

PIF-Net: Ill-Posed Prior Guided Multispectral and Hyperspectral Image Fusion via Invertible Mamba and Fusion-Aware LoRA

Baisong Li^{1,2}, Xingwang Wang^{1,2*}, Haixiao Xu^{1,2}

¹College of Computer Science and Technology, Jilin University

²Key Laboratory of Symbolic Computation and Knowledge Engineering of Ministry of Education, Jilin University
lbs23@mails.jlu.edu.cn, {xww, haixiao}@jlu.edu.cn

Abstract

The goal of multispectral and hyperspectral image fusion (MHIF) is to generate high-quality images that simultaneously possess rich spectral information and fine spatial details. However, due to the inherent trade-off between spectral and spatial information and the limited availability of observations, this task is fundamentally ill-posed. Previous studies have not effectively addressed the ill-posed nature caused by data misalignment. To tackle this challenge, we propose a fusion framework named PIF-Net, which explicitly incorporates ill-posed priors to effectively fuse multispectral images and hyperspectral images. To balance global spectral modeling with computational efficiency, we design a method based on an invertible Mamba architecture that maintains information consistency during feature transformation and fusion, ensuring stable gradient flow and process reversibility. Furthermore, we introduce a novel fusion module called the *Fusion-Aware Low-Rank Adaptation* module, which dynamically calibrates spectral and spatial features while keeping the model lightweight. Extensive experiments on multiple benchmark datasets demonstrate that PIF-Net achieves significantly better image restoration performance than current state-of-the-art methods while maintaining model efficiency.

Introduction

Multispectral and Hyperspectral Image Fusion (MHIF) aims to effectively combine the rich spatial details from multispectral images (MSI) with the continuous and abundant spectral information from hyperspectral images (HSI), generating high-quality images that possess both high spatial resolution and high spectral resolution. By integrating data from these two modalities, MHIF seeks to overcome the limitations of single sensors in either spatial or spectral resolution, enhancing spectral fidelity and spatial detail representation in the fused images. This, in turn, better supports vision-based downstream applications such as image classification (Liu et al. 2025; Hang et al. 2021; Cao et al. 2024), object detection (Qin et al. 2024; Liu et al. 2024a; He et al. 2023), and other practical tasks (Sethy et al. 2022; Bi et al. 2024; Bhargava et al. 2024).

*Corresponding Author

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

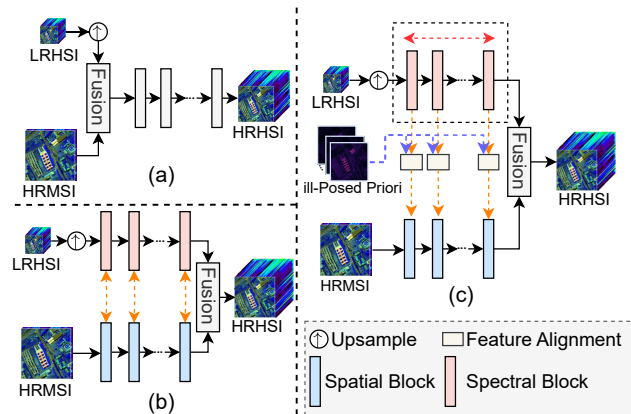


Figure 1: Illustration of three typical hyperspectral image fusion frameworks: (a) Single-branch framework: relies on a single pathway and is prone to alignment errors; (b) Dual-branch framework: models spectral and spatial features separately but lacks effective spatial transformation and collaborative mechanisms; (c) The proposed PIF-Net model: integrates invertible state-space modeling with fusion-aware LoRA, enabling bidirectional information flow in the frequency domain and robust cross-modal alignment.

Most fusion methods based on single-branch (as Figure 1(a)) (Hu et al. 2022b; Ma et al. 2023; Deng et al. 2023) or dual-branch (as Figure 1(b)) (Peng et al. 2023; Hu et al. 2022c) architectures, despite showing improvements over traditional model-driven approaches (He et al. 2016; Vivone et al. 2014; Aiazzi, Baronti, and Selva 2007), still suffer from two fundamental limitations: (1) *Significant differences in spectral characteristics, spatial resolution, and data distribution between MSIs and HSIs result in severe information misalignment and a substantial modality gap between the two sources. This constitutes a highly ill-posed problem that renders unified and effective feature modeling particularly challenging, thereby constraining further advances in fusion performance.* (2) *Existing fusion methods commonly fail to ensure reversible information transfer, leading to inevitable information loss during the fusion process. Such irreversible loss directly impairs spectral fidelity and the high-quality recovery of spatial details.*

To address the ill-posedness inherent in image restoration or image fusion, researchers have conducted extensive studies and explorations. For example, the *Structure-Aware Deep Network* (Fan et al. 2025; Lu, You, and Barnes 2018) enhances the model’s ability to capture geometric structures in images by introducing a structure-aware module, thereby significantly improving the visual quality of super-resolved images. The adoption of generative adversarial training (Creswell et al. 2018; Ma et al. 2020; Wang et al. 2018) further facilitates the generation of realistic high-resolution images with complex structures. Denoising diffusion models (Sun, Chen, and Messinger 2014; Saharia et al. 2022) have demonstrated excellent performance in detail restoration and noise suppression, producing more natural and realistic images while effectively mitigating the over-smoothing caused by the ill-posed nature of the problem. *However, these methods typically rely on general structural priors and mainly focus on low-level fixed patterns, making it difficult to capture the semantic shifts introduced during fusion.*

Furthermore, the introduction of invertible neural networks (INNs) (Dinh, Krueger, and Bengio 2014; Xiao et al. 2020) offers a novel perspective for image restoration tasks. Unlike traditional unidirectional mapping, invertible networks learn bidirectionally consistent mappings, establishing a reversible information flow between image degradation and restoration (Kingma and Dhariwal 2018; Lugmayr et al. 2020), thereby theoretically preventing information loss. This mechanism not only helps alleviate the ill-posedness in the restoration process but also significantly enhances the model’s capability to transfer information across different modalities, making it particularly suitable for the complex cross-modal modeling demands in multi-source image fusion. *Nevertheless, effectively leveraging the invertible structure to address the issue of information misalignment caused by modality inconsistencies remains a significant challenge.*

