


000 UNSUPERVISED MULTI-SENSOR SPECTRAL IMAGE 001 FUSION VIA FREQUENCY-SPATIAL RECIPROCAL- 002 VIEW LEARNING

006 **Anonymous authors**

007 Paper under double-blind review

010 ABSTRACT

013 The unsupervised fusion of multi-sensor spectral images is often limited by non-
014 absolute registration. This misalignment leads to significant differences between
015 the spectral shape of the fused image and the original hyperspectral signal. To ad-
016 dress this challenge, we propose the Frequency-Spatial Reciprocal-View Learning
017 (FSRVL) for unsupervised multi-sensor MSI-HSI fusion. 1) Feature Synchroni-
018 zation: weight-sharing convolutions are employed to process Low-Resolution
019 Hyperspectral Image (LR-HSI) and High-Resolution Multispectral Image (HR-
020 MSI) in the frequency domain jointly, achieving an information correspondence
021 between the two modalities with parameter transfer. 2) Frequency Recalibration:
022 sub-pixel information is assigned to learnable filters to adaptively refine spatial
023 distributions across various materials and promote the reestablishment of their
024 spectral characteristics. The advantages of FSRVL were demonstrated across var-
025 ious simulated and real-world scenarios, with experiments confirming that FSRVL
026 outperforms the baselines. *The source code will be linked here.*



037 Figure 1: When using traditional deep learning methods for unsupervised MSI-HSI fusion, it typ-
038 ically encounters the information mismatch between sensors. Our FSRVL can effectively increase
039 matched features to mitigate this issue.

041 1 INTRODUCTION

044 The unsupervised fusion of Hyperspectral Image (HSI) and Multispectral Image (MSI) is typically
045 guided by spectral unmixing, wherein the Low-Resolution Hyperspectral Image (LR-HSI) serves as
046 the source of Endmember Signal Dictionary (ESD), and the High-Resolution Multispectral Image
047 (HR-MSI) provides the Corresponding Subpixel Abundances (CSAs) (Yang et al., 2024a). A High-
048 Resolution Hyperspectral Image (HR-HSI) is then reconstructed by combining the two (Xie et al.,
049 2019). Despite the success of both deep learning-based and handcrafted methods, their practicality is
050 fundamentally limited by the assumption that ideal correspondence exists between LR-HSI's ESD
051 and HR-MSI's CSAs in the spatial domain. This assumption is rarely satisfied in multi-source
052 scenarios, severely hindering the quality and reliability of the fusion data (Qu et al., 2025).

053 Figure 1 highlights the key bottleneck in current deep learning-based unsupervised MSI-HSI fusion,
054 where most methods rely on the ideal assumption that endmember signals and abundance fractions

054 from HSI and MSI correspond perfectly. For example, the i -th endmember signal in the ESD of the
 055 HSI is assumed to represent the same material as the i -th endmember signal in the ESD of MSI,
 056 and that the abundance vectors obtained from both datasets are also corresponding. However, in
 057 practice, the MSI and HSI do not always correspond ideally, which might lead to low-quality fusion
 058 performance (Dong et al., 2024). Therefore, we develop the Frequency-Spatial Reciprocal-View
 059 Learning (FSRVL) to address the issue of information mismatch for unsupervised MSI-HSI fusion.
 060 Figure 1 adopts a CNN to acquire the results of classic deep learning.

061 HSIs are data cubes that contain dozens to hundreds of spectral bands (Zhang et al., 2024). Un-
 062 like other types of images, HSIs provide more detailed spectral data, which can reflect the intrinsic
 063 characteristics of various objects (Yang et al., 2024b). However, due to the limitations of imaging
 064 hardware, HSIs have to face a trade-off between spectral resolution and spatial resolution (Zeng
 065 et al., 2025; Wang et al., 2024). Specifically, increasing the spectral resolution leads to a decrease
 066 in spatial resolution, while enhancing spatial resolution causes a reduction in spectral resolution (Su
 067 et al., 2025). To improve the spatial resolution of a single image, one must either use more advanced
 068 sensors or process the image by supervised learning (Li et al., 2024). The former comes with enor-
 069 mous costs, while the latter relies on prior knowledge and some degree of manual expertise (Zhang
 070 et al., 2020; Qu et al., 2024). Conversely, MSIs have fewer spectral channels compared to HSIs, but
 071 offer higher spatial resolution and are generally more cost-effective to obtain (Zhu et al., 2025).

