

MMMamba: A Versatile Cross-Modal In Context Fusion Framework for Pan-Sharpening and Zero-Shot Image Enhancement

Yingying Wang^{1*}, Xuanhua He^{2*}, Chen Wu^{3*}, Jialing Huang¹, Suiyun Zhang⁴, Rui Liu⁴,
Xinghao Ding¹, Haoxuan Che^{4†}

¹ Key Laboratory of Multimedia Trusted Perception and Efficient Computing,
Ministry of Education of China, Xiamen University, China

²The Hong Kong University of Science and Technology

³University of Science and Technology of China

⁴Huawei Research

wangyingying7@stu.xmu.edu.cn, xhecd@connect.ust.hk, wuchen5x@mail.ustc.edu.cn, che.haoxuan@huawei.com

Abstract

Pan-sharpening aims to generate high-resolution multi-spectral (HRMS) images by integrating a high-resolution panchromatic (PAN) image with its corresponding low-resolution multispectral (MS) image. To achieve effective fusion, it is crucial to fully exploit the complementary information between the two modalities. Traditional CNN-based methods typically rely on channel-wise concatenation with fixed convolutional operators, which limits their adaptability to diverse spatial and spectral variations. While cross-attention mechanisms enable global interactions, they are computationally inefficient and may dilute fine-grained correspondences, making it difficult to capture complex semantic relationships. Recent advances in the Multimodal Diffusion Transformer (MMDiT) architecture have demonstrated impressive success in image generation and editing tasks. Unlike cross-attention, MMDiT employs in-context conditioning to facilitate more direct and efficient cross-modal information exchange. In this paper, we propose MMMamba, a cross-modal in-context fusion framework for pan-sharpening, with the flexibility to support image super-resolution in a zero-shot manner. Built upon the Mamba architecture, our design ensures linear computational complexity while maintaining strong cross-modal interaction capacity. Furthermore, we introduce a novel multimodal interleaved (MI) scanning mechanism that facilitates effective information exchange between the PAN and MS modalities. Extensive experiments demonstrate the superior performance of our method compared to existing state-of-the-art (SOTA) techniques across multiple tasks and benchmarks.

Code — <https://github.com/Gracewangyy/MMMamba>

Introduction

With the growing demand for high-quality satellite imagery in areas such as agriculture (Jenerowicz and Woroszkievicz 2016), urban planning (Aiazzi et al. 2003), and environmental monitoring (Sunuprpto, Danoedoro, and Ritohardoyo

2016), obtaining high-resolution multi-spectral (HRMS) data has become more critical than ever. However, the physical limitations of satellite sensors impede the direct acquisition of multi-spectral images that offer both fine spatial detail and rich spectral information. To address this issue, most satellites are equipped with two separate sensors: panchromatic (PAN) and multi-spectral (MS), each designed to capture complementary aspects. PAN images provide high spatial resolution but limited spectral coverage, while MS images offer rich spectral information at lower spatial resolutions. Pan-sharpening has therefore emerged as a practical and essential technique, aiming to fuse these two data sources into a single image that combines the spatial sharpness of PAN with the spectral fidelity of MS.

Early efforts in pan-sharpening were predominantly based on classical paradigms such as component substitution (CS) (Kwarteng and Chavez 1989), multi-resolution analysis (MRA) (Mallat 2002), and variational optimization (VO) (Ballester et al. 2006). These hand-crafted techniques relied on physical modeling and prior domain knowledge, limiting their ability to capture complex cross-modal relationships and yielding suboptimal results. The introduction of deep learning into the pan-sharpening field has led to significant improvements in both spatial resolution and spectral fidelity. A notable breakthrough was the pioneering PNN model (Masi et al. 2016), which demonstrated remarkable performance improvements over traditional approaches. Since then, the research community has witnessed rapid advancements with increasingly sophisticated neural network architectures (Wang et al. 2025a; Li et al. 2025). Based on varying fusion strategies, these approaches can be broadly categorized into channel concatenation-based methods, such as DIRFL (Lin et al. 2023) and HFEAN (Wang et al. 2023), PAN injection with multi-scale techniques like MSDDN (He et al. 2023) and WaveletNet (Zhang et al. 2024b), cross-attention methods exemplified by Panformer (Zhou, Liu, and Wang 2022) and CMINet (Wang et al. 2024), and gating-based approaches, including FAME (He et al. 2024) and Pan-Mamba (He et al. 2025).

Despite their progress, existing methods still exhibit cer-

*These authors contributed equally.

†Corresponding Author.

tain limitations that impede further performance improvements. CNN-based approaches typically rely on channel-wise concatenation, a static mechanism that lacks the adaptive flexibility to model the complex relationships between modalities. Transformer-based methods, while employing cross-attention and offering more dynamism, still have their drawbacks. First, they aggregate features through weighted averaging, which tends to smooth out the high-frequency spatial details crucial for preserving the integrity of the PAN image. Second, the information flows in only one direction, restricting the depth and richness of the interaction between modalities. Recent architectures, such as the Multimodal Diffusion Transformer (MMDiT) (Esser et al. 2024), have demonstrated significant success in multimodal interaction by adopting an in-context conditioning strategy (Tan et al. 2024; Labs et al. 2025). This approach discards traditional fusion modules like channel concatenation and cross-attention, instead concatenating tokens from all modalities into a single unified sequence, which is then jointly processed by self-attention, enabling deep and bidirectional interactions between all tokens. However, despite its advantages, directly employing this paradigm with Transformers is computationally prohibitive for image fusion due to the quadratic complexity of self-attention. Moreover, its direct application does not guarantee effective cross-modal interaction and integration in image fusion.

In this paper, we propose MMMamba, a novel cross-modal in-context fusion framework for pan-sharpening. Built upon the Mamba architecture, our design achieves linear computational complexity while maintaining strong cross-modal interaction capacity. To fully unleash the potential of in-context conditioning within our framework for pan-sharpening task, we introduce a specially designed multimodal interleaved (MI) scanning mechanism that facilitates effective information exchange between the PAN and MS modalities. This method arranges the input sequence so that corresponding PAN and MS tokens are spatially adjacent and can be scanned from different directions. A key advantage of this powerful and unified design is zero-shot task generalization: trained solely on pan-sharpening, MM-Mamba can perform MS image super-resolution by simply dropping the input PAN modality, without requiring retraining or fine-tuning. Extensive experiments across multiple benchmarks demonstrate that MMMamba consistently outperforms existing state-of-the-art (SOTA) methods both visually and quantitatively.

To summarize, this work offers the following key contributions:

- We propose MMMamba, a novel cross-modal in-context fusion framework for pan-sharpening. Built upon the Mamba architecture, it achieves linear complexity and enables bidirectional information flow, while also supporting zero-shot generalization to image super-resolution task.
- We are the first to explore the in-context conditioning paradigm in pan-sharpening, enabling deep and efficient cross-modal interactions among all tokens, thereby achieving superior multimodal image fusion results.

- We design a novel multimodal interleaved (MI) scanning mechanism that facilitates bidirectional information exchange by effectively exploiting complementary cues between PAN and MS modalities.
- Extensive experiments conducted on multiple benchmarks demonstrate that MMMamba consistently outperforms existing SOTA methods across various tasks.