To address the above issues, this paper proposes an advanced fusion framework guided by ill-posed priors, as shown in Figure 1(c), called PIF-Net. This framework ingeniously integrates the invertible visual state-space model Mamba with a lightweight *Fusion-Aware Multi-Head Low-Rank Adaptation* (FAM-LoRA) module, achieving high-fidelity and efficient multimodal image fusion. PIF-Net consists of two complementary branches: the frequency domain branch precisely separates high- and low-frequency components using a two-dimensional HWT, combined with the invertible Mamba module to enhance bidirectional information flow, ensuring lossless extraction of global spatial features and structural fidelity; the spatial domain branch leverages low-frequency global guidance and the lightweight LoRA module to finely extract rich spatial texture information, incorporating *Large-Kernel Attention* (LKA) (Wang et al. 2023; Liu et al. 2022; Peng et al. 2017) and *Squeeze-and-Excitation* (SE) (Hu, Shen, and Sun 2018) mechanisms to significantly improve cross-modal semantic modeling. Furthermore, PIF-Net introduces guided feature consistency loss and volumetric-aware loss to strengthen the consistency between the dual-modal features, ensuring high preci-

sion in modality alignment and feature fusion with semantic coherence. Extensive experimental results demonstrate that PIF-Net outperforms current state-of-the-art methods across multiple hyperspectral datasets, achieving outstanding fusion quality and computational efficiency. The main contributions of this work are summarized as follows:

- We propose PIF-Net, a novel fusion network combining invertible Mamba and fusion-aware LoRA, effectively tackling modality mismatch and ill-posedness in hyperspectral–multispectral image fusion for high-fidelity, efficient results.
- We design a dual-branch architecture: the frequency branch combines Haar wavelet decomposition and invertible Mamba to enable bidirectional information flow and structural preservation in the frequency domain; the spatial branch uses low-frequency guidance and a lightweight LoRA module to extract semantically rich and stable spatial features.
- Extensive experiments on several public hyperspectral datasets demonstrate that PIF-Net significantly outperforms existing state-of-the-art (SOTA) methods in terms of fusion quality and computational efficiency.

Related Work

Review of Existing Methods in MHIF

In the field of multispectral and hyperspectral image fusion (MHIF), traditional methods such as pansharpening (Wu et al. 2021; Aiazzi, Baronti, and Selva 2007) and multiscale decomposition (Yang et al. 2017) offer good interpretability but heavily rely on prior knowledge, making them prone to failure in complex scenarios. Bayesian approaches (Yang et al. 2018; Aiazzi et al. 2006) can effectively model the fusion process but require careful parameter tuning. In recent years, deep learning methods have significantly advanced MHIF. Early convolutional neural network (CNN)-based models, such as MHFnet (Xie et al. 2022) and HSRnet (Hu et al. 2022c), perform well under general conditions but struggle to capture long-range dependencies, leading to degraded performance in complex cases. To address this, Transformer-based Fusformer (Hu et al. 2022b) was introduced, which greatly improves fusion quality but at a higher computational cost. Other approaches like PSRT (Deng et al. 2023), U2Net (Peng et al. 2023), and 3DT-Net (Ma et al. 2023) bring innovations in architecture design but still face limitations in fusion capability, high training costs, or increased model complexity. SMGU-Net (Yan et al. 2025) attempts to combine traditional methods with deep learning to enhance cross-modal interactions, yet there remains substantial room for improvement in modeling complex nonlinear cross-modal relationships.

Invertible Neural Networks

Invertible Neural Networks (INNs) are a key component of normalizing flow models. Since NICE (Dinh, Krueger, and Bengio 2014) introduced additive coupling layers, methods like RealNVP (Dinh, Sohl-Dickstein, and Bengio 2017) and Glow (Kingma and Dhariwal 2018) have improved INNs,

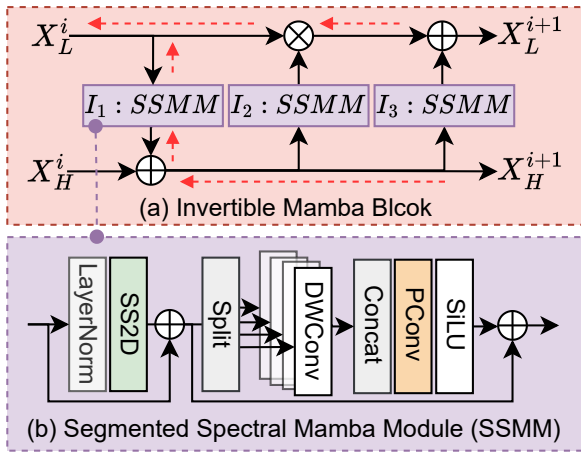


Figure 3: Illustration of the *Invertible Mamba Block*. The block takes low-frequency features X_L^i and high-frequency features X_H^i as input, and employs an affine coupling mechanism built on lightweight *Segmented Spectral Mamba Modules (SSMM)* to enable efficient bidirectional interaction and fusion of spectral information. *Split* and *Concat* denote channel-wise feature division and aggregation, respectively. SS2D refers to the *2D State Space Module* in VMamba (Liu et al. 2024c).

ture differences in an adaptive manner, enhancing the discriminative capacity of subsequent feature fusion. The computation is formulated as:

$$\begin{cases} C_R = \max_{c \in D} \bar{X}_s - \min_{c \in D} \bar{X}_s, \\ \mathbf{X}_R = \text{ReLU}(\text{Conv}(\mathbf{X}_s - \mathbf{Y}_s)) \cdot C_R, \\ \mathbf{X}_R = \text{ReLU}(\text{Conv}(\mathbf{X}_R)), \end{cases} \quad (2)$$

where \bar{X}_s denotes the mean feature across the channel dimension, and D is the set of all channels. The modulation factor C_R serves to amplify residual cues, guiding the network to focus on global structural variations, thereby enhancing the discriminability and robustness of the extracted prior features \mathbf{X}_R .