072 Nowadays, the fusion of an LR-HSI and an HR-MSI can reconstruct a new HR-HSI, which can
 073 achieve HSI super-resolution (Liu et al., 2022; Gao et al., 2023). In response to this demand, nu-
 074 merous HSI-MSI fusion methods have been developed, such as DFMff (Guo et al., 2023), SSG-
 075 SRL (Liu et al., 2025), and OTIAS (Deng et al., 2025). However, these methods are supervised,
 076 meaning their fusion performance heavily relies on manual sampling, limiting their scalability.
 077 Conversely, unsupervised HSI-MSI fusion is more capable of meeting the user demand for high-
 078 resolution HSI (HR-HSI). This has led to the development of several representative methods, in-
 079 cluding CUCaNet (Yao et al., 2020), DTDNML (Wang et al., 2025), and EDIP (Li et al., 2025).
 080 Currently, most unsupervised fusion approaches assume that the MSI and HSI imaging records un-
 081 der the same scene are identical, and they use spectral mixing models as the linkage between the two
 082 modalities. This manner preserves the spectral information from LR-HSI and the spatial details of
 083 HR-MSI. Note that this requires a perfect correspondence between the CSAs of MSI and the ESD
 084 of HSI. However, this ideal registration assumption is problematic to meet in practical multi-source
 085 data fusion, as distortions cannot be fully eliminated, inevitably leading to issues of information
 mismatch between the modalities.

086 To overcome the technical bottleneck for unsupervised multi-sensor fusion, we design a new ap-
 087 proach, named FSRVL. This novel deep learning framework transforms MSI and HSI into the fre-
 088 quency domain, enabling adaptive fusion by applying frequency filtering. Specifically, we adapt the
 089 discrete Fourier transform to process CSAs corresponding to ESD for both MSI and HSI modal-
 090 ties, where spatial details are encoded in the phase components and global structure is encoded
 091 in the amplitude. FSRVL can adaptively harmonize the CSAs in hidden layers for both MSI and
 092 HSI modalities, thereby improving information matching. ESDs of the two modes correspond to
 093 the weights of abundance features, and they will also become increasingly aligned during learning.
 094 Although this serves as an illustration, all pictures are derived from real experiments. As shown in
 095 Figure 1, FSRVL significantly increases the number of information matches between LR-HSI and
 096 HR-MSI, thereby substantially enhancing the fusion performance.

097 The main contributions of this paper are summarized as:

- 099 • FSRVL introduces the idea of deep frequency filtering to address the misalignment of
 100 modal information within the same scene during unsupervised MSI-HSI fusion, thereby
 101 enhancing performance.
- 102 • We analyze the issue of mismatch between the Endmember Signal Dictionary (ESD) of
 103 LR-HSI and the Corresponding Sub-pixel Abundances (CSAs) of HR-MSI, providing a
 104 novel direction for enhancing the effectiveness of unsupervised MSI-HSI fusion.
- 105 • This paper not only used a simulated scenario with Ground Truth to quantitatively evaluate
 106 FSRVL, but also validated it using unregistered multi-source datasets.

108 **2 RELATED WORKS**

109
110 **2.1 UNSUPERVISED MSI-HSI FUSION**

111 This paper defines the HR-MSI and LR-HSI using the notations $X \in \mathbb{R}^{WH \times g}$ and $Y \in \mathbb{R}^{wh \times G}$, re-
112 spectively. Here, (W, H) and (w, h) represent the spatial sizes of the HR-MSI and LR-HSI datasets,
113 and g and G denote the number of bands corresponding to the HR-MSI and LR-HSI, respectively.
114 Generally, they also should have the following relationship: $W \gg w$, $H \gg h$, and $G \gg g$ (Li et al.,
115 2025). When the Linear Spectral Mixing Model (LSMM) is employed, HR-MSI and LR-HSI can
116 be expressed as:
117

$$X = \mathbf{A}\hat{\mathbf{E}}, \quad Y = \hat{\mathbf{A}}\mathbf{E} \quad (1)$$

118 where $\hat{\mathbf{E}} \in \mathbb{R}^{e \times g}$ and $\mathbf{E} \in \mathbb{R}^{e \times G}$ denote ESDs of X and Y , respectively. In Eq. (1), $\mathbf{A} \in \mathbb{R}^{WH \times e}$
119 and $\hat{\mathbf{A}} \in \mathbb{R}^{wh \times e}$ represent the CSA matrices. They must satisfy the abundance non-negativity con-
120 straint (ANC) and the abundance sum-to-one constraint (ASC), ensuring the physical interpretation
121 of abundance as the fraction of material in a pixel (Deng et al., 2025). The deep learning-based unsu-
122 pervised fusion typically sets two modalities corresponding to LR-HSI and HR-MSI. One modality
123 focuses on obtaining ESD, while another concerns CSAs. We integrate \mathbf{A} and \mathbf{E} as:
124

$$\mathbf{Z} = \mathbf{A}\mathbf{E} \quad (2)$$

125 where $\mathbf{Z} \in \mathbb{R}^{WH \times G}$ is the reconstructed HR-HSI. The effectiveness of fusion in Eq. (2) depends on
126 whether the row vectors of \mathbf{A} correspond to the column vectors of \mathbf{E} , representing the same material.
127 This correspondence plays a crucial role in determining the quality of the fused data. Increasing
128 matching information between the two modalities is also the main motivation of this paper.
129

