

# Mamba-based Light Field Super-Resolution with Efficient Subspace Scanning

Ruisheng Gao , Zeyu Xiao , and Zhiwei Xiong  

University of Science and Technology of China, Hefei, China  
{grsmc4180, zeyuxiao}@mail.ustc.edu.cn, zwxiong@ustc.edu.cn

**Abstract.** Transformer-based methods have demonstrated impressive performance in 4D light field (LF) super-resolution by effectively modeling long-range spatial-angular correlations, but their quadratic complexity hinders the efficient processing of high resolution 4D inputs, resulting in slow inference speed and high memory cost. As a compromise, most prior work adopts a patch-based strategy, which fails to leverage the full information from the entire input LFs. The recently proposed selective state-space model, Mamba, has gained popularity for its efficient long-range sequence modeling. In this paper, we propose a Mamba-based Light Field Super-Resolution method, named MLFSR, by designing an efficient subspace scanning strategy. Specifically, we tokenize 4D LFs into subspace sequences and conduct bi-directional scanning on each subspace. Based on our scanning strategy, we then design the Mamba-based Global Interaction module to capture global information and the Spatial-Angular Modulator to complement local details. Additionally, we introduce a Transformer-to-Mamba loss to further enhance overall performance. Extensive experiments on public benchmarks demonstrate that MLFSR surpasses CNN-based models and rivals Transformer-based methods in performance while maintaining higher efficiency. With quicker inference speed and reduced memory demand, MLFSR facilitates full-image processing of high-resolution 4D LFs with enhanced performance. Our code is available at <https://github.com/RSGao/MLFSR>.

**Keywords:** Light field · Super-resolution · State space model.

## 1 Introduction

Light field (LF) imaging technique enables the capture of both intensity and direction of light rays, therefore finding wide applications in virtual reality and computational photography. However, the design of commercialized LF cameras (*e.g.*, Lytro) is mainly based on the micro-lens array. This design divides the main lens into sub-apertures, introducing a trade-off between the spatial and angular resolution of LFs. Therefore, LF Super-Resolution (LFSR) has become an important research topic to address these challenges.

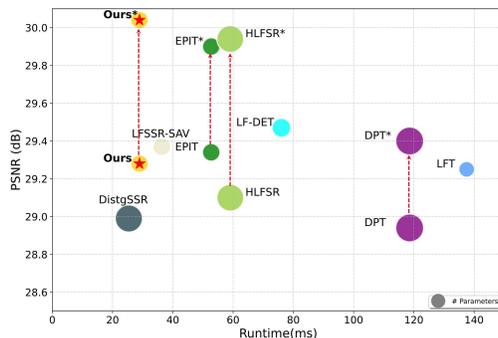
The spatial-angular redundancy in the 4D LF structure is vital for LFSR methods to exploit, compensating for missing high-frequency details in each

view. Deep learning techniques, including CNNs and Transformer-based architectures, have been widely adopted and proven effective for modeling 4D correlations within LFs. Directly processing entire 4D LF representations, such as Sub-Aperture Images (SAI) or Macro-Pixel Images (MacPI), poses optimization challenges during the training stage [25]. Consequently, prior methods extract features from 2D slices, including spatial slices, angular slices, and Epipolar Plane Images (EPIs) [35,34,5,31,32,20,21,8,37,39]. However, CNN-based methods, with their limited receptive fields, struggle to capture long-range spatial-angular correspondences, leading to inferior performance. Although Transformer-based methods effectively model non-local information, their attention mechanisms incur high memory costs and slow inference time.

Another critical aspect is the inference scheme. The drawbacks mentioned earlier have led prior works to adopt a patch-inference scheme, where input low-resolution (LR) LFs are divided into patches for individual processing, and the outputs are merged to form the final super-resolved result. This approach has two main issues: (1) For full-resolution LFs, the number of patches increases, requiring multiple forward passes through the network. (2) Each LF patch carries limited information, preventing the network from fully leveraging the entire input LF [7]. Therefore, there is an urgent need for efficient architectures that consume less memory and offer faster inference speeds to process full-resolution input LFs, thereby fully utilizing the whole input 4D LF structure.

The currently trending state space models (SSMs)[11,29,10] have proven effective in modeling long-range information with theoretically linear complexity. Among these, Mamba [10] introduces a selective mechanism for initializing structural parameters, enabling context-aware sequence modeling and boosting performance. Inspired by Mamba’s efficient long-range modeling capability, we explore its potential in handling the complex 4D structure of LFs.

In this paper, we attempt to leverage the efficient merits of Mamba to overcome the limitations of CNNs and Transformer-based methods. Specifically, we consider the inherent redundancy in the 4D LF structure and propose an efficient bi-directional subspace scanning scheme to effectively capture long-term spatial-angular correspondences. Based on the above scanning scheme, we propose a Mamba-based Global Interaction (MGI) module for the spatial-angular subspaces (SA-Mamba) and EPI subspaces (EPI-Mamba). Additionally, we propose a Spatial-Angular Modulator (SAM) to preserve the local structure of processed features. Combining the above designs, we devise a Mamba-based LFSR network termed MLFSR, which comprises interleaved MGI and SAM to efficiently model the 4D LF structure. In practice, there’s still some performance gap between MLFSR and Transformer-based methods due to the lossy state compression in SSMs. We therefore propose a Transformer-to-Mamba (T2M) distillation loss to align the non-local modelling ability between the attention mechanism and SSMs to improve the overall performance. As demonstrated in Figure 1, experimental results show that MLFSR outperforms CNN-based methods and benefits from less runtime and memory cost compared to Transformer-based counterparts while achieving competitive performance.



**Fig. 1.** Runtime (ms) and PSNR (dB) comparison. The runtime is calculated on an input LF of size  $5 \times 5 \times 32 \times 32$ . \* denotes full-resolution inference. MLFSR outperforms state-of-the-art CNN-based methods and has competitive performance compared to Transformer-based methods with less runtime and parameters. Full-resolution inference gives a further performance boost on available methods. Comparisons are performed for  $4 \times$  SR on the EPFL dataset.

The contributions of this paper are summarized below.

- We consider the spatial-angular redundancy that existed in 4D LFs and propose a bi-directional subspace scanning scheme to efficiently enable global information utilization.
- Based on the scanning scheme, we propose an MGI module and a local SAM to model complex 4D correlations efficiently, forming the MLFSR. We then introduce a T2M distillation loss to further boost the overall performance.
- Our proposed MLFSR outperforms CNN-based methods while comparable with Transformer-based counterparts while enjoying lower memory cost and less inference time, which enables full-resolution inference for further performance improvement.

