CutMIB: Boosting Light Field Super-Resolution via Multi-View Image Blending

Zeyu Xiao Yutong Liu Ruisheng Gao Zhiwei Xiong*
University of Science and Technology of China
{zeyuxiao,ustclyt,grsmc4180}@mail.ustc.edu.cn zwxiong@ustc.edu.cn

Abstract

Data augmentation (DA) is an efficient strategy for improving the performance of deep neural networks. Recent DA strategies have demonstrated utility in single image super-resolution (SR). Little research has, however, focused on the DA strategy for light field SR, in which multi-view information utilization is required. For the first time in light field SR, we propose a potent DA strategy called CutMIB to improve the performance of existing light field SR networks while keeping their structures unchanged. Specifically, CutMIB first cuts low-resolution (LR) patches from each view at the same location. Then CutMIB blends all LR patches to generate the blended patch and finally pastes the blended patch to the corresponding regions of high-resolution light field views, and vice versa. By doing so, CutMIB enables light field SR networks to learn from implicit geometric information during the training stage. Experimental results demonstrate that CutMIB can improve the reconstruction performance and the angular consistency of existing light field SR networks. We further verify the effectiveness of CutMIB on real-world light field SR and light field denoising. The implementation code is available at `https://github.com/zeyuxiao1997/CutMIB`.

1. Introduction

Light field cameras, which can record spatial and angular information of light rays, have rapidly become prominent imaging devices in virtual and augmented reality. Light fields are suitable for various applications, such as post-capture refocusing [35, 55], disparity estimation [52], and foreground occlusion removal [54, 69], thanks to the abundance of 4D spatial-angular information they contain. Commercialized light field cameras generally adopt micro-lens-array in front of the sensor, which poses an essential trade-off between the angular and spatial resolutions [29, 35]. Therefore, light field super-resolution (SR) has been an important and popular topic. Convolutional neural network

*Corresponding author.

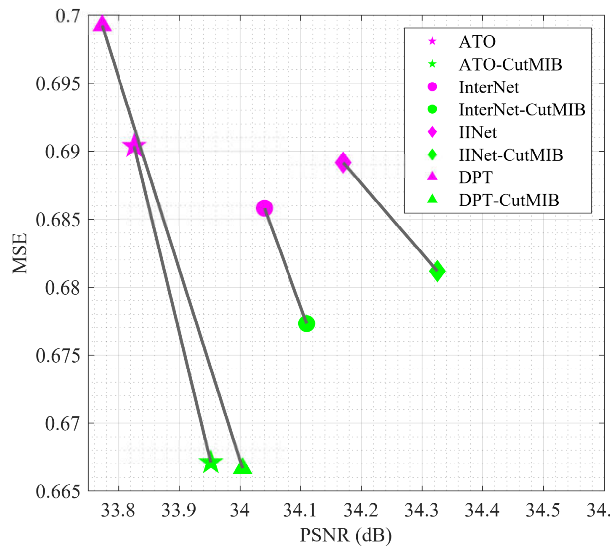

Figure 1. Comparisons on the reconstruction fidelity (PSNR, ↑) and the angular consistency (MSE, ↓) between light fields super-resolved through different methods. Following [9], we super-resolve the whole light field of the scene *Bicycle* from the HCI dataset to analyze the angular consistency of the super-resolved results in terms of disparity estimation using SPO [70]. Note that, CutMIB improves the values of PSNR and lowers the values of MSE by a large margin as compared to naïve light field SR methods (*e.g.*, ATO [27], InterNet [53], IINet [33], and DPT [47]).

(CNN) based and Transformer based methods have recently shown promising performance for light field SR [7, 8, 10, 26, 31, 33, 47, 52, 53, 56], outperforming traditional non-learning based methods [1, 38] with noticeable gains. This performance boost is obtained by training deep methods on external datasets. Few works have investigated data augmentation (DA) strategies for light field SR, which can improve the model performance without the need for additional training datasets given that obtaining these light field data is often time-consuming and expensive [19, 23, 36, 48].

DA has been well studied in high-level vision tasks (*e.g.*, image recognition, image classification, and semantic segmentation) for achieving better network performance and alleviating the overfitting problem [14, 44, 49, 64, 66, 71]. For example, as one of the pioneering strategies, Mixup [66]

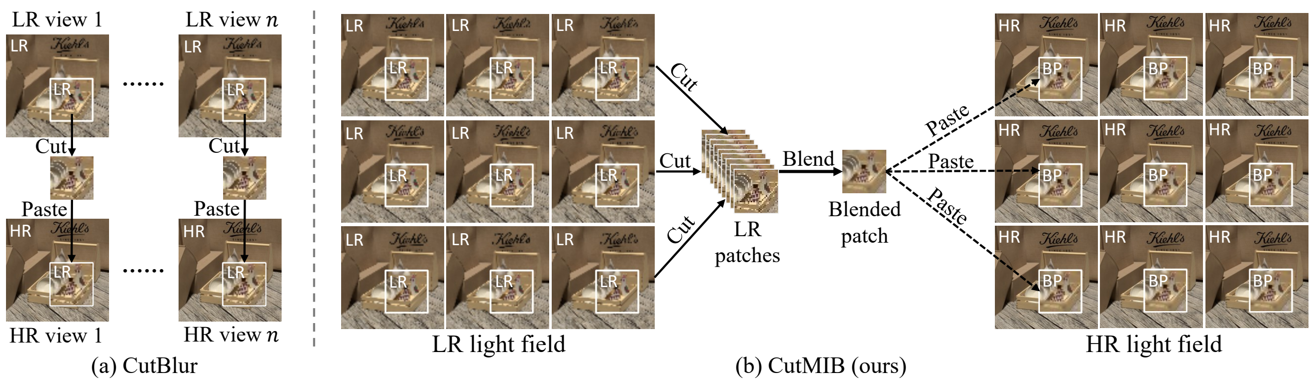

Figure 2. Illustrative examples of (a) CutBlur and (b) our proposed CutMIB. CutBlur generates augmented SAIs view-by-view via the "cutting-pasting" operation. CutMIB generates the augmented light field via the "cutting-blending-pasting" operation. The implicit geometric information can be utilized during the training stage.

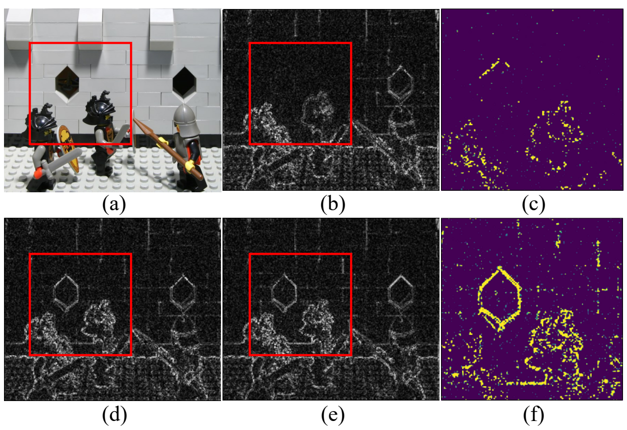

Figure 3. Analyzing CutBlur and CutMIB from a phase spectrum perspective. (a) The center view image in a 5×5 light field. The red rectangle denotes the area for the cutting and pasting operation. (b) The phase spectrum of the original LR center view image. (c) is the calculated residual map between (b) and (d). (d) The phase spectrum of the LR center view image with the pasted LR patch using CutBlur. We cut an HR patch from the HR center view image, and paste it to the LR center view image. (e) The phase spectrum of the LR center view image with the pasted blended patch using CutMIB. We cut all HR patches from the HR light field, blend them, and then paste the blended patch to the LR center view image. (f) is the calculated residual map between (b) and (e).