Related Work

Pan-Sharpener

Pan-sharpening can be categorized into conventional and deep learning-based approaches. Early studies predominantly relied on prior knowledge and handcrafted features, including Component Substitution (CS) (Kwarteng and Chavez 1989; Gillespie, Kahle, and Walker 1987), Multi-Resolution Analysis (MRA) (Schowengerdt 1980; Nunez et al. 2002), and Variational Optimization (VO) (Fasbender, Radoux, and Bogaert 2008; Ballester et al. 2006). While traditional approaches offered interpretability and computational efficiency, their limited capacity to model the complex and nonlinear correlations between PAN and MS modalities hindered their performance. The advent of deep learning has reshaped the landscape of pan-sharpening (Huang et al. 2023; Zhang et al. 2024a; Meng et al. 2025). PNN (Masi et al. 2016) first introduced a simple three-layer CNN that achieved promising results. This was followed by a surge of CNN-based sophisticated models, such as PanNet (Yang et al. 2017), HFEAN (Wang et al. 2023), BiMPan (Hou et al. 2023), PIF-Net (Meng et al. 2024). More recently, the integration of Transformer-based models, such as Panformer (Zhou, Liu, and Wang 2022), CMINet (Wang et al. 2024), LFormer (Hou et al. 2024), and FCSA (Wu et al. 2025), introduced self-attention mechanisms to capture long-range dependencies, significantly improving the modeling of spatial relationships.

State Space Model

State Space Models (SSMs) have recently emerged as a powerful alternative to CNNs and Transformers, owing to their long-range dependencies with linear computational complexity. S4 (Gu, Goel, and Ré 2021) introduced diagonal state-space parameterizations for efficient parallelization, and Mamba (Gu and Dao 2023) further incorporated a dynamic selection mechanism to enhance training and sequence modeling. Recent research has successfully adapted SSMs to the visual domain by reshaping images into sequential representations and integrating specialized scanning mechanisms. Specifically, Vmamba (Liu et al. 2024) and Vision Mamba (Zhu et al. 2024) employed directionally-aware scanning schemes to effectively model spatial structures, facilitating the integration of contextual information from various perspectives. LEVM (Cao et al. 2024) introduced a local-enhanced vision Mamba block tailored for image fusion tasks, which strengthened local spatial perception and improved the integration of spatial and spectral information. Pan-Mamba (He et al. 2025) is the pioneering work that introduces Mamba into pan-sharpening,

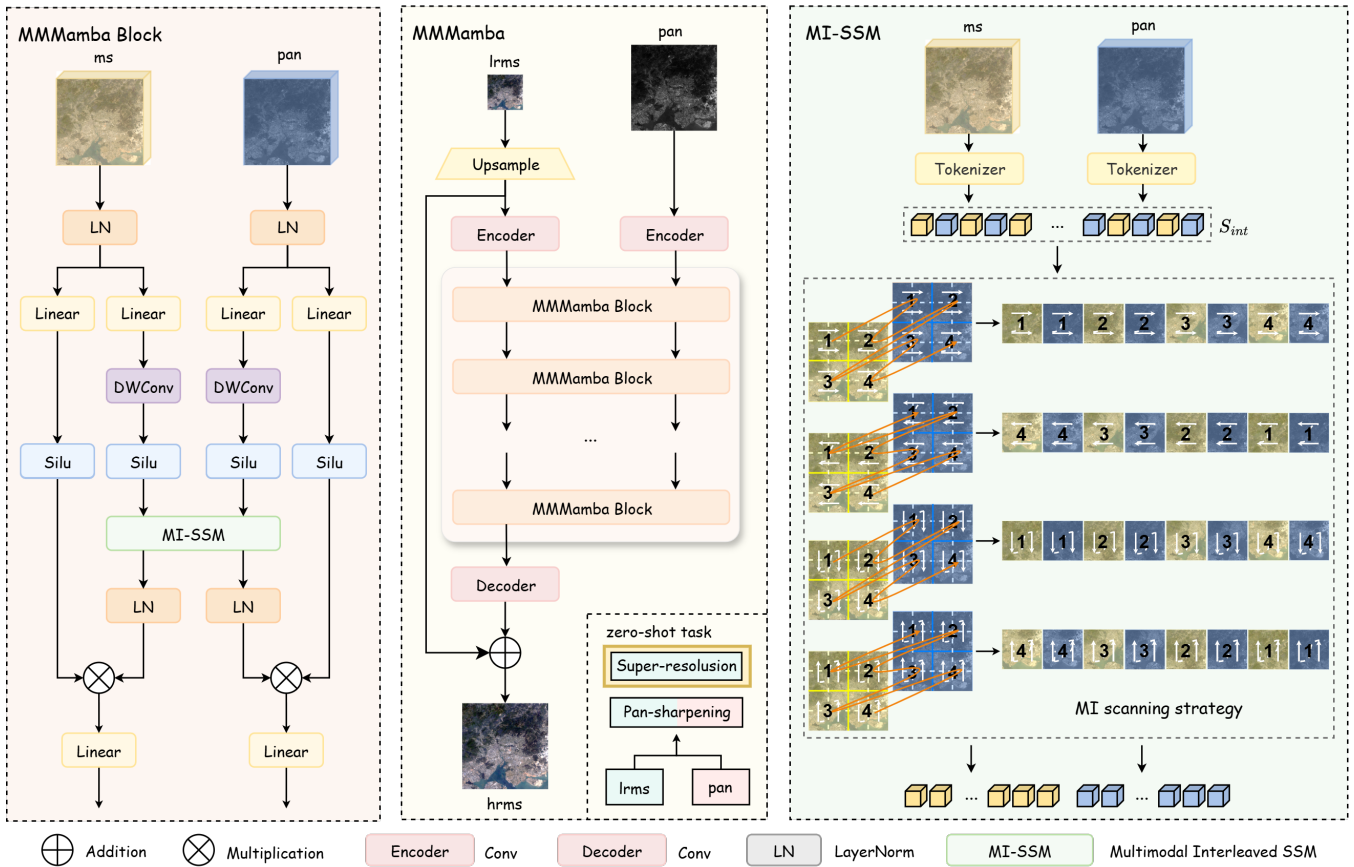


Figure 1: The overall framework of our proposed MMMamba, the first exploration of in-context conditioning paradigm in pan-sharpening. This framework enables bidirectional information flow between PAN and MS modalities and supports zero-shot generalization to task like image super-resolution. The proposed MI scanning strategy captures complementary information and facilitates effective cross-modal interaction.

effectively modeling long-range dependencies and cross-modal correlations for efficient global processing and superior spectral-spatial fusion. These approaches, although effective, are typically limited to a single task, such as image fusion or super-resolution, and cannot flexibly handle zero-shot generalization to other tasks. Moreover, the scanning strategies employed in these methods fail to facilitate efficient cross-modal information exchange, thereby constraining the quality of the fusion results.