## 2 Related Work

### 2.1 Light Field Super-Resolution

Representative traditional LFSR methods attempt to utilize depth information under optimization frameworks [3,2] or learning from dictionaries [6] to ensure angular consistency. With the advancement of deep learning techniques, CNN-based methods have gradually become the mainstream choice. The pioneer work LFCNN [42] uses stacked convolutional layers to process vertical, horizontal, and surrounding sub-aperture images. Following methods adopt 4D convolutions on SAIs to extract spatial-angular features [25]. Yeung et al. [41] further ease the computational burden by utilizing spatial-angular separable convolutions. To well-exploit the entangled 4D information, later works disentangle the LF into 2D

subspaces and design convolution variants tailored for each subspace [35,34,9]. HLFSR [31] further proposes to leverage multi-orientation epipolar information to boost the performance. However, due to the limited receptive field of CNN, non-local spatial-angular information cannot be well incorporated. Liu *et. al* [24] design a selective matching mechanism to explicitly extract 4D information in LFs with larger receptive fields. Recently, Transformer-based methods have emerged and achieved superior performance. By tokenizing 2D slices into sequences, the attention mechanism is adopted in spatial-angular [20] or EPI [21] subspaces to learn non-local spatial-angular correlations. Instead of network design, Xiao *et. al* [38] improves the SR performance from a data augmentation perspective. Different from prior works, we introduce the newly proposed Mamba to overcome the drawback of limited receptive field in CNN-based methods and the inferior efficiency of Transformer-based methods.

## 2.2 State Space Models

Originating from control theory, SSMs have garnered increasing attention due to their efficacy in long-term language modeling [11,29,10]. Unlike self-attention-based transformers, most SSMs capture long-range token interactions through linear recurrent processes, entailing  $\mathcal{O}(N)$  complexity theoretically. Mamba [10] improves the expressiveness of SSMs by introducing a selective mechanism, with its structural parameters adaptively learned from inputs. Motivated by its effectiveness, researchers have endeavored to adapt Mamba to the domain of computer vision. Vim [46] and VMamba [23] take pioneering efforts employing multidirectional scanning to overcome the causal nature of vanilla Mamba, inspiring subsequent work in various vision domains [13,40,27,16,19]. In this paper, we investigate the applicability of Mamba to LFs, to unlock its potential in efficient LF processing.

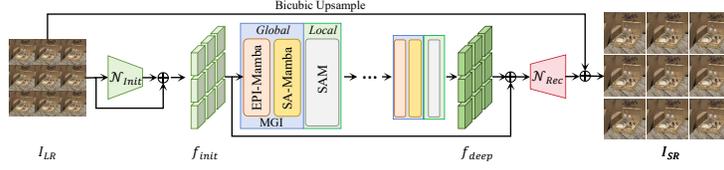
## 3 Method

### 3.1 Preliminaries

SSMs are inspired by the continuous control systems, which map input sequence  $x(t) \in \mathbb{R}^d$  to output sequence  $y(t) \in \mathbb{R}^d$  through hidden state  $h(t) \in \mathbb{R}^n$ . Without nonlinear transformation, SSMs set linear mappings  $\mathbf{A}$ ,  $\mathbf{B}$  and  $\mathbf{C}$  as transition parameters to update the hidden state over time and obtain the output. The mapping process is as follows,

$$\begin{aligned} h'(t) &= \mathbf{A}h(t) + \mathbf{B}x(t), \\ y(t) &= \mathbf{C}h(t). \end{aligned} \tag{1}$$

To enable deep network training, recent SSM such as the Structure State Space Sequence model (S4) [11] and Mamba [10] discretize the above process to a discretized version through an extra timescale parameter  $\mathbf{\Delta}$ . Then, continuous



**Fig. 2.** Overview of MLFSR. Initial features  $f_{init}$  are first extracted by encoder  $\mathcal{N}_{Init}$ , followed by alternate MGI modules and SAM to extract deep features  $f_{deep}$ . Finally, we obtain super-resolved results  $I_{SR}$  through the reconstruction module  $\mathcal{N}_{Rec}$ .

parameters  $\mathbf{A}$  and  $\mathbf{B}$  are transformed into the discretized ones  $\bar{\mathbf{A}}$  and  $\bar{\mathbf{B}}$  via the commonly used zero-order hold (ZOH) method,

$$\begin{aligned}\bar{\mathbf{A}} &= \exp(\Delta\mathbf{A}), \\ \bar{\mathbf{B}} &= (\Delta\mathbf{A})^{-1}(\exp(\Delta\mathbf{A}) - \mathbf{I})\Delta\mathbf{B},\end{aligned}\quad (2)$$

where  $\mathbf{I}$  denotes the identity matrix and  $(\cdot)^{-1}$  is the matrix inversion operation. Therefore, we now derive the discretized version of E.q. 1,

$$\begin{aligned}h_t &= \bar{\mathbf{A}}h_{t-1} + \bar{\mathbf{B}}x_t, \\ y_t &= \mathbf{C}h_t.\end{aligned}\quad (3)$$

The above recurrent form can be deviated into a convolutional form with global kernel  $\bar{\mathbf{K}}$ , which enables fast parallel training,

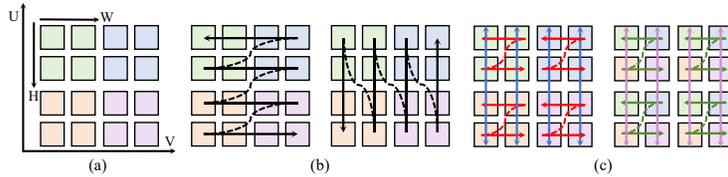
$$\begin{aligned}\bar{\mathbf{K}} &= (\mathbf{C}\bar{\mathbf{B}}, \mathbf{C}\bar{\mathbf{A}}\bar{\mathbf{B}}, \dots, \mathbf{C}\bar{\mathbf{A}}^{L-1}\bar{\mathbf{B}}), \\ \mathbf{y} &= \mathbf{x} * \bar{\mathbf{K}},\end{aligned}\quad (4)$$

where  $L$  denote the sequence length and  $*$  represent the convolution operator.

Different from S4, Mamba introduces a selective mechanism on  $\bar{\mathbf{B}}$  and  $\mathbf{C}$  to enable context-aware information filtering. Although it breaks the available convolutional form by introducing nonlinearity, a parallel scan algorithm [10] is proposed to speed up the process in E.q. 3 for faster training. In this paper, input LR LFs are tokenized as subspace sequences for efficient scanning.