Invertible Mamba Block

CNNs (Hu et al. 2022c; Peng et al. 2023) struggle to capture long-range dependencies. Transformers (Hu et al. 2022b) offer strong global modeling but suffer from quadratic complexity as feature size grows, reducing efficiency. To balance performance and cost, we adopt Mamba (Liu et al. 2024c; Gu and Dao 2023), a state-space model with linear complexity. Mamba captures global context efficiently and expressively. However, most Mamba-based models (Guo et al. 2024, 2025; Li, Wang, and Xu 2025) still rely on irreversible designs, which cause feature ambiguity and information loss in image restoration. To solve this, we propose the *Invertible Mamba Block*, which enables bidirectional interaction between low- and high-frequency spectral features (see Figure 3(a)).

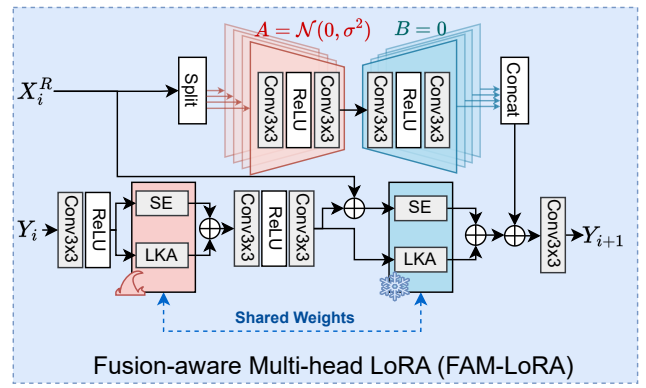


Figure 4: Illustration of the FAM-LoRA module that fuses the main spatial feature Y_i and auxiliary guidance X_i^R through channel transformation, LKA, SE attention, and multi-head LoRA, achieving efficient and accurate semantic fusion. The first part of LoRA’s parameters is initialized with a standard Gaussian distribution, while the second part is initialized to zero. The operations *Split* and *Concat* refer to feature division and aggregation along the channel dimension.

Invertible Mamba Block uses coupling transformations from invertible neural networks. Each transformation unit is built with our *Segmented Spectral Mamba Module (SSMM)*, as shown in Figure 3(b). This design supports global spectral modeling and multi-scale feature extraction, improving representation quality and diversity. We split the input into high-frequency X_H^i and low-frequency X_L^i components and update them with the following invertible mapping:

$$\begin{aligned} \mathbf{X}_H^{i+1} &= \mathbf{X}_H^i + I_1(\mathbf{X}_L^i), \\ \mathbf{X}_L^{i+1} &= \mathbf{X}_L^i \odot \exp[I_2(\mathbf{X}_H^i)] + I_3(\mathbf{X}_H^i). \end{aligned} \quad (3)$$

Here, \odot denotes element-wise multiplication. I_1, I_2 , and I_3 are SSMM-based subnetworks designed to model nonlinear interactions between different spectral components. Thanks to the invertible architecture, the module enables efficient information fusion during the forward pass. Notably, to ensure the flexibility of invertible blocks in fusion tasks, we employ a gradient-guided invertibility strategy that heuristically facilitates backward interaction of spectral features. Instead of using traditional rescaling schemes (Xiao et al. 2020), we design a tailored invertible loss function (see Eq. (4)).

Fusion-aware Multi-head LoRA

Inspired by the success of LoRA in large language models and vision tasks (Hu et al. 2022a; Liu et al. 2024b; Zhong et al. 2024; Talla et al. 2017), we propose a Fusion-Aware Multi-head Low-Rank Adaptation module (FAM-LoRA), as shown in Figure 4, which enables efficient semantic interaction during feature fusion with minimal parameters and computational overhead. Specifically, FAM-LoRA takes the main spatial feature Y_i and auxiliary guidance feature X_i^R as inputs. Y_i first passes through four consecutive 1×1 convolutions for channel compression and transformation,

