

EigenSR: Eigenimage-Bridged Pre-Trained RGB Learners for Single Hyperspectral Image Super-Resolution

Xi Su, Xiangfei Shen, Mingyang Wan, Jing Nie, Lihui Chen,
Haijun Liu*, Xichuan Zhou

Chongqing University
xisu@stu.cqu.edu.cn, haijun.liu@126.com

Abstract

Single hyperspectral image super-resolution (single-HSI-SR) aims to improve the resolution of a single input low-resolution HSI. Due to the bottleneck of data scarcity, the development of single-HSI-SR lags far behind that of RGB natural images. In recent years, research on RGB SR has shown that models pre-trained on large-scale benchmark datasets can greatly improve performance on unseen data, which may stand as a remedy for HSI. But how can we transfer the pre-trained RGB model to HSI, to overcome the data-scarcity bottleneck? Because of the significant difference in the channels between the pre-trained RGB model and the HSI, the model cannot focus on the correlation along the spectral dimension, thus limiting its ability to utilize on HSI. Inspired by the HSI spatial-spectral decoupling, we propose a new framework that first fine-tunes the pre-trained model with the spatial components (known as eigenimages), and then infers on unseen HSI using an iterative spectral regularization (ISR) to maintain the spectral correlation. The advantages of our method lie in: 1) we effectively inject the spatial texture processing capabilities of the pre-trained RGB model into HSI while keeping spectral fidelity, 2) learning in the spectral-decorrelated domain can improve the generalizability to spectral-agnostic data, and 3) our inference in the eigenimage domain naturally exploits the spectral low-rank property of HSI, thereby reducing the complexity. This work bridges the gap between pre-trained RGB models and HSI via eigenimages, addressing the issue of limited HSI training data, hence the name EigenSR. Extensive experiments show that EigenSR outperforms the state-of-the-art (SOTA) methods in both spatial and spectral metrics.

Code — <https://github.com/enter-i-username/EigenSR>

Introduction

Hyperspectral imaging sensors acquire three-dimensional hyperspectral images (HSIs) by capturing a vector with hundreds of contiguous spectral elements from each pixel within a given scene, offering the perception capacity of various types of real-world materials (Keshava and Mustard 2002; Bioucas-Dias et al. 2013). However, physical restrictions in imaging sensors lead to the final product of low spatial resolution HSI (Hong et al. 2021; Yokoya, Grohnfeldt, and

*Corresponding author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

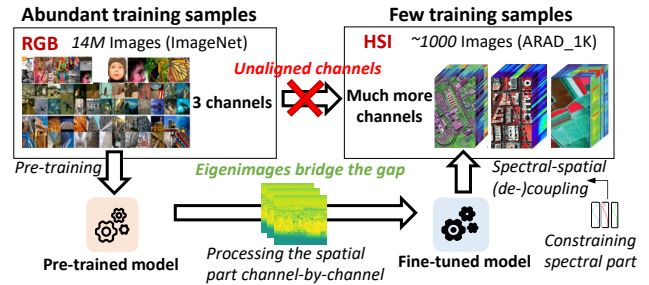


Figure 1: The motivation. The lack of HSI training data is the main bottleneck for single-HSI-SR. The channel differences between RGB and HSI make it challenging to directly transfer RGB models to HSI. Inspired by spectral-spatial decoupling, we leverage spatial part of HSI, namely eigenimages, to bridge the gap between pre-trained RGB models and HSI, and constrain the spectral part to keep spectral fidelity.

Chanussot 2017). To tackle this restriction, hyperspectral image super-resolution (HSI-SR) task (Vivone 2023) has received extensive attention by injecting abundant spatial details to HSI for spatial resolution improvement. Typically, HSI-SR can be categorized into two branches of methods: fusion-based (Xu et al. 2022; Peng et al. 2021; Xu et al. 2020) and single-image-based methods (Wan et al. 2022; Liang et al. 2023; Dong et al. 2021; Zhang et al. 2023). In this work, we focus on the single-HSI-SR task.

Similar to the definition of single RGB image SR (Dong et al. 2015; Zhang et al. 2018), single-HSI-SR aims to recover spatial details with only a single low-resolution (LR) image. With the booming of deep learning (Bai et al. 2024, 2025; Zhou et al. 2024a; Nie, Xie, and Sun 2024; Xie et al. 2023; Nie et al. 2023), and the development of large-scale benchmark datasets (Deng et al. 2009; Agustsson and Timofte 2017), single RGB image SR has reached a level of maturity, exemplified in the studies that leverage pre-training techniques (Chen et al. 2021; Li et al. 2021; Chen et al. 2023). Unfortunately, such an effective strategy cannot be applied to HSI due to the scarcity of data (Liu et al. 2022; Miao et al. 2022). Most existing single-HSI-SR methods can only be trained on a small number of indoor and outdoor images, or on low-resolution sub-images of remote sensing

data (Jiang et al. 2020; Zhang et al. 2023), thus restricting their capabilities. This scarcity of training data has become the major bottleneck in the field of single-HSI-SR.

Given the lack of a large-scale dataset for pre-training an HSI model, a natural question arises: *how can we leverage pre-trained RGB models to feed the data-hungry bottleneck of HSI?* Directly transferring the RGB model for HSI-SR is difficult, because the channels are not aligned (3 for RGB and much more for HSI) (Yuan, Zheng, and Lu 2017; Liang et al. 2023). Thus the RGB model cannot handle the correlation along the spectral dimension.

To address the issue, we borrow the idea of HSI spatial-spectral decoupling using the Singular Value Decomposition (SVD) (Zhuang and Ng 2023), and operate SR using pre-trained RGB models solely on the spatial components (known as eigenimages (Zhuang and Bioucas-Dias 2018)), while constraining the spectral correlations with the spectral components (called the spectral basis), as shown in Figure 1. Specifically, our method consists of two stages. First, we fine-tune the pre-trained RGB model to adapt to the eigenimage domain in a single-channel manner. Second, when the fine-tuned model is used to infer¹ on the unseen HSI, we decompose the LR HSI, perform SR on the eigenimages channel-by-channel, and reconstruct the SR HSI with the spectral basis. Additionally, an iterative spectral regularization (ISR) is proposed to enhance the spectral correlation by fully utilizing the LR HSI. The advantages lie in threefold: 1) we effectively inject the spatial texture processing capabilities learned from abundant RGB images into HSI while maintaining spectral correlation, 2) due to the learning in the spectral-decorrelated domain, the generalizability to spectral-agnostic data can be improved, and 3) our inference in the eigenimage domain naturally exploits the spectral low-rank property of HSI, thereby reducing the complexity. This work bridges the gap between pre-trained RGB models and HSI through eigenimages, addressing the issue of limited HSI training data, hence the name EigenSR. Extensive experiments show that EigenSR outperforms the SOTA methods in both spatial and spectral metrics.

Our main contributions can be summarized as follows:

- We propose a new framework for single-HSI-SR that first introduces large-scale pre-trained RGB models into HSI. It leverages pre-trained models learned from abundant RGB images to address the data-scarcity problem of HSI.
- Our method is based on the HSI spatial-spectral decoupling. We inject the spatial texture knowledge provided by the pre-trained model to the spatial part via fine-tuning, and regularize the spectral correlation when inferring on unseen data. This effectively leverages the pre-trained model while maintaining the spectral fidelity.
- Extensive experiments demonstrate that our method surpasses the current SOTA methods in both spatial and spectral metrics, validating the effectiveness of introducing the pre-trained RGB models.

