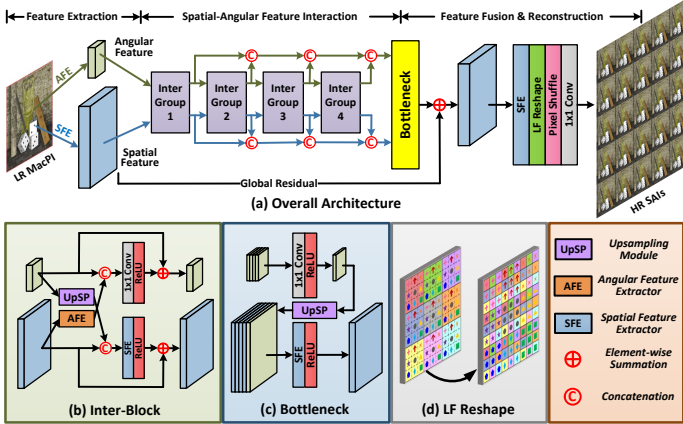


Fig. 4: An overview of our LF-InterNet. Angular and spatial features are first extracted from the input MacPI, and then fed to a series of Inter-Groups (which consists of several cascaded Inter-Blocks) to achieve spatial-angular interaction. After LF reshape and pixel shuffling, HR SAIs are generated.

to an SFE to incorporate the spatial and angular information. In this way, the complementary angular information can be used to guide spatial feature extraction. Simultaneously, the new angular feature is extracted from the input spatial feature by an AFE, and then concatenated with the input angular feature. The concatenated angular feature is further fed to a 1×1 convolution to integrate and update the angular information. Note that, the fused angular and spatial features are added with their input features to achieve local residual learning. In this paper, we cascade K Inter-Blocks in an Inter-Group, i.e., the output of an Inter-Block forms the input of its subsequent Inter-Block. In summary, the spatial-angular feature interaction can be formulated as

$$\mathcal{F}_{S,n}^{(k)} = H_S \left(\left[\mathcal{F}_{S,n}^{(k-1)}, \left(\mathcal{F}_{A,n}^{(k-1)} \right) \uparrow \right] \right) + \mathcal{F}_{S,n}^{(k-1)}, (k = 1, 2, \dots, K), \quad (5)$$

$$\mathcal{F}_{A,n}^{(k)} = H_{1 \times 1} \left(\left[\mathcal{F}_{A,n}^{(k-1)}, H_A \left(\mathcal{F}_{S,n}^{(k-1)} \right) \right] \right) + \mathcal{F}_{A,n}^{(k-1)}, (k = 1, 2, \dots, K), \quad (6)$$

where \uparrow represents the upsampling operation, $\mathcal{F}_{S,n}^{(k)}$ and $\mathcal{F}_{A,n}^{(k)}$ represent the output spatial and angular features of the k^{th} Inter-Block in the n^{th} Inter-Group, respectively.

Feature Fusion and Reconstruction The objective of this stage is to fuse the interacted features to reconstruct an HR SAI array. The fusion and reconstruction stage mainly consists of bottleneck fusion (as shown in Fig. 4(c)), LF reshape (as shown in Fig. 4(d)), pixel shuffling, and final reconstruction.

In the bottleneck, the concatenated angular features $[\mathcal{F}_{A,1}, \dots, \mathcal{F}_{A,N}] \in \mathbb{R}^{H \times W \times NC}$ are first fed to a 1×1 convolution and a ReLU layer to generate

a feature map $\mathcal{F}_A \in \mathbb{R}^{H \times W \times C}$. Then, the squeezed angular feature \mathcal{F}_A is up-sampled and concatenated with spatial features. The final fused feature $\mathcal{F}_{S,t}$ can be obtained as

$$\mathcal{F}_{S,t} = H_S([\mathcal{F}_{S,1}, \dots, \mathcal{F}_{S,N}, (\mathcal{F}_A) \uparrow]) + \mathcal{F}_{S,0}. \quad (7)$$