blends two images to generate an unseen training sample. The effectiveness of the DA strategy on light field SR has received very little attention. Instead, only geometric transformation strategies such as flipping and rotating are used in light field SR. Recently, Yoo *et al.* [60] propose CutBlur, a DA strategy for training a stronger single image SR model, in which a low-resolution (LR) patch is *cut* and *pasted* to the corresponding high-resolution (HR) image region, and vice versa. A straightforward way to utilize the DA strategy on light field SR is to perform CutBlur on each view in a light field and train single image SR networks view by view, as shown in Figure 2(a). However, the ignorance of the inherent correlation in the spatial-angular domain makes it suboptimal. We provide a visual observation using the phase spectrum since it contains rich texture information [46, 62] in Figure 3. Specifically, we use CutBlur on the LR center view (Figure 3(a)) in a 5×5 light field, cut an HR patch, and then paste it to the original LR image, and analyze the phase spectrum of the processed LR image. We can directly observe from the calculated residual map in Figure 3(c) that there is little additional information from the pasted HR patch using the CutBlur strategy. This encourages us to realize the need for a more effective strategy to exploit patches from multiple views.

Based on the aforementioned observation, we propose CutMIB, a novel DA strategy specifically designed for light field SR, as shown in Figure 2(b). Our CutMIB, which is inspired by CutBlur [60], first *cuts* LR patches from different views in an LR light field at the same position. The cut LR patches are then *blended* to generate the blended LR patch, which is then *pasted* to the corresponding areas of various HR light field views, and vice versa. Therefore, each augmented light field pair has partially blended LR and blended HR pixel distributions with a random ratio. By feeding the augmented training pairs into light field SR networks, these networks can not only learn "how" and "where" to super-resolve the LR light field (*i.e.*, benefit from the cutting-blending operation [60]), but also utilize the implicit geometric information in multi-view images, resulting in better performance and higher angular consistency among super-resolved light field views (*i.e.*, benefit from the blending operation [2, 5, 17]). Figure 3(f) illustrates that pasting the blended HR patch to the LR center view (Figure 3(a)) results in more additional details in the pasted area. This demonstrates that our CutMIB can more effectively use multi-view information in a light field.

Thanks to CutMIB, we can improve both the reconstruction quality and the angular consistency of light field SR results while maintaining the network structures unchanged (see Figure 1). Additionally, we verify the effectiveness of the proposed CutMIB on real-world light field SR and light field denoising tasks.

Contributions of this paper are summarized as follows:

(1) We propose a novel DA strategy, CutMIB, to improve the performance of existing light field SR networks. To our best knowledge, it is the first DA strategy for light field SR. Through the "cutting-blending-pasting" operation, CutMIB is designed to efficiently explore the geometric information in light fields during the training stage.

(2) Extensive experiments demonstrate CutMIB can boost the reconstruction fidelity and the angular consistency of existing typical light field SR methods.

(3) We verify the effectiveness of CutMIB on real-world light field SR and light field denoising tasks.

2. Related work

Light field super-resolution. Classic non-learning-based methods utilize projection and optimization techniques to super-resolve the LR observations, relying on geometric [30, 38] and mathematical [1, 57] modeling of the 4D light field structure. Due to their promising performance when trained on large external datasets, CNN-based methods now predominate light field SR. Yoon *et al.* [61] propose the first light field SR network LFCNN by reusing the SRCNN architecture [15] with multiple channels. After that, several CNN-based methods have been designed to exploit across-view redundancy in the 4D light field, either explicitly [26,51,68] or implicitly [34,52,53,56,59,63,72]. Transformer [4, 13, 45] based methods have recently made significant progress in light field SR [31, 32, 47]. Furthermore, Cheng *et al.* [9] address the domain gap issue in light field SR by applying a zero-shot learning framework, which learns the SR mapping function solely from the input. In contrast to these methods, we propose CutMIB to boost the performance of existing light field SR methods without increasing the inference time and changing architectures.

Data augmentation strategies in high-level vision. DA strategies in high-level vision tasks can be roughly divided into the following categories. (1) Geometric transformation, including horizontal flip, vertical flip, and rotation. (2) Photometric transformation, such as color jitter [40]. (3) Information-dropping strategies, including Cutout and random erasing [71] and Cutout [14]. This strategy primarily causes the loss or misunderstanding of spatial information between neighboring pixels. (4) Search-based strategies, including AutoAug [11] and RandAug [12]. The search-based strategy utilizes reinforcement learning to search from a pool of augmentation policies for an optimal combination. (5) Mixing-based strategies, Mixup [66] and CutMix [64]. Mixing-based augmentation employs multi-image information by creating mixed input images with soft labels for training. (6) Feature-level augmentation, such as [44, 49]. (7) GAN-based augmentation, such as [3]. In contrast to the abovementioned strategies, we propose CutMIB, which specializes in light field SR in low-level vision.

Data augmentation strategies in low-level vision. As a pioneering work along the line of DA for low-level vision, Timofte *et al.* [43] propose seven techniques to improve the performance of example-based single image SR, one of which is DA. Consistent improvements are gained across models and datasets using rotation and flipping operations. However, they only utilize simple geometric manipulations with traditional SR models [41, 42, 65] and an early deep method, SRCNN [15]. Feng *et al.* [16] analyze Mixup [66] in single image SR to suppress the model overfitting phenomenon. The paper most related to us is [60], in which a comprehensive analysis of the existing DA strategies is applied to the single image SR task, and a novel DA strategy called CutBlur is proposed. In addition, CutBlur shows promising results in image denoising and JPEG artifact removal. However, when directly applying CutBlur to light field SR, it achieves satisfactory results, neither in reconstruction accuracy nor angular consistency. In this paper, we draw inspiration from CutBlur and then design a novel DA strategy for light fields to fully exploit and utilize the multi-view information during the training stage.

3. Method

3.1. Problem Formulation

Following [51, 53, 56, 59, 63, 68], we convert the input light field from the RGB space to the YCbCr channel and only super-resolve the Y channel images, leaving Cb and Cr channel images being bicubic upscaled. Consequently, without considering the channel dimension, an LR light field can be formulated as a 4D tensor $\mathcal{L}^{LR} \in \mathbb{R}^{U\times V\times H\times W}$, where U and V represent angular dimensions, and H and W represent spatial dimensions. Specifically, an LR light field can be considered as a $U\times V$ array of LR sub-aperture images (SAIs), and the resolution of each LR SAI $\mathcal{I}_i^{LR}$ is $H \times W$, where $i \in [1, U \times V]$. The light field SR task aims at generating a super-resolved light field $\mathcal{L}^{SR} \in \mathbb{R}^{U\times V\times rH\times rW}$ and the resolution of each HR SAI $\mathcal{I}_i^{HR}$ is $rH \times rW$. r denotes the upsampling scale factor. In this paper, we set $r = 2$ and $r = 4$, *i.e.*, we verify the effectiveness of CutMIB on $\times 2$ and $\times 4$ light field SR tasks. The reconstructed HR light field is desired to be close to the ground-truth light field $\mathcal{L}^{HR} \in \mathbb{R}^{U\times V\times rH\times rW}$.