¹In this paper, we do not particularly differentiate between “test” and “infer”, both of which refer to performing single-image super-resolution on the unseen low-resolution HSI.

Related Work

Single RGB Image Super-Resolution

Single RGB image SR has been a long-standing topic of interest. Compared to convolutional neural networks (CNNs) for SR (Dong et al. 2015; Ledig et al. 2017; Wang et al. 2018), Transformer-based architectures (Liang et al. 2021; Chen et al. 2022, 2023; Hsu, Lee, and Chou 2024) have recently gained prominence in SR due to their ability to extract non-local spatial features and their scalability. Recent studies have found that self-supervised pre-training on large benchmark datasets, such as ImageNet (Deng et al. 2009), followed by fine-tuning on high-quality domain-specific datasets, such as DIV2K (Agustsson and Timofte 2017) and Flickr2K (Timofte et al. 2017), significantly enhances the Transformer’s power to process spatial information (Chen et al. 2021; Li et al. 2021; Chen et al. 2023). However, HSI lacks such large datasets, which hinders the application of pre-training. Therefore, we introduce pre-trained RGB models for HSI to address this data bottleneck.

Single HSI Super-Resolution

Single-HSI-SR not only involves processing spatial information but also focuses on the spectral correlation. Researchers used the 3D-CNN (Li, Wang, and Li 2020), group convolution (Jiang et al. 2020), spectrum and feature context (Wang, Li, and Li 2020), recurrent mechanism (Wang, Ma, and Jiang 2022), and Transformers (Chen, Zhang, and Zhang 2023; Zhang et al. 2023) to learn the spatial-spectral correlations in HSI in an end-to-end manner. Limited by the data, these models may not reach their full potential.

To address this issue, two interesting attempts (Yuan, Zheng, and Lu 2017; Liang et al. 2023) have found the effectiveness of leveraging RGB images for HSI-SR, inspired by transfer learning. They trained single-channel models on RGB images and inferred on HSI band-by-band to align the channels, followed by a non-negative matrix factorization (NMF) to ensure the spectral fidelity. However, the band-by-band operation in original HSI space and optimization of NMF may result in high time costs. Unlike them, we perform SR in the eigenimage domain, which more effectively and efficiently leverages large-scale pre-trained models.

Learning in the Eigenimage Domain

Learning in the eigenimage domain enables the independent processing of spectral or spatial components. Current denoising methods leverage the natural low-rank property of eigenimages and have achieved impressive results. For instance, off-the-shelf methods (Dabov et al. 2007) were inserted into eigenimages for efficient denoising and inpainting (Zhuang and Bioucas-Dias 2018). CNNs were used on eigenimages to remove mixed noise in HSI (Zhuang and Ng 2023). Deep spatial priors were learned exclusively in the eigenimage domain for HSI denoising (Peng et al. 2024). Inspired by the spectral-spatial decoupling, we perform single-HSI-SR on eigenimages in this work. We will demonstrate the effectiveness of eigenimages in the field of SR through theoretical analysis and experimental observations.

Methodology

Problem Formulation

Single-HSI-SR learns a SR model $f_{SR}(\cdot)$ on the training data, mapping the low-resolution HSI $\mathbf{Y}_{LR} \in \mathbb{R}^{L \times (N/\kappa^2)}$ to the high-resolution (HR) HSI $\mathbf{Y}_{HR} \in \mathbb{R}^{L \times N}$, where L , N , and κ represent the number of spectral bands, the number of pixels, and the spatial downsampling rate, respectively. Since there is no real dataset of LR-HR pairs, we follow the common practice that uses a downsampling operator to generate LR HSI from HR HSI by: $\mathbf{Y}_{LR} = \mathbf{Y}_{HR} \mathbf{D}$ where $\mathbf{D} \in \mathbb{R}^{N \times (N/\kappa^2)}$ is the linear downsampling operator. Specifically, the downsampling operator \mathbf{D} is the Bicubic function (Chen et al. 2023; Hsu, Lee, and Chou 2024).

Limited by the data scarcity, $f_{SR}(\cdot)$ cannot generalize well on unseen data. To address this limitation, we introduce the pre-trained RGB model into HSI. However, due to the difference in the image channels, RGB models cannot capture the spectral correlations in HSI. Next, we will explain how to overcome this using eigenimages.

SR with Eigenimages. Singular Value Decomposition (SVD) decomposes an HSI \mathbf{Y} into spectral and spatial components. This can be described as:

$$\{\mathbf{U}, \mathbf{\Sigma}, \mathbf{V}^T\} = \text{SVD}(\mathbf{Y}), \text{ or } \mathbf{Y} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T, \quad (1)$$

where $(\cdot)^T$ is the transpose operation; $\mathbf{U} \in \mathbb{R}^{L \times L}$ and $\mathbf{V} \in \mathbb{R}^{N \times N}$ are semiunitary matrices, satisfying $\mathbf{U}^T\mathbf{U} = \mathbf{I}$ and $\mathbf{V}^T\mathbf{V} = \mathbf{I}$; $\mathbf{\Sigma} = \text{diag}\{\sigma_1, \dots, \sigma_L\}$ is a diagonal matrix of singular values arranged in descending order. We combine $\mathbf{\Sigma}$ and \mathbf{V}^T to form $\mathbf{E} = \mathbf{\Sigma}\mathbf{V}^T$, and the rows $\mathbf{E}_{i,:}$ in \mathbf{E} are referred to as a set of *eigenimages* (Zhuang and Bioucas-Dias 2018). In this way, \mathbf{Y} can be decomposed into orthogonal basis vectors $\mathbf{U}_{:,i}$ representing spectral information and corresponding eigenimages $\mathbf{E}_{i,:}$ representing spatial information. Due to the semiunitary property of \mathbf{U} , eigenimages can also be represented as the projection form $\mathbf{E}_{i,:} = (\mathbf{U}_{:,i})^T \mathbf{Y}$.

We rewrite the SVD of \mathbf{Y} in the following form:

$$\mathbf{Y} = \sum_{i=1}^L \sigma_i \mathbf{U}_{:,i} (\mathbf{V}^T)_{i,:} = \sum_{i=1}^L \mathbf{U}_{:,i} \mathbf{E}_{i,:} \approx \sum_{i=1}^p \mathbf{U}_{:,i} \mathbf{E}_{i,:}, \quad (2)$$

where we have $p \ll L$. The rank- p approximation (2) holds for most real-world HSIs (Bioucas-Dias et al. 2012), as they may live in a low-dimensional space. Thus, \mathbf{Y} can be well reconstructed by: $\mathbf{Y} \approx \mathbf{U}_{:,1:p} ((\mathbf{U}_{:,1:p})^T \mathbf{Y})$.

Now, through the above analysis, we derive an efficient method for manipulating the HSI in the eigenimage domain rather than in the original space:

$$\begin{cases} \mathbf{U} = \text{SVD}(\mathbf{Y}), \\ \hat{\mathbf{Y}} = \mathbf{U}_{:,1:p} f_{SR}((\mathbf{U}_{:,1:p})^T \mathbf{Y}), \end{cases} \quad (3)$$

where we use only the \mathbf{U} matrix from the SVD for simplicity; $f_{SR}(\cdot)$ is the SR model that super-resolves the eigenimages. Specifically, we perform SR with the eigenimages channel-by-channel, to align the channel difference between the RGB model and eigenimages.