After feature fusion, we apply another SFE layer to extend the channel size of $\mathcal{F}_{S,t}$ to $\alpha^2 C$ for pixel shuffling [44]. However, since $\mathcal{F}_{S,t}$ is organized in the MacPI pattern, we apply LF reshape to convert $\mathcal{F}_{S,t}$ into an SAI array representation for pixel shuffling. To achieve LF reshape, we first extract pixels with the same angular coordinates in the MacPI feature, and then re-organize these pixels according to their spatial coordinates, which can be formulated as

$$\mathcal{I}_{SAIs}(x, y) = \mathcal{I}_{MacPI}(\xi, \eta), \quad (8)$$

where

$$x = H(\xi - 1) + \lfloor \xi/A \rfloor (1 - AH) + 1, \quad (9)$$

$$y = W(\eta - 1) + \lfloor \eta/A \rfloor (1 - AW) + 1. \quad (10)$$

Here, $x = 1, 2, \dots, AH$ and $y = 1, 2, \dots, AW$ denote the pixel coordinates in the output SAI arrays, ξ and η denote the corresponding coordinates in the input MacPI, $\lfloor \cdot \rfloor$ represents the round-down operation. The derivation of Eqs. (9) and (10) is presented in the supplemental material. Finally, a 1×1 convolution is applied to squeeze the number of feature channels to 1 for HR SAI reconstruction.

4 Experiments

In this section, we first introduce the datasets and our implementation details. Then we conduct ablation studies to investigate our network. Finally, we compare our LF-InterNet to several state-of-the-art LF image SR and SISR methods.

4.1 Datasets and Implementation Details

As listed in Table 1, we used 6 public LF datasets [42,45,46,47,48,49] in our experiments. All the LFs in the training and test sets have an angular resolution of 9×9 . In the training stage, we cropped each SAI into patches of size 64×64 , and then used bicubic downsampling with a factor of α ($\alpha = 2, 4$) to generate LR patches. The generated LR patches were re-organized into MacPIs to form the input of our network. The L_1 loss was used since it can produce good results and is robust to outliers [50]. We augmented the training data by 8 times using random flipping and 90-degree rotation. Note that, during each data augmentation, all SAIs need to be flipped and rotated along both spatial and angular directions to maintain their LF structures.

By default, we used the model with $N = 4$, $K = 4$, $C = 64$, and angular resolution of 5×5 for both $2 \times$ and $4 \times$ SR. We also investigated the performance of other branches of our LF-InterNet in Section 4.2. We used PSNR and SSIM

Table 1: Datasets used in our experiments.

	EPFL [45]	HCInew [42]	HCIold [46]	INRIA [47]	STFGantry [48]	STFlytro [49]
Training	70	20	10	35	9	250
Test	10	4	2	5	2	50

as quantitative metrics for performance evaluation. Note that, PSNR and SSIM were separately calculated on the Y channel of each SAI. To obtain the overall metric score for a dataset with M scenes (each with an angular resolution of $A \times A$), we first obtain the score for a scene by averaging its A^2 scores, and then obtain the overall score by averaging the scores of all M scenes.

Our LF-InterNet was implemented in PyTorch on a PC with an Nvidia RTX 2080Ti GPU. Our model was initialized using the Xavier method [51] and optimized using the Adam method [52]. The batch size was set to 12 and the learning rate was initially set to 5×10^{-4} and decreased by a factor of 0.5 for every 10 epochs. The training was stopped after 40 epochs.

4.2 Ablation Study

In this subsection, we compare the performance of our LF-InterNet with different architectures and angular resolutions to investigate the potential benefits introduced by different design choices.

Angular Information. We investigated the benefit of angular information by removing the angular path in LF-InterNet. That is, we only use SFE for LF image SR. Consequently, the network is identical to an SISR network, and can only incorporate spatial information within each SAI. As shown in Table 2, only using the spatial information, the network (i.e., LF-InterNet-SpatialOnly) achieves a PSNR of 29.98 and an SSIM of 0.897, which are significantly inferior to LF-InterNet. Therefore, the benefit of angular information to LF image SR is clearly demonstrated.

Spatial Information. To investigate the benefit introduced by spatial information, we changed the kernel size of all SFEs from 3×3 to 1×1 . In this case, the spatial information cannot be exploited and integrated by convolutions. As shown in Table 2, the performance of LF-InterNet-AngularOnly is even inferior to bicubic interpolation. That is because, neighborhood context in an image is highly significant in recovering details. It is clear that spatial information plays a major role in LF image SR, while angular information can only be used as a complementary part to spatial information but cannot be used alone.