Without loss of generality, a light field SR network aims to learn the mapping function $f(\cdot)$ of an LR light field $\mathcal{L}^{LR}$ to an HR light field $\mathcal{L}^{SR}$, which can be denoted as

$$\mathcal{L}^{SR} = f(\mathcal{L}^{LR}). \tag{1}$$

A light field SR network is optimized with a loss function,

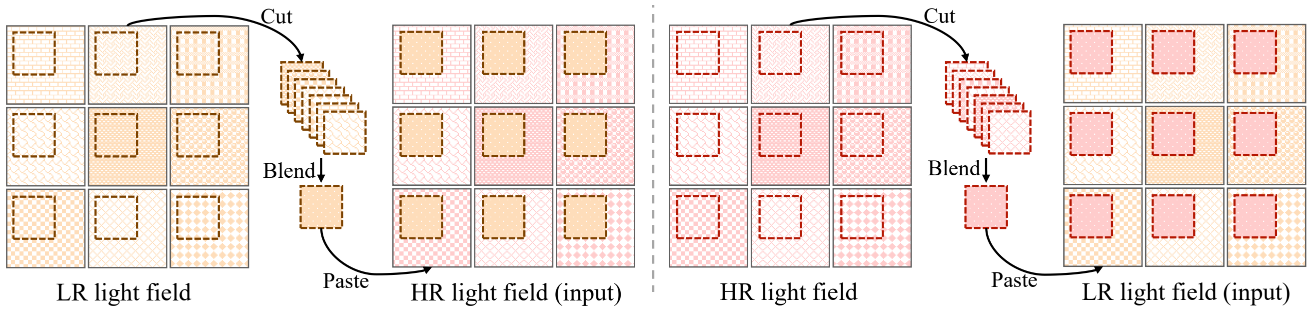

Figure 4. Schematic illustration of our proposed CutMIB strategy tailored for light field. Orange and pink rectangles represent the LR and HR light fields, respectively. Different textures in rectangles represent different views. Best viewed in color.

and L_1 is the most commonly used one. Given a training set $\{\mathcal{L}_i^{LR}, \mathcal{L}_i^{HR}\}_{i=1}^N$, which contains N LR input light fields and their HR counterparts. The goal of training light field SR network is to minimize the L_1 loss function

$$\begin{aligned} L(\Theta) &= \frac{1}{N}\sum_{i=1}^{N}\left\|\mathcal{L}_i^{SR} - \mathcal{L}_i^{HR}\right\|_1 \\ &= \frac{1}{N}\sum_{i=1}^{N}\left\|f\left(\mathcal{L}_i^{LR}\right) - \mathcal{L}_i^{HR}\right\|_1, \end{aligned} \tag{2}$$

where Θ is the parameters of the light field SR network.

3.2. CutMIB

As illustrated in Figure 2(b) and Figure 4, the goal of CutMIB is to generate a pair of new training samples by the cutting-blending-pasting operation:

(1) Cutting: cut the random regions in k SAIs from $\mathcal{L}^{LR}$ and get k LR patches $P^{LR} = \{p_i^{LR}\}_{i=1}^k$

$$p_i^{LR} = \boldsymbol{M} \odot \mathcal{I}_i^{LR}, i \in [1, k], \tag{3}$$

where $k \leq K$, and $\boldsymbol{M} \in \{0,1\}^{H\times W}$ denotes a binary mask indicating where to cut and $\odot$ denotes the element-wise multiplication operation. K indicates all views in an LF and $K = U \times V$.

(2) Blending: k LR patches are blended to get the blended LR patch P_{blend}^{LR}

$$P_{blend}^{LR} = \frac{1}{k}\sum_{i=1}^{k} p_i^{LR}. \tag{4}$$

(3) Pasting: the blended LR patch P_{blend}^{LR} are pasted to the corresponding k HR SAIs $\mathcal{I}_i^{HR}$ and corresponding region. Note that, due to the resolution mismatch of the blended LR patch and the corresponding region in HR SAIs, we utilize the bicubic kernel to upsample P_{blend}^{LR} with the factor of r and get $P_{blend}^{LR_r^{\uparrow}}$.

The HR blended patch and the generated LR SAIs can be achieved similarly. Based on the above operations, we get the augmented SAI as

$$\begin{aligned} \hat{\mathcal{I}}_i^{HR\to LR} &= P_{\text{blend}}^{HR^{\downarrow r}} + (\mathbf{1} - M) \odot \mathcal{I}_i^{LR}, \\ \hat{\mathcal{I}}_i^{LR\to HR} &= P_{\text{blend}}^{LR^{\uparrow r}} + (\mathbf{1} - M) \odot \mathcal{I}_i^{HR}, \end{aligned} \tag{5}$$

where $\mathbf{1}$ is a binary mask filled with ones. Therefore, a pair of new training samples $\{\hat{\mathcal{L}}^{HR\to LR}, \hat{\mathcal{L}}^{LR\to HR}\}$ can be generated by combining $\{\hat{\mathcal{I}}_i^{HR\to LR}, \hat{\mathcal{I}}_i^{LR\to HR}\}$ according to the view position.

3.3. Discussion: Variants of CutMIB

The cutting-blending-pasting operation allows CutMIB to exploit and utilize the multi-view information contained in light field pairs during the training stage. So naturally, we have two questions here: (1) which SAIs should be considered in CutMIB, and (2) how many SAIs should be considered in CutMIB? We design several variants to answer the above two questions, as shown in Figure 5.

Given a highly redundant light field, SAIs are stacked along four specific directions into different view stacks, reducing computational costs significantly while maintaining high performance in the light field SR task. These four specific directions, namely horizontal (θ_1, $\theta = 0°$), vertical (θ_2, $\theta = 90°$), main diagonal (θ_3, $\theta = 45°$) and antidiagonal (θ_4, $\theta = 135°$), contain the implicit geometric information along one direction of the light field. The variants of CutMIB consider SAIs along these four directions and their typical combinations to perform the cutting-blending-pasting operation. In addition, we consider the random selection of SAIs for CutMIB. Detailed results for the variants of CutMIB are shown in Section 4.2.

4. Experiments

4.1. Experimental Settings

Network structures. We adopt several advanced and typical light field SR networks to verify the effectiveness of our CutMIB strategy. We first consider CNN-based methods, including ATO [27], InterNet [31], IINet [33], and DistgSSR [52], in which well-designed CNN-based structures

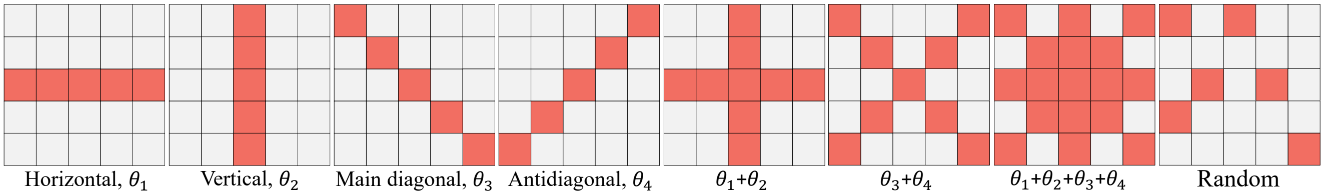

Figure 5. Variants of CutMIB on a light field with 5×5 angular resolution. Red rectangles denote the views in a light field utilized in CutMIB (*i.e.*, the cutting-blending-pasting operationsss). Gray rectangles denote the views trained without the CutMIB strategy.