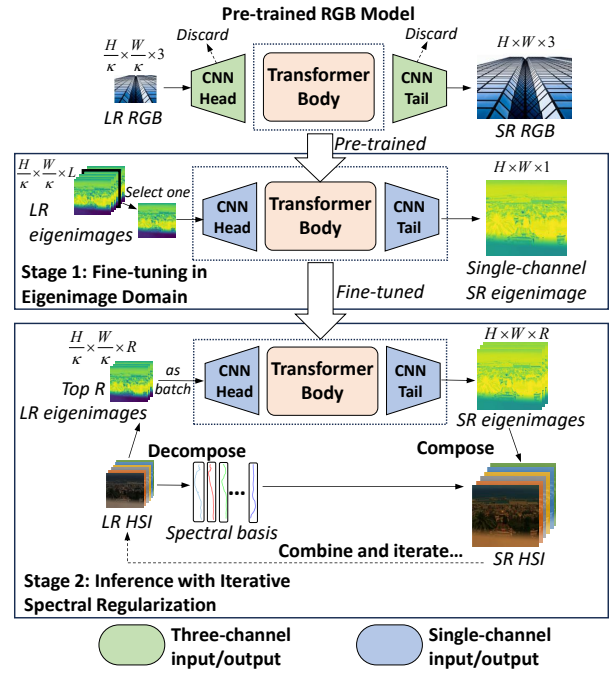


Figure 2: The flowchart of EigenSR. It consists of two stages. Stage 1 uses the pre-trained Transformer Body to fine-tune a single-channel model in the eigenimage domain. Stage 2 utilizes the fine-tuned model for inference with iterative spectral regularization on unseen LR HSI.

The advantages of SR with eigenimages are: 1) thanks to spectral decoupling, the SR model can focus solely on spatial information, while spectral correlations are stored in $\mathbf{U}_{:,1:p}$ to restore the original HSI space. This ensures that the super-resolved HSI maintains spectral fidelity; 2) eigenimages only contain the spatial information, which is more generalizable to spectral-agnostic unseen data; 3) processing p channels is more efficient than L in the original HSI space.

The Proposed EigenSR

Now, we will explain in detail how our EigenSR fully leverages the pre-trained RGB model to process HSI based on eigenimages. Our method consists of two stages, as shown in Figure 2. First, we fine-tune the pre-trained model using eigenimages. Second, when the model is applied to unseen data, we use Eq. (3) to perform SR and introduce an iterative spectral regularization (ISR). By repeatedly performing model inference, updating \mathbf{U} , and taking the weighted average of the SR results, we achieve better spectral fidelity.

Fine-Tuning the Pre-Trained Model in the Eigenimage Domain. Transformer-based architectures generally consist of a Transformer body and a CNN Head and Tail, with the Head and Tail serving as three-channel input and output layers specifically for RGB images. There are two reasons for fine-tuning: 1) to replace the input and output channels of the Head and Tail with one channel to align with the channels of eigenimages and focus solely on processing spatial details, and 2) to adapt to the data in the eigenimage domain.

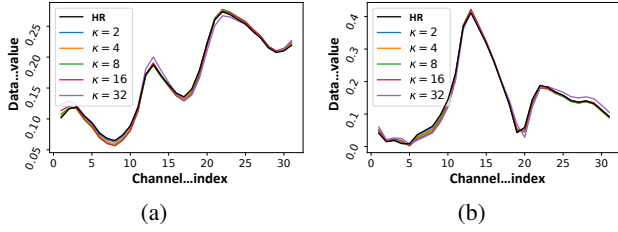


Figure 3: Absolute spectral basis of the CAVE-beads image at different downsampling rates. (a) $|\mathbf{U}_{:,1}|$ and (b) $|\mathbf{U}_{:,2}|$.

We list the key steps of the fine-tuning process below:

Step 1: Prepare in advance the training set of triplets $\{(\mathbf{Y}_{HR}, \mathbf{Y}_{LR}, \mathbf{U}_{HR})_n\}_n$, where each \mathbf{Y}_{HR} is the available HR HSI, $\mathbf{Y}_{LR} = \mathbf{Y}_{HR}\mathbf{D}$ is the generated LR HSI, and $\mathbf{U}_{HR} = \text{SVD}(\mathbf{Y}_{HR})$ is the spectral basis from \mathbf{Y}_{HR} ;

Step 2: In each epoch, randomly select one channel c , and obtain $(\mathbf{E}_{HR})_{c,:} = ((\mathbf{U}_{HR})_{:,c})^\top \mathbf{Y}_{HR}$ and $(\mathbf{E}_{LR})_{c,:} = ((\mathbf{U}_{HR})_{:,c})^\top \mathbf{Y}_{LR}$, by projecting both \mathbf{Y}_{HR} and \mathbf{Y}_{LR} onto the same eigenimage space using $((\mathbf{U}_{HR})_{:,c})^\top$.

Step 3: Feed the model with the LR eigenimage $(\mathbf{E}_{LR})_{c,:}$ and yield the super-resolved eigenimage $(\mathbf{E}_{SR})_{c,:}$, by $(\mathbf{E}_{SR})_{c,:} = f_{SR}((\mathbf{E}_{LR})_{c,:})$.

Step 4: Calculate the loss between $(\mathbf{E}_{SR})_{c,:}$ and $(\mathbf{E}_{HR})_{c,:}$, and perform backward propagation.

In this work, we use the pre-trained Image Processing Transformer (IPT) (Chen et al. 2021) as the pre-trained RGB model. The loss function is the L1 loss. For the settings of fine-tuning, please refer to the Experiments section.

Why Eigenimages Work? Here we illustrate the reason behind the effectiveness of learning eigenimages with the following theorem.

Theorem 1: For eigenimages $\mathbf{E} = \mathbf{U}^\top \mathbf{Y}$, we have:

1. If two column vectors $\mathbf{Y}_{:,i}$ and $\mathbf{Y}_{:,j}$ of \mathbf{Y} are identical, then their corresponding $\mathbf{E}_{:,i}$ and $\mathbf{E}_{:,j}$ are also identical.
2. If \mathbf{Y} is downsampled by a linear operator \mathbf{D} , then the same \mathbf{D} concurrently acts on \mathbf{E} .

Theorem 1 indicates that each channel $\mathbf{E}_{i,:}$ of \mathbf{E} preserves the spatial structure of the original space \mathbf{Y} , and that \mathbf{D} acts on each $\mathbf{E}_{i,:}$. This enables us to fine-tune a single-channel model for \mathbf{E} , which has already been pre-trained on RGB datasets, to learn the inverse mapping of the same degradation \mathbf{D} . The proof of Theorem 1 is given in the Appendix.

Inference with Iterative Spectral Regularization on Unseen Data.

The high-quality spectral basis \mathbf{U}_{HR} from \mathbf{Y}_{HR} can be used during fine-tuning. However, during inference on unseen data, only the LR HSI is available. We found that the dominant components of \mathbf{U} (e.g., $\mathbf{U}_{:,1}$ and $\mathbf{U}_{:,2}$) at different downsampling rates are not affected seriously, as shown in Figure 3. This is because most spectral information in \mathbf{Y}_{LR} is preserved despite the spatial degradation, even in the extreme case where $\kappa = 32$. Therefore, by assuming that the spectral information in \mathbf{Y}_{LR} and \mathbf{Y}_{HR} is close (Yuan,

Algorithm 1: Inference with iterative spectral regularization.

Input: Observed low-resolution HSI \mathbf{Y}_{LR} ; SR model $f_{SR}(\cdot)$; Rank R ; Iteration number N_{it} ; Combination weight constant λ .