### 3.2 Mamba-based Light Field Super-Resolution

**Overview.** Taking an input LR LF  $I_{LR} \in \mathbb{R}^{U \times V \times H \times W}$ , we aim to increase its spatial resolution by scale  $s$ , resulting in a super-resolved output  $I_{SR} \in \mathbb{R}^{U \times V \times sH \times sW}$ , where  $U, V$  denotes the angular resolution and  $H, W$  denotes the spatial resolution. As shown in Fig. 2, we first extract shallow features  $f_{init}$  from LR LF input  $I_{LR}$  using a convolutional-based encoder  $\mathcal{N}_{Init}$ . Then the shallow features are fed to alternate global-local 4D correspondence learning modules including Mamba-based Global Interaction (MGI) module and Spatial-Angular



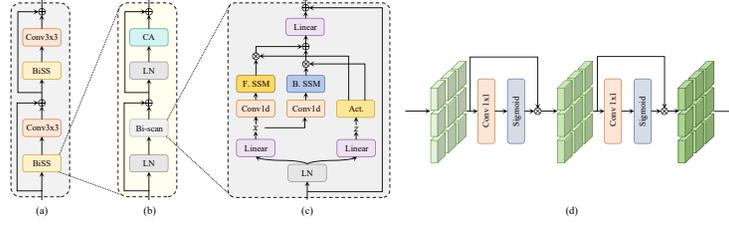
**Fig. 3.** A toy example ( $U = V = 2, H = W = 2$ ) on LF tokenization and scanning directions. (a) Sub-aperture images. (b) Whole 4D sequence with quad-directional scanning (SAI for example). (c) Subspace sequences with bi-directional scanning. Bi-directional spatial (red arrows) and angular (green arrows) scanning are complemented by bi-directional EPI (blue and purple arrows) scanning.

Modulator (SAM) to extract informative deep features  $f_{deep}$ . Finally,  $f_{deep}$  is upsampled by the reconstruction module  $\mathcal{N}_{Rec}$  to obtain the upsampled LF  $I_{SR}$ .

**Efficient Subspace Scanning.** As a high-dimensional vision modality, 4D LF has various representation forms such as SAIs, MacPIs, and subspace images. Therefore, it’s vital to select an appropriate way for effective and efficient scanning of specific representations. A straightforward way is to directly conduct 2D quad-directional scanning on the whole 4D representation (SAI or MacPI), as shown in Fig.3(b). In this way, the tokenized sequence has  $U \times V \times H \times W$  length and the efficient merit of Mamba on long sequence inputs can be fully exploited. Considering the causal nature of the original Mamba [10], quad-directional scanning [23] thus is required to fully benefit each token from each view. However, as it’s important in LFSR to well leverage the intra-view and inter-view correspondences, the mixed intra-inter view information cannot be well distinguished for each token. On the other hand, the compressed hidden state struggles to preserve prior spatial-angular information when the sequence gets longer. The trade-off between the long-range efficiency of Mamba and the limited state memory provides challenges to well adopting Mamba for LF inputs.

To overcome the above dilemma, we attempt to utilize the inherent structure redundancy that exists in LFs. Specifically, we notice the role of EPI representation, which provides rich information across the spatial and angular dimensions. When conducting bi-directional scanning on the EPI-H (or EPI-W) subspace, the vertical (or horizontal) of spatial and angular information can be incorporated, as can be seen in Fig.3(c). The quad-directional spatial (or angular) scanning can therefore be decomposed into two bi-directional scanning on spatial (or angular) dimensions, marked by red and green arrows) and EPI dimensions (marked by blue and purple arrows), respectively. In this way, by lowering the scanning length, the long-term memory issues can be well handled without sacrificing the complete 4D global information.

**Mamba-based Global Interaction (MGI) Module.** Based on the above bi-directional subspace scanning, we design a Mamba-based Global Interaction (MGI) module which consists of cascaded EPI-Mamba and SA-Mamba for effective non-local spatial-angular correspondence learning, as shown in Fig. 2.



**Fig. 4.** (a) The detailed structure of SA-Mamba/EPI-Mamba. Each SA-Mamba/EPI-Mamba includes two Bidirectional Subspace Scanning (BiSS) blocks. We omit the reshape operation for simplicity. (b) The BiSS block follows a typical Transformer-style design, including bidirectional scanning for token interaction and channel attention for channel mixing. (c) The details of bidirectional scanning. (d) The detailed structure of Spatial-Angular Modulator (SAM).

We take EPI-Mamba for example to elaborate its details since the structures of EPI-Mamba and SA-Mamba are the same except for the input form.

The detailed structures are illustrated in Fig. 4(a). Given an input feature  $f_l^{i-1}$  from the output of  $i-1$ -th SAM, we first flatten it into EPI-H token sequences  $T_h^i \in \mathbb{R}^{B \times V \times W \times U \times H \times C}$ , where  $C$  is the channel dimension and  $B$  is the batch size. Then, two successive Bidirectional Subspace Scanning (BiSS) blocks and convolution layers are set to enable non-local token interaction on the EPI-H and EPI-W subspace, which can be formulated as follows,

$$\begin{aligned} \hat{T}_h^i &= \text{Conv}(\text{BiSS}(T_h^i)) + T_h^i, \\ \hat{T}_w^i &= \text{Conv}(\text{BiSS}(T_w^i)) + T_w^i, \end{aligned} \quad (5)$$

where the EPI-W token sequences  $T_w^i \in \mathbb{R}^{B \times U \times H \times V \times W \times C}$  is reshaped from  $\hat{T}_h^i$ ,  $\hat{T}_h^i$  and  $\hat{T}_w^i$  are enhanced features.  $\hat{T}_w^i$  is then reshaped to  $f_g^i$  as the output of the  $i$ -th MGI module. The design of the BiSS block follows a typical Transformer-style structure, which is shown in Fig. 4(b). Input token sequences are first normalized by a LayerNorm [1] and fed into a Bidirectional Scanning (Bi-scan) block which is introduced in [46] (see Fig. 4(c)). Then, another LayerNorm and a Channel Attention (CA) layer [44] is used for channel mixing. We also add shortcut connections [14] to ease the optimization process. As for the SA-Mamba, we can simply replace EPI token sequences  $T_h^i$  and  $T_w^i$  with spatial and angular sequences  $T_s^i \in \mathbb{R}^{B \times U \times V \times H \times W \times C}$  and  $T_a^i \in \mathbb{R}^{B \times H \times W \times U \times V \times C}$ . In practice, we share the weights of two BiSS and convolutional layers in EPI-Mamba to share the intrinsic information along the epipolar line of LFs [21], which also leads to more efficient network architecture.

By combining the EPI-Mamba and SA-Mamba, complementary information between two directions of spatial-angular information and EPI information can be well incorporated, achieving efficient global 4D correspondence learning. The experimental results in Section 4.3 validate the effectiveness of each component.

**Spatial-Angular Modulator (SAM).** Although MGI enables the global receptive field to enable the network to utilize non-local spatial-angular information, the importance of locality in SR [4] is overlooked. To compensate for local spatial-angular information, we design a convolution-based Spatial-Angular Modulator (SAM) to learn input-adaptive weights for modulating the output of MGI. Inspired by pixel attention [45], we adopt a lightweight design of SAM with only two convolutional layers with kernel size 1, as shown in Fig 4(d).