Information Decoupling. To investigate the benefit of spatial-angular information decoupling, we stacked all SAIs along the channel dimension as input, and used 3×3 convolutions with a stride of 1 to extract both spatial and angular information from these stacked images. Note that, global and local residual learning was maintained in this variant to keep the overall network architecture unchanged. To achieve fair comparison, we adjusted the feature depths to keep the model size (i.e., LF-InterNet-SAcoupled_1) or computational complexity (i.e., LF-InterNet-SAcoupled_2) comparable to LF-InterNet. As shown in Ta-

Table 2: Comparative results achieved on the STFlytro dataset [49] by several variants of our LF-InterNet for $4\times$ SR. Note that, we carefully adjusted the feature depths of different variants to make their model size comparable. FLOPs are computed with an input MacPI of size 160×160 . The results of bicubic interpolation are listed as baselines.

Model	PSNR	SSIM	Params.	FLOPs
Bicubic	27.84	0.855	—	—
LF-InterNet-SpatialOnly	29.98	0.897	5.40M	134.7G
LF-InterNet-AngularOnly	26.57	0.823	5.43M	13.4G
LF-InterNet-SACoupled_1	31.11	0.918	5.42M	5.46G
LF-InterNet-SACoupled_2	31.17	0.919	50.8M	50.5G
LF-InterNet	31.65	0.925	5.23M	50.1G

Table 3: Comparative results achieved on the STFlytro dataset [49] by our LF-InterNet with different number of interactions for $4\times$ SR.

IG_1	IG_2	IG_3	IG_4	PSNR	SSIM
				29.84	0.894
✓				31.44	0.922
✓	✓			31.61	0.924
✓	✓	✓		31.66	0.925
✓	✓	✓	✓	31.84	0.927

ble 2, both two variants are inferior to LF-InterNet. That is, our LF-InterNet can well handle the 4D LF structure and achieve LF image SR much more efficiently by using the proposed spatial-angular feature decoupling mechanism.

Spatial-Angular Interaction. We investigated the benefits introduced by our spatial-angular interaction mechanism. Specifically, we canceled feature interaction in each Inter-Group by removing upsampling and AFE modules in each Inter-Block. In this case, spatial and angular features are processed separately. When all interactions are removed, these spatial and angular features can only be incorporated by the bottleneck block. Table 3 presents the results achieved by our LF-InterNet with different numbers of interactions. It can be observed that, without any feature interaction, our network achieves a very low reconstruction accuracy (i.e., 29.84 in PSNR and 0.894 in SSIM). As the number of interactions increases, the performance is steadily improved. This clearly demonstrates the effectiveness of our spatial-angular feature interaction mechanism.

Angular-to-Spatial Upsampling. To demonstrate the effectiveness of the pixel shuffling layer used in angular-to-spatial upsampling, we introduced two variants by replacing pixel shuffling with nearest upsampling and bilinear upsampling, respectively. It can be observed from Table 4 that LF-InterNet-bilinear achieves much lower PSNR and SSIM scores than LF-InterNet-nearest and LF-InterNet. That is because, bilinear interpolation introduces aliasing among macro-pixels during angular-to-spatial upsampling, resulting in ambiguities in spatial-angular feature interaction. In contrast, both nearest upsampling and pixel shuffling do not introduce aliasing and thus achieve improved performance. Moreover, since pixels in a macro-pixel can be unevenly distributed due to edges and occlusions in real scenes, pixel shuffling achieves a further im-

Table 4: Comparisons of different approaches for angular-to-spatial upsampling.

Model	Scale	PSNR	SSIM	Scale	PSNR	SSIM
LF-InterNet-nearest	2×	38.60	0.982	4×	31.65	0.925
LF-InterNet-bilinear	2×	37.67	0.976	4×	30.71	0.911
LF-InterNet	2×	38.81	0.983	4×	31.84	0.927

Table 5: Comparative results achieved on the STFflytro dataset [49] by our LF-InterNet with different angular resolutions for 2× and 4×SR.

AngRes	Scale	PSNR	SSIM	Scale	PSNR	SSIM
3×3	2×	37.95	0.980	4×	31.30	0.918
5×5	2×	38.81	0.983	4×	31.84	0.927
7×7	2×	39.05	0.984	4×	32.04	0.931
9×9	2×	39.08	0.985	4×	32.07	0.933

provement over nearest upsampling due to its discontinuity modeling capability within macro-pixels.

Angular Resolution. We analyze the performance of LF-InterNet with different angular resolution. Specifically, we extracted the central $A \times A$ ($A = 3, 5, 7, 9$) SAIs from the input LFs, and trained different models for both 2× and 4×SR. As shown in Table 5, the PSNR and SSIM values for both 2× and 4×SR are improved as the angular resolution is increased. That is because, additional views provide rich angular information for LF image SR. It is also notable that, the performance tends to be saturated when the angular resolution is further increased from 7×7 to 9×9. That is because, the complementary information provided by additional views is already sufficient. Since the angular information has been fully exploited for an angular resolution of 7×7, a further increase of views can only provide minor performance improvement.