132 **2.2 DEEP FREQUENCY FILTERING**

133 Deep Frequency Filtering (DFF) is a technique introduced to enhance the domain generalization
134 ability of deep neural networks by explicitly modulating the frequency components of learned fea-
135 tures (Lin et al., 2023). DFF and Frequency Domain Analysis (FDA) in image processing are closely
136 related, particularly when it comes to manipulating and enhancing features in the frequency domain.
137 For image processing, FDA's basic idea involves transforming an image from the spatial domain to
138 the frequency domain (Wang et al., 2020). DFF shares a common theoretical foundation with tradi-
139 tional FDA and applies this concept to deep learning to tackle domain generalization, improving the
140 model's performance by modulating its learned features in the frequency domain.
141

142 Following the research in (Lin et al., 2023), DFF adopts Fast Fourier Transform (FFT) (Cooley &
143 Tukey, 1965) to obtain the feature representations in the frequency domain. Let $\mathbf{F} \in \mathbb{R}^{H \times W \times d}$ be
144 the intermediate features, and the Discrete FFT (DFFT) is denoted as:
145

$$\mathbf{F}^{DFFT}(i, j) = \sum_{u=0}^{H-1} \sum_{v=0}^{W-1} \mathbf{F}(u, v) \cdot e^{-p2\pi(i \frac{u}{H} + j \frac{v}{W})} \quad (3)$$

146 where $\mathbf{F}^{DFFT} \in \mathbb{R}^{2d \times H \times (\lfloor \frac{W}{2} \rfloor + 1)}$ is the frequency representation of \mathbf{F} , (i, j) denotes the frequency
147 index, (u, v) defines the spatial index, and p is the imaginary unit. $\lfloor \cdot \rfloor$ rounds a number down to the
148 nearest integer less than.

149 Afterward, the real and imaginary parts of the frequency components are separated for deep feature
150 representation. The FFT decomposes the signal into its frequency components, which are repre-
151 sented as complex numbers. Subsequently, the Inverse DFFT (IDFFT) (Katznelson, 2004) is imple-
152 mented as:
153

$$\mathbf{F}(u, v) = \frac{1}{HW} \sum_{i=0}^{H-1} \sum_{j=0}^{W-1} \mathbf{F}^{DFFT}(i, j) \cdot e^{p2\pi(i \frac{u}{H} + j \frac{v}{W})} \quad (4)$$

154 By cascading multiple layers of processing, DFF can effectively isolate or amplify specific fre-
155 quency components, offering enhanced performance in applications in signal, image, audio, and
156 video processing (Jiang et al., 2024).
157




Figure 2: The overview of FSRVL. This work employs the discrete fast Fourier transform to transfer the features from the spatial domain to the frequency domain.

3 METHODOLOGY

3.1 OVERALL PIPELINE

Figure 2 illustrates the overall pipeline of FSRVL for MSI-HSI fusion. This work employs convolutions to achieve the multiscale representation of local correlations and adjust feature dimension. It adopts the DFFT to transfer the features from the spatial domain to the frequency domain, after aggregating the local correlations. The frequency recalibration in FSRVL will enhance information correspondence. As the unsupervised training progresses, the hidden layers generate CSAs with higher matching values, and their node weights are adjusted accordingly to provide an ESD that better matches the CSAs. Ultimately, the HR-HSI is reconstructed by the ESD and CSAs. *Feature Synchronization* and *Frequency Recalibration* are two key components of innovation. Despite the new method allowing their removal in the framework of MSI-HSI fusion, the models' performance is expected to drop significantly.

3.2 INITIAL REPRESENTATION

Given the high spatial resolution of HR-MSI, MSI-HSI fusion relies on it to provide spatial details. To enhance the retention of local relationships between pixels, we combine 2D convolutions and LeakyReLU (Maas et al., 2013) to achieve multiscale representation, which obtains $\mathbf{X}_{3 \times 3} \in \mathbb{R}^{WH \times d}$ and $\mathbf{X}_{5 \times 5} \in \mathbb{R}^{WH \times d}$. LeakyReLU can mitigate the dying neuron issue to ensure training stability. Corresponding to the left side in Figure 2, the multiscale embedding of HR-MSI is denoted as:

$$\mathbf{X}' = \mathbf{X}_{3 \times 3} \odot \mathbf{X}_{5 \times 5} \quad (5)$$

where \odot defines the Hadamard Product, $\mathbf{X}' \in \mathbb{R}^{WH \times d}$ denotes multiscale features.