The output feature  $f_g^i \in \mathbb{R}^{B \times U \times V \times H \times W \times C}$  from last MGI module is first reshaped to the spatial subspace and go through a single convolutional layer, followed by a sigmoid operation to generate the spatial attention map  $\text{Attn}_s$ . Then,  $f_g^i$  is modulated using  $\text{Attn}_s$  by element-wise multiplication,

$$\begin{aligned} \text{Attn}_s &= \text{Sigmoid}(\text{Conv}(f_g^i)), \\ \tilde{f}_g^i &= f_g^i \cdot \text{Attn}_s + f_g^i, \end{aligned} \quad (6)$$

where  $\tilde{f}_g^i$  represents spatially-modulated features. A similar process is conducted on the angular subspace to obtain the corresponding angular attention map  $\text{Attn}_a$  and we can obtain the output  $f_l^i$ ,

$$f_l^i = \tilde{f}_g^i \cdot \text{Attn}_a + \tilde{f}_g^i. \quad (7)$$

### 3.3 Learning Objective

We use  $L_1$  distance to measure the reconstruction loss between the final output and the ground-truth LF,

$$\mathcal{L}_{rec} = \|I_{SR} - I_{HR}\|_1. \quad (8)$$

However, only using the reconstruction loss results in sub-optimal performance compared to Transformer-based methods according to the experiment results. Considering the difference between the attention mechanism and SSMS, we can find that attention-based methods explicitly record relations between tokens into the attention map, while SSMSs implicitly compress non-local information in hidden memory. We therefore propose a Transformer-to-Mamba (T2M) distillation loss to align the deep features of the Transformer and Mamba, which is formulated as follows,

$$\mathcal{L}_{dist} = \|f_{deep} - f_{deep}^{Trans}\|_1. \quad (9)$$

In this work, we utilize EPIT [21] as the teacher network. The total loss function is the combination of the above two loss functions, adjusted by  $\lambda$ ,

$$\mathcal{L}_{total} = \mathcal{L}_{rec} + \lambda \mathcal{L}_{dist}. \quad (10)$$

## 4 Experiments

### 4.1 Experiment Settings

**Datasets.** Aligning with prior works, we use the BasicLFSR benchmark [33] for the training and evaluation of all methods on  $2\times$  and  $4\times$  scale. The benchmark

comprises five LF datasets (HCI-new [15], HCI-old [36], EPFL [28], INRIA [18] and STFGantry [30]) which include 144 training scenes and 23 test scenes with diverse contents and varying disparities. We extract the central  $5\times 5$  SAIs from each LF for training and testing.

**Implementation Details.** We implement our method using Pytorch [26] framework with four NVIDIA 3090 GPUs. The result of baseline methods is obtained from open-source codes. The number of MGI and SAM used in MLFSR is set to 3 and the state dimension is 8. The hyperparameter in  $\mathcal{L}_{total}$  is set to 0.1. The MLFSR is trained using the Adam optimizer with  $\beta_1 = 0.9$  and  $\beta_2 = 0.999$  with batch size 4 on each card. The initial learning rate is set to  $4e-4$  and decreases by a factor of 0.5 for every 15 epochs. Data augmentation including random flipping and rotation is applied in the training stage. The network is trained for 90 epochs with L1 loss only and finetuned for another 30 epochs with a combination of L1 loss and T2M distillation loss. We use widely adopted Peak Signal-to-Noise Ratio (PSNR) and Structure Similarity (SSIM) to evaluate the fidelity of the super-resolved results.

## 4.2 Comparison with State-of-the-Arts

We compare our method with 13 baseline methods, including 9 CNN-based methods [17,22,44,43,41,35,34,5,31] and 4 Transformer-based methods [20,32,8,21].

**Quantitative Results.** The quantitative results are shown in Table 1. We can notice that MLFSR outperforms CNN-based methods at both  $2\times$  and  $4\times$  scale. For example, MLFSR has a 0.33dB gain compared to DistgSSR [34] and HLFSR[31] on the EPFL testset at  $2\times$  scale. Compared to Transformer-based methods, MLFSR shows competitive performance on most testsets. When inference with full-resolution input, the performance further improves on most testsets at both  $2\times$  and  $4\times$  scale. Specifically, the  $4\times$  results even outperform all Transformer-based methods while maintaining lower number of parameters and less runtime.

**Efficiency Comparison.** We provide an in depth efficiency comparison respect to input resolution between MLFSR and current Transformer-based methods on two metrics: runtime and peak memory use. Table 2 clearly shows that with the increase of input resolution, the runtime of Transformer-based methods grows significantly. The increase of peak memory use in Table 2 shows a similar trend. Almost all baseline methods failed to infer at  $384\times 384$  resolution, and without memory optimization (`torch.nn.MultiheadAttention` interface), LFT cannot even handle  $5\times 5\times 64\times 64$  inputs. With optimized implementation, LFT enables inference at  $128\times 128$  resolution, which is slower than our method inference at  $384\times 384$  resolution. Although EPIT with optimization consumes less memory, the inference speed is still much slower compared to MLFSR. For example, at  $256\times 256$  and  $384\times 384$  resolution, EPIT is about  $3.0\times$  and  $4.2\times$  slower than our method, respectively. On the other hand, when input asymmetric resolution such as  $216\times 312$ , optimized implementation fails to reduce the memory consumption

**Table 1.** Quantitative results (PSNR / SSIM) on the five public benchmarks of LF super-resolution with the input angular resolution of  $5\times 5$ . We report the number of parameters (Params.) and runtime for efficiency comparison. Note that the runtime is calculated on an input LF of size  $5\times 5\times 32\times 32$ . Method marked with \* denotes full-resolution inference results. We mark the best, second best results in **bold** and with underline, respectively.