Table 1. Training InterNet [53] with different variants of CutMIB (SAIs at different typical positions and different number of SAIs).

Method	HCInew		HCIold		INRIA		STFgantry		EPFL		Average	
	PSNR	Δ	PSNR	Δ	PSNR	Δ	PSNR	Δ	PSNR	Δ	PSNR	Δ
InterNet	30.942	-	37.104	-	30.743	-	30.343	-	28.773	-	30.440	-
θ_1	30.978	**+0.036**	37.143	**+0.039**	30.793	**+0.050**	30.370	**+0.027**	28.838	**+0.065**	30.491	**+0.051**
θ_2	30.972	**+0.030**	37.141	**+0.037**	30.766	**+0.023**	30.368	**+0.025**	28.817	**+0.044**	30.474	**+0.034**
θ_3	30.977	**+0.035**	37.148	**+0.044**	30.793	**+0.050**	30.367	**+0.024**	28.840	**+0.067**	30.492	**+0.052**
θ_4	30.979	**+0.037**	37.152	**+0.048**	30.795	**+0.052**	30.368	**+0.025**	28.838	**+0.065**	30.492	**+0.052**
θ_1+θ_2	30.980	**+0.038**	37.150	**+0.046**	30.792	**+0.049**	30.375	**+0.032**	28.834	**+0.061**	30.490	**+0.050**
θ_3+θ_4	30.981	**+0.039**	37.147	**+0.043**	30.798	**+0.055**	30.369	**+0.026**	28.844	**+0.071**	30.495	**+0.055**
θ_1+θ_2+θ_3+θ_4	31.002	**+0.060**	37.177	**+0.073**	30.804	**+0.061**	30.445	**+0.102**	28.846	**+0.073**	30.510	**+0.070**
$k = 1$	30.942	0.000	37.111	**+0.007**	30.736	**-0.007**	30.338	**-0.005**	28.777	**+0.004**	30.440	0.000
$k = 10$	30.999	**+0.057**	37.173	**+0.069**	30.808	**+0.065**	30.441	**+0.098**	28.847	**+0.074**	30.510	**+0.070**
$k = 20$	31.006	**+0.064**	37.179	**+0.075**	30.808	**+0.065**	30.449	**+0.106**	28.852	**+0.079**	30.515	**+0.075**
InterNet	31.009	**+0.067**	37.184	**+0.079**	30.813	**+0.069**	30.460	**+0.117**	28.856	**+0.082**	30.519	**+0.080**

are proven effective on light field SR. Transformer-based method, DPT [47], is also adopted in our experiments.

Training settings and implementation details. We follow the same experimental setting as in [31,31,33,47,56] and retrain selected networks on the mixed datasets [22,28,37,58] from scratch based on their publicly available codes. In total, 144 scenes are used for training and 23 for testing. The training and testing LR light fields are generated by bicubic downsampling with MATLAB. We keep each method's training hyper-parameters (*e.g.*, learning rate and batch size) the same as reported in the original paper. We crop 320×320 patches for training light field SR methods. The spatial size of the patch used in CutMIB is randomly set to $16 \sim 72$. All experiments are conducted using PyTorch on two NVIDIA 1080Ti GPUs.

Inference settings. We evaluate our proposed CutMIB strategy and its variants on several benchmarks from BasicLFSR[1], including EPFL, HCInew, HCIold, INRIA, and STFgantry. We utilize PSNR on the Y channel to evaluate the performance of different methods. To quantitatively evaluate the results, we choose PSNR on the Y channel as the main metric.

4.2. Results of CutMIB Variants

To determine which SAIs and how many SAIs should be considered in CutMIB for better performance, we explore the impact of different variants of CutMIB on the performance of InterNet in terms of $\times 4$ light field SR.

Which SAIs should be considered in CutMIB? We show several typical views and their combinations in the upper part of Table 1. We can observe two phenomena from the table. (1) The views along specific directions do not matter in CutMIB, although the implicit geometric information of these views is practical in other tasks [20,25,39,68]. As can be seen from the table, various views along different directions remain the same as the final results. For instance, the difference between the average results of θ_1 and θ_3 is only 0.001dB. (2) Involving more views yields better results. As we can see, $\theta_1 + \theta_2 + \theta_3 + \theta_4$ gains over θ_1 with about 0.020dB. This is consistent with Section 1 analysis because more SAIs can provide more implicit geometric information during the training stage.

How many SAIs should be considered in CutMIB? As in the lower part of Table 1, the more views (*i.e.*, the larger the k), the better performance. This is consistent with the conclusion obtained above. Therefore, we utilize all views in CutMIB to achieve the best results.

4.3. Results of CutMIB

We first test our proposed CutMIB and CutBlur [60] on various benchmark datasets in Table 2. The table shows that as compared to CutBlur, the networks trained with CutMIB achieve greater reconstruction performance. Taking ATO as an example, we can see that training with CutMIB results in

[1]https://github.com/ZhengyuLiang24/BasicLFSR

Table 2. Quantitative comparison with existing light field SR methods. We show the PSNR (dB, ↑) results for $\times 2$ and $\times 4$ light field SR tasks on benchmark datasets. We compare the baseline methods with the methods trained with CutBlur and CutMIB.

