

M3SR: Multi-Scale Multi-Perceptual Mamba for Efficient Spectral Reconstruction

Yuze Zhang¹, Lingjie Li^{2*}, Qiuzhen Lin^{1*}, Zhong Ming¹, Fei Yu¹, Victor C. M. Leung¹

¹College of Computer Science and Software Engineering, Shenzhen University, Shenzhen, China

²School of Artificial Intelligence, Shenzhen Technology University, Shenzhen, China

{zhangyuze2021, qiuzhlin, mingz, yufei, vleung}@szu.edu.cn, lilingjie@sztu.edu.cn

Abstract

The Mamba architecture has been widely applied to various low-level vision tasks due to its exceptional adaptability and strong performance. Although the Mamba architecture has been adopted for spectral reconstruction, it still faces the following two challenges: (1) Single spatial perception limits the ability to fully understand and analyze hyperspectral images; (2) Single-scale feature extraction struggles to capture the complex structures and fine details present in hyperspectral images. To address these issues, we propose a multi-scale, multi-perceptual Mamba architecture for the spectral reconstruction task, called M3SR. Specifically, we design a multi-perceptual fusion block to enhance the ability of the model to comprehensively understand and analyze the input features. By integrating the multi-perceptual fusion block into a U-Net structure, M3SR can effectively extract and fuse global, intermediate, and local features, thereby enabling accurate reconstruction of hyperspectral images at multiple scales. Extensive quantitative and qualitative experiments demonstrate that the proposed M3SR outperforms existing state-of-the-art methods while incurring a lower computational cost.

Code — <https://github.com/zhangyuze/M3SR>

Introduction

Hyperspectral imaging uses narrow wavelength bands to gather detailed spatial and spectral data, providing a more comprehensive view of scenes compared to RGB images. This enhanced detail renders it valuable across various domains, including environmental monitoring (Lassalle 2021; Wang et al. 2024b), medical imaging (Leon et al. 2023; Liu et al. 2024b), and agriculture (Khan et al. 2022; Ram et al. 2024). Capturing hyperspectral images (HSIs) is a very time-consuming and complex process, but fortunately Spectral Reconstruction (SR) provides an efficient solution by generating HSIs from RGB images. Earlier SR methods, based on traditional techniques such as sparse dictionaries (Arad and Ben-Shahar 2016), Gaussian processes (Akhtar and Mian 2018), and low-rank representations (Gao et al. 2020), struggle to recognize complex patterns or adapt to different scenarios. In contrast, the advent

*Corresponding author.

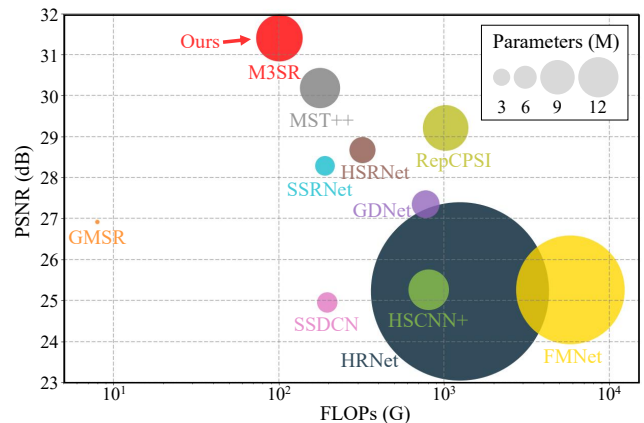


Figure 1: PSNR v.s. Parameters v.s. FLOPs comparison with existing spectral reconstruction methods. For an intuitive analysis, FLOPs and PSNR are represented by the horizontal and vertical axis, and the circle radius indicates parameters. The proposed M3SR achieves a better balance between PSNR and parameters and FLOPs.

of deep learning, particularly convolutional neural networks (CNNs), has greatly improved SR performance, although challenges remain in handling long-range connections and ensuring consistent spectral detail. Moreover, the Transformer with its multi-head self-attention mechanism effectively tackles these issues by capturing distant relationships and fine spectral variations, making it a powerful tool for SR. However, the computational requirements of the Transformer increase sharply with image size, posing efficiency challenges for HSIs with both detailed spatial resolution and broad spectral resolution.

Recently, the novel architecture Mamba (Gu and Dao 2023), based on the state space model (SSM), has quickly risen in popularity within natural language processing. Compared to the Transformer, Mamba exhibits superior efficiency in processing extended sequences, while consistently delivering robust performance during both the training and inference phases (Patro and Agneeswaran 2024). Inspired by this success, VMamba (Liu et al. 2024c) integrates SSM-based blocks into visual networks to enhance the learning of visual features. As a result, Mamba-based meth-

ods have gained wide acceptance for tasks such as semantic segmentation (Wan et al. 2024), object detection (Wang et al. 2024c), and HSI classification (Li et al. 2024b). While Mamba-based methods are effective in modelling long-range dependencies, they are fundamentally based on processing pixel sequences and have limitations in terms of feature perception capabilities. Therefore, the application of the Mamba architecture to SR tasks still faces the following two challenges: (1) Single spatial perception limits the ability to fully understand and analyze HSIs; (2) Single-scale feature extraction struggles to capture the complex structures and detailed information in HSIs.

To address these challenges, we design a Multi-scale Multi-perceptual Mamba architecture for SR, called M3SR. Specifically, we introduce a Multi-Perceptual Fusion (MPF) block to enhance the ability of the model to comprehensively understand and analyze the input features. Furthermore, by integrating the MPF block into the U-Net structure (Ronneberger, Fischer, and Brox 2015), M3SR is able to capture complex structures and detailed information in HSIs. The architecture of M3SR integrates multi-scale and multi-perceptual feature extraction capabilities to facilitate efficient SR by capturing and fusing both detailed and global features across spatial, frequency, and spectral dimensions at multiple scales. As shown in Fig. 1, the proposed M3SR strikes a suitable balance between reconstruction accuracy and computational cost. Finally, our M3SR model achieves superior performance compared to state-of-the-art (SOTA) super-resolution models on four public benchmark datasets, while using fewer parameters and requiring less computational resources. The main contributions of this paper are as follows.

- A multi-perceptual fusion (MPF) block is designed, which adaptively integrates spatial, frequency, and spectral information, significantly enhancing the richness and robustness of feature representation.
- A multi-scale, multi-perceptual Mamba architecture (M3SR) is proposed, which is capable of extracting perceptual features at global, intermediate, and local scales, and effectively fusing multi-scale information for efficient SR.
- Extensive quantitative and qualitative experiments demonstrate that our proposed M3SR achieves superior performance compared to state-of-the-art (SOTA) methods, with significantly fewer parameters and lower computational cost (FLOPs).