Method	Scale	Chikusei				Houston				PaviaU			
		PSNR↑	SSIM↑	SAM↓	ERGAS↓	PSNR↑	SSIM↑	SAM↓	ERGAS↓	PSNR↑	SSIM↑	SAM↓	ERGAS↓
Brovoy (Vivone et al. 2014)	×2	29.2704	0.8751	8.1522	11.0668	32.5312	0.8431	8.4829	5.7936	27.8390	0.7958	6.2207	6.0319
GSA (Aiazzi, Baronti, and Selva 2007)	×2	37.4039	0.9627	4.1328	7.0972	33.2998	0.8630	9.3221	5.4163	28.7031	0.7945	6.6719	6.4978
HSRnet (Hu et al. 2022c)	×2	53.0324	0.9964	1.9876	1.4615	40.8710	0.9805	3.4652	1.9588	40.9947	0.9685	3.6476	2.2422
Fusformer (Hu et al. 2022b)	×2	53.8104	0.9967	1.9414	1.4521	42.1020	0.9888	3.1807	1.9238	41.8442	0.9727	3.2919	2.0589
PSRT (Deng et al. 2023)	×2	53.8821	0.9970	1.8994	1.3537	42.3937	0.9901	2.8675	1.8847	41.9956	0.9790	2.2518	1.8971
U2Net (Peng et al. 2023)	×2	53.9657	0.9974	1.8824	1.2972	42.7243	0.9920	2.7825	1.8308	42.2347	0.9813	2.1596	1.4162
3DT-Net (Ma et al. 2023)	×2	53.9155	0.9988	<u>1.8222</u>	1.3098	43.1856	<u>0.9947</u>	<u>2.5811</u>	<u>1.7677</u>	41.2224	0.9807	2.3080	1.6433
SMGU-Net (Yan et al. 2025)	×2	52.7751	0.9985	1.8869	<u>1.2842</u>	43.1647	0.9946	2.6137	1.7902	39.3881	0.9775	2.8193	1.8930
PIF-Net	×2	54.0236	0.9990	1.7962	1.1804	43.4025	0.9963	2.4857	1.7498	42.8673	0.9892	1.7289	1.1392
Brovoy (Vivone et al. 2014)	×4	27.5923	0.8325	6.2724	15.1083	30.1662	0.7745	11.5137	7.9550	25.5351	0.7143	8.8247	7.9284
GSA (Aiazzi, Baronti, and Selva 2007)	×4	31.2757	0.8985	6.0018	10.3126	30.9217	0.7605	12.3469	8.5070	28.5528	0.6295	9.5707	10.0441
HSRnet (Hu et al. 2022c)	×4	49.3548	0.9901	2.4509	1.9242	40.4509	0.9835	3.8988	2.9216	37.2608	0.9540	3.7748	3.1791
Fusformer (Hu et al. 2022b)	×4	50.1466	0.9909	2.3864	1.8927	41.1338	0.9895	3.4892	2.3436	37.7936	0.9558	3.8403	2.9017
PSRT (Deng et al. 2023)	×4	50.5377	0.9919	2.2097	1.7832	41.5616	0.9882	3.2434	2.1695	38.2794	0.9645	3.3579	2.4509
U2Net (Peng et al. 2023)	×4	50.5061	0.9929	2.1852	1.7317	41.7714	0.9889	3.1039	1.9978	38.7681	0.9700	3.0176	2.2220
3DT-Net (Ma et al. 2023)	×4	48.1940	0.9961	2.4262	2.1546	42.6873	0.9931	3.0318	1.9952	38.1864	0.9744	<u>2.8405</u>	2.3481
SMGU-Net (Yan et al. 2025)	×4	<u>51.3382</u>	<u>0.9981</u>	<u>2.1080</u>	<u>1.5762</u>	<u>42.9130</u>	<u>0.9937</u>	<u>2.8873</u>	<u>1.8979</u>	<u>38.9201</u>	<u>0.9761</u>	<u>2.9800</u>	<u>1.9973</u>
PIF-Net	×4	51.6257	0.9983	2.0653	1.5401	43.1128	0.9947	2.6452	1.7624	39.8246	0.9845	2.4018	1.8126
Brovoy (Vivone et al. 2014)	×8	26.0978	0.7989	8.1166	18.1408	27.0096	0.7217	16.0689	11.4050	22.7245	0.6497	13.4666	10.7571
GSA (Aiazzi, Baronti, and Selva 2007)	×8	28.1606	0.8609	8.0270	13.4385	26.2180	0.6829	17.6147	11.9947	21.2674	0.5245	14.1969	14.3429
HSRnet (Hu et al. 2022c)	×8	47.7262	0.9902	2.6291	2.3439	30.6384	0.9550	6.0210	5.9721	30.1921	0.8956	6.0662	4.9901
Fusformer (Hu et al. 2022b)	×8	47.8910	0.9919	2.5806	2.2506	31.8052	0.9589	5.7651	5.4997	30.4891	0.9011	5.9990	4.9715
PSRT (Deng et al. 2023)	×8	48.6457	0.9942	2.5291	2.0439	32.6384	0.9650	5.1151	5.1257	31.7671	0.9189	5.6891	4.8965
U2Net (Peng et al. 2023)	×8	48.8911	0.9959	2.4006	1.9506	32.8052	0.9676	4.9910	4.9711	<u>31.9211</u>	0.9210	5.5451	4.6712
3DT-Net (Ma et al. 2023)	×8	44.9324	0.9936	2.9072	2.6482	34.1538	0.9688	5.0265	4.8583	29.2755	0.9302	5.8975	5.0295
SMGU-Net (Yan et al. 2025)	×8	<u>49.6582</u>	<u>0.9975</u>	<u>2.3260</u>	<u>1.7872</u>	<u>34.2202</u>	<u>0.9699</u>	<u>4.6268</u>	<u>4.8072</u>	29.7848	<u>0.9417</u>	<u>5.4906</u>	<u>4.4843</u>
PIF-Net	×8	50.0124	0.9979	2.2603	1.7328	34.2989	0.9706	4.3902	4.7218	32.2310	0.9521	5.4258	4.1726
Best Value	-	+∞	1	0	0	+∞	1	0	0	+∞	1	0	0

Table 1: Quantitative comparisons of different approaches were conducted on the Chikusei, Houston, and PaviaU test datasets. The **best** results are in bold, and the second-best results are underlined.

followed by LKA and SE modules to enhance spatial context awareness and channel-wise selectivity. The gradients of LKA and SE are then frozen, and \mathbf{X}_i^R is injected into the SE input to steer attention toward critical semantic regions, improving fusion accuracy. To capture multi-scale representations, the channel dimension is evenly split into four heads, each independently applying Low-Rank Adaptation.

Loss Function

To achieve efficient training and precise representation in PIF-Net, the loss function is composed of three parts, aiming to comprehensively enhance the fusion accuracy and feature consistency of the model, while ensuring the invertibility of the transformation process. The loss function is formulated as follows:

$$\mathcal{L} = \underbrace{\|\hat{\mathbf{Z}} - \mathbf{Z}\|_1}_{\mathcal{L}_1} - \lambda_{inv} \cdot \underbrace{\log \left| \det \left(\frac{\partial \hat{\mathbf{Z}}}{\partial \mathbf{X}} \right) \right|}_{\mathcal{L}_{inv}} + \lambda_{cos} \cdot \underbrace{\left(1 - \cos(\hat{\mathbf{Z}}, \mathbf{Z}) \right)}_{\mathcal{L}_{cos}} \quad (4)$$

First, the ℓ_1 fusion loss $\mathcal{L}_1 = \|\hat{\mathbf{Z}} - \mathbf{Z}\|_1$ precisely constrains the difference between the predicted feature $\hat{\mathbf{Z}}$ and the ground truth feature \mathbf{Z} , effectively reducing fusion error.

Second, the log-determinant of the Jacobian of the affine transformation, $\mathcal{L}_{inv} = \log \left| \det \left(\frac{\partial \hat{\mathbf{Z}}}{\partial \mathbf{X}} \right) \right|$, serves as an invertibility regularization term to ensure information preservation and numerical stability during the mapping process (Kingma and Dhariwal 2018; Lugmayr et al. 2020), encouraging the model to learn stable and invertible feature transformations. Finally, the cosine similarity loss

$\mathcal{L}_{cos} = 1 - \cos(\hat{\mathbf{Z}}, \mathbf{Z})$ enhances angular consistency between different feature representations, improving semantic feature fusion.