Output: The predicted SR HSI $\mathbf{Y}_{comb}^{(N_{it}+1)}$.

- 1: $\mathbf{Y}_{comb}^{(1)} = \mathbf{Y}_{LR}$;
- 2: **for** i **in** $\{1, \dots, N_{it}\}$ **do**
- 3: $\mathbf{U}^{(i)} = \text{SVD}(\mathbf{Y}_{comb}^{(i)})$;
- 4: $\mathbf{E}_{LR}^{(i)} = (\mathbf{U}_{:,1:R}^{(i)})^\top \mathbf{Y}_{LR}$;
The channel-by-channel processing in lines 5-7 can be executed in parallel within a batch.
- 5: **for** j **in** $\{1, \dots, R\}$ **do**
- 6: $(\mathbf{E}_{SR}^{(i)})_{j,:} = f_{SR}((\mathbf{E}_{LR}^{(i)})_{j,:})$;
- 7: **end for**
- 8: $\mathbf{Y}_{SR}^{(i)} = \mathbf{U}_{:,1:R}^{(i)} \mathbf{E}_{SR}^{(i)}$;
- 9: $\mathbf{Y}_{comb}^{(i+1)} = \lambda \mathbf{Y}_{SR}^{(i)} + (1 - \lambda) \mathbf{Y}_{comb}^{(i)}$; # We replace $\mathbf{Y}_{comb}^{(1)}$ here with Bicubic $\uparrow_\kappa(\mathbf{Y}_{comb}^{(1)})$.
- 10: **end for**

Zheng, and Lu 2017), we can extract the spectral information from \mathbf{Y}_{LR} to faithfully reconstruct \mathbf{Y}_{HR} . We thereby achieve an efficient method **EigenSR- α** , through:

$$\begin{cases} \mathbf{U} = \text{SVD}(\mathbf{Y}_{LR}), \\ \mathbf{Y}_{SR} = \mathbf{U}_{:,1:R} f_{SR}((\mathbf{U}_{:,1:R})^\top \mathbf{Y}_{LR}), \end{cases} \quad (4)$$

where \mathbf{Y}_{SR} is the SR HSI; $f_{SR}(\cdot)$ is the model fine-tuned on eigenimages that operates channel-by-channel; R is the rank hyperparameter balancing the reconstruction performance and inference efficiency. Although \mathbf{U} assists to maintain the spectral correlations in the original HSI space, it contains compressed and incomplete spectral information of \mathbf{Y}_{LR} . We further propose a new iterative spectral regularization (ISR) to fully utilize the spectral correlations of \mathbf{Y}_{LR} , by combining the result from the last iteration and then updating a new spectral basis as the following:

$$\begin{cases} \mathbf{Y}_{comb}^{(i+1)} = \lambda \mathbf{Y}_{SR}^{(i)} + (1 - \lambda) \mathbf{Y}_{comb}^{(i)}, \\ \mathbf{U}^{(i+1)} = \text{SVD}(\mathbf{Y}_{comb}^{(i+1)}), \end{cases} \quad (5)$$

where $\lambda \in (0, 1]$ is a weight constant; $\mathbf{Y}_{comb}^{(i)}$ is the combined result starting from \mathbf{Y}_{LR} . The detailed explanation is in Algorithm 1. We denote this version as **EigenSR- β** . Note that EigenSR- β takes more inference steps, but has increased spectral regularization. When we set λ and the iteration number N_{it} to 1, EigenSR- β degenerates to EigenSR- α .

Dataset	Type	#test/train	Sensor	Wavelength	Size
ARAD_1K	HSI	50/900	Specim IQ	400-700nm	482 × 512 × 31
CAVE	HSI	32/-	Apogee Alta U260	400-700nm	512 × 512 × 31
Harvard	HSI	50/-	Nuance FX	420-720nm	512 × 512 × 31
Pavia	HSI	5/-	ROSIS	430-860nm	192 × 192 × 102
DC Mall	HSI	8/-	Hydise	400-2400nm	128 × 128 × 191
Chikusei	HSI	48/-	HH-VNIR-C	363-1018nm	128 × 128 × 128
RESISC45	RGB	-/29250	-	Visible light	256 × 256 × 3

Table 1: Dataset information.

Method	Scale	ARAD_1K-Test			CAVE			Harvard		
		PSNR \uparrow	SSIM \uparrow	SAM \downarrow	PSNR \uparrow	SSIM \uparrow	SAM \downarrow	PSNR \uparrow	SSIM \uparrow	SAM \downarrow
Bicubic	$\times 2$	43.76	0.9824	0.878	40.11	0.9770	2.905	41.05	0.9545	2.314
DHSP (Sidorov and Yngve Hardeberg 2019)	$\times 2$	43.67	0.9804	2.016	42.40	0.9769	4.025	42.45	0.9604	2.496
SSPSR (Jiang et al. 2020)	$\times 2$	48.06	0.9913	0.826	42.95	0.9842	3.217	44.58	0.9750	2.331
SFCSR (Wang, Li, and Li 2020)	$\times 2$	48.70	0.9921	0.656	44.49	0.9865	2.612	45.56	0.9797	2.012
RFSR (Wang, Ma, and Jiang 2022)	$\times 2$	48.25	0.9914	0.894	43.73	0.9855	2.922	45.20	0.9783	2.086
TSBSR (Liang et al. 2023)	$\times 2$	48.12	0.9917	0.658	44.00	0.9860	2.654	43.20	0.9707	2.045
MSDformer (Chen, Zhang, and Zhang 2023)	$\times 2$	47.44	0.9908	1.011	42.53	0.9829	3.457	44.49	0.9748	2.321
ESSAformer (Zhang et al. 2023)	$\times 2$	48.97	0.9927	0.714	43.56	0.9857	2.909	44.93	0.9767	2.265
DKP (Yang et al. 2024)	$\times 2$	44.16	0.9827	1.957	42.96	0.9770	4.070	42.70	0.9650	2.599
EigenSR-α (Ours)	$\times 2$	49.09	0.9930	0.607	44.67	0.9869	2.609	45.65	0.9805	2.006
EigenSR-β (Ours)	$\times 2$	49.11	0.9930	0.600	44.97	0.9874	2.448	45.77	0.9819	1.924
Bicubic	$\times 4$	37.08	0.9280	1.641	34.38	0.9312	4.208	36.12	0.8811	2.741
DHSP (Sidorov and Yngve Hardeberg 2019)	$\times 4$	35.75	0.9257	1.707	36.26	0.9379	6.052	37.15	0.8911	2.922
SSPSR (Jiang et al. 2020)	$\times 4$	39.87	0.9524	1.713	37.11	0.9538	4.505	39.84	0.9287	2.716
SFCSR (Wang, Li, and Li 2020)	$\times 4$	40.06	0.9535	1.511	37.65	0.9562	3.989	39.97	0.9296	2.649
RFSR (Wang, Ma, and Jiang 2022)	$\times 4$	39.91	0.9526	1.823	37.39	0.9551	4.474	39.97	0.9299	2.665
TSBSR (Liang et al. 2023)	$\times 4$	39.81	0.9548	1.280	37.77	0.9583	3.715	38.31	0.9134	2.593
MSDformer (Chen, Zhang, and Zhang 2023)	$\times 4$	39.00	0.9437	2.134	36.03	0.9376	6.186	38.78	0.9135	2.893
ESSAformer (Zhang et al. 2023)	$\times 4$	40.40	0.9565	1.384	37.48	0.9576	4.144	40.04	0.9309	2.661
DKP (Yang et al. 2024)	$\times 4$	38.27	0.9343	1.824	36.48	0.9401	6.077	37.66	0.8999	3.005
EigenSR-α (Ours)	$\times 4$	40.28	0.9602	1.223	38.29	0.9619	3.576	40.18	0.9313	2.543
EigenSR-β (Ours)	$\times 4$	40.46	0.9605	1.182	38.36	0.9624	3.480	40.21	0.9315	2.524
Bicubic	$\times 8$	32.66	0.8453	2.613	30.07	0.8563	5.756	32.53	0.7932	3.222
DHSP (Sidorov and Yngve Hardeberg 2019)	$\times 8$	33.08	0.8409	5.083	31.95	0.8700	8.210	34.91	0.8501	3.343
SSPSR (Jiang et al. 2020)	$\times 8$	34.60	0.8770	2.646	32.60	0.8917	6.120	36.18	0.8616	3.090
SFCSR (Wang, Li, and Li 2020)	$\times 8$	34.59	0.8760	2.484	32.75	0.8926	5.514	36.21	0.8613	3.052
RFSR (Wang, Ma, and Jiang 2022)	$\times 8$	34.51	0.8738	3.222	32.53	0.8894	7.170	36.17	0.8607	3.192
TSBSR (Liang et al. 2023)	$\times 8$	34.55	0.8810	2.310	32.95	0.9000	5.158	34.43	0.8309	3.087
MSDformer (Chen, Zhang, and Zhang 2023)	$\times 8$	34.26	0.8705	3.018	31.04	0.8476	8.851	35.11	0.8415	3.750
ESSAformer (Zhang et al. 2023)	$\times 8$	34.71	0.8831	2.363	33.00	0.9005	5.670	36.27	0.8651	3.043
DKP (Yang et al. 2024)	$\times 8$	33.69	0.8517	5.136	32.23	0.8769	8.424	34.52	0.8479	3.563
EigenSR-α (Ours)	$\times 8$	34.62	0.8814	2.140	33.08	0.9006	5.195	36.31	0.8689	3.009
EigenSR-β (Ours)	$\times 8$	34.75	0.8835	2.086	33.13	0.9008	4.945	36.37	0.8696	2.975