Related Work

Spectral Reconstruction from RGB Images

SR technology aims to reconstruct HSIs from RGB images, tackling challenges such as high acquisition costs and low spatial resolution. Existing SR methods can be categorized into traditional methods, CNN-based methods, and Transformer-based methods. Traditional methods use statistical priors for SR, such as sparse dictionaries (Arad and Ben-Shahar 2016), Gaussian processes (Akhtar and Mian 2018), and low-rank representations (Gao et al. 2020), but

they are limited in reconstructing high-quality HSI (Wu et al. 2023a). CNN-based methods use CNNs to map RGB to HSI, with models such as HSCNN+ (Shi et al. 2018) and HRNet (Zhao et al. 2020) delivering strong results, although they struggle to capture long-range dependencies. Transformer-based methods (e.g., MST++ (Cai et al. 2022a) and ESSAformer (Zhang et al. 2023)) excel in long-range correlation modelling, but the high computational cost limits their application. Therefore, we design M3SR based on Mamba, which can capture long-range dependencies with linear complexity to achieve efficient SR.

State Space Models for Low-level Vision

Recently, SSMs are increasingly applied to low-level vision tasks, with remarkable results (Zhang et al. 2024; Wu et al. 2024; Weng et al. 2025). MambaIR (Guo et al. 2025) first introduces SSMs to image restoration and provides a simple baseline for related research. Subsequently, MambaFormerSR (Zhi, Fan, and Shi 2024) combines SSMs with the Transformer architecture and demonstrates superior performance in super-resolution tasks for remote sensing images. SRMamba-T (Liu et al. 2025) further innovates by designing a multi-directional selective scanning module and a feature fusion module, achieving an excellent balance between computational efficiency and performance in single-image super-resolution tasks.

However, these methods represent only the initial exploration of SSMs in low-level vision tasks. Their single-scale feature extraction mechanisms struggle to fully capture the complex structures and detailed information in images. To address this, U-net is integrated with SSMs due to its multi-scale feature extraction capabilities, which leads to significant advances in several low-level vision tasks. For example, in dehazing (Zheng and Wu 2024), deblurring (Liu et al. 2024a; Kong et al. 2024), and image restoration (Deng and Gu 2024; Sepehri, Fabian, and Soltanolkotabi 2024), the combination of SSMs and U-net shows excellent performance. Meanwhile, thanks to the advantages of frequency domain analysis in localization and multi-resolution capabilities, the integration of SSMs with frequency domain analysis also shows excellent results in several low-level vision tasks. For example, the combination of SSMs and frequency domain analysis achieves remarkable success in super-resolution (Xiao et al. 2024), deraining (Li et al. 2024a) and magnetic resonance imaging reconstruction (Zou et al. 2024).

Nevertheless, the task of reconstructing HSIs from RGB images using SSMs still faces numerous challenges. GMSR (Wang et al. 2024a) is the first method to combine SSMs with spatial and spectral gradients for SR tasks. While GMSR significantly improves computational efficiency, it relies solely on SSMs to extract spatial information, limiting the ability of SSMs to comprehensively understand and analyze images. Additionally, its single-scale feature extraction mechanism struggles to balance global semantic information with local details. Therefore, we propose a multi-scale, multi-perceptual Mamba architecture for SR tasks, aiming to enable comprehensive understanding and analysis of images, as well as to effectively capture complex structures and

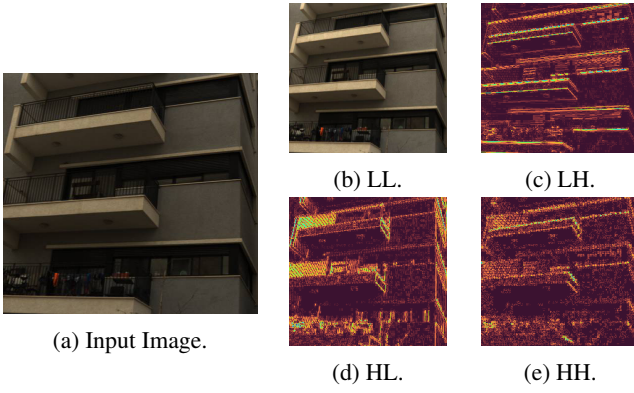


Figure 2: Input image and sub-bands after discrete wavelet transform decomposition.

fine-grained details.

Preliminaries

Discrete Wavelet Transform: The DWT is a highly effective method of image processing, allowing detailed analysis in both the time and frequency domains to accurately capture localized image features (Latif, Abdulredha, and Hassan 2024). By applying the Inverse Discrete Wavelet Transform (IDWT) to the transformed signal, the original image can be accurately reconstructed using the wavelet coefficients. The DWT decomposes an input image $I \in \mathbb{R}^{H \times W \times C}$ into four distinct sub-bands: I_{LL} , I_{LH} , I_{HL} , and I_{HH} . This decomposition can be expressed as:

$$I_{LL}, I_{LH}, I_{HL}, I_{HH} = DWT(I), \quad (1)$$

where I_{LL} represents a low-frequency approximation suitable for image scaling, while I_{LH} , I_{HL} , and I_{HH} emphasize high-frequency details such as edges in different orientations. As shown in Fig. 2, after DWT decomposition, the resulting sub-bands contain texture information at different resolutions. By performing analysis and processing of these sub-bands, key texture features in the image can be effectively extracted, leading to more efficient and accurate SR.

State Space Model: SSM is a framework designed to describe the state representation of a sequence at each time step and predict its next state based on given inputs. It operates by mapping a one-dimensional input function or sequence $x(t) \in \mathbb{R}$ to an output $y(t) \in \mathbb{R}$ through a hidden state representation $h(t) \in \mathbb{R}^N$. This transformation can be mathematically formulated using a linear ordinary differential equation as shown below:

$$h'(t) = \mathbf{A}h(t) + \mathbf{B}x(t), \quad (2)$$

$$y(t) = \mathbf{C}h(t). \quad (3)$$

where $\mathbf{A} \in \mathbb{R}^{N \times N}$ represents the state matrix, $\mathbf{B} \in \mathbb{R}^{N \times 1}$ and $\mathbf{C} \in \mathbb{R}^{1 \times N}$ denote the projection parameters.