Overall, the loss function balances fusion fidelity (\mathcal{L}_1), semantic consistency (\mathcal{L}_{cos}), and transformation invertibility (\mathcal{L}_{inv}) under multi-objective constraints, enabling the model to learn efficiently and robustly within complex high-dimensional feature spaces. In practice, we set the weights as $\lambda_{inv} = 0.01$ and $\lambda_{cos} = 0.1$ to appropriately balance the influence of each term, ensuring the model achieves not only superior fusion performance but also stable mappings and consistent semantic representations.

Experiments

Experimental Settings

Dataset and Evaluation. We conduct experiments on three widely used hyperspectral datasets: Chikusei (Yokoya and Iwasaki 2016), PaviaU, and Houston. The Chikusei image (2517 × 2335, 128 bands) is split by using a 1000 × 2000 top-left region for training and the remaining area, divided into 680 × 680 patches, for testing. PaviaU (610 × 340, 103 bands) uses the top 340 × 340 region for testing and the rest for training. For Houston (349 × 1905, 144 bands), the left 349 × 349 area is used for testing. Following (Wei et al. 2015), HRMSI are simulated, and LRHSI are obtained via Gaussian blur (kernel size 3 × 3, std 0.5) and downsampling. Training samples are cropped into 64 × 64 ground-truth patches paired with corresponding LRHSI (16 × 16) and HRMSI patches.

Metrics. For quantitative evaluation, we adopt four commonly-used metrics: Peak Signal-to-Noise Ratio (PSNR), Spectral Angle Mapper (SAM) (Yuhas, Goetz, and Boardman 1992), Relative Global Dimensional Synthesis Error (ERGAS) (Wald 2002), and Structural Similarity Index (SSIM) (Wang et al. 2004).

Training Details. Our model is implemented in PyTorch and trained on an NVIDIA A30 GPU. The hidden dimension D of PIF-Net is set to 64, and the number of feature extraction blocks L is 4. We adopt the AdamW optimizer (Kingma and Ba 2015) with an initial learning rate of 1×10^{-4} , which is halved every 200 epochs. The batch size is 8, and the total training spans 500 epochs. For all comparison methods, losses follow their official default implementations.

Ablation Study

Effects of Ill-Posed Residual Prior. To investigate the influence of the ill-posed residual prior on model performance, we introduce a weighting parameter β to control the strength of this prior. A series of ablation experiments were conducted on the PaviaU $\times 4$ dataset to evaluate the effect of varying β values. The results, presented in Table 2, reveal a clear trend in model performance as β changes. Specifically, the ill-posed residual prior significantly contributes to improving both model stability and fusion accuracy. The best performance is achieved when $\beta = 1$. In contrast, setting β to 0.4 or 0.8 results in no notable improvement and may even slightly degrade performance. Moreover, when $\beta = 0$, indicating the model does not utilize the residual prior, performance is at its lowest, further underscoring the critical role of this prior.

β	0	0.4	0.8	1
PSNR	36.9541 dB	38.0012 dB	38.9129 dB	39.8246 dB

Table 2: Impact of varying β values on PSNR performance on the PaviaU $\times 4$ dataset, demonstrating how the weighting of the ill-posed residual prior influences fusion quality. The **best** results are in bold.

Effects of the Invertible Mamba Block. To assess the impact of the *Invertible Mamba Block* on fusion performance, we performed an ablation study on the PaviaU $\times 4$ dataset, comparing the model’s performance with and without this block. As shown in Table 3, incorporating the Mamba Block leads to a notable improvement in fusion quality. Specifically, the PSNR increases from 37.1781 dB to 39.8246 dB, while SSIM improves from 0.9637 to 0.9845. These results highlight the effectiveness of the Mamba Block in enhancing feature representation and boosting the fidelity and quality of the fused images.

Structure	w/o Mamba Block	w/ Mamba Block
PSNR / SSIM	37.1781 dB / 0.9637	39.8246 dB / 0.9845

Table 3: Performance impact of the *Invertible Mamba Block* on the PaviaU $\times 4$ dataset. The **best** results are in bold.

Effects of FAM-LoRA. We also investigate the influence of the FAM-LoRA module using a similar ablation study on the PaviaU $\times 4$ dataset. As reported in Table 4, the inclusion of FAM-LoRA considerably enhances the model’s accuracy. Without this module, the PSNR and SSIM are 36.9211 dB and 0.9607, respectively. When FAM-LoRA is integrated, these metrics rise substantially to 39.8246 dB and 0.9845. This demonstrates that FAM-LoRA effectively facilitates feature aggregation, leading to superior fine-detail fusion and overall image quality.

Structure	w/o FAM-LoRA	w/ FAM-LoRA
PSNR / SSIM	36.9211 dB / 0.9607	39.8246 dB / 0.9845

Table 4: Performance impact of the FAM-LoRA module on the PaviaU $\times 4$ dataset. The **best** results are in bold.

Ablation Study on Loss Components. To evaluate the impact of the invertibility regularization term \mathcal{L}_{inv} and the cosine similarity loss \mathcal{L}_{cos} on the overall loss function, we conducted an ablation study. The fusion loss \mathcal{L}_1 is always included as the fundamental loss. Table 5 shows the PSNR and SSIM results of the model on the benchmark dataset under different loss combinations. When only \mathcal{L}_1 is used, the model achieves a PSNR of 37.4521 dB and an SSIM of 0.9187. Adding \mathcal{L}_{inv} significantly improves performance, with PSNR reaching 38.6794 dB and SSIM 0.9342. Adding \mathcal{L}_{cos} also enhances the results, with PSNR of 38.1158 dB and SSIM of 0.9278. The best performance is obtained when both \mathcal{L}_{inv} and \mathcal{L}_{cos} are included, achieving a PSNR of 39.8246 dB and SSIM of 0.9845. These results indicate that the two auxiliary loss terms effectively improve fusion accuracy and semantic consistency, confirming their important roles in PIF-Net.

\mathcal{L}_1	\mathcal{L}_{inv}	\mathcal{L}_{cos}	PSNR (dB)	SSIM
✓	✗	✗	37.4521	0.9187
✓	✓	✗	38.6794	0.9342
✓	✗	✓	38.1158	0.9278
✓	✓	✓	39.8246	0.9845

Table 5: Ablation study on loss components (\mathcal{L}_1 loss always included) conducted on the PaviaU $\times 4$ dataset. The **best** results are in bold.