4.3 Comparison to the State-of-the-arts

We compare our method to six SISR methods [19,20,21,22,23,24] and five LF image SR methods [12,14,16,13,17]. Bicubic interpolation was used as baselines.

Quantitative Results. Quantitative results in Table 6 demonstrate the state-of-the-art performance of our LF-InterNet on all the 6 test datasets. Thanks to the use of angular information, our method achieves an improvement of 1.54 dB (2×SR) and 1.00 dB (4×SR) in PSNR over the powerful SISR method RCAN [21]. Moreover, our LF-InterNet can achieve a comparable PSNR and SSIM scores as compared to the most recent LF image SR method LF-ATO [17].

Qualitative Results. Qualitative results of 2×/4×SR are shown in Fig. 5, with more visual comparisons being provided in our supplemental material. Our LF-InterNet can well preserve the textures and details (e.g., the horizontal stripes in the scene *HCInew_origami*) in the super-resolved images. In contrast, state-of-the-art SISR methods RCAN [21] and SAN [22] produce oversmoothed images with poor details. The visual superiority of our method is more obvious for 4×SR. That is because, the input LR images are severely degraded by the down-sampling operation, and the process of 4×SR is highly ill-posed. In such cases,

Table 6: PSNR/SSIM values achieved by different methods for 2× and 4×SR. The best results are in **red** and the second best results are in **blue**.

Method	Scale	EPFL	HCInew	HCInold	INRIA	STFGantry	STFlytro
Bicubic	2×	29.50/0.935	31.69/0.934	37.46/0.978	31.10/0.956	30.82/0.947	33.02/0.950
VDSR [19]	2×	32.01/0.959	34.37/0.956	40.34/0.985	33.80/0.972	35.80/0.980	35.91/0.970
EDSR [20]	2×	32.86/0.965	35.02/0.961	41.11/0.988	34.61/0.977	37.08/0.985	36.87/0.975
RCAN [21]	2×	33.46/0.967	35.56/0.963	41.59/0.989	35.18/0.978	38.18/0.988	37.32/0.977
SAN [22]	2×	33.36/0.967	35.51/0.963	41.47/0.989	35.15/0.978	37.98/0.987	37.26/0.976
LFBM5D [12]	2×	31.15/0.955	33.72/0.955	39.62/0.985	32.85/0.969	33.55/0.972	35.01/0.966
GB [14]	2×	31.22/0.959	35.25/0.969	40.21/0.988	32.76/0.972	35.44/0.983	35.04/0.956
resLF [13]	2×	33.22/0.969	35.79/0.969	42.30/0.991	34.86/0.979	36.28/0.985	35.80/0.970
LFSSR [16]	2×	34.15/0.973	36.98/0.974	43.29/0.993	35.76/0.982	37.67/0.989	37.57/0.978
LF-ATO [17]	2×	34.49/0.976	37.28/0.977	43.76/0.994	36.21/0.984	39.06/0.992	38.27/0.982
LF-InterNet	2×	34.76/0.976	37.20/0.976	44.65/0.995	36.64/0.984	38.48/0.991	38.81/0.983
Bicubic	4×	25.14/0.831	27.61/0.851	32.42/0.934	26.82/0.886	25.93/0.843	27.84/0.855
VDSR [19]	4×	26.82/0.869	29.12/0.876	34.01/0.943	28.87/0.914	28.31/0.893	29.17/0.880
EDSR [20]	4×	27.82/0.892	29.94/0.893	35.53/0.957	29.86/0.931	29.43/0.921	30.29/0.903
RCAN [21]	4×	28.31/0.899	30.25/0.896	35.89/0.959	30.36/0.936	30.25/0.934	30.66/0.909
SAN [22]	4×	28.30/0.899	30.25/0.898	35.88/0.960	30.29/0.936	30.30/0.933	30.71/0.909
SRGAN [23]	4×	26.85/0.870	28.95/0.873	34.03/0.942	28.85/0.916	28.19/0.898	29.28/0.883
ESRGAN [24]	4×	25.59/0.836	26.96/0.819	33.53/0.933	27.54/0.880	28.00/0.905	27.09/0.826
LFBM5D [12]	4×	26.61/0.869	29.13/0.882	34.23/0.951	28.49/0.914	28.30/0.900	29.07/0.881
GB [14]	4×	26.02/0.863	28.92/0.884	33.74/0.950	27.73/0.909	28.11/0.901	28.37/0.873
resLF [13]	4×	27.86/0.899	30.37/0.907	36.12/0.966	29.72/0.936	29.64/0.927	28.94/0.891
LFSSR [16]	4×	29.16/0.915	30.88/0.913	36.90/0.970	31.03/0.944	30.14/0.937	31.21/0.919
LF-ATO [17]	4×	29.16/0.917	31.08/0.917	37.23/0.971	31.21/0.950	30.78/0.944	30.98/0.918
LF-InterNet	4×	29.52/0.917	31.01/0.917	37.23/0.972	31.65/0.950	30.44/0.941	31.84/0.927