To make the continuous system compatible with deep learning, the zero-order hold (ZOH) technique is employed,

using a time-scale parameter Δ to facilitate discretization. This approach simplifies the conversion of the continuous parameters \mathbf{A} and \mathbf{B} into their discrete counterparts $\bar{\mathbf{A}}$ and $\bar{\mathbf{B}}$. After discretization, the discretized SSM system can be formulated as follows:

$$h_t = \bar{\mathbf{A}}h_{t-1} + \bar{\mathbf{B}}x_t, \quad (4)$$

$$y_t = \mathbf{C}h_t. \quad (5)$$

To enhance the computational efficiency and scalability, the convolution operation is harnessed to expedite the linear recurrence process outlined above. Consequently, the ultimate output can be synthesized as

$$\bar{\mathbf{K}} = (\mathbf{C}\bar{\mathbf{B}}, \mathbf{C}\bar{\mathbf{A}}\bar{\mathbf{B}}, \dots, \mathbf{C}\bar{\mathbf{A}}^{L-1}\bar{\mathbf{B}}), \quad (6)$$

$$y = x * \bar{\mathbf{K}}, \quad (7)$$

where L is the length of the input sequence x , and $\bar{\mathbf{K}} \in \mathbb{R}^L$ represents a structured convolutional kernel.

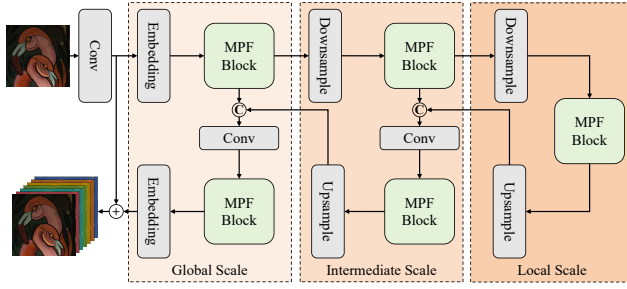
To tackle the limitation of SSMs in capturing contextual information, Mamba introduces a selection mechanism and proposes the selective SSM (S6) to achieve dynamic interactions between sequential states. Additionally, Mamba introduces a hardware-aware algorithm that computes the model recurrently with a scan, ensuring both effectiveness and efficiency in capturing global contextual information.

Methodology

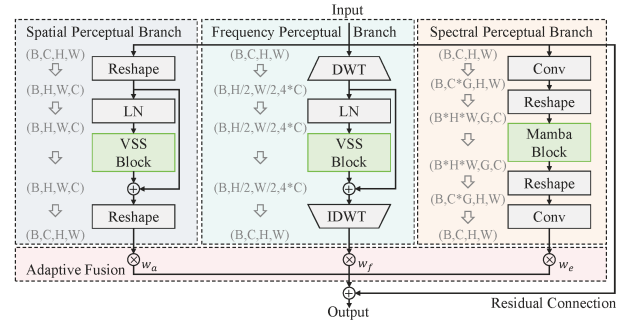
The Overall Architecture

The architecture of M3SR is illustrated in Fig. 3a. Based on the U-net structure, M3SR uses an encoder-decoder structure to efficiently extract and fuse multi-scale features. The architecture progressively extracts multi-scale semantic features by downsampling operations and restores spatial resolution by upsampling operations, forming a symmetric feature processing workflow. The multi-scale feature extraction mechanism of M3SR is divided into three scales: global scale, intermediate scale and local scale. The global scale captures the overall structure, the intermediate scale focuses on contextual information of the HSI, while the local scale concentrates on recovering detailed texture information.

The MPF block is designed as the fundamental block of M3SR. As shown in Fig. 3b, the MPF block achieves feature extraction across multiple perceptual domains through three parallel branches. The first branch focuses on spatial perception, capturing the spatial structure and geometric features. The second branch uses DWT for frequency perception, extracting multiple frequency features. The third branch targets spectral perception, enhancing the ability of the MPF block to model spectral features. Through this multi-perceptual feature extraction and fusion mechanism, M3SR efficiently integrates spatial, frequency and spectral feature information. It maintains the overall structural consistency of the image while accurately recovering detailed texture and high-frequency information, ultimately achieving high-quality reconstruction of HSIs.

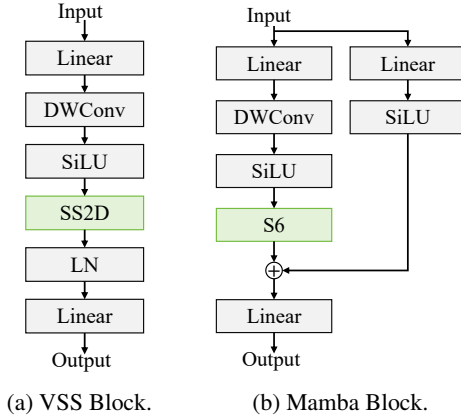


(a) The architecture of our M3SR.



(b) The structure of the MPF Block.

Figure 3: The architecture of our M3SR and the structure of the multi-perceptual fusion (MPF) block.



(a) VSS Block.

(b) Mamba Block.

Figure 4: The structures of VSS block and Mamba block.

Multi-perceptual Fusion Block

As a core component, the MPF block is proposed to overcome the limitations of single perception. By exploiting the perception of multi-dimensional information, including spatial, frequency, and spectral domains, more accurate and comprehensive SR is achieved. As shown in Fig. 3b, there are three main branches, spatial perceptual, frequency perceptual, and spectral perceptual branches, which are introduced as follows.

Spatial Perceptual Branch: In this branch, we employ the Mamba architecture as the foundational component to develop a spatial perceptual branch capable of modeling long-range dependencies with linear computational efficiency. Although the core mechanism of Mamba and its parallel selective scanning process is highly effective, it is primarily designed for one-dimensional sequential data processing. Therefore, the 2D Selective-Scan (SS2D) technique from VMamba is used to capture two-dimensional spatial information. Specifically, the SS2D technique transforms a 2D image into a 1D sequence and conducts scans along four distinct directions to extract long-range dependencies within each sequence. These sequences are subsequently aggregated through summation and reconfigured into the original 2D format. The structure of the VSS block is presented

in Fig. 4a, formulated as:

$$VSS(x) = Lin(LN(SS2D(x'))), \quad (8)$$

$$x' = SiLU(DWConv(Lin(x))), \quad (9)$$

where x is the input, Lin is the linear layer, $SiLU$ is the SiLU activation function, $DWConv$ is the depth-wise convolution and $SS2D$ is the 2D Selective-Scan module.