Methodology

Problem Formulation

Pan-sharpening seeks to fuse the complementary information between the multispectral (MS) image $I_{lms} \in \mathbb{R}^{H/s \times W/s \times C}$ and the panchromatic (PAN) image $I_p \in \mathbb{R}^{H \times W \times 1}$ in order to produce the high-resolution multispectral (HRMS) image $I_{hrms} \in \mathbb{R}^{H \times W \times C}$. Here, H , W , and C represent the image height, width, and number of spectral channels, respectively, and s defines the spatial resolution ratio between I_{lms} and I_{hrms} , which is set to 4. The overall architecture of MMMamba is shown in Figure 1.

Network Architecture

Given the upsampled LRMS image $I_{lms} \in \mathbb{R}^{H \times W \times C}$ and PAN image $I_p \in \mathbb{R}^{H \times W \times 1}$, both inputs are first passed through separate gated convolutional encoders (Rao et al. 2022), denoted as E_φ^{ms} and E_φ^p , to extract shallow features from their respective modalities, resulting in $F_{ms} \in \mathbb{R}^{B \times C \times H \times W}$ and $F_p \in \mathbb{R}^{B \times C \times H \times W}$:

$$F_{ms} = E_\varphi^{ms}(I_{lms}), \quad (1)$$

$$F_p = E_\varphi^p(I_p). \quad (2)$$

MMMamba Blocks The shallow features F_{ms} and F_p , derived from the MS and PAN modalities, are then independently processed through a series of MMMamba blocks, which enable deep cross-modal interaction and efficient in-context fusion.

Specifically, F_{ms} and F_p first undergo layer normalization, followed by a linear projection to transform the feature dimensions. The outputs are denoted as $F_{ms}^{ln} \in \mathbb{R}^{B \times H \times W \times C}$ and $F_p^{ln} \in \mathbb{R}^{B \times H \times W \times C}$:

$$F_{ms}^{ln} = \text{Linear}(\text{LN}(F_{ms})), \quad (3)$$

$$F_p^{ln} = \text{Linear}(\text{LN}(F_p)). \quad (4)$$

Next, these normalized and projected features are processed by depth-wise convolutional layers (DWConv), and then activated using the sigmoid linear unit (SiLU) function, yielding $F_{ms}^{silu} \in \mathbb{R}^{B \times C \times H \times W}$ and $F_p^{silu} \in \mathbb{R}^{B \times C \times H \times W}$:

$$F_{ms}^{silu} = \text{SiLU}(\text{DWConv}(F_{ms}^{ln})), \quad (5)$$

$$F_p^{silu} = \text{SiLU}(\text{DWConv}(F_p^{ln})). \quad (6)$$

Multimodal Interleaved (MI) SSM The multimodal interleaved scanning operation, denoted as MI-Scan(\cdot), is applied to enable effective cross-modal information exchange and to capture complementary characteristics between the MS and PAN modalities. The details of MI-SSM are illustrated in the right part of Figure 1.

Tokenization Initially, the features F_{ms}^{silu} and F_p^{silu} from the MS and PAN modalities are tokenized into non-overlapping patches. These patches are then interleaved in four predefined directions: “left-to-right and up-to-down” (“ltr_utt”), “up-to-down and left-to-right” (“utd_ltr”), “right-to-left and down-to-up” (“rtl_dtu”), and “down-to-up and right-to-left” (“dtu_rtl”):

$$T_{ms}^{\text{ltr_utd}}, T_p^{\text{ltr_utd}} = \text{Tokenize}(F_{ms}^{silu}, F_p^{silu})_{\text{ltr_utd}}, \quad (7)$$

$$T_{ms}^{\text{utd_ltr}}, T_p^{\text{utd_ltr}} = \text{Tokenize}(F_{ms}^{silu}, F_p^{silu})_{\text{utd_ltr}}, \quad (8)$$

$$T_{ms}^{\text{rtl_dtu}}, T_p^{\text{rtl_dtu}} = \text{Tokenize}(F_{ms}^{silu}, F_p^{silu})_{\text{rtl_dtu}}, \quad (9)$$

$$T_{ms}^{\text{dtu_rtl}}, T_p^{\text{dtu_rtl}} = \text{Tokenize}(F_{ms}^{silu}, F_p^{silu})_{\text{dtu_rtl}}, \quad (10)$$

where $T_{ms}^k, T_p^k \in \mathbb{R}^{B \times C \times H_g \times W_g \times s \times s}$, and $k \in \{\text{ltr_utd}, \text{utd_ltr}, \text{rtl_dtu}, \text{dtu_rtl}\}$. Here, H_g and W_g denote the number of rows and columns in the patch grid, respectively, with $H_g = H/s$ and $W_g = W/s$, and s represents the spatial size of each patch.

For each direction, patch-wise interleaving is performed to form a fused sequence of patches from both the MS and PAN modalities:

$$S_{int}^{\text{ltr_utd}} = \text{Interleave}(T_{ms}^{\text{ltr_utd}}, T_p^{\text{ltr_utd}}), \quad (11)$$

$$S_{int}^{\text{utd_ltr}} = \text{Interleave}(T_{ms}^{\text{utd_ltr}}, T_p^{\text{utd_ltr}}), \quad (12)$$

$$S_{int}^{\text{rtl_dtu}} = \text{Interleave}(T_{ms}^{\text{rtl_dtu}}, T_p^{\text{rtl_dtu}}), \quad (13)$$

$$S_{int}^{\text{dtu_rtl}} = \text{Interleave}(T_{ms}^{\text{dtu_rtl}}, T_p^{\text{dtu_rtl}}). \quad (14)$$

The sequences from all four directions are then concatenated to generate the interleaved sequence:

$$S_{int} = \text{Concat}(S_{int}^{\text{ltr_utd}}, S_{int}^{\text{utd_ltr}}, S_{int}^{\text{rtl_dtu}}, S_{int}^{\text{dtu_rtl}}), \quad (15)$$

where $S_{int} \in \mathbb{R}^{B \times 4 \times C \times L}$, with $L = 2 \times H \times W$.

MI Scanning Strategy The MI scanning strategy first splits the S_{int} into sequence of four directions $S_{int}^{\text{ltr_utd}}, S_{int}^{\text{utd_ltr}}, S_{int}^{\text{rtl_dtu}}, S_{int}^{\text{dtu_rtl}}$. Each sequence is reshaped into $\mathbb{R}^{B \times C \times H_g \times W_g \times 2 \times s^2}$:

$$S_{int}^{\text{ltr_utd}}, S_{int}^{\text{utd_ltr}}, S_{int}^{\text{rtl_dtu}}, S_{int}^{\text{dtu_rtl}} = \text{Split}(S_{int}). \quad (16)$$

The sequences are then split into two parts to perform cross-modal MI scanning:

$$S_{int.1}^{\text{ltr_utd}}, S_{int.2}^{\text{ltr_utd}} = \text{Split}(S_{int}^{\text{ltr_utd}}), \quad (17)$$

$$S_{int.1}^{\text{utd_ltr}}, S_{int.2}^{\text{utd_ltr}} = \text{Split}(S_{int}^{\text{utd_ltr}}), \quad (18)$$

$$S_{int.1}^{\text{rtl_dtu}}, S_{int.2}^{\text{rtl_dtu}} = \text{Split}(S_{int}^{\text{rtl_dtu}}), \quad (19)$$

$$S_{int.1}^{\text{dtu_rtl}}, S_{int.2}^{\text{dtu_rtl}} = \text{Split}(S_{int}^{\text{dtu_rtl}}), \quad (20)$$

where $S_{int.1}^k, S_{int.2}^k \in \mathbb{R}^{B \times C \times H_g \times W_g \times s \times s}$, and $k \in \{\text{ltr_utd}, \text{utd_ltr}, \text{rtl_dtu}, \text{dtu_rtl}\}$.