Method	Scale	Params.(M)	Time (ms)	HCI-new	HCI-old	EPFL	INRIA	STFGantry
Bicubic	2x	-	-	31.89/.9356	37.69/.9785	29.74/.9376	31.33/.9577	31.06/.9498
VDSR [17]	2x	0.66	29.50	34.37/.9561	40.61/.9867	32.50/.9598	34.43/.9741	35.54/.9789
EDSR [22]	2x	38.62	137.25	34.83/.9592	41.01/.9874	33.09/.9629	34.97/.9764	36.29/.9818
RCAN [44]	2x	15.31	1586.50	34.98/.9603	41.05/.9875	33.16/.9634	35.01/.9769	36.33/.9831
resLF [43]	2x	7.98	101.89	36.69/.9739	43.42/.9932	33.62/.9706	35.39/.9804	38.36/.9904
LFSSR [41]	2x	0.89	3.98	36.81/.9749	43.81/.9938	33.68/.9744	35.28/.9832	37.95/.9898
LF-InterNet [35]	2x	5.04	18.08	37.28/.9763	44.45/.9946	34.14/.9760	35.80/.9843	38.72/.9909
DistgSSR [34]	2x	3.53	24.76	37.96/.9796	44.94/.9949	34.81/.9787	36.59/.9859	40.40/.9942
LFSAV [5]	2x	1.22	10.57	37.43/.9776	44.22/.9942	34.62/.9772	36.36/.9849	38.69/.9914
HLFSR [31]	2x	3.45	54.29	38.03/.9798	44.86/.9949	34.71/.9779	36.59/.9856	40.49/.9943
LFT [20]	2x	1.11	132.37	37.84/.9791	44.52/.9945	34.80/.9781	36.59/.9855	40.51/.9941
DPT [32]	2x	3.73	117.97	37.35/.9771	44.31/.9943	34.48/.9758	36.40/.9843	39.52/.9926
LF-DET [8]	2x	1.59	73.97	<b>38.31/.9807</b>	<u>44.99/.9950</u>	<u>35.26/.9797</u>	<u>36.95/.9864</u>	<u>41.76/.9955</u>
EPIT [21]	2x	1.42	51.58	<u>38.23/.9810</u>	<b>45.08/.9949</b>	34.83/.9775	36.67/.9853	<b>42.17/.9957</b>
MLFSR (Ours)	2x	1.36	27.81	38.14/.9803	44.90/.9950	35.22/.9801	36.92/.9865	40.98/.9949
MLFSR* (Ours)	2x	1.36	27.81	38.15/.9802	44.97/.9950	<b>36.01/.9798</b>	<b>38.65/.9864</b>	41.03/.9947
Bicubic	4x	-	-	27.61/.8517	32.42/.9344	25.14/.8324	26.82/.8867	25.93/.8452
VDSR [17]	4x	0.66	35.75	29.31/.8823	34.81/.9515	27.25/.8777	29.19/.9204	28.51/.9009
EDSR [22]	4x	38.89	147.00	29.60/.8869	35.18/.9536	27.84/.8854	29.66/.9257	28.70/.9072
RCAN [44]	4x	15.36	1716.05	29.63/.8886	35.20/.9548	27.88/.8863	29.76/.9276	28.90/.9131
resLF [43]	4x	8.65	104.23	30.73/.9107	36.71/.9682	28.27/.9035	30.34/.9412	30.19/.9372
LFSSR [41]	4x	1.77	22.37	30.72/.9145	36.70/.9696	28.27/.9118	30.31/.9467	30.15/.9426
LF-InterNet [35]	4x	5.48	19.65	30.98/.9161	37.11/.9716	28.67/.9162	30.64/.9491	30.53/.9409
DistgSSR [34]	4x	3.58	25.40	31.38/.9217	37.56/.9732	28.99/.9195	30.99/.9519	31.65/.9535
LFSAV [5]	4x	1.54	36.36	31.45/.9217	37.50/.9721	29.37/.9223	31.27/.9531	31.36/.9505
HLFSR [31]	4x	3.48	59.08	31.43/.9226	37.72/.9738	29.10/.9212	31.17/.9531	31.41/.9524
LFT [20]	4x	1.16	137.53	31.46/.9218	37.63/.9735	29.25/.9210	31.20/.9524	31.86/.9548
DPT [32]	4x	3.78	118.62	31.20/.9188	37.41/.9721	28.94/.9170	30.96/.9503	31.15/.9488
LF-DET [8]	4x	1.69	76.05	<u>31.56/.9235</u>	<u>37.84/.9744</u>	<u>29.47/.9230</u>	<u>31.39/.9534</u>	<u>32.14/.9573</u>
EPIT [21]	4x	1.47	52.75	31.51/.9231	37.68/.9737	29.34/.9197	31.37/.9526	<u>32.18/.9571</u>
MLFSR (Ours)	4x	1.41	28.94	<u>31.56/.9235</u>	<u>37.83/.9745</u>	29.28/.9218	31.24/.9531	32.03/.9567
MLFSR* (Ours)	4x	1.41	28.94	<b>31.59/.9237</b>	<b>37.92/.9747</b>	<b>30.04/.9221</b>	<b>32.21/.9534</b>	<b>32.21/.9575</b>

**Table 2.** Efficiency comparison against state-of-the-art Transformer-based methods on 2x scale. We take runtime (ms) and peak memory use (GB) on different input spatial resolutions to evaluate the efficiency. † denotes methods without optimized implementation and OOM denotes out of memory on a single RTX 3090 GPU. Blanked entries correspond to results unable to be reported.

Method	32x32	64x64	128x128	256x256	216x312	384x384
DPT [32]	117.97/0.26	130.08/0.57	429.94/1.67	4539.51/8.60	4759.83/8.99	OOM
LF-DET [8]	73.97/0.35	387.88/3.58	OOM	OOM	OOM	OOM
LFT† [20]	146.34/1.79	OOM	OOM	OOM	OOM	OOM
LFT [20]	132.37/0.22	778.57/0.69	8828.29/3.51	-	-	-
EPIT† [21]	53.21/0.38	164.94/2.51	1019.11/18.38	OOM	OOM	OOM
EPIT [21]	51.58/0.13	123.61/0.47	752.24/1.81	5518.96/7.19	OOM	17852.41/16.16
MLFSR(Ours)	27.81/0.18	105.19/0.67	455.30/2.63	1855.72/10.47	1935.71/10.77	4277.76/23.55

which largely limits its practical use. The above results indicate that our method demonstrates significant efficiency advantages as the input resolution increases.

**Table 3.** Full-resolution inference (denoted by \*) quantitative results (PSNR / SSIM) on the five public benchmarks of LF super-resolution with the input angular resolution of  $5\times 5$ . We mark the best, second best results in **bold** and with underline, respectively. OOM denotes out of memory on a single RTX 3090 GPU.