Since LR-HSI is utilized to provide ESD, we do not need to pay attention to its spatial information via a multiscale representation. This part utilizes 1×1 convolutions to adjust feature dimensions. Likewise, we integrate the 1×1 convolutions and LeakyReLU to implement the LR-HSI embedding $\mathbf{Y}' \in \mathbb{R}^{wh \times d}$, and the process is stacked to obtain $\mathbf{Y}'' \in \mathbb{R}^{wh \times d}$. Additionally, we utilize 1×1 convolutions and LeakyReLU to further process \mathbf{X}' to acquire $\mathbf{X}'' \in \mathbb{R}^{WH \times d}$, which enables consistent modeling of multi-source representation by sharing weights.

Considering that the CSAs reflect the distribution of materials and preserve their spatial information, this work applies Softmax to ensure the ANC and the ASC:

$$\begin{aligned} \mathbf{H} &= \text{Softmax}(\text{Concat}(\mathbf{X}', \mathbf{X}'')\mathbf{W}_0) \\ \widehat{\mathbf{H}} &= \text{Softmax}(\text{Concat}(\mathbf{Y}', \mathbf{Y}'')\mathbf{W}_1) \end{aligned} \quad (6)$$

where $\mathbf{H} \in \mathbb{R}^{WH \times d}$ and $\widehat{\mathbf{H}} \in \mathbb{R}^{wh \times d}$ are the initial representation of CSAs, $\mathbf{W}_0 \in \mathbb{R}^{2d \times d}$ and $\mathbf{W}_1 \in \mathbb{R}^{2d \times d}$ are two learnable weight matrices.

3.3 FEATURE SYNCHRONIZATION

We represent each column vector in \mathbf{H} or $\widehat{\mathbf{H}}$ by hidden layer nodes, and the edge matrices connecting to the next layer are denoted as \mathbf{W} or $\widehat{\mathbf{W}}$. Each row vector in \mathbf{W} or $\widehat{\mathbf{W}}$ corresponds to the edge of a node. To prepare for frequency domain operations, we reshape \mathbf{H} and $\widehat{\mathbf{H}}$ as the tensors $\mathbf{H} \in \mathbb{R}^{W \times H \times d}$ and $\widehat{\mathbf{H}} \in \mathbb{R}^{w \times h \times d}$. Then, we bilinearly interpolate $\widehat{\mathbf{H}}$ to match the size of \mathbf{H} .

Specifically, let $\widehat{\mathbf{H}}_c$ be the c -th component of $\widehat{\mathbf{H}}$, and \mathbf{H}_c is the c -th component of \mathbf{H} , $c = 1, 2, \dots, d$. We map the indices (U, V) of \mathbf{H}_c to floating coordinates (u, v) in $\widehat{\mathbf{H}}_c$, obtaining $u = \frac{U}{W-1} \cdot (w-1)$, $v = \frac{V}{H-1} \cdot (h-1)$. The bilinear interpolation of $\widehat{\mathbf{H}}$ is written as:

$$\begin{aligned} \mathbf{H}'_c[U, V, k] &= (1 - \delta_u)(1 - \delta_v) \cdot \widehat{\mathbf{H}}_c[u_1, v_1, k] \\ &\quad + \delta_u(1 - \delta_v) \cdot \widehat{\mathbf{H}}_c[u_2, v_1, k] \\ &\quad + \delta_v(1 - \delta_u) \cdot \widehat{\mathbf{H}}_c[u_1, v_2, k] \\ &\quad + \delta_u\delta_v \cdot \widehat{\mathbf{H}}_c[u_2, v_2, k] \end{aligned} \quad (7)$$

where k be the interpolated value of the channel, $u_1 = \lfloor u \rfloor$, $u_2 = u_1 + 1$, $v_1 = \lfloor v \rfloor$, $v_2 = v_1 + 1$, $\delta_u = u - u_1$, and $\delta_v = v - v_1$. The tensor $\mathbf{H}' \in \mathbb{R}^{W \times H \times d}$ consists of d matrices using Equ. (7). For convenience, we treat the DFFT in Equ. (3) as a function $\mathcal{F}_{DFFT}(\cdot)$. Our FSRVL applies the DFFT to transfer features from the spatial domain to the frequency domain:

$$\begin{aligned} [\text{Re}(\mathbf{H}), \text{Im}(\mathbf{H})] &= \mathcal{F}_{DFFT}(\mathbf{H}) \\ [\text{Re}(\mathbf{H}'), \text{Im}(\mathbf{H}')] &= \mathcal{F}_{DFFT}(\mathbf{H}') \end{aligned} \quad (8)$$