Table 2: Quantitative comparison with the SOTA methods with reference image. ARAD_1K-Train is used for training. The best results are marked in **bold**, and the second and third are underlined.

Experiments

Experimental Settings

Datasets. We list the dataset information in Table 1. The ARAD_1K (He et al. 2023), CAVE (Yasuma et al. 2010), and Harvard (Chakrabarti and Zickler 2011) datasets contain high-quality HSIs of indoor and outdoor scenes, each with 31 spectral bands, collected using three different sensors. The ARAD_1K dataset includes 950 HSIs, with the first 900 used as the training set (denoted as ARAD_1K-Train) and the remaining 50 used as the test set (denoted as ARAD_1K-Test). The CAVE and Harvard datasets contain only a few dozen images, and we used them as test sets.

The remote sensing datasets Pavia², DC Mall³, and Chikusei (Yokoya and Iwasaki 2016) contain more than 100 channels, but have low spatial resolutions. The RESISC45⁴ includes HR RGB images of remote sensing scenes. We used the RESISC45 dataset as the training set, and the Pavia, DC Mall, and Chikusei datasets as test sets to conduct blind testing on the three datasets without the HR reference image.

Metrics. We used the peak signal-to-noise ratio (PSNR) and structural similarity (SSIM) to evaluate the spatial quality of SR HSI, and the spectral angle mapper (SAM) for the spectral quality. When testing without reference image, a no-reference quality measurement score (NRQMS) (Yang et al. 2017) was used to jointly assess the spectral-spatial quality.

Comparative Methods. Several methods were compared, including three unsupervised methods Bicubic, DHSP

(Sidorov and Yngve Hardeberg 2019), and DKP (Yang et al. 2024), and six supervised methods SSPSR (Jiang et al. 2020), SFCSR (Wang, Li, and Li 2020), RFSR (Wang, Ma, and Jiang 2022), TSBSR (Liang et al. 2023), MSDformer (Chen, Zhang, and Zhang 2023), and ESSAformer (Zhang et al. 2023). Among them, TSBSR transfers RGB models to HSI and uses NMF as the spectral constraint. Since it only has remote sensing models, we replaced its model with the fine-tuned IPT on original HSIs to adapt to various scenes.

Implementation Details. We used two NVIDIA 3090 cards to run the algorithms in PyTorch. We adopted the low-rank adaptation (LoRA) (Hu et al. 2021) for parameter efficient fine-tuning, applying LoRA parameters in parallel with all q and v parameters in the pre-trained Transformer (Zhou et al. 2024b), and set the low-rank hyperparameter r to 4 as suggested in the original paper. Thus, our trainable parameters consist only of LoRA and the Head and Tail. We fine-tuned on the ARAD_1K-Train dataset for 2000 epochs and on the RESISC45 dataset for 50 epochs, considering the number of training samples. We used the Adam optimizer with a learning rate of 0.001 and the batch size was 64.

When testing on unseen data, we set R to 50% of the number of channels L . For EigenSR- β , the number of iterations N_{it} was set to 5, and the constant λ was empirically set to 0.8 for SR $\times 2$, and 0.4 for SR $\times 4$ and $\times 8$, respectively.

Comparisons with the SOTA Methods

Quantitative Comparisons. First, we conducted comparisons for SR at $\times 2$, $\times 4$, and $\times 8$ scales with HR reference images. The LR images were obtained by downsampling the HR reference images. The supervised methods were