Subsequently, the spatial perceptual branch is designed based on the VSS block to improve the spatial feature extraction. The input feature \mathbf{F}_{in} is first reshaped to obtain the expanded feature \mathbf{F}_a^1 . This feature is then processed by the VSS block, which extracts spatial information from the HSI. Finally, the features are reshaped to generate the output feature \mathbf{F}_a^2 , as defined below:

$$\mathbf{F}_a^1 = Reshape(\mathbf{F}_{in}), \quad (10)$$

$$\mathbf{F}_a^2 = Reshape(Concat(VSS(LN(\mathbf{F}_a^1)), \mathbf{F}_a^1)), \quad (11)$$

where LN is the layer norm, and VSS is the VSS block.

Frequency Perceptual Branch: To comprehensively and accurately extract texture features from HSIs, a frequency perceptual branch is designed to fully leverage information from different frequency domains. Specifically, DWT is used to decompose the input feature \mathbf{F}_{in} into sub-bands, resulting in one low-frequency component and three high-frequency components to separate the local details from the overall information. Then, the VSS block is used to perform feature learning on these frequency components to capture global contextual information from the low-frequency component while extracting local detail features from the high-frequency components. Finally, Inverse Discrete Wavelet Transform (IDWT) is used to reconstruct the processed frequency components into the original representation \mathbf{F}_a^2 , as defined below:

$$\mathbf{F}_f^1 = DWT(\mathbf{F}_{in}), \quad (12)$$

$$\mathbf{F}_f^2 = IDWT(Concat(VSS(LN(\mathbf{F}_f^1)), \mathbf{F}_f^1)), \quad (13)$$

where DWT is the Discrete Wavelet Transform, and $IDWT$ is the Inverse Discrete Wavelet Transform.

Spectral Perceptual Branch: Unlike traditional vision systems that rely on RGB channels for image capture, HSIs offer a wider spectral range and finer spectral resolution. The spectral data associated with each pixel can be interpreted as a continuous sequence with approximate continuity across wavelengths. Consequently, the spectral perceptual branch is built on the basis of the Mamba block to effectively model the dependencies within the spectral data and extract distinctive features that improve the analysis. The structure of the Mamba block, illustrated in Fig. 4b, is mathematically defined as:

$$Mamba(x) = Lin(x' + x''), \quad (14)$$

$$x' = SiLU(Lin(x)), \quad (15)$$

$$x'' = S6(SiLU(DWConv(Lin(x)))), \quad (16)$$

where x is input and $S6$ is the selective SSM ($S6$).

Using this Mamba block, the spectral perceptual branch can capture relationships across the spectral dimension. For an input feature \mathbf{F}_{in} , the channel dimension C is first expanded to $C \times G$ by a convolutional operation, and the features are segmented into G groups by reshaping. The Mamba block is then used to characterize the interactions and obtain \mathbf{F}_e^1 , refining the spectral information in the process. The refined spectral feature is processed by convolution and reshaping operations to obtain \mathbf{F}_e^2 , as defined below:

$$\mathbf{F}_e^1 = Mamba(Reshape(Conv(\mathbf{F}_{in}))), \quad (17)$$

$$\mathbf{F}_e^2 = Conv(Reshape(\mathbf{F}_e^1)), \quad (18)$$

where $Mamba$ is the Mamba block.

Adaptive Fusion: Finally, the multi-perceptual features extracted from the spatial perceptual, frequency perceptual, and spectral perceptual branches are adaptively fused to achieve an accurate reconstruction of the HSI. For this purpose, three adaptive weights are used to evaluate the importance of the spatial, frequency, and spectral perceptual features and fuse them with a weighted manner. In addition, we incorporate a residual connection to mitigate overfitting that may occur during training. The fusion process is expressed as follows:

$$\mathbf{F}_{out} = \omega_a * \mathbf{F}_a^2 + \omega_f * \mathbf{F}_f^2 + \omega_e * \mathbf{F}_e^2 + \mathbf{F}_{in}, \quad (19)$$

where ω_a , ω_f , and ω_e denote the fusion weights of spatial, frequency, and spectral features, respectively. Note that ω_a , ω_f , and ω_e are randomly initialized and updated through back propagation to determine the final fusion weights.

Loss Function

The M3SR network optimization employs the mean absolute error (MAE) as the loss function, striking a balance between simplicity and effectiveness. Despite not using a more complex loss strategy, the network shows strong convergence during training, demonstrating its ability to capture fine image detail while maintaining structural integrity. Let the reconstructed HSI be $Z \in \mathbb{R}^{H \times W \times C}$ and the reference HSI be $\hat{Z} \in \mathbb{R}^{H \times W \times C}$. The loss function L is defined as:

$$L = \frac{1}{H \times W \times C} \sum_{i=1}^H \sum_{j=1}^W \sum_{k=1}^C |Z_{i,j,k} - \hat{Z}_{i,j,k}|, \quad (20)$$

where H , W , and C denote the height, width, and spectral channels of the image, respectively.

Experiments

Dataset and Evaluative Metrics

To evaluate the performance of M3SR, we conducted extensive experiments using four well-known public SR datasets: NTIRE2022 (Arad et al. 2022), NTIRE2020-Clean, NTIRE2020-Realworld (Arad et al. 2020), and CAVE (Yasuma et al. 2010). Specifically, the NTIRE2022 dataset contains 900 RGB-HSI training pairs, 50 RGB-HSI validation pairs, and 50 unpaired RGB test images. Each of the NTIRE2020-Clean and NTIRE2020-Realworld datasets contains 450 RGB-HSI training pairs, 10 validation pairs and 20 unpaired test images. The HSIs in these three datasets have a resolution of 482×512 and 31 spectral bands covering 400-700 nm. The CAVE dataset consists of 32 RGB HSI pairs, each with a resolution of 512×512 and 31 spectral bands spanning 400-700 nm.

To ensure fair comparisons, we reorganize the datasets into training, validation and test sets. Since the original test images of NTIRE2022 and NTIRE2020 are no longer available after the competitions, we adjust the splits as follows: for NTIRE2022, the training set (900 pairs) remains unchanged. The validation set of 50 pairs is split into 20 pairs for validation and 30 pairs for testing. For NTIRE2020-Clean and NTIRE2020-Realworld, the test set is formed by taking 10 pairs from the validation set and 20 pairs randomly selected from the training set. From the remaining training pairs, 30 pairs are randomly selected for validation and the rest remain in the training set. For CAVE, we assign three pairs (face, fake_and_real_beers, watercolors) for validation, three pairs (balloons, oil_painting, fake_and_real_strawberries) for testing, and the remaining pairs for training. Following the method in (Wu et al. 2023b), evaluation metrics include root mean square error (RMSE), peak signal-to-noise ratio (PSNR), spectral angle mapper (SAM), and mean structural similarity (MSSIM). RMSE and PSNR measure numerical accuracy, SAM evaluates spectral fidelity, and MSSIM assesses spatial structure recovery. Note that the symbols \uparrow and \downarrow in tables are used to indicate that larger or smaller values of the metric are better, respectively.