Method	HCInew		HCIold		INRIA		STFgantry		EPFL		Average	
$\times 2$	PSNR	Δ	PSNR	Δ	PSNR	Δ	PSNR	Δ	PSNR	Δ	PSNR	Δ
ATO	37.170	-	43.956	-	36.133	-	39.456	-	34.228	-	36.454	-
ATO	37.236	+0.066	44.210	+0.254	36.196	+0.063	39.609	+0.153	34.294	+0.066	36.543	+0.089
ATO	37.324	+0.154	44.233	+0.277	36.350	+0.217	39.929	+0.472	34.442	+0.214	36.686	+0.232
InterNet	37.072	-	44.290	-	35.671	-	38.169	-	33.904	-	36.113	-
InterNet	37.236	+0.164	44.546	+0.256	36.104	+0.433	38.557	+0.387	34.320	+0.416	36.473	+0.360
InterNet	37.320	+0.248	44.590	+0.300	36.043	+0.372	39.004	+0.834	34.300	+0.397	36.508	+0.395
IINet	37.690	-	44.664	-	36.536	-	39.595	-	34.667	-	36.897	-
IINet	37.749	+0.059	44.727	+0.063	36.510	-0.026	39.780	+0.186	34.689	+0.021	36.933	+0.036
IINet	37.836	+0.146	44.746	+0.082	36.519	-0.018	40.264	+0.669	34.755	+0.087	37.022	+0.125
DPT	37.288	-	44.057	-	36.381	-	39.342	-	34.480	-	36.637	-
DPT	37.353	+0.066	44.274	+0.217	36.403	+0.022	39.455	+0.113	34.485	+0.005	36.684	+0.047
DPT	37.471	+0.183	44.340	+0.283	36.476	+0.095	39.738	+0.396	34.562	+0.082	36.784	+0.147
DistgSSR	37.956	-	44.917	-	36.579	-	40.360	-	34.802	-	37.100	-
DistgSSR	37.952	-0.004	44.898	-0.019	36.569	-0.010	40.335	-0.025	34.792	-0.001	37.089	-0.011
DistgSSR	37.967	+0.011	44.919	+0.002	36.575	-0.004	40.380	+0.020	34.811	+0.009	37.107	+0.007
Method	HCInew		HCIold		INRIA		STFgantry		EPFL		Average	
$\times 4$	PSNR	Δ	PSNR	Δ	PSNR	Δ	PSNR	Δ	PSNR	Δ	PSNR	Δ
ATO	30.813	-	36.893	-	30.677	-	30.573	-	28.515	-	30.292	-
ATO	30.896	+0.083	36.991	+0.098	30.798	+0.121	30.585	+0.012	28.612	+0.097	30.385	+0.093
ATO	30.950	+0.138	37.050	+0.157	30.829	+0.152	30.771	+0.198	28.631	+0.116	30.430	+0.138
InterNet	30.942	-	37.104	-	30.743	-	30.343	-	28.773	-	30.440	-
InterNet	30.980	+0.039	37.149	+0.045	30.800	+0.056	30.371	+0.028	28.845	+0.072	30.496	+0.056
InterNet	31.009	+0.067	37.184	+0.079	30.813	+0.069	30.460	+0.117	28.856	+0.082	30.519	+0.080
IINet	31.313	-	37.547	-	31.086	-	31.198	-	29.005	-	30.792	-
IINet	31.357	+0.044	37.595	+0.047	31.026	-0.060	31.300	+0.102	29.046	+0.041	30.818	+0.025
IINet	31.422	+0.109	37.613	+0.066	31.089	+0.003	31.450	+0.252	29.106	+0.102	30.884	+0.092
DPT	31.135	-	37.212	-	30.924	-	31.060	-	28.881	-	30.631	-
DPT	31.177	+0.043	37.364	+0.151	30.959	+0.034	31.092	+0.032	28.940	+0.059	30.688	+0.057
DPT	31.279	+0.144	37.385	+0.173	31.072	+0.147	31.334	+0.274	29.029	+0.148	30.791	+0.160
DistgSSR	31.410	-	37.588	-	31.015	-	31.635	-	29.015	-	30.840	-
DistgSSR	31.418	+0.008	37.599	+0.011	31.022	+0.007	31.638	+0.003	29.023	+0.008	30.847	+0.008
DistgSSR	31.445	+0.035	37.618	+0.030	31.030	+0.015	31.681	+0.046	29.034	+0.019	30.864	+0.024

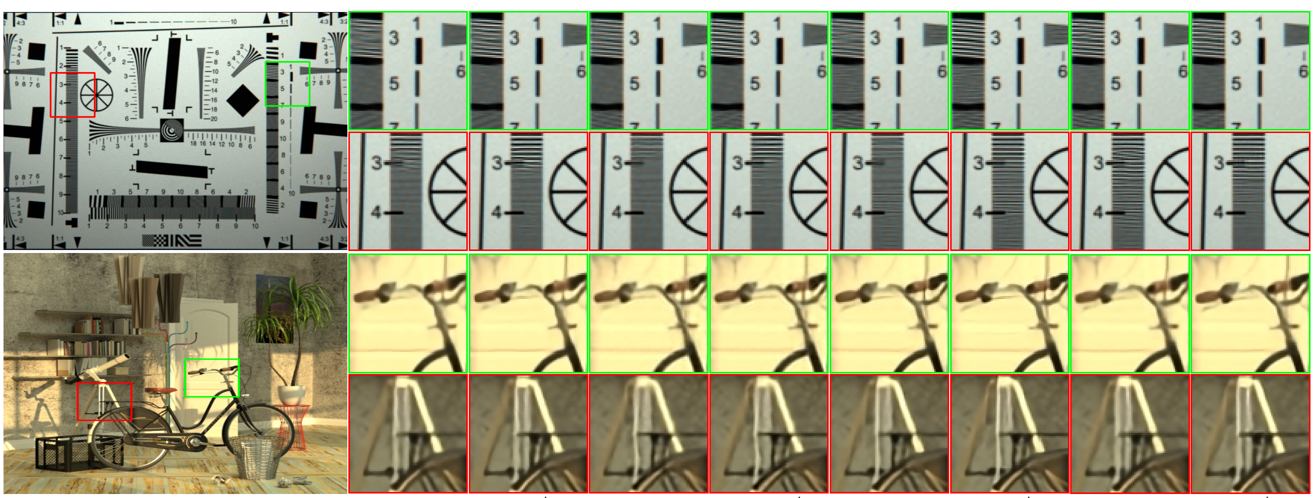

Figure 6. Visual comparisons of different models trained without and with the proposed CutMIB on $\times 2$ and $\times 4$ light field SR. † means the networks are trained with the CutMIB. Please zoom in for better visualization and best viewed on the screen.

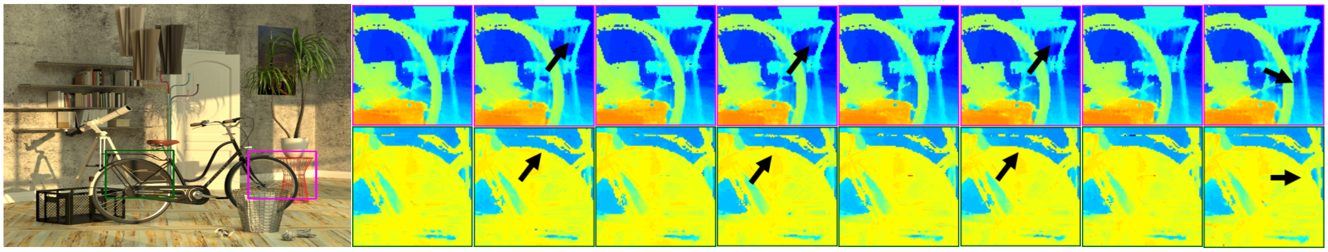

Figure 7. Disparity estimation results achieved by SPO [70] using $\times 2$ light field SR results generated by different methods. Please zoom in for better visualization and best viewed on the screen.

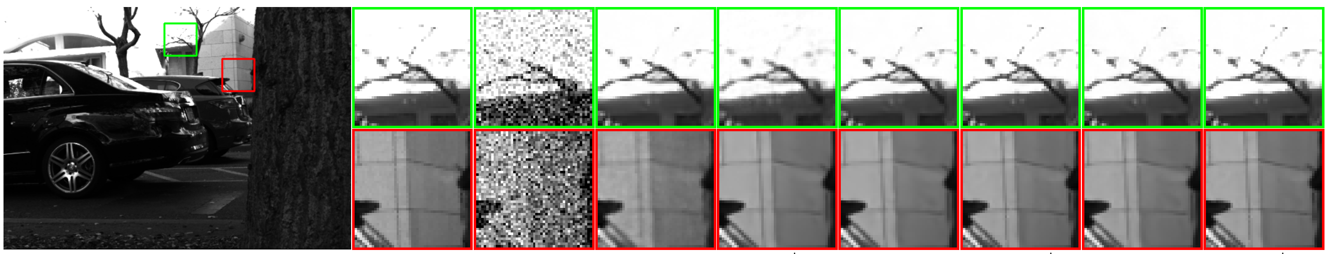

Figure 8. Visual comparisons (view coordinates: (4, 4) in an 8×8 light field) of different models trained without and with the proposed CutMIB on light field denoising under the setting of $\sigma = 50$. Please zoom in for better visualization and best viewed on the screen.