Table 7: Comparisons of the number of parameters (#Params.) and FLOPs for 2× and 4×SR. Note that, the FLOPs is calculated on an input LF with a size of 5×5×32×32, and the PSNR and SSIM scores are averaged over the 6 test datasets [45,42,46,47,48,49] in Table 6.

Method	Scale	#Params.	FLOPs(G)	PSNR/SSIM	Scale	#Params.	FLOPs(G)	PSNR/SSIM
RCAN [21]	2×	15.44M	15.71×25	36.88/0.977	4×	15.59M	16.34×25	30.95/0.922
SAN [22]	2×	15.71M	16.05×25	36.79/0.977	4×	15.86M	16.67×25	31.96/0.923
resLF [13]	2×	6.35M	37.06	36.38/0.977	4×	6.79M	39.70	30.08/0.916
LFSSR [16]	2×	0.81M	25.70	37.57/0.982	4×	1.61M	128.44	31.55/0.933
LF-ATO [17]	2×	1.51M	597.66	38.18/0.984	4×	1.66M	686.99	31.74/0.937
LF-InterNet_32	2×	1.20M	11.87	37.88/0.983	4×	1.31M	12.53	31.57/0.933
LF-InterNet_64	2×	4.80M	47.46	38.42/0.984	4×	5.23M	50.10	31.95/0.937

some perceptual-oriented methods (e.g., SRGAN [23] and ESRGAN [24]) use spatial information only to hallucinate missing details, resulting in ambiguous and even fake textures (e.g., wheel in scene *STFlytro_buildings*). In contrast, our method can use complementary angular information among different views to produce more faithful results.

Efficiency. We compare our LF-InterNet to several competitive methods [21,22,13,16,17] in terms of the number of parameters and FLOPs. As shown in Table 7, our LF-InterNet achieves superior SR performance with reasonable number of parameters and FLOPs. Note that, although LF-ATO has very small model sizes (i.e., 1.51M for 2×SR and 1.66M for 4×SR), its FLOPs are very high since it uses the *All-to-One* strategy to separately super-resolve individual views in a sequence. In contrast, our method (i.e., LF-InterNet_64) super-resolves all views within a single inference, and achieves comparable or even



Fig. 5: Visual results of $2\times/4\times$ SR.

Fig. 6: Comparative results (i.e., PSNR values) achieved on each perspective of scene *HCInew_MonasRoom*. Here, 7×7 input views are used to perform $2\times$ SR. We use standard deviation (Std) to represent their uniformity. Our LF-InterNet achieves high reconstruction quality with a relatively balanced distribution.

better performance than LF-ATO with significantly lower FLOPs. It is worth noting that, even the feature depth of our model is halved to 32, our method (i.e., LF-InterNet₃₂) can still achieve promising PSNR/SSIM scores, which are comparable to LFSSR and higher than RCAN, SAN, and resLF. The above comparisons clearly demonstrate the high efficiency of our network architecture.

Performance w.r.t. Perspectives. Since our LF-InterNet can super-resolve all SAIs in an LF, we further investigate the reconstruction quality with respect to different perspectives. We followed [13] to use the central 7×7 views of scene *HCInew_MonasRoom* to perform $2\times$ SR, and used PSNR for performance evaluation. Note that, due to the changing perspectives, the contents of different SAIs are not identical, resulting in inherent PSNR variations. Therefore, we evaluate this variation by using RCAN to perform SISR on each SAI. Results are reported and visualized in Fig. 6. Since resLF uses part of views to super-resolve different perspectives, the reconstruction qualities of resLF for non-central views are relatively low. In contrast, LFSSR, LF-ATO and our LF-InterNet can use the angular information from all input views to super-resolve each view, and thus achieve a relatively balanced distribution (i.e., lower Std scores) among different perspectives. The reconstruction quality (i.e., PSNR scores) of LF-InterNet is higher than those of LFSSR and LF-ATO on this scene.