Next, these sequences are divided into multiple local windows. For each local window, selective scanning is first applied to $S_{int.1}^{\text{ltr_utd}}$ using the “ltr_utt” scanning direction. After completing this, the scanning is transferred to the corresponding local window of the $S_{int.2}^{\text{ltr_utd}}$, where the same selective scanning is executed. Once finished, the process returns to the next local window of $S_{int.1}^{\text{ltr_utd}}$ and repeats the same procedure. This alternating scanning continues for all local windows:

$$S_{mi1}^{\text{ltr_utd}}, S_{mi2}^{\text{ltr_utd}} = \text{MI-Scan}(S_{int.1}^{\text{ltr_utd}}, S_{int.2}^{\text{ltr_utd}}), \quad (21)$$

where $S_{mi1}^{\text{ltr_utd}}, S_{mi2}^{\text{ltr_utd}} \in \mathbb{R}^{B \times C \times L'}$, and $L' = H \times W$.

The scanning strategy then proceeds with three additional directions, “utd_ltr”, “rtl_dtu”, and “dtu_rtl”. Such multi-directional scanning approach enhances cross-modal interaction and enables better exploitation of complementary information:

$$S_{mi1}^{\text{utd_ltr}}, S_{mi2}^{\text{utd_ltr}} = \text{MI-Scan}(S_{int.1}^{\text{utd_ltr}}, S_{int.2}^{\text{utd_ltr}}), \quad (22)$$

$$S_{mi1}^{\text{rtl_dtu}}, S_{mi2}^{\text{rtl_dtu}} = \text{MI-Scan}(S_{int.1}^{\text{rtl_dtu}}, S_{int.2}^{\text{rtl_dtu}}), \quad (23)$$

$$S_{mi1}^{\text{dtu_rtl}}, S_{mi2}^{\text{dtu_rtl}} = \text{MI-Scan}(S_{int.1}^{\text{dtu_rtl}}, S_{int.2}^{\text{dtu_rtl}}). \quad (24)$$

The outputs of the MI-SSM are computed by summing the results of the four directional scans:

$$S_{mi1}^{\text{out}} = S_{mi1}^{\text{ltr_utd}} + S_{mi1}^{\text{utd_ltr}} + S_{mi1}^{\text{rtl_dtu}} + S_{mi1}^{\text{dtu_rtl}}, \quad (25)$$

$$S_{mi2}^{\text{out}} = S_{mi2}^{\text{ltr_utd}} + S_{mi2}^{\text{utd_ltr}} + S_{mi2}^{\text{rtl_dtu}} + S_{mi2}^{\text{dtu_rtl}}, \quad (26)$$

where $S_{mi1}^{\text{out}}, S_{mi2}^{\text{out}} \in \mathbb{R}^{B \times C \times L}$.

The output features from the MI-SSM, S_{mi1}^{out} and S_{mi2}^{out} , are subsequently combined with the SiLU-activated projections of the normalized F_{ms}^{ln} and F_p^{ln} respectively through element-wise multiplication and summation:

$$F_{ms}^{mm} = \text{LN}(S_{mi1}^{\text{out}}) \odot \text{SiLU}(F_{ms}^{ln}), \quad (27)$$

$$F_p^{mm} = \text{LN}(S_{mi2}^{\text{out}}) \odot \text{SiLU}(F_p^{ln}), \quad (28)$$

where $F_{ms}^{mm}, F_p^{mm} \in \mathbb{R}^{B \times C \times L}$.

These features are then passed through linear projections and reshaped to produce $F_{ms}^{\text{out}}, F_p^{\text{out}} \in \mathbb{R}^{B \times C \times H \times W}$, delivering the final output of the MMMamba block:

$$F_{ms}^{\text{out}}, F_p^{\text{out}} = \text{Linear}(F_{ms}^{mm}), \text{Linear}(F_p^{mm}). \quad (29)$$

The resulting outputs are forwarded to the subsequent MMMamba block, which progressively refines the multimodal representations and enriches cross-modal feature interactions, effectively exploiting complementary cues between modalities and enabling efficient in-context fusion.

Afterward, a convolutional decoder D_φ is applied to the output of the last MMMamba block to generate the final MS feature F_{ms}^{final} :

$$F_{ms}^{\text{final}} = D_\varphi(F_{ms}^{\text{out_last}}), \quad (30)$$

where $F_{ms}^{\text{out_last}}$ denotes the output of the last MMMamba block.

Finally, the HRMS result is obtained by adding F_{ms}^{final} to the upsampled LRMS image $I_{ms} \in \mathbb{R}^{H \times W \times C}$:

$$I_{hms} = F_{ms}^{\text{final}} + I_{ms}. \quad (31)$$

Loss Function

We employ the \mathcal{L}_1 as the loss function (Zhao et al. 2016). The predicted HRMS image is denoted by I_{hms} , and the corresponding ground truth is defined by I_{gt} . The loss can be expressed as:

$$\mathcal{L} = \|I_{gt} - I_{hms}\|_1. \quad (32)$$

Experiments

Datasets and Benchmark

We conducted experiments using data from three satellites: WorldView-II (WV2), GaoFen2 (GF2), and WorldView-III (WV3). These datasets provide a variety of resolutions and scenes, including industrial areas and natural landscapes from WV2, mountains and rivers from GF2, and urban environments from WV3. As ground truth was not available, we generated all test datasets at a reduced resolution according to the Wald protocol. We compared our proposed model against several traditional methods, specifically GF-PCA (Liao et al. 2015), LRTCfPan (Wu et al. 2023), Brovey (Gillespie, Kahle, and Walker 1987), IHS (Haydn 1982), and SFIM (Liu 2000), as well as recent deep learning-based methods, including SRPPNN (Cai and Huang 2020), INN-former (Zhou et al. 2022), FAME (He et al. 2024), SFINet++ (Zhou et al. 2024), WaveletNet (Zhang et al. 2024b), Pan-Mamba (He et al. 2025), and CFLIHPs (Wang et al. 2025b). The performance was quantitatively evaluated using a combination of full-reference and no-reference metrics. The full-reference metrics were Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index (SSIM), Spectral Angle Mapper (SAM), and the relative dimensionless global error in synthesis (ERGAS). The no-reference metrics were the spatial distortion index (D_S), the spectral distortion index (D_λ), and the Quality with No Reference (QNR) index.