where $\text{Re}(\mathbf{H})$ and $\text{Re}(\mathbf{H}')$ are the real parts of complex numbers, $\text{Im}(\mathbf{H})$ and $\text{Im}(\mathbf{H}')$ are the corresponding imaginary part. Considering that both the real part and imaginary parts are in Tensor form, to facilitate information interaction across feature channels, we apply Fully Connected (FC) layers to process the real and imaginary parts:

$$\mathbf{F}^{Re} = \text{FC}(\text{Re}(\mathbf{H})), \quad \mathbf{F}^{Im} = \text{FC}(\text{Im}(\mathbf{H})) \quad (9)$$

where $\mathbf{F}^{Re} \in \mathbb{R}^{W \times H \times d}$ and $\mathbf{F}^{Im} \in \mathbb{R}^{W \times H \times d}$ represent the frequency embedding with HR-MSI. Similarly, we also adopt Equ. (9) to process $\text{Re}(\mathbf{H}')$ and $\text{Im}(\mathbf{H}')$ to acquire the frequency embedding with LR-HSI, obtaining $\mathbf{F}'^{Re} \in \mathbb{R}^{W \times H \times d}$ and $\mathbf{F}'^{Im} \in \mathbb{R}^{W \times H \times d}$. The bottom left of Figure 2 provides the schematic of the feature synchronization.

3.4 FREQUENCY RECALIBRATION

After conducting the DFFT, if spectral convolution updates a value in the spectral domain, all original nodes will be affected globally (Huang et al., 2025). Therefore, we can still employ 2D convolutions as the learnable filters to capture the global frequency features. In the stream with HR-MSI, the process with the c -th channel can be written as:

$$\mathbf{Q}_c^{Re} = \text{Con}^{m \times m}(\mathbf{F}_c^{Re}), \quad \mathbf{Q}_c^{Im} = \text{Con}^{m \times m}(\mathbf{F}_c^{Im}) \quad (10)$$

where $\mathbf{Q}_c^{Re} \in \mathbb{R}^{W \times H}$ and $\mathbf{Q}_c^{Im} \in \mathbb{R}^{W \times H}$ are the Global Frequency Feature (GFF) matrices of the c -th channel, m defines the size of the convolution kernel. These GFF matrices with all channels can be combined into the feature tensor of real and imaginary parts: $\mathbf{Q}^{Re} \in \mathbb{R}^{W \times H \times d}$ and $\mathbf{Q}^{Im} \in \mathbb{R}^{W \times H \times d}$.

For the embedding of LR-HSI, we adopt the same strategy in Equ. (10) to process \mathbf{F}'^{Re} and \mathbf{F}'^{Im} for calculating the GFF matrices. Like the above processing, the GFF matrices are integrated as the tensors of real and imaginary parts: $\mathbf{Q}'^{Re} \in \mathbb{R}^{W \times H \times d}$ and $\mathbf{Q}'^{Im} \in \mathbb{R}^{W \times H \times d}$. We employ Batch Normalization (BatchNorm) (Ioffe & Szegedy, 2015) to reduce internal covariate shift, thereby stabilizing the training. BatchNorm will generate Recalibrated Feature Embeddings (RFEs) in the frequency domain: $\mathbf{B}'^{Re} \in \mathbb{R}^{W \times H \times d}$, $\mathbf{B}'^{Im} \in \mathbb{R}^{W \times H \times d}$, $\mathbf{B}''^{Re} \in \mathbb{R}^{W \times H \times d}$, and $\mathbf{B}''^{Im} \in \mathbb{R}^{W \times H \times d}$. By the IDFFT in Equ. (4), these features are transformed back to the spatial domain:

$$\begin{aligned} \mathbf{S} &= \text{Con}^{1 \times 1} (\mathcal{F}_{IDFFT}[\mathbf{B}'^{Re}, \mathbf{B}'^{Im}]), \\ \mathbf{S}' &= \text{Con}^{1 \times 1} (\mathcal{F}_{IDFFT}[\mathbf{B}''^{Re}, \mathbf{B}''^{Im}]) \end{aligned} \quad (11)$$

where $\mathbf{S} \in \mathbb{R}^{W \times H \times e}$ and $\mathbf{S}' \in \mathbb{R}^{W \times H \times e}$ are the RFEs in the spatial domain, $\mathcal{F}_{IDFFT}[\cdot, \cdot]$ expresses implementation of the IDFFT. Moreover, we use a bilinear downsample to address \mathbf{S}' , thereby restoring the spatial size of LR-HSI to obtain $\mathbf{S}'' \in \mathbb{R}^{w \times h \times e}$. Let each component of \mathbf{S} and \mathbf{S}'' correspond to an abundance node, and its weight can represent a vector in the ESD. Consequently, ESDs associated with \mathbf{S} and \mathbf{S}'' will also undergo continuous updates as iterations. The bottom right of Figure 2 illustrates the architectures of the frequency recalibration.