Implementation Details

During the training process, RGB images are linearly rescaled to $[0, 1]$, and 128×128 RGB and HSI sample pairs are cropped. The batch size is set to 32, and the Adam optimizer with $\beta_1 = 0.9$ and $\beta_2 = 0.999$ is used for parameter optimization. The learning rate is initialized to 0.0004 and follows a cosine annealing schedule over 100 epochs. Training data augmentation includes random rotation and flipping. The models are implemented using the PyTorch

Methods	Params (M)	FLOPs (G)	NTIRE2022				CAVE			
			RMSE(↓)	PSNR(↑)	SAM(↓)	MSSIM(↑)	RMSE(↓)	PSNR(↑)	SAM(↓)	MSSIM(↑)
HSCNN+ (Shi et al. 2018, CVPR)	1.642	807.998	0.058	25.2622	5.9151	0.8576	0.0227	33.8117	6.9895	0.9718
FMNet (Zhang et al. 2020, AAAI)	11.793	5819.522	0.0586	25.2593	6.8735	0.863	0.0247	33.1439	8.5393	0.9706
HRNet (Zhao et al. 2020, CVPR)	31.705	1249.005	0.0577	25.2223	6.055	0.8772	0.0205	34.531	6.6962	0.9787
GDNet (Zhu et al. 2021, TCI)	0.772	774.456	0.0455	27.3483	8.1031	0.8841	0.0291	31.3349	8.6479	0.9685
HSRNet (He et al. 2021, TNNLS)	0.683	321.078	0.0438	28.6688	6.4454	0.9111	0.0227	33.3828	7.8339	0.9775
SSDCN (Chen, Zheng, and Lu 2021, TGRSL)	0.415	196.525	0.0623	24.9525	7.9265	0.8501	0.0218	34.2542	7.1912	0.9778
MST++ (Cai et al. 2022b, CVPR)	1.62	177.737	0.035	30.1822	6.0892	0.9308	0.0205	34.6534	6.6725	0.9826
RepCPSI (Wu et al. 2023b, TGRS)	2.078	1023.235	0.0409	29.2083	6.4858	0.9181	0.0229	33.6155	7.4338	0.9772
SSRNet (Dian et al. 2023, TNNLS)	0.388	190.479	0.0465	28.2843	6.1205	0.9118	0.0252	32.7655	7.4848	0.9767
GMSR (Wang et al. 2024a, arXiv)	0.019	7.988	0.0492	26.9156	8.2357	0.8798	0.0206	34.5797	7.418	0.9769
M3SR [Ours]	2.166	100.924	0.0343	31.3995	6.6199	0.9351	0.0184	35.6063	6.7887	0.9814

Table 1: Reconstruction results on the NTIRE2022 and CAVE datasets. The **champion** and **second place** for each metric are marked in underlined bold and bold, respectively.

Methods	Params (M)	FLOPs (G)	NTIRE2020-Clean				NTIRE2020-realworld			
			RMSE(↓)	PSNR(↑)	SAM(↓)	MSSIM(↑)	RMSE(↓)	PSNR(↑)	SAM(↓)	MSSIM(↑)
HSCNN+ (Shi et al. 2018, CVPR)	1.642	807.998	0.0274	32.7659	3.628	0.9555	0.0281	31.5774	4.3754	0.9427
FMNet (Zhang et al. 2020, AAAI)	11.793	5819.522	0.0241	33.634	3.6259	0.9757	0.0247	32.8718	4.2207	0.9583
HRNet (Zhao et al. 2020, CVPR)	31.705	1249.005	0.0309	32.0203	3.5785	0.9617	0.0282	32.1507	4.1579	0.9527
GDNet (Zhu et al. 2021, TCI)	0.772	774.456	0.0196	36.1221	3.9366	0.9777	0.0192	35.3121	4.687	0.9611
HSRNet (He et al. 2021, TNNLS)	0.683	321.078	0.0198	37.1743	3.6409	0.9841	0.0213	34.5468	4.6183	0.9592
SSDCN (Chen, Zheng, and Lu 2021, TGRSL)	0.415	196.525	0.0216	35.7702	3.9704	0.9777	0.0199	35.0647	4.3352	0.96
MST++ (Cai et al. 2022b, CVPR)	1.62	177.737	0.0198	36.324	3.3886	0.9833	0.0185	35.6348	4.3617	0.9636
RepCPSI (Wu et al. 2023b, TGRS)	2.078	1023.235	0.0212	36.2077	3.6616	0.9802	0.0203	34.9522	4.0163	0.9625
SSRNet (Dian et al. 2023, TNNLS)	0.388	190.479	0.0239	35.0742	3.6473	0.9757	0.021	35.0337	4.2956	0.9591
GMSR (Wang et al. 2024a, arXiv)	0.019	7.988	0.0239	33.9653	3.9759	0.9644	0.0278	31.9018	4.5457	0.9456
M3SR [Ours]	2.166	100.924	0.0196	37.7088	3.2842	0.9841	0.0171	36.3524	4.2843	0.964

Table 2: Reconstruction results on the NTIRE2020-Clean and NTIRE2020-Realworld datasets. The **champion** and **second place** for each metric are marked in underlined bold and bold, respectively.

framework and trained on a single NVIDIA 4090 GPU. In our M3SR implementation, the number of groups is $G = 4$.

Comparisons With SOTA Methods

To demonstrate the effectiveness of the M3SR, we compare the proposed M3SR with ten state-of-the-art (SOTA) methods on the four benchmarks, including HSCNN+ (Shi et al. 2018), FMNet (Zhang et al. 2020), HRNet (Zhao et al. 2020), GDNet (Zhu et al. 2021), HSRNet (He et al. 2021), SSDCN (Chen, Zheng, and Lu 2021), MST++ (Cai et al. 2022b), RepCPSI (Wu et al. 2023b), SSRNet (Dian et al. 2023) and GMSR (Wang et al. 2024a).