Implement Details

We implemented the model in PyTorch and conducted all training on a single Nvidia V100 GPU. For optimization, we used the Adam optimizer with a gradient clipping norm of 4.0 to ensure training stability. The learning rate was initialized to 5×10^{-4} and adjusted using a cosine decay schedule, which reduced it to 5×10^{-8} by the final epoch. To account for variations in data volume, we trained the model for 200 epochs on the WorldView-II dataset and 500 epochs on both the GaoFen2 and WorldView-III datasets.

Comparison with SOTA Methods

Evaluation on Reduced-Resolution Scene Table 1 summarizes the quantitative results of MMMamba in comparison with existing methods across three benchmark datasets, demonstrating its superior performance over existing SOTA techniques across multiple evaluation metrics. In particular, our approach achieves notable gains in PSNR, outperforming the CFLIHPs by 0.40 dB and 0.61 dB on the WV2 and GF2 datasets, respectively. Figure 2 presents qualitative results from the WV3 dataset. The residual plots produced by our method exhibit the lowest brightness, reflecting a high

degree of consistency with the ground truth. Additionally, our approach yields sharper edges and more accurate spectral details, further emphasizing its advantage over competing methods.

Evaluation on Full-Resolution Scene We further conducted a full-resolution evaluation under real-world conditions to assess the generalization capability of our method. This experiment was carried out on the full-resolution GF2 (FGF2) datasets, where no-reference quality metrics were employed due to the absence of ground truth references. The FGF2 dataset was utilized in its original form without any downsampling, providing a testing environment that closely replicates real-world image degradation. As summarized in Table 2, our method consistently outperforms other approaches across all three metrics, demonstrating its strong generalization performance in real-world scenarios.

Zero-Shot Task Generalization

To evaluate MMMamba’s zero-shot generalization capabilities, we tested it on MS image super-resolution. Although trained exclusively on pan-sharpening, MMMamba can perform this task without any retraining or fine-tuning. By leveraging its in-context fusion mechanism, it adapts by simply omitting one input modality—performing super-resolution when given only the MS image.

We also compared our approach with other deep learning models. Since these models cannot inherently work with a single input, we had to adapt them. For the super-resolution task, we fed the MS image into both the PAN and MS encoders during inference.

The qualitative results for the zero-shot super-resolution task are presented in Figure 3. As illustrated in the figure, our proposed method generates visually compelling results, successfully reconstructing finer details and sharper edges. In contrast, the outcomes from the adapted SFINet++ and Pan-Mamba methods appear comparatively blurry, with a noticeable loss of textural information.

The quantitative metrics, summarized in Table 4, provide further evidence of our model’s effectiveness. Our approach consistently outperforms the compared methods across all evaluation criteria. Notably, our model achieves a PSNR of 36.49 dB and an SSIM of 0.9114. These scores not only surpass those of the other deep learning-based methods, SFINet++ and Pan-Mamba, but also exceed the performance of the traditional Bicubic interpolation. Furthermore, our method yields the lowest error values with a SAM of 0.0299 and an ERGAS of 1.5515, indicating superior spectral and radiometric fidelity in the reconstructed images.

Ablation Experiments

We conducted ablation studies on the WV2 dataset to validate core components, as presented in Table 3.

Ablation on Core Paradigm Our analysis confirms the effectiveness of the core design (M1). We replaced the Mamba operator with a self-attention (SA) based block to compare their effectiveness. For a fair comparison under a similar computational budget, we built a computationally matched SA by incorporating sequence downsampling

| Methods | WorldView-II | | | | GaoFen2 | | | | Worldview-III | | | |
|--|-----------------|-----------------|------------------|--------------------|-----------------|-----------------|------------------|--------------------|-----------------|-----------------|------------------|--------------------|
| | PSNR \uparrow | SSIM \uparrow | SAM \downarrow | ERGAS \downarrow | PSNR \uparrow | SSIM \uparrow | SAM \downarrow | ERGAS \downarrow | PSNR \uparrow | SSIM \uparrow | SAM \downarrow | ERGAS \downarrow |
| IHS (Haydn 1982) | 35.2962 | 0.9027 | 0.0461 | 2.0278 | 38.1754 | 0.9100 | 0.0243 | 1.5336 | 22.5579 | 0.5354 | 0.1266 | 8.3616 |
| Brovey (Gillespie, Kahle, and Walker 1987) | 35.8646 | 0.9216 | 0.0403 | 1.8238 | 37.7974 | 0.9026 | 0.0218 | 1.3720 | 22.5060 | 0.5466 | 0.1159 | 8.2331 |
| SFIM (Liu 2000) | 34.1297 | 0.8975 | 0.0439 | 2.3449 | 36.9060 | 0.8882 | 0.0318 | 1.7398 | 21.8212 | 0.5457 | 0.1208 | 8.9730 |
| GFPCA (Liao et al. 2015) | 34.5581 | 0.9038 | 0.0488 | 2.1411 | 37.9443 | 0.9204 | 0.0314 | 1.5604 | 22.3344 | 0.4826 | 0.1294 | 8.3964 |
| LRTCFPan (Wu et al. 2023) | 34.7756 | 0.9112 | 0.0426 | 2.0075 | 36.9253 | 0.8946 | 0.0332 | 1.7060 | 22.1574 | 0.5735 | 0.1380 | 8.6796 |
| SRPPNN (Cai and Huang 2020) | 41.4538 | 0.9679 | 0.0233 | 0.9899 | 47.1998 | 0.9877 | 0.0106 | 0.5586 | 30.4346 | 0.9202 | 0.0770 | 3.1553 |
| INNformer (Zhou et al. 2022) | 41.6903 | 0.9704 | 0.0227 | 0.9514 | 47.3528 | 0.9893 | 0.0102 | 0.5479 | 30.5365 | 0.9225 | 0.0747 | 3.0997 |
| FAME (He et al. 2024) | 42.0262 | 0.9723 | 0.0215 | 0.9172 | 47.6721 | 0.9898 | 0.0098 | 0.5242 | 30.9903 | 0.9287 | 0.0697 | 2.9531 |
| WaveletNet (Zhang et al. 2024b) | 41.9131 | 0.9715 | 0.0220 | 0.9274 | 47.5907 | 0.9894 | 0.0099 | 0.5310 | 30.9139 | 0.9279 | 0.0710 | 2.9770 |
| SFINet++ (Zhou et al. 2024) | 41.8115 | 0.9731 | 0.0220 | 0.9489 | 47.5344 | 0.9906 | 0.0100 | 0.5356 | 30.7665 | 0.9261 | 0.0732 | 3.0217 |
| Pan-Mamba (He et al. 2025) | 42.2354 | 0.9729 | 0.0212 | 0.8975 | 47.6453 | 0.9894 | 0.0103 | 0.5286 | 31.1740 | 0.9302 | 0.0698 | 2.8910 |
| CFLIHPs (Wang et al. 2025b) | 41.9077 | 0.9712 | 0.0220 | 0.9284 | 47.3824 | 0.9892 | 0.0102 | 0.5409 | 30.8341 | 0.9269 | 0.0737 | 2.9980 |
| Ours | 42.3120 | 0.9733 | 0.0209 | 0.8888 | 47.9932 | 0.9902 | 0.0098 | 0.5126 | 31.2311 | 0.9305 | 0.0687 | 2.8950 |