3.5 IMAGE RECONSTRUCTION

Mathematically, $\widehat{\mathbf{E}}$ and $\widehat{\mathbf{A}}$ in Equ. (1) can be further decomposed as:

$$\widehat{\mathbf{E}} = \mathbf{EP}, \quad \widehat{\mathbf{A}} = \mathbf{RA} \quad (12)$$

where $\mathbf{P} \in \mathbb{R}^{G \times g}$ and $\mathbf{R} \in \mathbb{R}^{wh \times WH}$ are the blurring factor matrices with the spectral and spatial degradations. Such degradations in the spatial and spectral dimensions are often modeled using the spectral response function and point spread function, respectively. We treat the spectral signals \mathbf{E} and degraded spectral signals $\widehat{\mathbf{E}}$ as learnable weight matrices. According to Equ. (2), we can reconstruct the HR-HSI:

$$\mathbf{Z} = \text{Reshape}(\mathbf{S})\mathbf{E}, \quad (13)$$

where $\text{Reshape}(\cdot)$ can reshape a tensor to a matrix, $\text{Reshape}(\mathbf{S}) = \mathbf{A}$. Following the research in (Liu et al., 2022), if \mathbf{Z} exists, \mathbf{X} and \mathbf{Y} can be deduced in reverse:

$$\widetilde{\mathbf{Y}} = \text{Conv}^{\phi \times \psi}(\mathbf{Z}), \quad \widetilde{\mathbf{X}} = \mathbf{X}\widetilde{\mathbf{E}} \quad (14)$$

where $\text{Conv}^{\phi \times \psi}$ is a learnable downsampling convolutional layer, ensuring that the downsampled HR-HSI has the same spatial resolution as the LR-HSI. Here, ϕ and ψ are two scale factors, $\phi = \lceil H/h \rceil$, $\psi = \lceil W/w \rceil$. The symbol $\lceil \cdot \rceil$ rounds a number up to the nearest integer more than. As the spectral degradation matrix, $\widetilde{\mathbf{E}}$ can be computed between the downsampled HR-MSI \mathbf{X} and LR-HSI \mathbf{Y} , and optimized offline prior to model training:

$$\widetilde{\mathbf{E}} \leftarrow \text{argmin}_{\mathbf{W}} \|\text{Downsample}(\mathbf{X}) - \mathbf{Y}\mathbf{W}\|_2. \quad (15)$$

where \mathbf{W} is a learnable weight matrix. In FSRVL, the loss function is written as:

$$\begin{aligned} \ell_{\text{FSRVL}} &= \alpha \left(\|\mathbf{Y} - \widehat{\mathbf{Y}}\| + \|\mathbf{X} - \widehat{\mathbf{X}}\| \right) \\ &\quad + \beta \left(\|\mathbf{Y} - \widetilde{\mathbf{Y}}\| + \|\mathbf{X} - \widetilde{\mathbf{X}}\| \right) \end{aligned} \quad (16)$$

where α and β are the balance coefficients.

4 EXPERIMENTS


Table 1 lists five datasets for evaluating fusion performances: the CAVE-Toy, CAVE-Face (Yasuma et al., 2010), TG-1, GF2-GF5, and SZUTree (Long et al., 2024). The CAVE-Toy, CAVE-Face, and TG-1 are used for simulated scenarios. Simulated data provide *Ground Truth (GT)* for quantitative evaluation, but cannot fully reflect real cross-sensor conditions. GF2-GF5 and SZUTree are real cross-sensor datasets. FSRVL employs Adam (Kingma & Ba, 2015) as the optimizer. The environment for implementing FSRVL includes PyTorch 2.4.1, Python 3.9, and CUDA 12.1. The experiments are conducted on a computing server with an NVIDIA RTX 4090 GPU.

Table 1: Dataset types and acquisition platforms.

TYPE	DATASET NAME	PLATFORM	SOURCE	GROUND TRUTH
Registered Data	CAVE-Toy	Ground Camera	Single Sensor	Yes
	CAVE-Face	Ground Camera	Single Sensor	Yes
	TG-1	Satellite	Cross-Sensor	Yes
Unregistered Data	GF2-GF5	Satellite	Cross-Sensor	No
	SZUTree	UAV	Cross-Sensor	No

Table 2: Left: *Initial* represents the model without the deep frequency domain filter, *FSRVL-S* represents the model without Feature Synchronization, and *FSRVL-R* represents the model without Feature Recalibration. Right: Features of spatial and frequency domain views

	Initial	FSRVL-S	FSRVL-R	FSRVL
SAM \downarrow	0.94	0.69	0.71	0.64
PSNR \uparrow	45.04	48.05	47.64	48.35
ERGAS \downarrow	0.22	0.18	0.17	0.15
SSIM \uparrow	0.98	0.98	0.99	0.99
Para (M)	0.98	0.98	0.99	0.99
Times (H)	0.19	0.50	0.59	0.63