Method	Scale	Params.(M)	Time (ms)	HCI-new	HCI-old	EPFL	INRIA	STFGantry
LFSAV* [5]	2×	1.22	10.57	37.44/.9777	44.32/.9942	<u>35.65</u> /.9773	<u>38.16</u> /.9850	38.74/.9914
HLFSR* [31]	2×	3.45	54.29	37.97/.9798	<u>45.01</u> /. <u>9949</u>	35.16/. <u>9779</u>	37.05/. <u>9856</u>	40.53/.9944
DPT* [32]	2×	3.73	117.97	37.29/.9770	OOM	34.48/.9758	36.19/.9843	39.38/.9928
EPIT* [21]	2×	1.42	51.58	<b>38.25</b> /. <b>9810</b>	<u>45.12</u> /. <u>9949</u>	OOM	OOM	<b>42.22</b> /. <b>9957</b>
MLFSR* (Ours)	2×	1.36	27.81	<u>38.15</u> /. <u>9802</u>	44.97/. <b>9950</b>	<b>36.01</b> /. <b>9798</b>	<b>38.65</b> /. <b>9864</b>	<u>41.03</u> /. <u>9947</u>
LFSAV* [5]	4×	1.54	36.36	31.47/.9219	37.52/.9722	<u>29.99</u> /. <b>9226</b>	32.17/. <b>9534</b>	31.41/.9506
HLFSR* [31]	4×	3.48	59.08	31.46/.9229	<u>37.77</u> /. <u>9739</u>	29.94/.9217	<u>32.20</u> /. <b>9534</b>	31.53/.9529
DPT* [32]	4×	3.78	118.62	31.11/.9188	37.20/.9721	29.40/.9171	31.38/.9505	31.00/.9489
EPIT* [21]	4×	1.47	52.75	<u>31.53</u> /. <u>9232</u>	37.74/.9738	29.90/.9200	32.16/. <u>9529</u>	<b>32.26</b> /. <u>9573</u>
MLFSR* (Ours)	4×	1.41	28.94	<b>31.59</b> /. <b>9237</b>	<b>37.92</b> /. <b>9747</b>	<b>30.04</b> /. <u>9221</u>	<b>32.21</b> /. <b>9534</b>	<u>32.21</u> /. <b>9575</b>

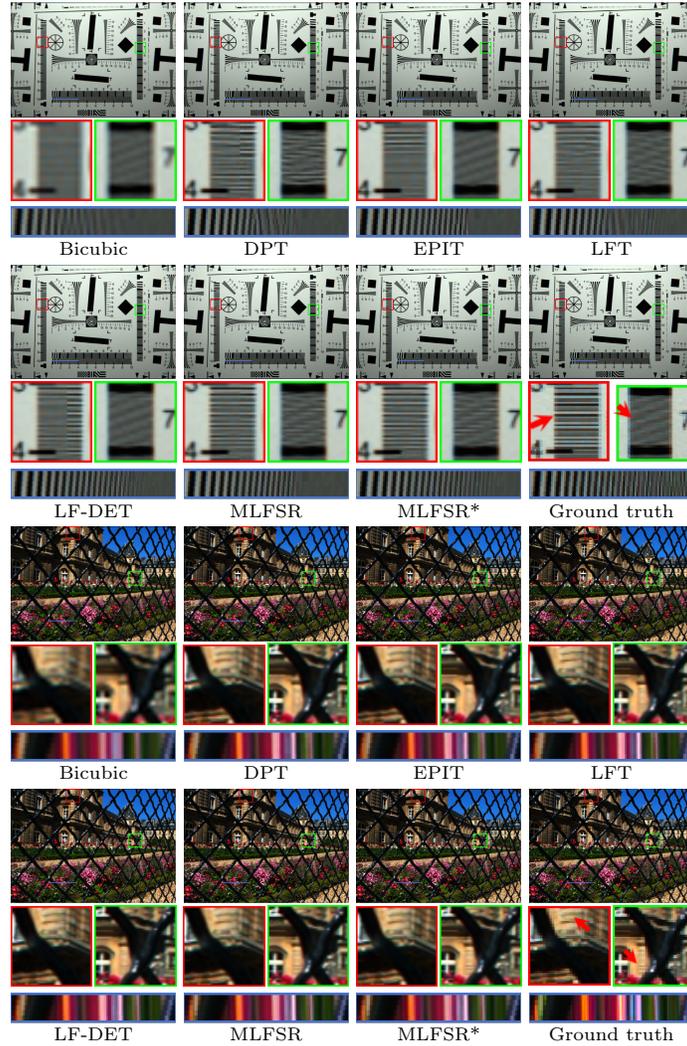
**Full-Resolution Inference Results.** We also provide the full-resolution inference results of available top-performing methods in Table 3 for a more comprehensive evaluation. With extra information from larger resolution inputs, the performance of baseline method improves in most case. Our method still achieves state-of-the-art performance under the full-resolution inference setting with lower parameters and latency.

**Qualitative Results.** We provide visual results on  $2\times$  and  $4\times$  scale in Fig. 5 and Fig. 6, respectively. For example, in the ISO scene in Fig. 5, our results manages to recover the horizontal line. Note that the result of DPT and LF-DET contains vertical line artifacts due to the patch inference scheme, which shows the superiority of whole-image inference. The details on the background wall in the below scene are also recovered by our method. In Fig. 6, our method successfully produces clear number 3 in the ISO scene while other baselines tend to generate blurry results and artifacts. The letters in the tarot card scene generated by our method have less distorted edges and are more close to the ground truth.

### 4.3 Ablation Studies

In this section, we perform ablation studies on  $4\times$  SR to validate the effectiveness of our designs including macro components and scanning methods.

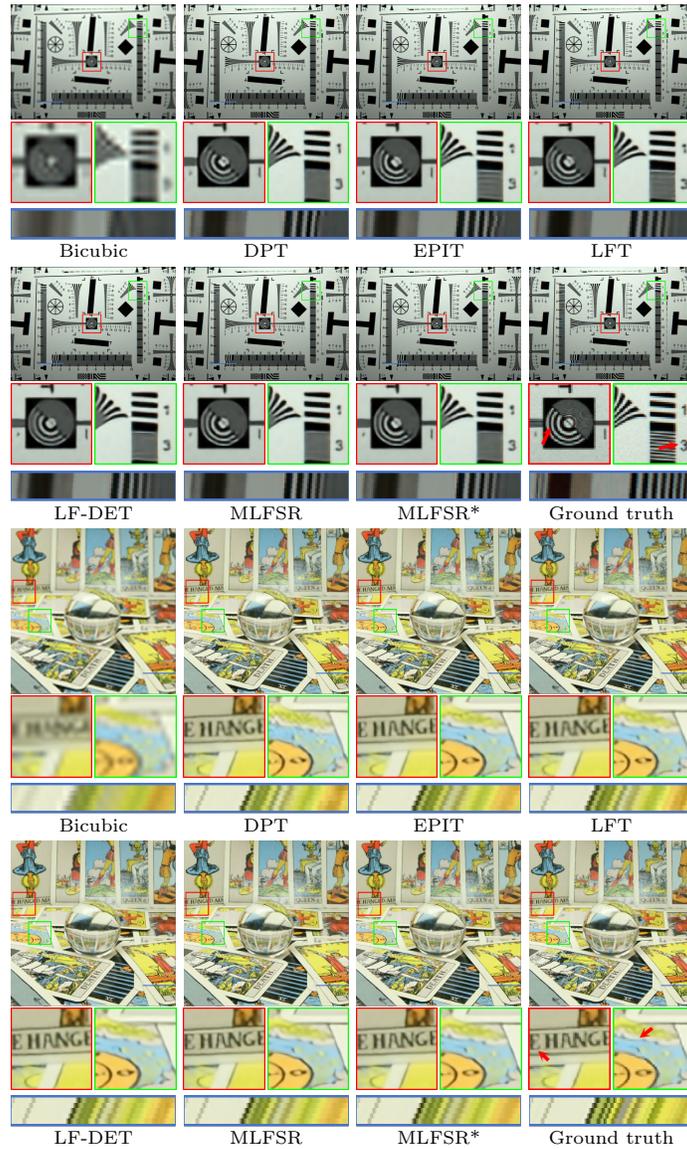
**Macro Design.** We first consider the macro design of our method, *i.e.*, the effect of SA-Mamba, EPI-Mamba, SAM, and the T2M distillation loss. The quantitative results are shown in Table 4. By removing either SA-Mamba or EPI-Mamba in MLFSR, the performance dropped by 0.14dB and 0.18dB on the average PSNR respectively, showing the importance of complementary information between subspaces. On the other hand, without the locality introduced by SAM, the performance also dropped by 0.11dB, indicating the important role of local information in the SR task. Finally, the T2M loss further boosts the performance by 0.06dB. We also use the Local Attribution Map (LAM) [12] to explore the role of each component, as shown in Fig. 7. Without EPI-Mamba or