²https://ehu.eus/ccwintco/index.php?title=Hyperspectral_Remote_Sensing_Scenes

³<http://lesun.weebly.com/hyperspectral-data-set.html>

⁴<https://tensorflow.google.cn/datasets/catalog/resisc45>

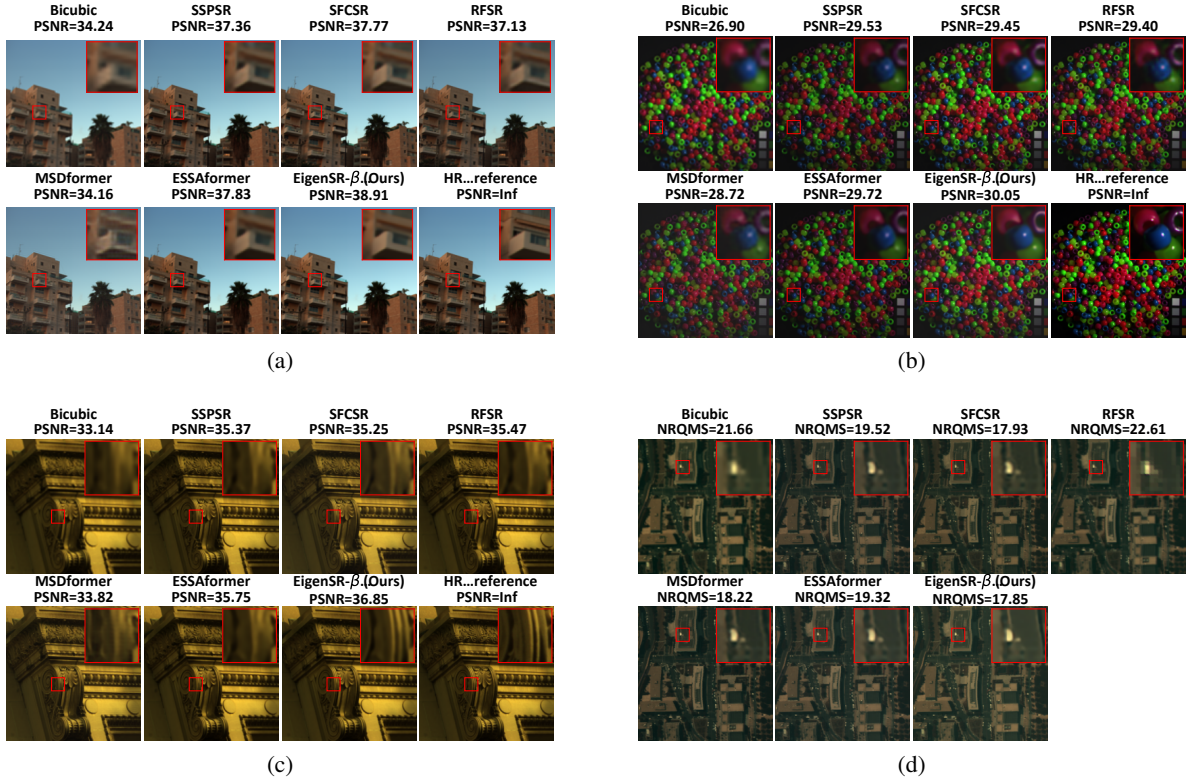


Figure 4: Visual results of SR $\times 4$. (a) ARAD_1K-Test. (b) CAVE. (c) Harvard. (d) DC Mall. We used band numbers (27, 14, 6) to generate the pseudo-color images for ARAD_1K, CAVE, and Harvard datasets, and band numbers (32, 14, 7) for DC Mall.

trained on ARAD_1K-Train, and all methods were tested on ARAD_1K-Test, CAVE, and Harvard datasets. We used the PSNR, SSIM, and SAM to evaluate the SR quality. The results are shown in Table 2. Our EigenSR- α and EigenSR- β models achieved second and first rankings in both spatial and spectral metrics, respectively, in almost all cases. Since our models were fine-tuned on ARAD_1K-Train, they also performed well on the CAVE and Harvard datasets, which were captured by different sensors, demonstrating their generalization ability to spectral-agnostic unseen data.

Second, we conducted blind SR tests at scale $\times 4$ without reference images on the LR remote sensing datasets Pavia, DC Mall, and Chikusei. We trained on the HR remote sensing RGB dataset RESISC45, due to the lack of HR remote sensing HSI for training. For those HSI-specific methods, we trained a 3-channel model and performed SR on the test HSI by processing adjacent three channels at a time. In contrast, since our single-channel model does not face channel alignment issues, we directly fine-tuned on the eigenimages of RGB images and tested on the eigenimages of HSI. We used the no-reference metric NRQMS to jointly evaluate the spatial and spectral quality, and recorded the inference time (in seconds), as shown in Table 3. Our models can provide SR HSI with better spatial and spectral quality. Meanwhile, EigenSR- α consumes less inference time due to its efficiency in the eigenimage domain, while EigenSR- β still

Method	Pavia		DC Mall		Chikusei	
	NRQMS \downarrow	Time(s)	NRQMS \downarrow	Time(s)	NRQMS \downarrow	Time(s)
Bicubic	22.42	-	22.26	-	24.70	-
DHSP (2019)	19.78	949.44	18.93	724.10	22.84	482.28
SSPSR (2020)	19.55	5.24	19.34	13.48	21.57	9.83
SFCRS (2020)	19.19	5.30	18.39	23.76	20.75	17.21
RFSR (2022)	22.90	3.66	22.80	10.30	24.70	7.76
TSBSR (2023)	19.83	211.86	19.43	159.28	21.22	107.91
MSDformer (2023)	19.33	8.88	18.33	28.52	20.72	20.31
ESSAformer (2023)	20.04	21.34	19.65	64.92	22.12	44.50
DKP (2024)	20.05	1205.31	19.89	886.52	22.98	450.40
EigenSR- α (Ours)	19.17	4.79	18.05	4.76	20.58	3.36
EigenSR- β (Ours)	19.13	24.54	17.98	24.36	20.42	17.16

Table 3: Quantitative comparison without reference image at scale $\times 4$. RESISC45 is used for training. The best results are marked in **bold**, and the second and third are underlined.

keeps an acceptable inference time among other methods.

Qualitative Comparisons. We visually present several SR test results from a spatial view, in Figure 4. Overall, our method can provide clearer and sharper SR results. For example, the stripes in the Harvard dataset and the small white target in the DC Mall dataset. This is because our method effectively introduces the model pre-trained on a large amount of RGB data, addressing the HSI data bottleneck, and enhancing spatial processing capabilities.

On the other hand, we compared the spectra of SR HSI in Figure 5. It can be seen that our method provides spec-

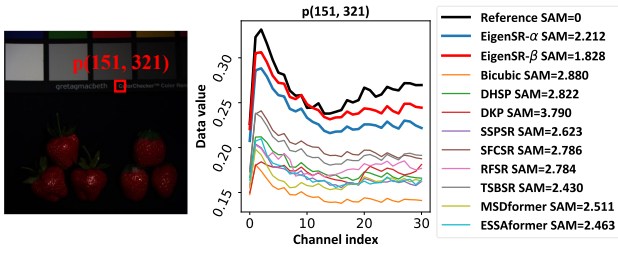


Figure 5: SR $\times 4$ comparison in the spectral dimension on CAVE dataset. We denote the spatial coordinate as $p(151, 321)$. Our method can provide more accurate spectral curves.

Setting (explanation)	PT	NMF	Eigen	ISR	PSNR \uparrow	SSIM \uparrow	SAM \downarrow	Time(s)
① (IPT baseline)	✓				39.96	0.9556	1.268	3.19
② (IPT w/ NMF)	✓	✓			39.81	0.9548	1.280	36.17
③ (EigenSR- α)	✓		✓		40.28	0.9602	1.223	1.13
④ (EigenSR- α w/o PT)			✓		39.74	0.9538	1.310	1.13
⑤ (EigenSR- β)	✓		✓	✓	40.46	0.9605	1.182	4.71

Table 4: Break-down ablation to improve the performance.

How to fine-tune	PSNR \uparrow	SSIM \uparrow	SAM \downarrow	#trainable params. (M)
Only Head & Tail	39.92	0.9551	1.250	0.78
Head & Tail + LoRA	40.46	0.9605	1.182	1.11
Full parameter	40.60	0.9610	1.175	115.68

Table 5: Different fine-tuning architectures.

tral curves closer to the reference spectral curve. This is because we preserved the spectral information during the SR of eigenimages and recovered it during reconstruction, and further constrained it through ISR. These resulted in more accurate spectral estimates.

Ablation Studies

We conducted ablation studies on the ARAD_1K dataset. The models were trained on ARAD_1K-Train and tested on ARAD_1K-Test at a scale factor of $\times 4$.

Break-down Ablation. We validated each component of our method by breaking them down, as shown in Table 4. We set the pre-trained (PT) model IPT followed by fine-tuning on HSIs, as the baseline (①). The additional NMF spectral constraint (Liang et al. 2023) was applied during inference (②). Differently, EigenSR- α (③) fine-tuned the pre-trained model in the eigenimage domain (Eigen). We removed the pre-trained weights (④) from EigenSR- α and trained the model from scratch instead of fine-tuning. The ISR was further added to EigenSR- α as the EigenSR- β (⑤).