- **CAVE Datasets:** The CAVE-Toy ($512 \times 512 \times 31$) was downsampled into the HR-MSI ($512 \times 512 \times 3$) and the LR-HSI ($64 \times 64 \times 31$), and the CAVE-Face follows the same process.
- **TG-1 Dataset:** The original data ($240 \times 240 \times 54$) was from TianGong01 satellite and downsampled into the HR-MSI ($240 \times 240 \times 8$) and the LR-HSI ($30 \times 30 \times 54$).
- **GF2-GF5 dataset:** GaoFen02 satellite provides the GF2 data-HR-MSI ($360 \times 360 \times 4$), while GaoFen05 satellite collects the LR-HSI ($60 \times 60 \times 330$).
- **SZUTree Dataset:** It was collected by an UAV with a RGB camera and a SPECIM FX10 sensor, yields the HR-MSI ($300 \times 300 \times 3$) and the LR-HSI ($150 \times 150 \times 98$).

Fusion performance is evaluated using five standard metrics: *Spectral Angle Mapper (SAM \downarrow)*, *Peak Signal-to-Noise Ratio (PSNR \uparrow)*, *Erreur Relative Globale Adimensionnelle de Synthese (ERGAS \downarrow)*, *Structural Similarity Index (SSIM \uparrow)*, *Mean Relative Absolute Error (MRAE \downarrow)*. Here, \uparrow indicates that higher values denote better results, while \downarrow signifies lower values are preferable.




Figure 3: Left compares performances when configuring different hyperparameters. Right lists the classification maps on the SZTree data using FSRVL.

378
379

4.1 ABLATION STUDY AND HYPERPARAMETER ANALYSIS

380
381
382
383
384
385
386
387

Table 2 provides the ablation study on the TG-1 dataset that evaluates the contribution of FSRVL’s three modules: Initialization, Feature Synchronization, and Frequency Recalibration. As shown in Table 2, each module plays a significant role in improving fusion performance. Figure 3 illustrates the impact of configuring different hyperparameters on fusion performance. Tests on the CAVE-Face dataset indicate that the convolution kernel in Equ. (10) is 7×7 , with $\alpha = 100$ and $\beta = 10$ yielding the best results, and this setting is used throughout all experiments. Figure 3 also illustrates the classification maps of the SZUTree dataset generated by FSRVL, demonstrating that the fused HR-HSI can facilitate more accurate land-cover classification.

388
389
390

4.2 COMPARISON WITH BASELINES

391
392
393
394

FSRVL is compared with baselines, such as CUCaNet (Yao et al., 2020), SSG-SRL (Liu et al., 2025), DFMFf (Guo et al., 2023), CS2DIPs (Fang et al., 2024), DTDNML (Wang et al., 2025), EDIP (Li et al., 2025), and OTIAS (Deng et al., 2025). Therein, CUCaNet, CS2DIPs, DTDNML, and EDIP are unsupervised approaches, whereas SSG-SRL, DFMFf, and OTIAS are supervised methods.

395
396
397
398
399
400

Figure 4 compares FSRVL, EDIP, and OTIAS on the CAVE dataset, showing that FSRVL yields spectral results closer to the GTs. Figure 5 and Table 3 show that FSRVL outperforms baselines on the TG-1 dataset. The CAVE and TG-1 datasets are both simulated with ideal conditions where data across different modalities are perfectly aligned, enabling unsupervised methods to leverage their full potential. In contrast, supervised methods exhibit slightly inferior fusion performance due to their reliance on sampling.

401
402
403
404
405

CS2DIPs, DTDNML, and EDIP are excluded from tests on GF2-GF5 and SZUTree because their implementations require strictly aligned data. Figure 6 shows that FSRVL best preserves the spectral/color features of LR-HSI and the spatial details of HR-MSI. We employ a simple classifier with a fully connected layer followed by a Softmax in the downstream.


406
407
408
409
410
411
412
413
414415
416

Figure 4: Comparing fusion results on the CAVE datasets, where the red curves represent the Ground Truth, and the blue dotted lines represent the estimated spectra.

417
418
419
420
421

Table 3: Quantitative evaluation of different methods on the TG-1 dataset. The **best** and second-best values are highlighted.