Table 3. Quantitative comparison with existing light field denoising methods. We show the PSNR (dB, ↑) results for $\sigma = 10$, $\sigma = 20$, and $\sigma = 50$ light field denoising task. We compare the baseline methods with the methods trained with CutMIB .

Method	$\sigma = 10$		$\sigma = 20$		$\sigma = 50$	
	PSNR	Δ	PSNR	Δ	PSNR	Δ
ATO	42.509	-	39.957	-	36.057	-
ATO	42.612	**+0.103**	40.051	**+0.094**	36.202	**+0.145**
InterNet	42.706	-	39.997	-	35.920	-
InterNet	42.776	**+0.070**	40.112	**+0.115**	36.152	**+0.232**
IINet	43.164	-	40.481	-	36.525	-
IINet	43.177	**+0.013**	40.613	**+0.132**	36.606	**+0.081**

a 0.319 dB gain compared with training using CutBlur does on the STFGantry dataset.

Table 2 also shows quantitative comparisons of the methods trained without and with CutMIB on benchmark datasets. The table shows that existing light field SR networks trained with CutMIB outperform the corresponding baselines by a considerable margin. Specifically, on $\times 4$ light field SR of HCInew, ATO trained with CutMIB achieves 30.950 dB (PSNR), while the same network trained without CutMIB only gets 30.813 dB (PSNR). Also, as shown in Figure 1, methods trained with our CutMIB obtain results with better angular consistency.

Qualitative results on $\times 2$ and $\times 4$ light field SR are presented in Figure 6. It is clear that the networks trained with CutMIB provide better qualitative results than their baselines, with more accurate details and fewer blurs (such as the area of the bicycle handlebar and the telescope mount). Since HR and angular-consistent light fields are beneficial to disparity estimation, we apply SPO [70] to estimate disparity maps of the super-resolved light field images to perform disparity estimation. As can be seen in Figure 7, the disparity maps estimated by the results generated from the networks trained with CutMIB have sharper edges and more accurate results, indicating the effectiveness of our CutMIB.

4.4. Extensions: Applications of CutMIB

Light field denoising. We validate the effectiveness of CutMIB on the light field denoising task. We follow the same settings in [6, 21], and generate the datasets for training and testing based on Stanford Lytro Light Field Archive. We center-crop each scene and set the angular resolution of each light field equal to 8×8. Zero-mean Gaussian noise with a standard variance of $\sigma = 10$, $\sigma = 20$, and $\sigma = 50$ are synthesized for each training data and testing data. We retrain ATO, InterNet, and IINet on the generated datasets under three different settings. We change the input angular number from 5 to 8, and remove the upsampling operation in each network for light field denoising. The average PSNR results between the denoised light fields and ground-truth ones are used to evaluate different methods quantitatively in Table 3. As can be seen in the table, all the methods trained with our proposed CutMIB achieve higher reconstruction fidelity than their baseline versions, validating the advantage of our CutMIB. We also show the qualitative results in Figure 8 at the level of $\sigma = 50$. Networks trained with CutMIB can generate more realistic and accurate details, especially in the area of branches and walls.

Real-world light field | Real-world | ATO | ATO† | InterNet | InterNet† | IINet | IINet† | DPT | DPT†

Figure 9. Visual comparisons of different methods on $\times 4$ real-world light field SR. Please zoom in for better visualization.

Table 4. Quantitative comparison with $\times 4$ light field SR methods under different isotropic Gaussian kernels on EPFL. We compare the baseline methods with the methods trained with CutMIB.

Method	$k = 1.8$		$k = 2.5$		$k = 3.2$	
	PSNR	Δ	PSNR	Δ	PSNR	Δ
ATO	26.776	-	25.260	-	24.243	-
ATO	26.851	**+0.075**	25.289	**+0.029**	24.263	**+0.020**
IINet	26.884	-	25.299	-	24.252	-
IINet	26.898	**+0.014**	25.308	**+0.009**	24.265	**+0.013**
InterNet	26.887	-	25.310	-	24.267	-
InterNet	26.937	**+0.050**	25.324	**+0.014**	24.274	**+0.007**
DPT	26.781	-	25.200	-	24.174	-
DPT	26.781	0.000	25.210	**+0.010**	24.182	**+0.008**

Light field SR under isotropic gaussian kernels. Existing light field SR networks are trained and evaluated on simulated datasets that assume simple and uniform degradation (*i.e.*, bicubic degradation). Degradations in real applications are much more complicated. Here we evaluate light field SR networks on light fields degraded using isotropic gaussian kernels to measure the generalizability. The degradation process of a light field can be denoted as [50]

$$\mathcal{I}_i^{LR} = (\mathcal{I}_i^{HR} \otimes k_i) \downarrow_r + n_i, \tag{6}$$

where $\otimes$ represents the convolution operation, and $\downarrow_r$ represents downsampling by a factor r. n_i is the real-world noise. Following the setting in [18, 24], the kernel size of k_i is set as 21, and the kernel widths are set to 1.8, 2.5, and 3.2 for evaluation. Note that this work focuses on not designing a blind light field SR network; we set each view to have the same degradation kernel and do not account for noise. As can be seen in Table 4, all results drop a lot. This is mainly because we do not explicitly address the degradation kernel mismatch issue. Still, networks trained with CutMIB generate results with higher PSNR values.

Real-world light field SR. We also conduct experiments to prove that networks trained with CutMIB can generalize well to real-world light fields. We choose a real-world light field from the HFUT dataset [67], and we super-resolve this light field directly using the baseline networks and their corresponding versions trained with CutMIB. As can be seen in Figure 9, networks trained with CutMIB produce visually more promising results with clearer details.

Table 5. PSNR (dB) comparison of different data augmentation strategies in light field SR.

DA strategy	Cutout (8px)	Mixup	Blend	RGB permute
InterNet	Average result: 30.440			
$k = 5$	+0.049	+0.043	+0.045	+0.042
$k = 10$	+0.026	+0.046	+0.043	+0.038
$k = 20$	+0.031	+0.050	+0.049	+0.011
$k = 25$	+0.001	+0.025	+0.045	+0.009

4.5. Discussion: Other DA Strategies

In this section, we evaluate the performance of InterNet [53] trained with different typical DA strategies using k random views. (1) Cutout [14]: we randomly erase 8 pixels in k views of a light field. (2) Mixup [66]: we blend two images on k views to generate an unseen training sample. (3) Blend: we blend image with vector $\boldsymbol{v} = (v1, v2, v3)$, where $v_i \sim$ Unif$(0.6, 1)$. (4) RGB permute: we randomly permute RGB channels. Results are shown in Figure 5. It can be seen that these DA strategies can improve the performance of InterNet, although they are not designed for the task of light field SR. We can also find that, performance gains of these DA strategies are not as significant as CutMIB, demonstrating CutMIB can make better use of multi-view information.

5. Conclusion

In this work, we propose a novel data augmentation strategy, *i.e.*, CutMIB, for light field SR. Our CutMIB is able to train better light field SR networks without changing their structures or post-processing operations. We demonstrate the effectiveness and versatility of our proposed CutMIB on the light field SR task, which can achieve improved reconstruction quality and better angular consistency. We also verify the effectiveness of CutMIB on light field denoising and real-world light field SR. In future work, we will explore other data augmentation strategies for light field SR to further improve the performance of existing networks. In addition, we will extend CutMIB to other high-dimensional data processing tasks.