**Fig. 5.** Visual comparisons of different methods on  $2\times$  SR (view coordinates: (2, 2)). Please zoom in for better visualization and best viewing on screen.

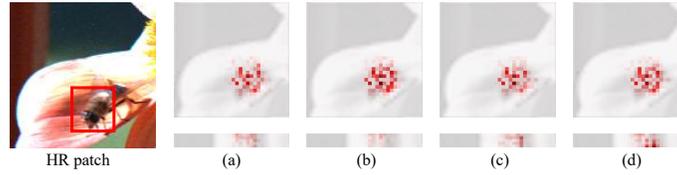
SA-Mamba, non-local spatial-angular information is less utilized. On the other hand, SAM focuses more on local information to complement more fine details. The above results validate the effectiveness of our core designs.

**Scanning Methods.** We explore the scanning methods of LFs by altering the LF tokenization way and the scanning directions. Specifically, we introduce three variants: (a) 4D sequence with quad-directional scanning. (b) 4D sequence with bi-directional scanning. (c) Subspace sequence with quad-directional scanning



**Fig. 6.** Visual comparisons of different methods on  $4\times$  SR (view coordinates: (2, 2)). Please zoom in for better visualization and best viewing on screen.

to compare with our scanning method. The quantitative results are shown in Table 5. When tokenizing the whole 4D representation (we use SAI here) for scanning, the quantitative results are much lower than the subspace tokenization. Even the scanning directions are increased to 4 (quad-scanning), there is no clear



**Fig. 7.** LAM visualization (Central view & EPI) of the macro-design ablation study. (a) w/o EPI-Mamba (b) w/o SA-Mamba (c) w/o SAM (d) Ours.

**Table 4.** Ablations on the macro design. The best results are marked in **bold**. All methods are under full-resolution inference scheme.

EPI-Mamba	SA-Mamba	SAM	T2M Loss	Avg. PSNR/SSIM
✗	✓	✓	✓	32.65/.9457
✓	✗	✓	✓	32.60/.9454
✓	✓	✗	✓	32.68/.9456
✓	✓	✓	✗	32.73/.9460
✓	✓	✓	✓	<b>32.79/.9463</b>

**Table 5.** Ablations on the scanning methods. We explore the configuration of LF tokenization and scanning directions. The runtime is calculated on an input LF of size  $5 \times 5 \times 32 \times 32$ . The best results are marked in **bold**. All methods are under full-resolution inference scheme.

Tokenization	Scanning	Params.(M)	Time(ms)	Avg. PSNR/SSIM
Whole	4	1.5	40.76	31.85/.9354
	2	1.4	31.58	31.72/.9343
Subspace	4	1.5	48.99	<b>32.73/.9460</b>
	2	1.4	28.94	32.79/.9463

performance improvement compared to our scanning method, which shows the redundancy of quad-directional scanning.

## 5 Conclusion

In this paper, we propose a Mamba-based LFSR method named MLFSR to ease the high memory consumption and latency brought by Transformer-based methods. Specifically, we utilize the inherent structure redundancy existing in LFs and propose an efficient subspace scanning method. Based on it, we design a Mamba-based Global Interaction module to model global spatial-angular correlations. Furthermore, a Spatial-Angular Modulator is proposed to complement local information. The overall performance is further boosted by a Transformer-to-Mamba distillation loss. Extensive experiments show that MLFSR has a clear advantage in efficiency while achieving state-of-art performance. With much less memory use, MLFSR facilitates further performance improvement with full-resolution inference.

**Acknowledgments.** We acknowledge funding from National Natural Science Foundation of China under Grants 62131003 and 62021001.

## References

1. Ba, J.L., Kiros, J.R., Hinton, G.E.: Layer normalization. arXiv preprint arXiv:1607.06450 (2016)
2. 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), 972–986 (2011)
3. Bishop, T.E., Zanetti, S., Favaro, P.: Light field superresolution. In: 2009 IEEE International Conference on Computational Photography (ICCP). pp. 1–9. IEEE (2009)
4. Chen, Z., Zhang, Y., Gu, J., Kong, L., Yang, X., Yu, F.: Dual aggregation transformer for image super-resolution. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 12312–12321 (2023)
5. Cheng, Z., Liu, Y., Xiong, Z.: Spatial-angular versatile convolution for light field reconstruction. *IEEE Transactions on Computational Imaging* **8**, 1131–1144 (2022)
6. Cho, D., Lee, M., Kim, S., Tai, Y.W.: Modeling the calibration pipeline of the lytro camera for high quality light-field image reconstruction. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 3280–3287 (2013)
7. Chu, X., Chen, L., Chen, C., Lu, X.: Improving image restoration by revisiting global information aggregation. In: European Conference on Computer Vision. pp. 53–71. Springer (2022)
8. Cong, R., Sheng, H., Yang, D., Cui, Z., Chen, R.: Exploiting spatial and angular correlations with deep efficient transformers for light field image super-resolution. *IEEE Transactions on Multimedia* (2023)
9. Gao, R., Liu, Y., Xiao, Z., Xiong, Z.: Diffusion-based light field synthesis. In: ECCVW (2024)
10. Gu, A., Dao, T.: Mamba: Linear-time sequence modeling with selective state spaces. arXiv preprint arXiv:2312.00752 (2023)
11. Gu, A., Goel, K., Re, C.: Efficiently modeling long sequences with structured state spaces. In: The Tenth International Conference on Learning Representations (2022)
12. Gu, J., Dong, C.: Interpreting super-resolution networks with local attribution maps. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9199–9208 (2021)
13. Guo, H., Li, J., Dai, T., Ouyang, Z., Ren, X., Xia, S.T.: Mambair: A simple baseline for image restoration with state-space model. In: ECCV (2024)
14. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
15. Honauer, K., Johannsen, O., Kondermann, D., Goldluecke, B.: A dataset and evaluation methodology for depth estimation on 4d light fields. In: Computer Vision–ACCV 2016: 13th Asian Conference on Computer Vision, Taipei, Taiwan, November 20–24, 2016, Revised Selected Papers, Part III 13. pp. 19–34. Springer (2017)
16. Huang, T., Pei, X., You, S., Wang, F., Qian, C., Xu, C.: Localmamba: Visual state space model with windowed selective scan. arXiv preprint arXiv:2403.09338 (2024)
17. Kim, J., Lee, J.K., Lee, K.M.: Accurate image super-resolution using very deep convolutional networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1646–1654 (2016)
18. Le Pendu, M., Jiang, X., Guillemot, C.: Light field inpainting propagation via low rank matrix completion. *IEEE Transactions on Image Processing* **27**(4), 1981–1993 (2018)