It can be found that NMF leads to a performance decrease compared to the baseline while incurring a non-negligible time cost during inference. This is probably because the NMF optimization has the non-convex nature, making it difficult to adapt to various situations. In contrast, EigenSR- α outperforms both ① and ② while maintaining efficient inference, which validates its effectiveness of learning the eigenimages. Leveraging the RGB pre-trained weights results in an improvement of ③ over ④. Moreover, despite introducing a longer but still acceptable inference time, EigenSR- β

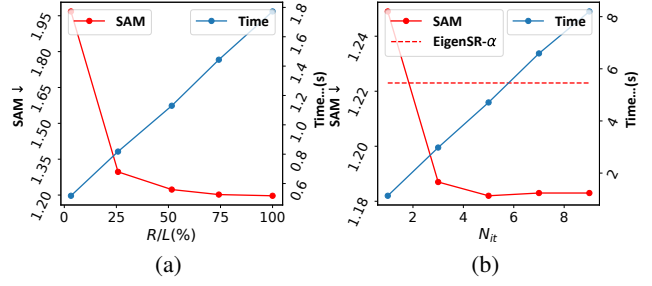


Figure 6: Effects of inference hyperparameters. (a) Rank R . (b) Iteration number N_{it} ($\lambda = 0.4$ for EigenSR- β).

further outperforms EigenSR- α in terms of performance.

Fine-Tuning Architectures. We studied three different fine-tuning architectures in Table 5: only fine-tuning the Head & Tail, fine-tuning the Head & Tail along with the LoRA embedded in the Transformer, and full parameter fine-tuning. It can be observed that full parameter fine-tuning leads to the best performance but imposes a heavy burden in terms of parameters. Only using the Head & Tail for fine-tuning does not achieve optimal performance. Incorporating LoRA allows the Transformer to better adapt to the eigenimage domain while keeping parameter efficiency.

Inference Hyperparameters. The rank R and the iteration number N_{it} balance the SR performance and time costs. Figure 6 studies the effects of the two hyperparameters. It can be observed that as the two hyperparameters increase, the performance gradually improves and eventually saturates, while the inference time grows linearly. Considering the balance between performance and time, we suggest a reasonable configuration: $R/L = 50\%$ and $N_{it} = 5$.

Conclusion

This work presents EigenSR for single-HSI-SR, a method that fully leverages the spatial capabilities of pre-trained RGB models while maintaining spectral correlations, to address the data bottleneck of HSI. It is based on the concept of spatial-spectral decoupling via SVD and consists of two stages. First, we fine-tune a single-channel model on the spatial components, i.e., the eigenimages, allowing the RGB model channels to align with HSI and adapt to the eigenimage domain. Then, during inference on unseen data, we decompose the LR HSI, perform channel-wise SR on its eigenimages, and reconstruct the SR HSI using the spectral components, i.e., the spectral basis, to maintain spectral correlations. Furthermore, we propose ISR to fully exploit the LR HSI and enhance spectral fidelity. We only introduce the pre-trained model in the spectral-agnostic spatial components, which enhances the model’s generalization ability. Additionally, by utilizing a small number of low-rank eigenimages, we reduce the inference complexity. Extensive experiments demonstrate that our method outperforms SOTA approaches, proving the effectiveness of introducing pre-trained RGB models to address the data scarcity in HSI.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (U2133211, 62301092, 62301093). Xi Su expresses special gratitude to Dr. Shuanghao Bai for providing valuable feedback on the writing of this paper.

References

- Agustsson, E.; and Timofte, R. 2017. NTIRE 2017 Challenge on Single Image Super-Resolution: Dataset and Study. In *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2017, Honolulu, HI, USA, July 21-26, 2017*, 1122–1131. IEEE Computer Society.
- Bai, S.; Zhang, M.; Zhou, W.; Huang, S.; Luan, Z.; Wang, D.; and Chen, B. 2024. Prompt-based distribution alignment for unsupervised domain adaptation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 729–737.
- Bai, S.; Zhang, Y.; Zhou, W.; Luan, Z.; and Chen, B. 2025. Soft prompt generation for domain generalization. In *European Conference on Computer Vision*, 434–450. Springer.
- Bioucas-Dias, J. M.; Plaza, A.; Camps-Valls, G.; Scheunders, P.; Nasrabadi, N.; and Chanussot, J. 2013. Hyperspectral remote sensing data analysis and future challenges. *IEEE Geosci. Remote Sens. Mag.*, 1(2): 6–36.
- Bioucas-Dias, J. M.; Plaza, A.; Dobigeon, N.; Parente, M.; Du, Q.; Gader, P.; and Chanussot, J. 2012. Hyperspectral unmixing overview: Geometrical, statistical, and sparse regression-based approaches. *IEEE journal of selected topics in applied earth observations and remote sensing*, 5(2): 354–379.
- Chakrabarti, A.; and Zickler, T. 2011. Statistics of real-world hyperspectral images. In *CVPR 2011*, 193–200. IEEE.
- Chen, H.; Wang, Y.; Guo, T.; Xu, C.; Deng, Y.; Liu, Z.; Ma, S.; Xu, C.; Xu, C.; and Gao, W. 2021. Pre-trained image processing transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12299–12310.
- Chen, S.; Zhang, L.; and Zhang, L. 2023. MSDformer: Multi-scale Deformable Transformer for Hyperspectral Image Super-Resolution. *IEEE Transactions on Geoscience and Remote Sensing*.
- Chen, X.; Wang, X.; Zhou, J.; Qiao, Y.; and Dong, C. 2023. Activating more pixels in image super-resolution transformer. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 22367–22377.
- Chen, Z.; Zhang, Y.; Gu, J.; Kong, L.; Yuan, X.; et al. 2022. Cross aggregation transformer for image restoration. *Advances in Neural Information Processing Systems*, 35: 25478–25490.
- Dabov, K.; Foi, A.; Katkovnik, V.; and Egiazarian, K. 2007. Image denoising by sparse 3-D transform-domain collaborative filtering. *IEEE Transactions on image processing*, 16(8): 2080–2095.
- Deng, J.; Dong, W.; Socher, R.; Li, L.; Li, K.; and Fei-Fei, L. 2009. ImageNet: A large-scale hierarchical image database. In *2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2009), 20-25 June 2009, Miami, Florida, USA*, 248–255. IEEE Computer Society.
- Dong, C.; Loy, C. C.; He, K.; and Tang, X. 2015. Image super-resolution using deep convolutional networks. *IEEE transactions on pattern analysis and machine intelligence*, 38(2): 295–307.
- Dong, W.; Zhou, C.; Wu, F.; Wu, J.; Shi, G.; and Li, X. 2021. Model-Guided Deep Hyperspectral Image Super-Resolution. *IEEE Trans. Image Process.*, 30: 5754–5768.
- He, J.; Yuan, Q.; Li, J.; Xiao, Y.; Liu, D.; Shen, H.; and Zhang, L. 2023. Spectral super-resolution meets deep learning: Achievements and challenges. *Information Fusion*, 101812.
- Hong, D.; He, W.; Yokoya, N.; Yao, J.; Gao, L.; Zhang, L.; Chanussot, J.; and Zhu, X. 2021. Interpretable hyperspectral artificial intelligence: When nonconvex modeling meets hyperspectral remote sensing. *IEEE Geosci. Remote Sens. Mag.*, 9(2): 52–87.
- Hsu, C.-C.; Lee, C.-M.; and Chou, Y.-S. 2024. DRCT: Saving Image Super-Resolution Away from Information Bottleneck. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 6133–6142.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Jiang, J.; Sun, H.; Liu, X.; and Ma, J. 2020. Learning spatial-spectral prior for super-resolution of hyperspectral imagery. *IEEE Transactions on Computational Imaging*, 6: 1082–1096.
- Keshava, N.; and Mustard, J. F. 2002. Spectral unmixing. *IEEE Signal Process. Mag.*, 19(1): 44–57.
- Ledig, C.; Theis, L.; Huszár, F.; Caballero, J.; Cunningham, A.; Acosta, A.; Aitken, A.; Tejani, A.; Totz, J.; Wang, Z.; et al. 2017. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4681–4690.
- Li, Q.; Wang, Q.; and Li, X. 2020. Mixed 2D/3D convolutional network for hyperspectral image super-resolution. *Remote sensing*, 12(10): 1660.
- Li, W.; Lu, X.; Qian, S.; Lu, J.; Zhang, X.; and Jia, J. 2021. On efficient transformer-based image pre-training for low-level vision. *arXiv preprint arXiv:2112.10175*.
- Liang, J.; Cao, J.; Sun, G.; Zhang, K.; Van Gool, L.; and Timofte, R. 2021. Swinir: Image restoration using swin transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, 1833–1844.
- Liang, Z.; Wang, S.; Zhang, T.; and Fu, Y. 2023. Blind Super-Resolution of Single Remotely Sensed Hyperspectral Image. *IEEE Trans. Geosci. Remote. Sens.*, 61: 1–14.
- Liu, Z.; Wang, X.; Zhong, Y.; Shu, M.; and Sun, C. 2022. SiamHYPER: Learning a Hyperspectral Object Tracker From an RGB-Based Tracker. *IEEE Transactions on Image Processing*, 31: 7116–7129.