Table 1 presents performance comparisons between the proposed M3SR and state-of-the-art (SOTA) SR methods on the NTIRE2022 and CAVE datasets. It shows that our method M3SR achieves the best performance in terms of RMSE, PSNR, and MSSIM on the NTIRE2022 dataset, and the best RMSE and PSNR on the CAVE dataset. In addition, Table 1 shows the parameters and computational cost of each method using a 482×512 RGB input, demonstrating that M3SR strikes an excellent balance between computational cost and model performance. Similarly, Table 2 shows the performance on the NTIRE2020-Clean and NTIRE2020-Realworld datasets. Our method M3SR achieves the best performance in RMSE, PSNR, SAM and MSSIM on the NTIRE2020-Clean dataset and the best RMSE, PSNR and MSSIM on the NTIRE2020-Realworld dataset. These results demonstrate that the proposed M3SR provides superior reconstruction performance at a significantly lower computational cost.

Method	SPA	FRE	SPE	RMSE(↓)	PSNR(↑)	SAM(↓)	MSSIM(↑)
M3SR-V1		✓	✓	0.0381	30.4932	12.5459	0.8827
M3SR-V2	✓		✓	0.0365	30.5882	6.52	0.9315
M3SR-V3	✓	✓		0.0369	30.3648	6.2789	0.9241
M3SR	✓	✓	✓	0.0343	31.3995	6.6199	0.9351

Table 3: Results of ablation experiments on effects of the spatial perceptual branch (SPA), frequency perceptual branch (FRE), and the spectral perceptual branch (SPE). The best scores are highlighted in **underline bold**.

To further illustrate the superiority of M3SR, we randomly select a spectral band and generated a heatmap of the absolute error between the reconstructed image and the input RGB image for the NTIRE2022 dataset. Fig. 5 provides a spatial comparison of the reconstruction performance of the different methods. The bluer the color, the higher the reconstruction accuracy. Obviously, M3SR performs well in reconstructing ARAD_1K_0907 (630nm) for the NTIRE2022 dataset. The observations confirm that M3SR achieves satisfactory reconstruction results, consistent with the quantitative analysis.

Ablation Studies

To validate the effectiveness of the MPF block, ablation studies are provided in this section. Parametric explorations are also provided in order to balance model performance and computational cost. Note that these ablation studies are performed on the NTIRE2022 dataset.

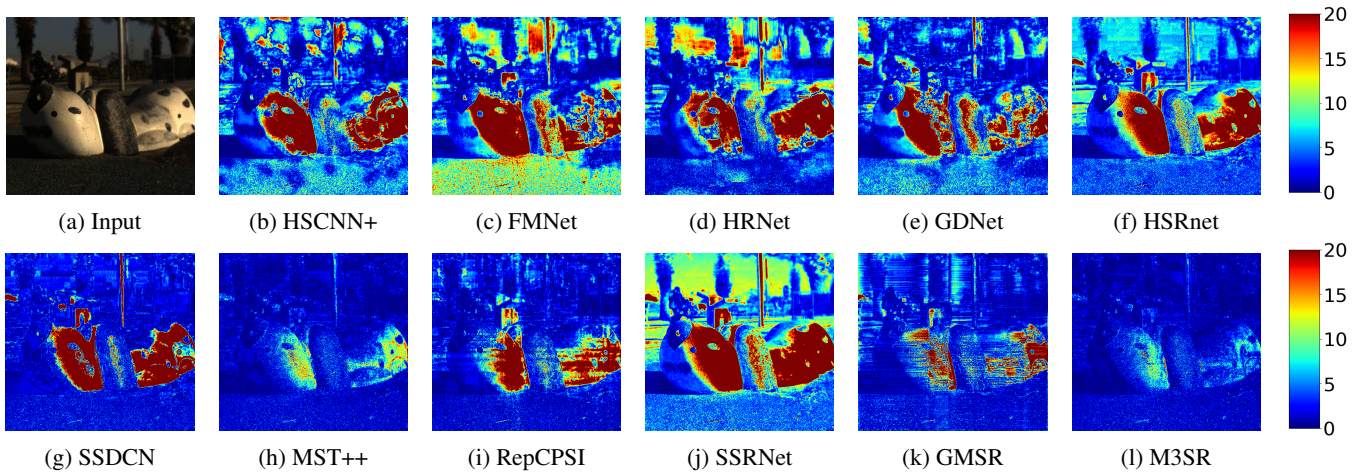


Figure 5: Comparison of reconstruction results on the NTIRE2022 dataset using input image ARAD_1K_0907 (630 nm).

Effects of Multi-Perceptual Fusion Block This section provides an ablation study to verify the effectiveness of the spatial perceptual, frequency perceptual and spectral perceptual branches. The ablation experimental results of M3SR and its three variants on the NTIRE2022 dataset are presented in Table 3.

Investigation of the Spatial Perceptual Branch: The variant M3SR-V1 is designed by removing the spatial perceptual branch to validate its effectiveness. The ablation results show that the absence of the spatial perceptual branch significantly weakens the ability to reconstruct spatial information, as evidenced by a marked worsening of RMSE and PSNR, together with a sharp increase in SAM. This indicates that the spatial perceptual branch plays a critical role in the reconstruction of spatial details and structural information.

Investigation of the Frequency Perceptual Branch: The variant M3SR-V2 is designed by removing the frequency perceptual branch to validate its effectiveness. The ablation results show that the absence of the frequency perceptual branch has a significant impact on the model performance, with RMSE, PSNR, and MSSIM all deteriorating. This indicates that the frequency perceptual branch is essential for reconstructing high frequency details and texture information in images.

Investigation of the Spectral Perceptual Branch: The variant M3SR-V3 is designed by removing the spectral perceptual branch to validate its effectiveness. The ablation results show that the absence of the spectral perceptual branch affects the model performance to some extent, with RMSE, PSNR, and MSSIM all deteriorating. This demonstrates that the spectral perceptual branch plays an important role in the reconstruction of spectral information in images.

In summary, these ablation results highlight the importance of all three branches of the MPF block, further validating the effectiveness of the M3SR architecture.