19. Li, K., Li, X., Wang, Y., He, Y., Wang, Y., Wang, L., Qiao, Y.: Videomamba: State space model for efficient video understanding. In: ECCV (2024)
20. Liang, Z., Wang, Y., Wang, L., Yang, J., Zhou, S.: Light field image super-resolution with transformers. *IEEE Signal Processing Letters* **29**, 563–567 (2022)
21. Liang, Z., Wang, Y., Wang, L., Yang, J., Zhou, S., Guo, Y.: Learning non-local spatial-angular correlation for light field image super-resolution. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 12376–12386 (2023)
22. 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. pp. 136–144 (2017)
23. Liu, Y., Tian, Y., Zhao, Y., Yu, H., Xie, L., Wang, Y., Ye, Q., Liu, Y.: Vmamba: Visual state space model. *Advances in Neural Information Processing Systems* (2024)
24. Liu, Y., Cheng, Z., Xiao, Z., Xiong, Z.: Light field super-resolution using decoupled selective matching. *IEEE Transactions on Circuits and Systems for Video Technology* (2023)
25. Meng, N., So, H.K.H., Sun, X., Lam, E.Y.: High-dimensional dense residual convolutional neural network for light field reconstruction. *IEEE transactions on pattern analysis and machine intelligence* **43**(3), 873–886 (2019)
26. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems* **32** (2019)
27. Peng, S., Zhu, X., Deng, H., Lei, Z., Deng, L.J.: Fusionmamba: Efficient image fusion with state space model. *arXiv preprint arXiv:2404.07932* (2024)
28. Rerabek, M., Ebrahimi, T.: New light field image dataset. In: 8th International Conference on Quality of Multimedia Experience (QoMEX) (2016)
29. Smith, J.T., Warrington, A., Linderman, S.: Simplified state space layers for sequence modeling. In: The Eleventh International Conference on Learning Representations (2023)
30. Vaish, V., Adams, A.: The (new) stanford light field archive. *Computer Graphics Laboratory, Stanford University* **6**(7) (2008)
31. Van Duong, V., Huu, T.N., Yim, J., Jeon, B.: Light field image super-resolution network via joint spatial-angular and epipolar information. *IEEE Transactions on Computational Imaging* **9**, 350–366 (2023)
32. Wang, S., Zhou, T., Lu, Y., Di, H.: Detail-preserving transformer for light field image super-resolution. In: Proceedings of the AAAI conference on artificial intelligence. vol. 36, pp. 2522–2530 (2022)
33. Wang, Y., Wang, L., Liang, Z., Yang, J., Timofte, R., Guo, Y., Jin, K., Wei, Z., Yang, A., Guo, S., Gao, M., Zhou, X., Duong, V.V., Huu, T.N., Yim, J., Jeon, B., Liu, Y., Cheng, Z., Xiao, Z., Xu, R., Xiong, Z., Liu, G., Jin, M., Yue, H., Yang, J., Gao, C., Zhang, S., Chang, S., Lin, Y., Chao, W., Wang, X., Wang, G., Duan, F., Xia, W., Wang, Y., Xia, P., Wang, S., Lu, Y., Cong, R., Sheng, H., Yang, D., Chen, R., Wang, S., Cui, Z., Chen, Y., Lu, Y., Cai, D., An, P., Salem, A., Ibrahim, H., Yagoub, B., Kang, H.S., Zeng, Z., Wu, H.: Ntire 2023 challenge on light field image super-resolution: Dataset, methods and results. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW) (2023)
34. Wang, Y., Wang, L., Wu, G., Yang, J., An, W., Yu, J., Guo, Y.: Disentangling light fields for super-resolution and disparity estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **45**(1), 425–443 (2022)

35. Wang, Y., Wang, L., Yang, J., An, W., Yu, J., Guo, Y.: Spatial-angular interaction for light field image super-resolution. In: *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIII* 16. pp. 290–308. Springer (2020)
36. Wanner, S., Meister, S., Goldluecke, B.: Datasets and benchmarks for densely sampled 4d light fields. In: *VMV*. vol. 13, pp. 225–226 (2013)
37. Xiao, Z., Cheng, Z., Xiong, Z.: Space-time super-resolution for light field videos. *IEEE Transactions on Image Processing* (2023)
38. Xiao, Z., Liu, Y., Gao, R., Xiong, Z.: Cutmib: Boosting light field super-resolution via multi-view image blending. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 1672–1682 (2023)
39. Xiao, Z., Xiong, Z.: Incorporating degradation estimation in light field spatial super-resolution. *arXiv preprint arXiv:2405.07012* (2024)
40. Xing, Z., Ye, T., Yang, Y., Liu, G., Zhu, L.: Segmamba: Long-range sequential modeling mamba for 3d medical image segmentation. *arXiv preprint arXiv:2401.13560* (2024)
41. 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), 2319–2330 (2018)
42. 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*. pp. 24–32 (2015)
43. Zhang, S., Lin, Y., Sheng, H.: Residual networks for light field image super-resolution. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 11046–11055 (2019)
44. 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)*. pp. 286–301 (2018)
45. Zhao, H., Kong, X., He, J., Qiao, Y., Dong, C.: Efficient image super-resolution using pixel attention. In: *Computer Vision–ECCV 2020 Workshops: Glasgow, UK, August 23–28, 2020, Proceedings, Part III* 16. pp. 56–72. Springer (2020)
46. Zhu, L., Liao, B., Zhang, Q., Wang, X., Liu, W., Wang, X.: Vision mamba: Efficient visual representation learning with bidirectional state space model. In: *Forty-first International Conference on Machine Learning* (2024)