- Miao, Y.-C.; Zhao, X.-L.; Fu, X.; Wang, J.-L.; and Zheng, Y.-B. 2022. Hyperspectral Denoising Using Unsupervised Disentangled Spatiospectral Deep Priors. *IEEE Transactions on Geoscience and Remote Sensing*, 60: 1–16.
- Nie, J.; Xie, J.; Cao, J.; and Pang, Y. 2023. Context and detail interaction network for stereo rain streak and raindrop removal. *Neural Networks*, 166: 215–224.
- Nie, J.; Xie, J.; and Sun, H. 2024. Remote Sensing Image Dehazing via a Local Context-Enriched Transformer. *Remote Sensing*, 16(8): 1422.
- Peng, J.; Wang, H.; Cao, X.; Zhao, Q.; Yao, J.; Zhang, H.; and Meng, D. 2024. Learnable Representative Coefficient Image Denoiser for Hyperspectral Image. *IEEE Transactions on Geoscience and Remote Sensing*, 62: 1–16.
- Peng, Y.; Li, W.; Luo, X.; and Du, J. 2021. Hyperspectral image superresolution using global gradient sparse and non-local low-rank tensor decomposition with hyper-laplacian prior. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 14: 5453–5469.
- Sidorov, O.; and Yngve Hardeberg, J. 2019. Deep hyperspectral prior: Single-image denoising, inpainting, super-resolution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, 0–0.
- Timofte, R.; Agustsson, E.; Van Gool, L.; Yang, M.-H.; and Zhang, L. 2017. Ntire 2017 challenge on single image super-resolution: Methods and results. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 114–125.
- Vivone, G. 2023. Multispectral and hyperspectral image fusion in remote sensing: A survey. *Information Fusion*, 89: 405–417.
- Wan, W.; Zhang, B.; Vella, M.; Mota, J. F. C.; and Chen, W. 2022. Robust RGB-Guided Super-Resolution of Hyperspectral Images via TV^3 Minimization. *IEEE Signal Process. Lett.*, 29: 957–961.
- Wang, Q.; Li, Q.; and Li, X. 2020. Hyperspectral image superresolution using spectrum and feature context. *IEEE Transactions on Industrial Electronics*, 68(11): 11276–11285.
- Wang, X.; Ma, J.; and Jiang, J. 2022. Hyperspectral Image Super-Resolution via Recurrent Feedback Embedding and Spatial-Spectral Consistency Regularization. *IEEE Transactions on Geoscience and Remote Sensing*, 60: 1–13.
- 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.
- Xie, J.; Nie, J.; Ding, B.; Yu, M.; and Cao, J. 2023. Cross-modal Local Calibration and Global Context Modeling Network for RGB-Infrared Remote Sensing Object Detection. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*.
- Xu, T.; Huang, T.; Deng, L.; and Yokoya, N. 2022. An Iterative Regularization Method Based on Tensor Subspace Representation for Hyperspectral Image Super-Resolution. *IEEE Trans. Geosci. Remote. Sens.*, 60: 1–16.
- Xu, T.; Huang, T.; Deng, L.; Zhao, X.; and Huang, J. 2020. Hyperspectral Image Superresolution Using Unidirectional Total Variation With Tucker Decomposition. *IEEE J. Sel. Top. Appl. Earth Obs. Remote. Sens.*, 13: 4381–4398.
- Yang, J.; Zhao, Y.-Q.; Yi, C.; and Chan, J. C.-W. 2017. No-reference hyperspectral image quality assessment via quality-sensitive features learning. *Remote Sensing*, 9(4): 305.
- Yang, Z.; Xia, J.; Li, S.; Huang, X.; Zhang, S.; Liu, Z.; Fu, Y.; and Liu, Y. 2024. A Dynamic Kernel Prior Model for Unsupervised Blind Image Super-Resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 26046–26056.
- 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.
- Yokoya, N.; Grohnfeldt, C.; and Chanussot, J. 2017. Hyperspectral and multispectral data fusion: A comparative review of the recent literature. *IEEE Geosci. Remote Sens. Mag.*, 5(2): 29–56.
- 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: 5.
- Yuan, Y.; Zheng, X.; and Lu, X. 2017. Hyperspectral Image Superresolution by Transfer Learning. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 10(5): 1963–1974.
- Zhang, M.; Zhang, C.; Zhang, Q.; Guo, J.; Gao, X.; and Zhang, J. 2023. Essformer: Efficient transformer for hyperspectral image super-resolution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 23073–23084.
- Zhang, Y.; Li, K.; Li, K.; Wang, L.; Zhong, B.; and Fu, Y. 2018. Image Super-Resolution Using Very Deep Residual Channel Attention Networks. In Ferrari, V.; Hebert, M.; Sminchisescu, C.; and Weiss, Y., eds., *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part VII*, volume 11211 of *Lecture Notes in Computer Science*, 294–310. Springer.
- Zhou, W.; Bai, S.; Zhao, Q.; and Chen, B. 2024a. Revisiting the Adversarial Robustness of Vision Language Models: a Multimodal Perspective. *arXiv preprint arXiv:2404.19287*.
- Zhou, X.; Liang, F.; Chen, L.; Liu, H.; Song, Q.; Vivone, G.; and Chanussot, J. 2024b. MeSAM: Multiscale Enhanced Segment Anything Model for Optical Remote Sensing Images. *IEEE Transactions on Geoscience and Remote Sensing*.
- Zhuang, L.; and Bioucas-Dias, J. M. 2018. Fast Hyperspectral Image Denoising and Inpainting Based on Low-Rank and Sparse Representations. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 11(3): 730–742.
- Zhuang, L.; and Ng, M. K. 2023. FastHyMix: Fast and Parameter-Free Hyperspectral Image Mixed Noise Removal. *IEEE Transactions on Neural Networks and Learning Systems*, 34(8): 4702–4716.