Table 8: Comparative results achieved on the UCSD dataset for 2× and 4×SR.

Method	Scale	PSNR/SSIM	Scale	PSNR/SSIM
RCAN [21]	2×	41.63/0.983	4×	36.49/0.955
SAN [22]	2×	41.56/0.983	4×	36.57/0.956
resLF [13]	2×	41.29/0.982	4×	35.89/0.953
LFSSR [16]	2×	41.55/0.984	4×	36.77/0.957
LF-ATO [17]	2×	41.80/0.985	4×	36.95/0.959
LF-InterNet	2×	42.36/0.985	4×	37.12/0.960




Fig. 7: Visual results achieved by different methods under real-world degradation.

Generalization to Unseen Scenarios. We evaluate the generalization capability of different methods by testing them on a novel and unseen real-world dataset (i.e., the UCSD dataset [53]). Note that, all methods have not been trained or fine-tuned on the UCSD dataset. Results in Table 8 show that our LF-InterNet outperforms the state-of-the-art methods [21,22,13,16,17], which demonstrates the generalization capability of our method to unseen scenarios.

Performance Under Real-World Degradation. We compare the performance of different methods under real-world degradation by directly applying them to LFs in the STFlytro dataset [49]. As shown in Fig. 7, our method produces images with faithful details and less artifacts. Since the LF structure keeps unchanged under both bicubic and real-world degradation, our method can learn to incorporate spatial and angular information from training LFs using the proposed spatial-angular interaction mechanism. It is also demonstrated that our method can be easily applied to LF cameras to generate high-quality images.

5 Conclusion and Future Work

In this paper, we proposed a deep convolutional network LF-InterNet for LF image SR. We first introduce an approach to extract and decouple spatial and angular features, and then design a feature interaction mechanism to incorporate spatial and angular information. Experimental results have demonstrated the superiority of our LF-InterNet over state-of-the-art methods. Since the spatial-angular interaction mechanism is a generic framework and can process LFs in an elegant and efficient manner, we will apply LF-InterNet to LF angular SR [8,9,10,11] and joint spatial-angular SR [54,55] as our future work.

Acknowledgement This work was supported by the National Natural Science Foundation of China (No. 61972435, 61602499).

References

1. Wang, Y., Yang, J., Guo, Y., Xiao, C., An, W.: Selective light field refocusing for camera arrays using bokeh rendering and superresolution. *IEEE Signal Processing Letters* **26**(1) (2018) 204–208
2. Shin, C., Jeon, H.G., Yoon, Y., So Kweon, I., Joo Kim, S.: Epinet: A fully-convolutional neural network using epipolar geometry for depth from light field images. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. (2018) 4748–4757
3. Wang, T., Piao, Y., Li, X., Zhang, L., Lu, H.: Deep learning for light field saliency detection. In: *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*. (2019) 8838–8848
4. Zhang, M., Li, J., WEI, J., Piao, Y., Lu, H.: Memory-oriented decoder for light field salient object detection. In: *Advances in Neural Information Processing Systems*. (2019) 896–906
5. Wang, Y., Wu, T., Yang, J., Wang, L., An, W., Guo, Y.: Deocnet: Learning to see through foreground occlusions in light fields. In: *Winter Conference on Applications of Computer Vision (WACV)*. (2020)
6. Wilburn, B., Joshi, N., Vaish, V., Talvala, E.V., Antunez, E., Barth, A., Adams, A., Horowitz, M., Levoy, M.: High performance imaging using large camera arrays. In: *ACM Transactions on Graphics*. Volume 24., ACM (2005) 765–776
7. Venkataraman, K., Lelescu, D., Duparré, J., McMahon, A., Molina, G., Chatterjee, P., Mullis, R., Nayar, S.: Picam: An ultra-thin high performance monolithic camera array. *ACM Transactions on Graphics* **32**(6) (2013) 166
8. Wu, G., Zhao, M., Wang, L., Dai, Q., Chai, T., Liu, Y.: Light field reconstruction using deep convolutional network on epi. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. (2017) 6319–6327
9. Wu, G., Liu, Y., Dai, Q., Chai, T.: Learning sheared epi structure for light field reconstruction. *IEEE Transactions on Image Processing* **28**(7) (2019) 3261–3273
10. Jin, J., Hou, J., Yuan, H., Kwong, S.: Learning light field angular super-resolution via a geometry-aware network. In: *AAAI Conference on Artificial Intelligence*. (2020)
11. Shi, J., Jiang, X., Guillemot, C.: Learning fused pixel and feature-based view reconstructions for light fields. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. (2020)
12. Alain, M., Smolic, A.: Light field super-resolution via lfbm5d sparse coding. In: *Proceedings of the IEEE International Conference on Image Processing (ICIP)*. (2018) 2501–2505
13. Zhang, S., Lin, Y., Sheng, H.: Residual networks for light field image super-resolution. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. (2019) 11046–11055
14. Rossi, M., Frossard, P.: Geometry-consistent light field super-resolution via graph-based regularization. *IEEE Transactions on Image Processing* **27**(9) (2018) 4207–4218
15. Wang, Y., Liu, F., Zhang, K., Hou, G., Sun, Z., Tan, T.: Lfnet: A novel bidirectional recurrent convolutional neural network for light-field image super-resolution. *IEEE Transactions on Image Processing* **27**(9) (2018) 4274–4286
16. Yeung, H.W.F., Hou, J., Chen, X., Chen, J., Chen, Z., Chung, Y.Y.: Light field spatial super-resolution using deep efficient spatial-angular separable convolution. *IEEE Transactions on Image Processing* **28**(5) (2018) 2319–2330