Parametric Analysis The M3SR relies on a key parameter, the number of groups G used in the spectral perceptual branch, and parametric analysis is conducted to determine an

Group	RMSE(↓)	PSNR(↑)	SAM(↓)	MSSIM(↑)	Params (M)	FLOPs (G)
2	0.0351	31.2306	7.7846	0.9202	2.066	91.32
4	0.0343	31.3995	6.6199	0.9351	2.166	100.924
8	0.0372	30.6399	7.348	0.9238	2.368	120.133
16	0.0361	30.7574	11.81	0.8793	2.77	158.551

Table 4: Results of ablation experiments on the number of groups. The best scores are highlighted in **underline bold**.

effective setting for this parameter. We evaluate the impact of the number of groups G on model performance by testing values of 2, 4, 8, and 16. As shown in Table 4, the model with $G = 4$ achieves the best results across most metrics, including RMSE, PSNR, SAM, and MSSIM. Although the model with $G = 2$ is lighter in terms of parameter count and computational complexity, its performance metrics are slightly lower compared to the model with $G = 4$. The models with $G = 8$ and $G = 16$ show competitive performance on some metrics, but the number of parameters and computational demands increase considerably as G grows. Therefore, considering both model performance and computational efficiency, setting $G = 4$ provides a favorable balance and is selected as the preferred configuration.

Conclusion

This paper presents a multi-scale, multi-perceptual Mamba architecture for SR tasks (M3SR) that addresses the limitations of traditional Mamba. The innovative multi-perceptual fusion (MPF) block consists of spatial perceptual, frequency perceptual, and spectral perceptual branches that fully analyze images and capture complex structures. By integrating the MPF block into the U-Net structure, M3SR can efficiently capture and fuse global, intermediate, and local features for superior SR. Extensive experiments confirm that M3SR not only outperforms existing SOTA SR methods, but also significantly reduces computational cost, making it an efficient solution for SR tasks.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (NSFC) under Grants 62506238, 62376163, 62173236, 62271324 and 62231020; in part by the Guangdong Regional Joint Foundation Key Project under Grant 2022B1515120076; in part by the Scientific Research Capacity Enhancement Program for Key Construction Disciplines in Guangdong Province under Grant 2024ZDJS063; in part by Shenzhen Technology University School-level Research Project under Grant 20251061020002; and in part by Shenzhen Science and Technology Program under Grant ZDSYS20220527171400002.

References

- Akhtar, N.; and Mian, A. 2018. Hyperspectral recovery from RGB images using Gaussian processes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(1): 100–113.
- Arad, B.; and Ben-Shahar, O. 2016. Sparse recovery of hyperspectral signal from natural RGB images. In *European Conference on Computer Vision*, 19–34. Springer.
- Arad, B.; Timofte, R.; Ben-Shahar, O.; Lin, Y.-T.; and Finlayson, G. D. 2020. Ntire 2020 challenge on spectral reconstruction from an rgb image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 446–447.
- Arad, B.; Timofte, R.; Yahel, R.; Morag, N.; Bernat, A.; Cai, Y.; Lin, J.; Lin, Z.; Wang, H.; Zhang, Y.; et al. 2022. Ntire 2022 spectral recovery challenge and data set. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 863–881.
- Cai, Y.; Lin, J.; Hu, X.; Wang, H.; Yuan, X.; Zhang, Y.; Timofte, R.; and Van Gool, L. 2022a. Mask-guided spectral-wise transformer for efficient hyperspectral image reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17502–17511.
- Cai, Y.; Lin, J.; Lin, Z.; Wang, H.; Zhang, Y.; Pfister, H.; Timofte, R.; and Van Gool, L. 2022b. MST++: Multi-stage spectral-wise transformer for efficient spectral reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 745–755.
- Chen, W.; Zheng, X.; and Lu, X. 2021. Semisupervised spectral degradation constrained network for spectral super-resolution. *IEEE Geoscience and Remote Sensing Letters*, 19: 1–5.
- Deng, R.; and Gu, T. 2024. Cu-mamba: Selective state space models with channel learning for image restoration. In *2024 IEEE 7th International Conference on Multimedia Information Processing and Retrieval (MIPR)*, 328–334. IEEE.
- Dian, R.; Shan, T.; He, W.; and Liu, H. 2023. Spectral super-resolution via model-guided cross-fusion network. *IEEE Transactions on Neural Networks and Learning Systems*, 35(7): 10059–10070.
- Gao, L.; Hong, D.; Yao, J.; Zhang, B.; Gamba, P.; and Chanussot, J. 2020. Spectral superresolution of multispectral imagery with joint sparse and low-rank learning. *IEEE Transactions on Geoscience and Remote Sensing*, 59(3): 2269–2280.
- Gu, A.; and Dao, T. 2023. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*.
- Guo, H.; Li, J.; Dai, T.; Ouyang, Z.; Ren, X.; and Xia, S.-T. 2025. Mambair: A simple baseline for image restoration with state-space model. In *European Conference on Computer Vision*, 222–241. Springer.
- He, J.; Li, J.; Yuan, Q.; Shen, H.; and Zhang, L. 2021. Spectral response function-guided deep optimization-driven network for spectral super-resolution. *IEEE Transactions on Neural Networks and Learning Systems*, 33(9): 4213–4227.
- Khan, A.; Vibhute, A. D.; Mali, S.; and Patil, C. H. 2022. A systematic review on hyperspectral imaging technology with a machine and deep learning methodology for agricultural applications. *Ecological Informatics*, 69: 101678.
- Kong, L.; Dong, J.; Yang, M.-H.; and Pan, J. 2024. Efficient visual state space model for image deblurring. *arXiv preprint arXiv:2405.14343*.
- Lassalle, G. 2021. Monitoring natural and anthropogenic plant stressors by hyperspectral remote sensing: Recommendations and guidelines based on a meta-review. *Science of the Total Environment*, 788: 147758.
- Latif, I. H.; Abdulredha, S. H.; and Hassan, S. K. A. 2024. Discrete Wavelet Transform-Based Image Processing: A Review. *Al-Nahrain Journal of Science*, 27(3): 109–125.
- Leon, R.; Fabelo, H.; Ortega, S.; Cruz-Guerrero, I. A.; Campos-Delgado, D. U.; Szolna, A.; Piñeiro, J. F.; Espino, C.; O’Shanahan, A. J.; Hernandez, M.; et al. 2023. Hyperspectral imaging benchmark based on machine learning for intraoperative brain tumour detection. *npj Precision Oncology*, 7(1): 119.
- Li, D.; Liu, Y.; Fu, X.; Xu, S.; and Zha, Z.-J. 2024a. Fourier-mamba: Fourier learning integration with state space models for image deraining. *arXiv preprint arXiv:2405.19450*.
- Li, Y.; Luo, Y.; Zhang, L.; Wang, Z.; and Du, B. 2024b. MambaHSI: Spatial-Spectral Mamba for Hyperspectral Image Classification. *IEEE Transactions on Geoscience and Remote Sensing*, 62: 1–16.
- Liu, C.; Zhang, D.; Lu, G.; Yin, W.; Wang, J.; and Luo, G. 2025. SRMamba-T: Exploring the hybrid Mamba-Transformer network for single image super-resolution. *Neurocomputing*, 129488.
- Liu, H.; Liu, C.; Xu, J.; Jiang, P.; and Lu, M. 2024a. XYScanNet: An Interpretable State Space Model for Perceptual Image Deblurring. *arXiv preprint arXiv:2412.10338*.
- Liu, J.; Zhang, H.; Tian, J.-H.; Su, Y.; Chen, Y.; and Wang, Y. 2024b. R2D2-GAN: Robust Dual Discriminator Generative Adversarial Network for Microscopy Hyperspectral Image Super-Resolution. *IEEE Transactions on Medical Imaging*, 43(11): 4064–4074.
- Liu, Y.; Tian, Y.; Zhao, Y.; Yu, H.; Xie, L.; Wang, Y.; Ye, Q.; and Liu, Y. 2024c. VMamba: Visual State Space Model. *arXiv preprint arXiv:2401.10166*.