Comparison with State-of-the-Art Methods

Quantitative Results. We comprehensively evaluate PIF-Net against several SOTA methods, with results in Table 1 showing its clear superiority across all metrics. PIF-Net consistently achieves the highest PSNR and SSIM and the lowest SAM and ERGAS across datasets and scales. Notably, on challenging $\times 4$ and $\times 8$ tasks, PIF-Net significantly outperforms competitors—for instance, on PaviaU $\times 4$, it surpasses the second-best by 0.9045 dB PSNR, and on Chikusei $\times 8$, it sets a new benchmark with 50.0124 dB PSNR. These results validate that the integration of the Invertible Mamba Block and FAM-LoRA effectively captures spectral-spatial correlations and preserves high-frequency details, empowering PIF-Net’s superior performance, especially at large scaling factors

Visual Comparison. We evaluate the computational efficiency of the models on the PaviaU dataset for the $\times 4$ super-resolution task, measuring inference latency on an NVIDIA A30 GPU under identical implementation settings. As shown in Figure 6, both the number of parameters and the inference latency are presented on a logarithmic scale for clearer comparison across models. PIF-Net achieves a favorable trade-off among model size, inference speed, and reconstruction accuracy. With only 1.73 million parameters and a latency of 9.3 ms, PIF-Net achieves the highest PSNR and SSIM among all compared methods, demonstrating superior efficiency and strong potential for real-time applications.




Figure 5: The pseudo-color images, corresponding $\times 4$ super-resolution results, and SAM error maps generated by all comparative models on three datasets: (1) the sixth test area of the Chikusei dataset (rows 1–2, RGB bands: R=101, G=40, B=10), (2) a test area from the PaviaU dataset (rows 3–4, RGB bands: R=20, G=30, B=40), and (3) a test area from the Houston dataset (rows 5–6, RGB bands: R=10, G=76, B=2).

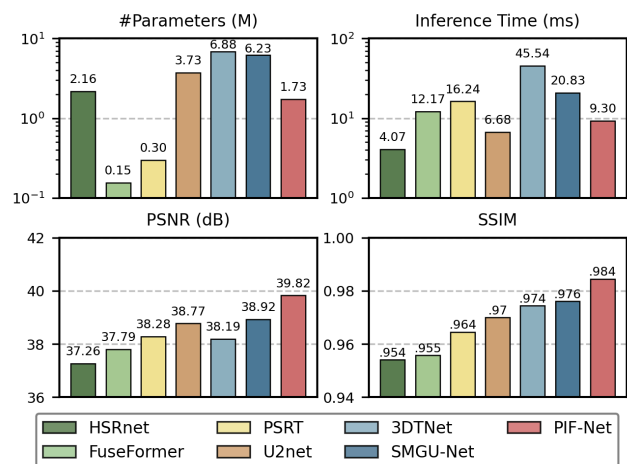


Figure 6: Comparison of model efficiency on PaviaU for $\times 4$ super-resolution. Both the number of parameters (in millions) and inference latency (in milliseconds, measured on an NVIDIA A30 GPU) are displayed on logarithmic scales.

Efficiency Comparison. We evaluate the computational efficiency of all competing methods on the PaviaU dataset for the $\times 4$

super-resolution task, measuring inference latency on an NVIDIA A30 GPU under identical implementation and batch settings. As illustrated in Figure 6, both model size (in millions of parameters) and inference latency (in milliseconds) are displayed on logarithmic scales to better highlight relative differences. PIF-Net achieves a compelling balance among model complexity, speed, and reconstruction fidelity. With only 1.73M parameters and a latency of 9.3 ms, it outperforms all other methods in terms of both PSNR and SSIM, demonstrating SOTA accuracy and high computational efficiency, which makes it well suited for practical and real-time fusion scenarios.

Conclusion

This paper proposes PIF-Net, an innovative MHIF framework. Guided by ill-posed priors, the framework employs a dual-branch architecture that integrates HWT with Invertible Mamba Blocks, achieving high-fidelity information transformation and stable transmission, while ensuring the invertibility of feature transformations and the stability of gradient flow. Meanwhile, by combining a spatial-domain fusion strategy based on low-frequency priors with a lightweight LoRA fusion module, it effectively enhances the dynamic calibration of spatial details and spectral features. Furthermore, we innovatively design a guided feature consistency loss that improves semantic alignment and representation across domains. Extensive experimental results demonstrate that PIF-Net achieves SOTA performance on three public datasets.

Acknowledgments

This research was supported by the National Key Research and Development Program of China (No. 2023YFB4502304).