17. Jin, J., Hou, J., Chen, J., Kwong, S.: Light field spatial super-resolution via deep combinatorial geometry embedding and structural consistency regularization. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2020)
18. Dong, C., Loy, C.C., He, K., Tang, X.: Learning a deep convolutional network for image super-resolution. In: Proceedings of the European Conference on Computer Vision (ECCV). (2014) 184–199
19. Kim, J., Kwon Lee, J., Mu Lee, K.: Accurate image super-resolution using very deep convolutional networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2016) 1646–1654
20. Lim, B., Son, S., Kim, H., Nah, S., Mu Lee, K.: Enhanced deep residual networks for single image super-resolution. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW). (2017) 136–144
21. Zhang, Y., Li, K., Li, K., Wang, L., Zhong, B., Fu, Y.: Image super-resolution using very deep residual channel attention networks. In: Proceedings of the European Conference on Computer Vision (ECCV). (2018) 286–301
22. Dai, T., Cai, J., Zhang, Y., Xia, S.T., Zhang, L.: Second-order attention network for single image super-resolution. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2019) 11065–11074
23. Ledig, C., Theis, L., Huszár, F., Caballero, J., Cunningham, A., Acosta, A., Aitken, A., Tejani, A., Totz, J., Wang, Z., et al.: Photo-realistic single image super-resolution using a generative adversarial network. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2017) 4681–4690
24. Wang, X., Yu, K., Wu, S., Gu, J., Liu, Y., Dong, C., Qiao, Y., Change Loy, C.: Esrgan: Enhanced super-resolution generative adversarial networks. In: Proceedings of the European Conference on Computer Vision (ECCV). (2018) 0–0
25. Yoon, Y., Jeon, H.G., Yoo, D., Lee, J.Y., So Kweon, I.: Learning a deep convolutional network for light-field image super-resolution. In: Proceedings of the IEEE International Conference on Computer Vision Workshops (ICCVW). (2015) 24–32
26. Yoon, Y., Jeon, H.G., Yoo, D., Lee, J.Y., Kweon, I.S.: Light-field image super-resolution using convolutional neural network. *IEEE Signal Processing Letters* **24**(6) (2017) 848–852
27. Yuan, Y., Cao, Z., Su, L.: Light-field image superresolution using a combined deep cnn based on epi. *IEEE Signal Processing Letters* **25**(9) (2018) 1359–1363
28. Wang, Z., Chen, J., Hoi, S.C.: Deep learning for image super-resolution: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2020)
29. Anwar, S., Khan, S., Barnes, N.: A deep journey into super-resolution: A survey. *ACM Computing Surveys* **53**(3) (2020) 1–34
30. Yang, W., Zhang, X., Tian, Y., Wang, W., Xue, J.H., Liao, Q.: Deep learning for single image super-resolution: A brief review. *IEEE Transactions on Multimedia* (2019)
31. Timofte, R., De Smet, V., Van Gool, L.: Anchored neighborhood regression for fast example-based super-resolution. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV). (2013) 1920–1927
32. Jianchao, Y., John, W., Thomas, H., Yi, M.: Image super-resolution via sparse representation. *IEEE Transactions on Image Processing* **19**(11) (2010) 2861–2873
33. Zeyde, R., Elad, M., Protter, M.: On single image scale-up using sparse-representations. In: Proceedings of the International Conference on Curves and Surfaces. (2010) 711–730