Acknowledgement

This work was supported in part by the National Natural Science Foundation of China under Grants 62131003 and 62021001.

References

[1] Martin Alain and Aljosa Smolic. Light field super-resolution via lfbm5d sparse coding. In *ICIP*, 2018. 1, 3

[2] Vivek Boominathan, Kaushik Mitra, and Ashok Veeraraghavan. Improving resolution and depth-of-field of light field cameras using a hybrid imaging system. In *ICCP*, 2014. 2

[3] Christopher Bowles, Liang Chen, Ricardo Guerrero, Paul Bentley, Roger Gunn, Alexander Hammers, David Alexander Dickie, Maria Valdés Hernández, Joanna Wardlaw, and Daniel Rueckert. Gan augmentation: Augmenting training data using generative adversarial networks. *arXiv preprint arXiv:1810.10863*, 2018. 3

[4] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *NeurIPS*, 2020. 3

[5] Gaurav Chaurasia, Sylvain Duchene, Olga Sorkine-Hornung, and George Drettakis. Depth synthesis and local warps for plausible image-based navigation. *ACM Transactions on Graphics (TOG)*, 32(3):1–12, 2013. 2

[6] Jie Chen, Junhui Hou, and Lap-Pui Chau. Light field denoising via anisotropic parallax analysis in a cnn framework. *IEEE Signal Processing Letters*, 25(9):1403–1407, 2018. 7

[7] Zhen Cheng, Yutong Liu, and Zhiwei Xiong. Spatial-angular versatile convolution for light field reconstruction. *IEEE Transactions on Computational Imaging*, 8:1131–1144, 2022. 1

[8] Zhen Cheng, Zhiwei Xiong, Chang Chen, and Dong Liu. Light field super-resolution: a benchmark. In *CVPRW*, 2019. 1

[9] Zhen Cheng, Zhiwei Xiong, Chang Chen, Dong Liu, and Zheng-Jun Zha. Light field super-resolution with zero-shot learning. In *CVPR*, 2021. 1, 3

[10] Zhen Cheng, Zhiwei Xiong, and Dong Liu. Light field super-resolution by jointly exploiting internal and external similarities. *IEEE Transactions on Circuits and Systems for Video Technology*, 30(8):2604–2616, 2019. 1

[11] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation policies from data. *arXiv preprint arXiv:1805.09501*, 2018. 3

[12] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated data augmentation with a reduced search space. In *CVPRW*, 2020. 3

[13] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *NAACL*, 2018. 3

[14] Terrance DeVries and Graham W Taylor. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017. 1, 3, 8

[15] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Learning a deep convolutional network for image super-resolution. In *ECCV*, 2014. 3

[16] Ruicheng Feng, Jinjin Gu, Yu Qiao, and Chao Dong. Suppressing model overfitting for image super-resolution networks. In *CVPR Workshops*, 2019. 3

[17] Todor G Georgiev, Ke Colin Zheng, Brian Curless, David Salesin, Shree K Nayar, and Chintan Intwala. Spatio-angular resolution tradeoffs in integral photography. *Rendering Techniques*, 2006(263-272):21, 2006. 2

[18] Jinjin Gu, Hannan Lu, Wangmeng Zuo, and Chao Dong. Blind super-resolution with iterative kernel correction. In *CVPR*, 2019. 8

[19] Laurent Guillo, Xiaoran Jiang, Gauthier Lafruit, and Christine Guillemot. *Light field video dataset captured by a R8 Raytrix camera (with disparity maps)*. PhD thesis, INTERNATIONAL ORGANISATION FOR STANDARDISATION ISO/IEC JTC1/SC29/WG1 & WG11, 2018. 1

[20] Bichuan Guo, Jiangtao Wen, and Yuxing Han. Deep material recognition in light-fields via disentanglement of spatial and angular information. In *ECCV*, 2020. 5

[21] Mantang Guo, Junhui Hou, Jing Jin, Jie Chen, and Lap-Pui Chau. Deep spatial-angular regularization for light field imaging, denoising, and super-resolution. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. 7

[22] Katrin Honauer, Ole Johannsen, Daniel Kondermann, and Bastian Goldluecke. A dataset and evaluation methodology for depth estimation on 4d light fields. In *ACCV*, 2016. 5

[23] Xinjue Hu, Chenchen Wang, Yuxuan Pan, Yunming Liu, Yumei Wang, Yu Liu, Lin Zhang, and Shervin Shirmohammadi. 4dlfvd: A 4d light field video dataset. In *Proceedings of the 12th ACM Multimedia Systems Conference*, MMSys '21, pages 287–292, 2021. 1

[24] Yan Huang, Shang Li, Liang Wang, Tieniu Tan, et al. Unfolding the alternating optimization for blind super resolution. *NeurIPS*, 2020. 8

[25] Zhicong Huang, Xuemei Hu, Zhou Xue, Weizhu Xu, and Tao Yue. Fast light-field disparity estimation with multi-disparity-scale cost aggregation. In *ICCV*, 2021. 5

[26] Jing Jin, Junhui Hou, Jie Chen, and Sam Kwong. Light field spatial super-resolution via deep combinatorial geometry embedding and structural consistency regularization. In *CVPR*, 2020. 1, 3

[27] Jing Jin, Junhui Hou, Hui Yuan, and Sam Kwong. Learning light field angular super-resolution via a geometry-aware network. In *AAAI*, 2020. 1, 4

[28] Mikael Le Pendu, Xiaoran Jiang, and Christine Guillemot. Light field inpainting propagation via low rank matrix completion. *IEEE Transactions on Image Processing*, 27(4):1981–1993, 2018. 5

[29] Anat Levin, William T Freeman, and Frédo Durand. Understanding camera trade-offs through a bayesian analysis of light field projections. In *ECCV*, 2008. 1

[30] Chia-Kai Liang and Ravi Ramamoorthi. A light transport framework for lenslet light field cameras. *ACM Transactions on Graphics (TOG)*, 34(2):1–19, 2015. 3

[31] Zhengyu Liang, Yingqian Wang, Longguang Wang, Jungang Yang, and Shilin Zhou. Light field image super-resolution with transformers. *IEEE Signal Processing Letters*, 2022. 1, 3, 4, 5

[32] Zhengyu Liang, Yingqian Wang, Longguang Wang, Jungang Yang, Shilin Zhou, and Yulan Guo. Learning non-

local spatial-angular correlation for light field image super-resolution. *arXiv preprint arXiv:2302.08058*, 2023. 3

[33] Gaosheng Liu, Huanjing Yue, Jiamin Wu, and Jingyu Yang. Intra-inter view interaction network for light field image super-resolution. *IEEE Transactions on Multimedia*, 2021. 1, 4, 5

[34] Nan Meng, Hayden K-H So, Xing Sun, and Edmund Y Lam. High-dimensional dense residual convolutional neural network for light field reconstruction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(3):873–886, 2019. 3

[35] Ren Ng, Marc Levoy, Mathieu Brédif, Gene Duval, Mark Horowitz, and Pat Hanrahan. *Light field photography with a hand-held plenoptic camera*. PhD thesis, Stanford University, 2005. 1