References

- Aiazzi, B.; Alparone, L.; Baronti, S.; Garzelli, A.; and Selva, M. 2006. MTF-tailored multiscale fusion of high-resolution MS and Pan imagery. *Photogrammetric Engineering & Remote Sensing*, 72(5): 591–596.
- Aiazzi, B.; Baronti, S.; and Selva, M. 2007. Improving Component Substitution Pansharpening Through Multivariate Regression of MS +Pan Data. *IEEE Transactions on Geoscience and Remote Sensing*, 3230–3239.
- Arduzzone, L.; Lüth, C.; Kruse, J.; Rother, C.; and Köthe, U. 2019. Guided image generation with conditional invertible neural networks. *arXiv preprint arXiv:1907.02392*.
- Bhargava, A.; Sachdeva, A.; Sharma, K.; Alsharif, M. H.; Uthansakul, P.; and Uthansakul, M. 2024. Hyperspectral imaging and its applications: A review. *Heliyon*, 10(12): e33208.
- Bi, X.; Wang, J.; Xue, B.; He, C.; Liu, F.; Chen, H.; Lin, L. L.; Dong, B.; Li, B.; Jin, C.; et al. 2024. SERSomes for metabolic phenotyping and prostate cancer diagnosis. *Cell Reports Medicine*, 5(6).
- Cao, X.; Yu, J.; Xu, R.; Wei, J.; and Jiao, L. 2024. Mask-Enhanced Contrastive Learning for Hyperspectral Image Classification. *IEEE Transactions on Geoscience and Remote Sensing*, 62: 1–15.
- Creswell, A.; White, T.; Dumoulin, V.; Arulkumaran, K.; Sengupta, B.; and Bharath, A. A. 2018. Generative adversarial networks: An overview. *IEEE signal processing magazine*, 35(1): 53–65.
- Deng, S.-Q.; Deng, L.-J.; Wu, X.; Ran, R.; Hong, D.; and Vivone, G. 2023. PSRT: Pyramid Shuffle-and-Reshuffle Transformer for Multispectral and Hyperspectral Image Fusion. *IEEE Transactions on Geoscience and Remote Sensing*, 1–15.
- Dinh, L.; Krueger, D.; and Bengio, Y. 2014. NICE: Non-linear Independent Components Estimation. *Computer Science*.
- Dinh, L.; Sohl-Dickstein, J.; and Bengio, S. 2017. Density estimation using Real NVP. In *International Conference on Learning Representations*.
- Fan, W.; Wang, Y.; Wang, C.; Zhang, Y.; Wang, W.; and Zhou, D. 2025. Semantic-Guided Global-Local Collaborative Networks for Lightweight Image Super-Resolution. *arXiv:2503.16056*.
- Gu, A.; and Dao, T. 2023. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*.
- Guo, H.; Guo, Y.; Zha, Y.; Zhang, Y.; Li, W.; Dai, T.; Xia, S.-T.; and Li, Y. 2025. Mambairv2: Attentive state space restoration. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 28124–28133.
- Guo, H.; Li, J.; Dai, T.; Ouyang, Z.; Ren, X.; and Xia, S.-T. 2024. Mambair: A simple baseline for image restoration with state-space model. In *European conference on computer vision*, 222–241. Springer.
- Hang, R.; Li, Z.; Liu, Q.; Ghamisi, P.; and Bhattacharyya, S. S. 2021. Hyperspectral Image Classification With Attention-Aided CNNs. *IEEE Transactions on Geoscience and Remote Sensing*, 2281–2293.
- He, S.; Zhou, H.; Wang, Y.; Cao, W.; and Han, Z. 2016. Super-resolution reconstruction of hyperspectral images via low rank tensor modeling and total variation regularization. In *2016 IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, 6962–6965. IEEE.
- He, X.; Tang, C.; Liu, X.; Zhang, W.; Sun, K.; and Xu, J. 2023. Object Detection in Hyperspectral Image via Unified Spectral–Spatial Feature Aggregation. *IEEE Transactions on Geoscience and Remote Sensing*, 61: 1–13.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W.; et al. 2022a. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2): 3.
- Hu, J.; Shen, L.; and Sun, G. 2018. Squeeze-and-Excitation Networks. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7132–7141.
- Hu, J.-F.; Huang, T.-Z.; Deng, L.-J.; Dou, H.-X.; Hong, D.; and Vivone, G. 2022b. Fusformer: A Transformer-Based Fusion Network for Hyperspectral Image Super-Resolution. *IEEE Geoscience and Remote Sensing Letters*, 1–5.
- Hu, J.-F.; Huang, T.-Z.; Deng, L.-J.; Jiang, T.-X.; Vivone, G.; and Chanussot, J. 2022c. Hyperspectral Image Super-Resolution via Deep Spatospectral Attention Convolutional Neural Networks. *IEEE Transactions on Neural Networks and Learning Systems*, 7251–7265.
- Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Kingma, D. P.; and Dhariwal, P. 2018. Glow: Generative flow with invertible 1x1 convolutions. *Advances in neural information processing systems*, 31.
- Li, B.; Wang, X.; and Xu, H. 2025. HSRMamba: Efficient Wavelet Stripe State Space Model for Hyperspectral Image Super-Resolution. *arXiv preprint arXiv:2505.11062*.
- Liao, R.; Li, F.; Wei, Y.; Shi, Z.; Zhang, L.; Bai, H.; and Wang, M. 2025. Prompt to Restore, Restore to Prompt: Cyclic Prompting for Universal Adverse Weather Removal. *arXiv preprint arXiv:2503.09013*.
- Liu, P.; Xu, T.; Chen, H.; Zhou, S.; Qin, H.; and Li, J. 2024a. Spectrum-Driven Mixed-Frequency Network for Hyperspectral Salient Object Detection. *IEEE Transactions on Multimedia*, 26: 5296–5310.
- Liu, P.; Xu, T.; Wang, J.; Chen, H.; Bai, H.; and Li, J. 2025. Dual-Stage Hyperspectral Image Classification Model with Spectral Super-token. In *European Conference on Computer Vision*, 368–386. Springer.
- Liu, S.-Y.; Wang, C.-Y.; Yin, H.; Molchanov, P.; Wang, Y.-C. F.; Cheng, K.-T.; and Chen, M.-H. 2024b. Dora: Weight-decomposed low-rank adaptation. In *Forty-first International Conference on Machine Learning*.
- Liu, Y.; Tian, Y.; Zhao, Y.; Yu, H.; Xie, L.; Wang, Y.; Ye, Q.; Jiao, J.; and Liu, Y. 2024c. Vmamba: Visual state space model. *Advances in neural information processing systems*, 37: 103031–103063.
- Liu, Z.; Mao, H.; Wu, C.-Y.; Feichtenhofer, C.; Darrell, T.; and Xie, S. 2022. A ConvNet for the 2020s. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Lu, K.; You, S.; and Barnes, N. 2018. Deep texture and structure aware filtering network for image smoothing. In *Proceedings of the European conference on computer vision (ECCV)*, 217–233.
- Lugmayr, A.; Danelljan, M.; Van Gool, L.; and Timofte, R. 2020. SrfLOW: Learning the super-resolution space with normalizing flow. In *Computer vision—ECCV 2020: 16th European conference, Glasgow, UK, August 23–28, 2020, proceedings, part v 16*, 715–732. Springer.