34. Zhang, Y., Tian, Y., Kong, Y., Zhong, B., Fu, Y.: Residual dense network for image super-resolution. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2018) 2472–2481
35. Bishop, T.E., Favaro, P.: The light field camera: Extended depth of field, aliasing, and superresolution. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **34**(5) (2011) 972–986
36. Wanner, S., Goldluecke, B.: Variational light field analysis for disparity estimation and super-resolution. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **36**(3) (2013) 606–619
37. Farrugia, R.A., Galea, C., Guillemot, C.: Super resolution of light field images using linear subspace projection of patch-volumes. *IEEE Journal of Selected Topics in Signal Processing* **11**(7) (2017) 1058–1071
38. Egiiazarian, K., Katkovnik, V.: Single image super-resolution via bm3d sparse coding. In: European Signal Processing Conference (EUSIPCO). (2015) 2849–2853
39. Huang, Y., Wang, W., Wang, L.: Bidirectional recurrent convolutional networks for multi-frame super-resolution. In: Advances in Neural Information Processing Systems (NeurIPS). (2015) 235–243
40. Williem, Park, I., Lee, K.M.: Robust light field depth estimation using occlusion-noise aware data costs. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **40**(10) (2018) 2484–2497
41. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-cam: Visual explanations from deep networks via gradient-based localization. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV). (2017) 618–626
42. Honauer, K., Johannsen, O., Kondermann, D., Goldluecke, B.: A dataset and evaluation methodology for depth estimation on 4d light fields. In: Proceedings of the Asian Conference on Computer Vision (ACCV). (2016) 19–34
43. Park, I.K., Lee, K.M., et al.: Robust light field depth estimation using occlusion-noise aware data costs. *IEEE transactions on pattern analysis and machine intelligence* **40**(10) (2017) 2484–2497
44. Shi, W., Caballero, J., Huszár, F., Totz, J., Aitken, A.P., Bishop, R., Rueckert, D., Wang, Z.: Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2016) 1874–1883
45. Rerabek, M., Ebrahimi, T.: New light field image dataset. In: Proceedings of the International Conference on Quality of Multimedia Experience (QoMEX). (2016)
46. Wanner, S., Meister, S., Goldluecke, B.: Datasets and benchmarks for densely sampled 4d light fields. In: Vision, Modelling and Visualization (VMV). Volume 13., Citeseer (2013) 225–226
47. Le Pendu, M., Jiang, X., Guillemot, C.: Light field inpainting propagation via low rank matrix completion. *IEEE Transactions on Image Processing* **27**(4) (2018) 1981–1993
48. Vaish, V., Adams, A.: The (new) stanford light field archive. *Computer Graphics Laboratory, Stanford University* **6**(7) (2008)
49. Raj, A.S., Lowney, M., Shah, R., Wetzstein, G.: Stanford lytro light field archive (2016)
50. Anagun, Y., Isik, S., Seke, E.: Srlibrary: Comparing different loss functions for super-resolution over various convolutional architectures. *Journal of Visual Communication and Image Representation* **61** (2019) 178–187

51. Glorot, X., Bengio, Y.: Understanding the difficulty of training deep feedforward neural networks. In: Proceedings of the International Conference on Artificial Intelligence and Statistics. (2010) 249–256
52. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. Proceedings of the International Conference on Learning and Representation (ICLR) (2015)
53. Wang, T.C., Zhu, J.Y., Hiroaki, E., Chandraker, M., Efros, A.A., Ramamoorthi, R.: A 4d light-field dataset and cnn architectures for material recognition. In: Proceedings of the European Conference on Computer Vision (ECCV). (2016) 121–138
54. Meng, N., So, H.K.H., Sun, X., Lam, E.: High-dimensional dense residual convolutional neural network for light field reconstruction. IEEE Transactions on Pattern Analysis and Machine Intelligence (2019)
55. Meng, N., Wu, X., Liu, J., Lam, E.Y.: High-order residual network for light field super-resolution. AAAI Conference on Artificial Intelligence (2020)