Table 1: Quantitative comparison on three datasets. The best results are highlighted in **bold**. \uparrow signifies better performance with larger values, while \downarrow indicates improved performance with smaller values.

| Metrics | IHS | Brovey | SFIM | GFPCA | LRTCFPan | SRPPNN | INNformer | FAME | WaveletNet | SFINet++ | Pan-Mamba | CFLIHPs | Ours |
|-----------------------|--------|--------|--------|--------|----------|--------|-----------|--------|------------|----------|-----------|---------|---------------|
| $D_\lambda\downarrow$ | 0.0770 | 0.1378 | 0.0822 | 0.0914 | 0.1170 | 0.0767 | 0.0782 | 0.0674 | 0.0700 | 0.0673 | 0.0652 | 0.0678 | 0.0656 |
| $D_S\downarrow$ | 0.2985 | 0.2605 | 0.1121 | 0.1635 | 0.2024 | 0.1162 | 0.1253 | 0.1121 | 0.1063 | 0.1108 | 0.1129 | 0.1170 | 0.1113 |
| QNR \uparrow | 0.6485 | 0.6390 | 0.8214 | 0.7615 | 0.7063 | 0.8173 | 0.8073 | 0.8291 | 0.8327 | 0.8471 | 0.8306 | 0.8287 | 0.8312 |

Table 2: Evaluation of our method on real-world full-resolution scenes from the GF2 dataset.

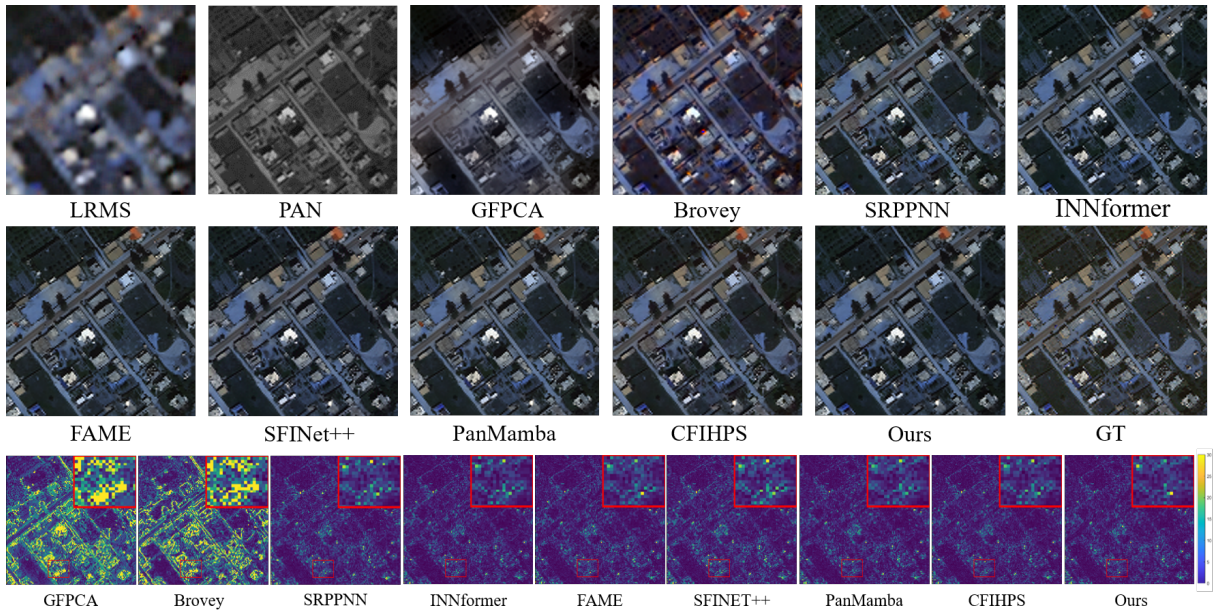


Figure 2: Visual comparison of all methods on WV3. The last row visualizes the MSE residues between the pan-sharpening results and the ground truth.

(via pixel shuffling) and a linear projection before the attention mechanism. Replacing the **Mamba operator** with this self-attention module degrades performance, underscoring Mamba’s linear ($O(N)$) efficiency for this task (He et al. 2025).

The proposed **in-context fusion** is also crucial: substituting it with naive channel-wise concatenation (M2) prevents

effective cross-modal interaction, causing the PSNR to drop to 41.28 dB.

The necessity of our **interleaved design** is demonstrated by replacing it with sequential token concatenation (M3), which caused a huge performance decrease. The interleaved approach places tokens from corresponding spatial positions of each modality adjacent to one another, enabling direct and

| ID | Methods / Variant | Performance Metrics | | | | Efficiency Metrics | |
|---|---|---------------------|-----------------|------------------|--------------------|-------------------------|------------------------|
| | | PSNR \uparrow | SSIM \uparrow | SAM \downarrow | ERGAS \downarrow | Params (M) \downarrow | FLOPs (G) \downarrow |
| <i>Ablation on Core Paradigm & Backbone</i> | | | | | | | |
| M0 | MMMamba (Full Model / Baseline) | 42.3120 | 0.9733 | 0.0209 | 0.8888 | 0.2453 | 5.0616 |
| M1 | w/o Mamba (use Transformer) | 41.3995 | 0.9675 | 0.0235 | 0.9862 | 0.3684 | 3.9206 |
| M2 | w/o In-context Fusion (use Channel Concat.) | 41.2898 | 0.9672 | 0.0237 | 0.9976 | 0.2704 | 5.0275 |
| M3 | w/o Interleaving (use Sequential Concat.) | 36.4702 | 0.9107 | 0.0302 | 1.5550 | 0.2432 | 5.0275 |
| <i>Ablation on Scanning Strategy</i> | | | | | | | |
| M4 | w/o Multi-direction (use 1-way Scan) | 42.0965 | 0.9723 | 0.0214 | 0.0989 | 0.2432 | 5.0275 |
| M5 | w/o Local Scan (use Global Scan) | 42.1998 | 0.9729 | 0.0211 | 0.8972 | 0.2432 | 5.0277 |

Table 3: Ablation study of the MMMamba model on the WV2 dataset. ‘ \uparrow ’ indicates that higher is better, while ‘ \downarrow ’ indicates that lower is better. **Bold** marks the best result in each column. All models are trained and evaluated under identical settings.




Figure 3: The visual comparison of the zero-shot image super-resolution results on the WV2 dataset.

efficient information exchange. Conversely, sequential concatenation separates these corresponding tokens, causing severe information decay as the signal propagates over a long distance within the Mamba state.

Ablation on Scanning Strategy

Benefits of Multi-Directional Scanning To validate multi-directional scanning, we simplified it to a single, unidirectional scan (M4). The results show a performance drop with negligible change in computational cost. This confirms that aggregating contextual information from multiple directions allows the model to build a more comprehensive and robust feature representation, essential for 2D spatial data.

Effectiveness of the Local Window Scan We compared our local window scan against a standard global scan (M5), where the latter showed a slight decline in performance (PSNR drops to 42.19 dB). This experiment demonstrates that our modification successfully introduces a crucial inductive bias of locality into the Mamba operator.