[36] Jiayong Peng, Zhiwei Xiong, Yicheng Wang, Yueyi Zhang, and Dong Liu. Zero-shot depth estimation from light field using a convolutional neural network. *IEEE Transactions Computational Imaging*, 6:682–696, 2020. 1

[37] Martin Rerabek and Touradj Ebrahimi. New light field image dataset. In *QoMEX*, 2016. 5

[38] Mattia Rossi and Pascal Frossard. Graph-based light field super-resolution. In *MMSP*, 2017. 1, 3

[39] Changha Shin, Hae-Gon Jeon, Youngjin Yoon, In So Kweon, and Seon Joo Kim. Epinet: A fully-convolutional neural network using epipolar geometry for depth from light field images. In *CVPR*, 2018. 5

[40] Ryo Takahashi, Takashi Matsubara, and Kuniaki Uehara. Ricap: Random image cropping and patching data augmentation for deep cnns. In Jun Zhu and Ichiro Takeuchi, editors, *Proceedings of The 10th Asian Conference on Machine Learning*, volume 95 of *Proceedings of Machine Learning Research*, pages 786–798. PMLR, 14–16 Nov 2018. 3

[41] Radu Timofte, Vincent De Smet, and Luc Van Gool. Anchored neighborhood regression for fast example-based super-resolution. In *ICCV*, 2013. 3

[42] Radu Timofte, Vincent De Smet, and Luc Van Gool. A+: Adjusted anchored neighborhood regression for fast super-resolution. In *ACCV*, 2014. 3

[43] Radu Timofte, Rasmus Rothe, and Luc Van Gool. Seven ways to improve example-based single image super resolution. In *CVPR*, 2016. 3

[44] Paul Upchurch, Jacob Gardner, Geoff Pleiss, Robert Pless, Noah Snavely, Kavita Bala, and Kilian Weinberger. Deep feature interpolation for image content changes. In *CVPR*, 2017. 1, 3

[45] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *NeurIPS*, 2017. 3

[46] Le Wang and Shengmei Zhao. Super resolution ghost imaging based on fourier spectrum acquisition. *Optics and Lasers in Engineering*, 139:106473, 2021. 2

[47] Shunzhou Wang, Tianfei Zhou, Yao Lu, and Huijun Di. Detail-preserving transformer for light field image super-resolution. *AAAI*, 2022. 1, 3, 5

[48] Ting-Chun Wang, Jun-Yan Zhu, Nima Khademi Kalantari, Alexei A Efros, and Ravi Ramamoorthi. Light field video capture using a learning-based hybrid imaging system. *ACM Transactions Graphics*, 36(4):1–13, 2017. 1

[49] Yulin Wang, Gao Huang, Shiji Song, Xuran Pan, Yitong Xia, and Cheng Wu. Regularizing deep networks with semantic data augmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. 1, 3

[50] Yingqian Wang, Zhengyu Liang, Longguang Wang, Jungang Yang, Wei An, and Yulan Guo. Learning a degradation-adaptive network for light field image super-resolution. *arXiv preprint arXiv:2206.06214*, 2022. 8

[51] 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 Transactions on Image Processing*, 27(9):4274–4286, 2018. 3

[52] Yingqian Wang, Longguang Wang, Gaochang Wu, Jungang Yang, Wei An, Jingyi Yu, and Yulan Guo. Disentangling light fields for super-resolution and disparity estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. 1, 3, 4

[53] Yingqian Wang, Longguang Wang, Jungang Yang, Wei An, Jingyi Yu, and Yulan Guo. Spatial-angular interaction for light field image super-resolution. In *ECCV*, 2020. 1, 3, 5, 8

[54] Yingqian Wang, Tianhao Wu, Jungang Yang, Longguang Wang, Wei An, and Yulan Guo. Deoccnet: Learning to see through foreground occlusions in light fields. In *WACV*, 2020. 1

[55] Yingqian Wang, Jungang Yang, Yulan Guo, Chao Xiao, and Wei An. Selective light field refocusing for camera arrays using bokeh rendering and superresolution. *IEEE Signal Processing Letters*, 26(1):204–208, 2018. 1

[56] Yingqian Wang, Jungang Yang, Longguang Wang, Xinyi Ying, Tianhao Wu, Wei An, and Yulan Guo. Light field image super-resolution using deformable convolution. *IEEE Transactions on Image Processing*, 30:1057–1071, 2020. 1, 3, 5

[57] Sven Wanner and Bastian Goldluecke. Variational light field analysis for disparity estimation and super-resolution. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(3):606–619, 2013. 3

[58] Sven Wanner, Stephan Meister, and Bastian Goldluecke. Datasets and benchmarks for densely sampled 4d light fields. In *VMV*, volume 13, pages 225–226, 2013. 5

[59] Henry Wing Fung Yeung, Junhui Hou, Xiaoming Chen, Jie Chen, Zhibo Chen, and Yuk Ying Chung. Light field spatial super-resolution using deep efficient spatial-angular separable convolution. *IEEE Transactions on Image Processing*, 28(5):2319–2330, 2018. 3

[60] Jaejun Yoo, Namhyuk Ahn, and Kyung-Ah Sohn. Rethinking data augmentation for image super-resolution: A comprehensive analysis and a new strategy. In *CVPR*, 2020. 2, 3, 5

[61] 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 *ICCVW*, 2015. 3

[62] Hu Yu, Naishan Zheng, Man Zhou, Jie Huang, Zeyu Xiao, and Feng Zhao. Frequency and spatial dual guidance for image dehazing. In *ECCV*, 2022. 2

[63] Yan Yuan, Ziqi Cao, and Lijuan Su. Light-field image super-resolution using a combined deep cnn based on epi. *IEEE Signal Processing Letters*, 25(9):1359–1363, 2018. 3

[64] Sangdoo Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *ICCV*, 2019. 1, 3

[65] Roman Zeyde, Michael Elad, and Matan Protter. On single image scale-up using sparse-representations. In *International conference on curves and surfaces*, pages 711–730, 2010. 3

[66] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *ICLR*, 2018. 1, 3, 8

[67] Jun Zhang, Yamei Liu, Shengping Zhang, Ronald Poppe, and Meng Wang. Light field saliency detection with deep convolutional networks. *IEEE Transactions on Image Processing*, 29:4421–4434, 2020. 8

[68] Shuo Zhang, Youfang Lin, and Hao Sheng. Residual networks for light field image super-resolution. In *CVPR*, 2019. 3, 5

[69] Shuo Zhang, Zeqi Shen, and Youfang Lin. Removing foreground occlusions in light field using micro-lens dynamic filter. In *IJCAI*, 2021. 1

[70] Shuo Zhang, Hao Sheng, Chao Li, Jun Zhang, and Zhang Xiong. Robust depth estimation for light field via spinning parallelogram operator. *Computer Vision and Image Understanding*, 145:148–159, 2016. 1, 7

[71] Zhun Zhong, Liang Zheng, Guoliang Kang, Shaozi Li, and Yi Yang. Random erasing data augmentation. In *AAAI*, 2020. 1, 3

[72] Hao Zhu, Mantang Guo, Hongdong Li, Qing Wang, and Antonio Robles-Kelly. Revisiting spatio-angular trade-off in light field cameras and extended applications in super-resolution. *IEEE Transactions on Visualization and Computer Graphics*, 27(6):3019–3033, 2019. 3