- Patro, B. N.; and Agneeswaran, V. S. 2024. Mamba-360: Survey of state space models as transformer alternative for long sequence modelling: Methods, applications, and challenges. *arXiv preprint arXiv:2404.16112*.
- Ram, B. G.; Oduor, P.; Igathinathane, C.; Howatt, K.; and Sun, X. 2024. A systematic review of hyperspectral imaging in precision agriculture: Analysis of its current state and future prospects. *Computers and Electronics in Agriculture*, 222: 109037.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention*, 234–241. Springer.
- Sepehri, M. S.; Fabian, Z.; and Soltanolkotabi, M. 2024. Serpent: Scalable and efficient image restoration via multi-scale structured state space models. *arXiv preprint arXiv:2403.17902*.
- Shi, Z.; Chen, C.; Xiong, Z.; Liu, D.; and Wu, F. 2018. HSCNN+: Advanced cnn-based hyperspectral recovery from rgb images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 939–947.
- Wan, Z.; Zhang, P.; Wang, Y.; Yong, S.; Stepputtis, S.; Sycara, K.; and Xie, Y. 2024. Sigma: Siamese mamba network for multi-modal semantic segmentation. *arXiv preprint arXiv:2404.04256*.
- Wang, X.; Huang, Z.; Zhang, S.; Zhu, J.; Gamba, P.; and Feng, L. 2024a. GMSR: gradient-guided mamba for spectral reconstruction from RGB images. *arXiv preprint arXiv:2405.07777*.
- Wang, Y.; Zou, B.; Chai, L.; Lin, Z.; Feng, H.; Tang, Y.; Tian, R.; Tu, Y.; Zhang, B.; and Zou, H. 2024b. Monitoring of soil heavy metals based on hyperspectral remote sensing: A review. *Earth-Science Reviews*, 104814.
- Wang, Z.; Li, C.; Xu, H.; and Zhu, X. 2024c. Mamba YOLO: SSMs-based YOLO for object detection. *arXiv preprint arXiv:2406.05835*.
- Weng, J.; Yan, Z.; Tai, Y.; Qian, J.; Yang, J.; and Li, J. 2025. Mamballie: Implicit retinex-aware low light enhancement with global-then-local state space. *Advances in Neural Information Processing Systems*, 37: 27440–27462.
- Wu, C.; Li, J.; Song, R.; Li, Y.; and Du, Q. 2023a. HPRN: Holistic prior-embedded relation network for spectral super-resolution. *IEEE Transactions on Neural Networks and Learning Systems*, 35(8): 11409–11423.
- Wu, C.; Li, J.; Song, R.; Li, Y.; and Du, Q. 2023b. RepCPSI: Coordinate-preserving proximity spectral interaction network with reparameterization for lightweight spectral super-resolution. *IEEE Transactions on Geoscience and Remote Sensing*, 61: 1–13.
- Wu, H.; Yang, Y.; Xu, H.; Wang, W.; Zhou, J.; and Zhu, L. 2024. Rainmamba: Enhanced locality learning with state space models for video deraining. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 7881–7890.
- Xiao, Y.; Yuan, Q.; Jiang, K.; Chen, Y.; Zhang, Q.; and Lin, C.-W. 2024. Frequency-assisted mamba for remote sensing image super-resolution. *IEEE Transactions on Multimedia*.
- Yasuma, F.; Mitsunaga, T.; Iso, D.; and Nayar, S. K. 2010. Generalized assorted pixel camera: postcapture control of resolution, dynamic range, and spectrum. *IEEE Transactions on Image Processing*, 19(9): 2241–2253.
- Zhang, L.; Lang, Z.; Wang, P.; Wei, W.; Liao, S.; Shao, L.; and Zhang, Y. 2020. Pixel-aware deep function-mixture network for spectral super-resolution. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 12821–12828.
- Zhang, M.; Li, L.; Shi, W.; Guo, J.; Li, Y.; and Gao, X. 2024. VmambaSCI: Dynamic Deep Unfolding Network with Mamba for Compressive Spectral Imaging. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 6549–6558.
- Zhang, M.; Zhang, C.; Zhang, Q.; Guo, J.; Gao, X.; and Zhang, J. 2023. ESSAformer: Efficient transformer for hyperspectral image super-resolution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 23073–23084.
- Zhao, Y.; Po, L.-M.; Yan, Q.; Liu, W.; and Lin, T. 2020. Hierarchical regression network for spectral reconstruction from RGB images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 422–423.
- Zheng, Z.; and Wu, C. 2024. U-shaped vision mamba for single image dehazing. *arXiv preprint arXiv:2402.04139*.
- Zhi, R.; Fan, X.; and Shi, J. 2024. MambaFormerSR: A Lightweight Model for Remote-Sensing Image Super-Resolution. *IEEE Geoscience and Remote Sensing Letters*, 21: 1–5.
- Zhu, Z.; Liu, H.; Hou, J.; Jia, S.; and Zhang, Q. 2021. Deep amended gradient descent for efficient spectral reconstruction from single RGB images. *IEEE Transactions on Computational Imaging*, 7: 1176–1188.
- Zou, J.; Liu, L.; Chen, Q.; Wang, S.; Hu, Z.; Xing, X.; and Qin, J. 2024. MMR-Mamba: Multi-Modal MRI Reconstruction with Mamba and Spatial-Frequency Information Fusion. *arXiv preprint arXiv:2406.18950*.