Computational Efficiency

We have evaluated the FLOPs and the number of parameters of our proposed method, along with other comparative meth-

| Methods | PSNR \uparrow | SSIM \uparrow | SAM \downarrow | ERGAS \downarrow |
|-----------|-----------------|-----------------|------------------|--------------------|
| Bicubic | 34.0869 | 0.8726 | 0.0397 | 2.1202 |
| SFNet++ | 33.3047 | 0.8679 | 0.0439 | 2.3105 |
| Pan-Mamba | 30.5913 | 0.7656 | 0.0524 | 3.1224 |
| Ours | 36.4892 | 0.9114 | 0.0299 | 1.5515 |

Table 4: Comparison results on the WV2 dataset for zero-shot image super-resolution evaluation.

| Methods | FLOPs (G) | Params (M) |
|------------|-----------|------------|
| SRPPNN | 21.1059 | 1.7114 |
| INNformer | 1.3079 | 0.0706 |
| FAME | 9.4093 | 0.5766 |
| WaveletNet | 7.770 | 1.3230 |
| SFNet++ | 1.3112 | 0.0848 |
| Pan-Mamba | 3.0088 | 0.1827 |
| CFLIHPs | 6.4500 | 0.1314 |
| Ours | 5.0616 | 0.2453 |

Table 5: The comparison of computational efficiency.

ods, on PAN images with a resolution of 128×128 and MS images with a resolution 32×32 on a single Nvidia V100 GPU. The results of this evaluation are presented in Table 5. Our proposed method demonstrates 5.0616 G FLOPs and 0.2453 M parameters.

Conclusion

In conclusion, we present MMMamba, a novel cross-modal in-context fusion framework that pioneers the exploration of the in-context conditioning paradigm in the pan-sharpening domain. Built on the Mamba architecture, MMMamba achieves linear computational complexity and enables efficient bidirectional information flow between PAN and MS modalities. To further strengthen multimodal interactions, we design a multimodal interleaved scanning mechanism that effectively captures complementary characteristics across modalities. Our framework also demonstrates strong generalization capabilities, including zero-shot adaptation to image super-resolution task. Extensive experiments across multiple benchmark datasets consistently validate the superiority of MMMamba over existing SOTA methods.

Acknowledgments

This work was supported by the National Natural Science Foundation of China under Grant 82272071, 62271430, 82172073, and 52105126.

References

- Aiazzi, B.; Alparone, L.; Baronti, S.; Garzelli, A.; and Selva, M. 2003. An MTF-based spectral distortion minimizing model for pan-sharpening of very high resolution multispectral images of urban areas. In *2003 2nd GRSS/ISPRS Joint Workshop on Remote Sensing and Data Fusion over Urban Areas*, 90–94. IEEE.
- Ballester, C.; Caselles, V.; Igual, L.; Verdera, J.; and Rougé, B. 2006. A variational model for P+ XS image fusion. *International Journal of Computer Vision*, 69(1): 43–58.
- Cai, J.; and Huang, B. 2020. Super-resolution-guided progressive pansharpening based on a deep convolutional neural network. *IEEE Transactions on Geoscience and Remote Sensing*, 59(6): 5206–5220.
- Cao, Z.; Wu, X.; Deng, L.-J.; and Zhong, Y. 2024. A novel state space model with local enhancement and state sharing for image fusion. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 1235–1244.
- Esser, P.; Kulal, S.; Blattmann, A.; Entezari, R.; Müller, J.; Saini, H.; Levi, Y.; Lorenz, D.; Sauer, A.; Boesel, F.; et al. 2024. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first international conference on machine learning*.
- Fasbender, D.; Radoux, J.; and Bogaert, P. 2008. Bayesian data fusion for adaptable image pansharpening. *IEEE Transactions on Geoscience and Remote Sensing*, 46(6): 1847–1857.
- Gillespie, A. R.; Kahle, A. B.; and Walker, R. E. 1987. Color enhancement of highly correlated images. II. Channel ratio and “chromaticity” transformation techniques. *Remote Sensing of Environment*, 22(3): 343–365.
- Gu, A.; and Dao, T. 2023. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*.
- Gu, A.; Goel, K.; and Ré, C. 2021. Efficiently modeling long sequences with structured state spaces. *arXiv preprint arXiv:2111.00396*.
- Haydn, R. 1982. Application of the IHS color transform to the processing of multisensor data and image enhancement. In *Proc. of the International Symposium on Remote Sensing of Arid and Semi-Arid Lands, Cairo, Egypt, 1982*.
- He, X.; Cao, K.; Zhang, J.; Yan, K.; Wang, Y.; Li, R.; Xie, C.; Hong, D.; and Zhou, M. 2025. Pan-mamba: Effective pan-sharpening with state space model. *Information Fusion*, 115: 102779.
- He, X.; Yan, K.; Li, R.; Xie, C.; Zhang, J.; and Zhou, M. 2024. Frequency-adaptive pan-sharpening with mixture of experts. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 2121–2129.
- He, X.; Yan, K.; Zhang, J.; Li, R.; Xie, C.; Zhou, M.; and Hong, D. 2023. Multiscale dual-domain guidance network for pan-sharpening. *IEEE Transactions on Geoscience and Remote Sensing*, 61: 1–13.
- Hou, J.; Cao, Q.; Ran, R.; Liu, C.; Li, J.; and Deng, L.-j. 2023. Bidomain modeling paradigm for pansharpening. In *Proceedings of the 31st ACM international conference on multimedia*, 347–357.
- Hou, J.; Cao, Z.; Zheng, N.; Li, X.; Chen, X.; Liu, X.; Cong, X.; Hong, D.; and Zhou, M. 2024. Linearly-evolved transformer for pan-sharpening. In *Proceedings of the 32nd ACM international conference on multimedia*, 1486–1494.
- Huang, J.; Meng, G.; Wang, Y.; Lin, Y.; Huang, Y.; and Ding, X. 2023. Dp-innet: Dual-path implicit neural network for spatial and spectral features fusion in pan-sharpening. In *Chinese Conference on Pattern Recognition and Computer Vision (PRCV)*, 268–279. Springer.
- Jenerowicz, A.; and Woroszkiewicz, M. 2016. The pansharpening of satellite and UAV imagery for agricultural applications. In *Remote Sensing for Agriculture, Ecosystems, and Hydrology XVIII*, volume 9998, 565–575. SPIE.
- Kwarteng, P.; and Chavez, A. 1989. Extracting spectral contrast in Landsat Thematic Mapper image data using selective principal component analysis. *Photogramm. Eng. Remote Sens*, 55(1): 339–348.
- Labs, B. F.; Batifol, S.; Blattmann, A.; Boesel, F.; Consul, S.; Diagne, C.; Dockhorn, T.; English, J.; English, Z.; Esser, P.; et al. 2025. FLUX. 1 Kontext: Flow Matching for In-Context Image Generation and Editing in Latent Space. *arXiv preprint arXiv:2506.15742*.
- Li, X.; He, X.; Hu, T.; Zhang, J.; Zhou, M.; Xie, C.; Wang, Y.; and Huang, B. 2025. Freq-RWKV: Granularity-Aware Spatial-Frequency Synergy via Dual-Domain Recurrent Scanning for Pan-sharpening. In *Proceedings of the 33rd ACM International Conference on Multimedia*, 1890–1899.
- Liao, W.; Huang, X.; Van Coillie, F.; Thoonen, G.; Pižurica, A.; Scheunders, P.; and Phillips, W. 2015. Two-stage fusion of thermal hyperspectral and visible RGB image by PCA and guided filter. In *2015 7th Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing (WHISPERS)*, 1–4. Ieee.
- Lin, Y.; Fu, Z.; Meng, G.; Wang, Y.; Dong, Y.; Fan, L.; Yu, H.; and Ding, X. 2023. Domain-irrelevant feature learning for generalizable pan-sharpening. In *Proceedings of the 31st ACM International Conference on Multimedia*, 3287–3296.
- Liu, J. 2000. Smoothing filter-based intensity modulation: A spectral preserve image fusion technique for improving spatial details. *International Journal of remote sensing*, 21(18): 3461–3472.
- Liu, Y.; Tian, Y.; Zhao, Y.; Yu, H.; Xie, L.; Wang, Y.; Ye, Q.; Jiao, J.; and Liu, Y. 2024. Vmamba: Visual state space model. *Advances in neural information processing systems*, 37: 103031–103063.
- Mallat, S. G. 2002. A theory for multiresolution signal decomposition: the wavelet representation. *IEEE transactions*