- Ma, J.; Yu, W.; Chen, C.; Liang, P.; Guo, X.; and Jiang, J. 2020. Pan-GAN: An unsupervised pan-sharpening method for remote sensing image fusion. *Information Fusion*, 62: 110–120.
- Ma, Q.; Jiang, J.; Liu, X.; and Ma, J. 2023. Learning a 3D-CNN and Transformer prior for hyperspectral image super-resolution. *Information Fusion*, 100: 101907.
- Peng, C.; Zhang, X.; Yu, G.; Luo, G.; and Sun, J. 2017. Large Kernel Matters – Improve Semantic Segmentation by Global Convolutional Network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Peng, S.; Guo, C.; Wu, X.; and Deng, L.-J. 2023. U2Net: A General Framework with Spatial-Spectral-Integrated Double U-Net for Image Fusion. In *Proceedings of the 31st ACM International Conference on Multimedia*, 3219–3227. ISBN 9798400701085.
- Qin, H.; Xu, T.; Liu, P.; Xu, J.; and Li, J. 2024. DMSSN: Distilled Mixed Spectral–Spatial Network for Hyperspectral Salient Object Detection. *IEEE Transactions on Geoscience and Remote Sensing*, 62: 1–18.
- Saharia, C.; Ho, J.; Chan, W.; Salimans, T.; Fleet, D. J.; and Norouzi, M. 2022. Image super-resolution via iterative refinement. *IEEE transactions on pattern analysis and machine intelligence*, 45(4): 4713–4726.
- Sethy, P. K.; Pandey, C.; Sahu, Y. K.; and Behera, S. K. 2022. Hyperspectral imagery applications for precision agriculture—a systemic survey. *Multimedia Tools and Applications*, 81: 3005–3038.
- Sun, W.; Chen, B.; and Messinger, D. W. 2014. Nearest-neighbor diffusion-based pan-sharpening algorithm for spectral images. *Optical Engineering*, 53(1): 013107–013107.
- Talla, V.; Hessar, M.; Kellogg, B.; Najafi, A.; Smith, J. R.; and Gollakota, S. 2017. Lora backscatter: Enabling the vision of ubiquitous connectivity. *Proceedings of the ACM on interactive, mobile, wearable and ubiquitous technologies*, 1(3): 1–24.
- Vivone, G.; Alparone, L.; Chanussot, J.; Dalla Mura, M.; Garzelli, A.; Licciardi, G. A.; Restaino, R.; and Wald, L. 2014. A critical comparison among pansharpening algorithms. *IEEE Transactions on Geoscience and Remote Sensing*, 53(5): 2565–2586.
- Wald, L. 2002. *Data fusion: definitions and architectures: fusion of images of different spatial resolutions*. Presses des MINES.
- Wang, W.; Dai, J.; Chen, Z.; Huang, Z.; Li, Z.; Zhu, X.; Hu, X.; Lu, T.; Lu, L.; Li, H.; et al. 2023. Internimage: Exploring large-scale vision foundation models with deformable convolutions. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 14408–14419.
- Wang, X.; Yu, K.; Wu, S.; Gu, J.; Liu, Y.; Dong, C.; Qiao, Y.; and Change Loy, C. 2018. Esrgan: Enhanced super-resolution generative adversarial networks. In *Proceedings of the European conference on computer vision (ECCV) workshops*, 0–0.
- Wang, Z.; Bovik, A.; Sheikh, H.; and Simoncelli, E. 2004. Image Quality Assessment: From Error Visibility to Structural Similarity. *IEEE Transactions on Image Processing*, 600–612.
- Wei, Q.; Bioucas-Dias, J.; Dobigeon, N.; and Tourneret, J.-Y. 2015. Hyperspectral and Multispectral Image Fusion Based on a Sparse Representation. *IEEE Transactions on Geoscience and Remote Sensing*, 3658–3668.
- Wu, L.; Yin, Y.; Jiang, X.; and Cheng, T. 2021. Pan-sharpening based on multi-objective decision for multi-band remote sensing images. *Pattern Recognition*, 118: 108022.
- Xiao, M.; Zheng, S.; Liu, C.; Wang, Y.; He, D.; Ke, G.; Bian, J.; Lin, Z.; and Liu, T.-Y. 2020. Invertible image rescaling. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16*, 126–144. Springer.
- Xie, Q.; Zhou, M.; Zhao, Q.; Xu, Z.; and Meng, D. 2022. MHF-Net: An Interpretable Deep Network for Multispectral and Hyperspectral Image Fusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(3): 1457–1473.
- Yan, J.; Zhang, K.; Sun, Q.; Ge, C.; Wan, W.; Sun, J.; and Zhang, H. 2025. Spatial–spectral unfolding network with mutual guidance for multispectral and hyperspectral image fusion. *Pattern Recognition*, 161: 111277.
- Yang, Y.; Wan, W.; Huang, S.; Lin, P.; and Que, Y. 2017. A novel pan-sharpening framework based on matting model and multiscale transform. *Remote Sensing*, 9(4): 391.
- Yang, Y.; Wu, L.; Huang, S.; Sun, J.; Wan, W.; and Wu, J. 2018. Compensation details-based injection model for remote sensing image fusion. *IEEE Geoscience and Remote Sensing Letters*, 15(5): 734–738.
- Yi, Q.; Li, J.; Dai, Q.; Fang, F.; Zhang, G.; and Zeng, T. 2021. Structure-preserving deraining with residue channel prior guidance. In *Proceedings of the IEEE/CVF international conference on computer vision*, 4238–4247.
- Yokoya, N.; and Iwasaki, A. 2016. Airborne hyperspectral data over Chikusei. *Space Appl. Lab., Univ. Tokyo, Tokyo, Japan, Tech. Rep. SAL-2016-05-27*, 5.
- Yuhas, R. H.; Goetz, A. F.; and Boardman, J. W. 1992. Discrimination among semi-arid landscape endmembers using the spectral angle mapper (SAM) algorithm. In *JPL, Summaries of the Third Annual JPL Airborne Geoscience Workshop. Volume 1: AVIRIS Workshop*.
- Zhang, K. A.; Cuesta-Infante, A.; and Veeramachaneni, K. 2019. SteganoGAN: High Capacity Image Steganography with GANs. *arXiv preprint arXiv:1901.03892*.
- Zhao, Z.; Bai, H.; Zhang, J.; Zhang, Y.; Xu, S.; Lin, Z.; Timofte, R.; and Van Gool, L. 2023. Cddfuse: Correlation-driven dual-branch feature decomposition for multi-modality image fusion. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5906–5916.
- Zhong, Z.; Tang, Z.; He, T.; Fang, H.; and Yuan, C. 2024. Convolution meets lora: Parameter efficient finetuning for segment anything model. *arXiv preprint arXiv:2401.17868*.