- on pattern analysis and machine intelligence, 11(7): 674–693.
- Masi, G.; Cozzolino, D.; Verdoliva, L.; and Scarpa, G. 2016. Pansharpening by convolutional neural networks. *Remote Sensing*, 8(7): 594.
- Meng, G.; Huang, J.; Tu, J.; Wang, Y.; Lin, Y.; Tu, X.; Huang, Y.; and Ding, X. 2025. Accelerated Diffusion via High-Low Frequency Decomposition for Pan-Sharpener. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 6117–6125.
- Meng, G.; Huang, J.; Wang, Y.; Fu, Z.; Ding, X.; and Huang, Y. 2024. Progressive high-frequency reconstruction for pansharpening with implicit neural representation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 4189–4197.
- Nunez, J.; Otazu, X.; Fors, O.; Prades, A.; Pala, V.; and Arbiol, R. 2002. Multiresolution-based image fusion with additive wavelet decomposition. *IEEE Transactions on Geoscience and Remote Sensing*, 37(3): 1204–1211.
- Rao, Y.; Zhao, W.; Tang, Y.; Zhou, J.; Lim, S. N.; and Lu, J. 2022. Hornet: Efficient high-order spatial interactions with recursive gated convolutions. *Advances in Neural Information Processing Systems*, 35: 10353–10366.
- Schowengerdt, R. A. 1980. Reconstruction of multispatial, multispectral image data using spatial frequency content. *Photogrammetric Engineering and Remote Sensing*, 46(10): 1325–1334.
- Sunuprpto, H.; Danoedoro, P.; and Ritohardoyo, S. 2016. Evaluation of pan-sharpening method: applied to artisanal gold mining monitoring in Gunung Pani Forest area. *Procedia Environmental Sciences*, 33: 230–238.
- Tan, Z.; Liu, S.; Yang, X.; Xue, Q.; and Wang, X. 2024. Ominicontrol: Minimal and universal control for diffusion transformer. *arXiv preprint arXiv:2411.15098*.
- Wang, Y.; He, X.; Dong, Y.; Lin, Y.; Huang, Y.; and Ding, X. 2024. Cross-modality interaction network for pansharpening. *IEEE Transactions on Geoscience and Remote Sensing*, 62: 1–16.
- Wang, Y.; Lin, Y.; He, X.; Zheng, H.; Yan, K.; Fan, L.; Huang, Y.; and Ding, X. 2025a. Learning Diffusion High-Quality Priors for Pan-Sharpener: A Two-Stage Approach With Time-Aware Adapter Fine-Tuning. *IEEE Transactions on Geoscience and Remote Sensing*.
- Wang, Y.; Lin, Y.; Meng, G.; Fu, Z.; Dong, Y.; Fan, L.; Yu, H.; Ding, X.; and Huang, Y. 2023. Learning high-frequency feature enhancement and alignment for pansharpening. In *Proceedings of the 31st ACM International Conference on Multimedia*, 358–367.
- Wang, Y.; Zheng, H.; Li, F.; Lin, Y.; Fan, L.; He, X.; Huang, Y.; and Ding, X. 2025b. Towards Generalizable Pan-sharpening: Conditional Flow-based Learning Guided by Implicit High-frequency Priors. *IEEE Transactions on Geoscience and Remote Sensing*.
- Wu, X.; Cao, Z.-H.; Huang, T.-Z.; Deng, L.-J.; Chanussot, J.; and Vivone, G. 2025. Fully-connected transformer for multi-source image fusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(3): 2071–2088.
- Wu, Z.-C.; Huang, T.-Z.; Deng, L.-J.; Huang, J.; Chanussot, J.; and Vivone, G. 2023. LRTCFFPan: Low-rank tensor completion based framework for pansharpening. *IEEE Transactions on Image Processing*, 32: 1640–1655.
- Yang, J.; Fu, X.; Hu, Y.; Huang, Y.; Ding, X.; and Paisley, J. 2017. PanNet: A deep network architecture for pansharpening. In *Proceedings of the IEEE international conference on computer vision*, 5449–5457.
- Zhang, J.; Cao, K.; Yan, K.; Lin, Y.; He, X.; Wang, Y.; Li, R.; Xie, C.; Zhang, J.; and Zhou, M. 2024a. Frequency decoupled domain-irrelevant feature learning for pan-sharpening. *IEEE Transactions on Circuits and Systems for Video Technology*.
- Zhang, J.; He, X.; Yan, K.; Cao, K.; Li, R.; Xie, C.; Zhou, M.; and Hong, D. 2024b. Pan-sharpening with wavelet-enhanced high-frequency information. *IEEE Transactions on Geoscience and Remote Sensing*, 62: 1–14.
- Zhao, H.; Gallo, O.; Frosio, I.; and Kautz, J. 2016. Loss functions for image restoration with neural networks. *IEEE Transactions on computational imaging*, 3(1): 47–57.
- Zhou, H.; Liu, Q.; and Wang, Y. 2022. PanFormer: A transformer based model for pan-sharpening. In *2022 IEEE international conference on multimedia and expo (ICME)*, 1–6. IEEE.
- Zhou, M.; Huang, J.; Fang, Y.; Fu, X.; and Liu, A. 2022. Pan-sharpening with customized transformer and invertible neural network. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, 3553–3561.
- Zhou, M.; Huang, J.; Yan, K.; Hong, D.; Jia, X.; Chanussot, J.; and Li, C. 2024. A general spatial-frequency learning framework for multimodal image fusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Zhu, L.; Liao, B.; Zhang, Q.; Wang, X.; Liu, W.; and Wang, X. 2024. Vision mamba: Efficient visual representation learning with bidirectional state space model. *arXiv preprint arXiv:2401.09417*.