422
423
424
425
426
427
428
429
430
431

Method	CUCaNet	SSG-SRL	DFMFf	CS2DIPs	DTDNML	EDIP	OTIAS	FSRVL
Published Year	2020	2025	2023	2024	2025	2025	2025	
TG-1	SAM \downarrow	1.21	1.14	2.23	1.15	0.99	<u>0.91</u>	0.89
	PSNR \uparrow	40.35	44.90	44.40	44.08	44.47	45.45	<u>45.71</u>
	ERGAS \downarrow	0.40	0.29	0.51	0.29	0.27	<u>1.44</u>	0.24
	SSIM \uparrow	0.95	0.96	0.95	0.96	<u>0.97</u>	<u>0.97</u>	0.99
	Para (M)	1.78	3.36	4.82	0.30	0.71	12.25	2.14
	Times (H)	0.54	0.93	0.87	0.10	0.39	0.61	0.66




Figure 5: Comparing fusion results on the TG-1 dataset. The top row displays the RGB compositions (R:28, G:18, B:5) of the fused HR-HSIs and the pixels' spectra. The second and third rows present the SAM heat map and the MRAE heat map.




Figure 6: Reconstructed HR-HSIs were obtained by different methods using the unregistered datasets. The two datasets captured the same area within a short time. GTs of spectra are from the LR-HSI.

5 CONCLUSION

This work presents our FSRVL, a novel framework for cross-sensor fusion of multispectral and hyperspectral images. We designed the Feature Synchronization module and the Frequency Recalibration module to adaptively modulate frequency features, effectively mitigating spectral-spatial information mismatches between HR-MSI and LR-HSI. Comprehensive experiments on both simulated and real-world datasets validate its robustness and demonstrate superior performance compared to existing methods in spectral image fusion. In future work, we plan to extend FSRVL to handle weakly aligned data and explore its applications in real-time onboard processing. Additionally, I acknowledge that this paper has been revised and polished with the assistance of ChatGPT to improve its clarity and readability. Upon acceptance of the paper, we will make the experimental datasets and demos available for download to ensure reproducibility of the study. Further details are shown in Appendix.

486 REFERENCES
487

- 488 James W Cooley and John W Tukey. An algorithm for the machine calculation of complex fourier
489 series. *Mathematics of computation*, 19(90):297–301, 1965.
- 490 Shangqi Deng, Jun Ma, Liang-Jian Deng, and Ping Wei. Otias: Octree implicit adaptive sampling
491 for multispectral and hyperspectral image fusion. In *Proceedings of the 39th AAAI Conference on
492 Artificial Intelligence*, pp. 2708–2716, Apr. 2025.
- 493 W. Dong, Y. Xu, J. Qu, and S. Hou. Learning multi-modal cross-scale deformable transformer
494 network for unregistered hyperspectral image super-resolution. In *Proceedings of the AAAI Con-
495 ference on Artificial Intelligence (AAAI)*, volume 38, pp. 1573–1581, 2024.
- 496 Y. Fang, Y. Liu, C.-Y. Chi, Z. Long, and C. Zhu. Cs2dips: Unsupervised hsi super-resolution using
497 coupled spatial and spectral dips. *IEEE Transactions on Image Processing*, 33:3090–3101, 2024.
- 498 L. Gao, J. Li, K. Zheng, and X. Jia. Enhanced autoencoders with attention-embedded degra-
500 dation learning for unsupervised hyperspectral image super-resolution. *IEEE Transactions on Geo-
501 science and Remote Sensing*, 61:1–17, 2023.
- 502 A. Guo, R. Dian, and S. Li. A deep framework for hyperspectral image fusion between different
503 satellites. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(7):7939–7954,
504 2023.
- 505 X. Huang, Y. Dong, Y. Zhang, and B. Du. Single-source frequency transform for cross-scene classi-
506 fication of hyperspectral image. *IEEE Transactions on Image Processing*, 34:3000–3012, 2025.
- 507 S. Ioffe and C. Szegedy. Batch normalization: accelerating deep network training by reducing
508 internal covariate shift. In *Proceedings of the 32nd International Conference on International
509 Conference on Machine Learning (ICML)*, pp. 448–456, 2015.
- 510 Z. Jiang, Y. Ding, J. Zhao, Y. Tian, S. Han, S. Ozharar, T. Wang, and J. M. Moore. Seeing the
511 vibration from fiber-optic cables: Rain intensity monitoring using deep frequency filtering. In
512 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)
513 workshop*, pp. 3017–3026, June 2024.
- 514 Yitzhak Katznelson. *An introduction to harmonic analysis*. Cambridge University Press, 2004.
- 515 D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *Proceedings of the 3rd
516 International Conference on Learning Representations (ICLR)*, 2015.
- 517 J. Li, K. Zheng, L. Gao, Zh. Han, Z. Li, and J. Chanussot. Enhanced deep image prior for unsu-
518 pervised hyperspectral image super-resolution. *IEEE Transactions on Geoscience and Remote
519 Sensing*, 63:1–18, 2025.
- 520 K. Li, L. Van Gool, and D. Dai. Test-time training for hyperspectral image super-resolution. *IEEE
521 Transactions on Pattern Analysis and Machine Intelligence*, pp. 1–12, 2024.
- 522 Shiqi Lin, Zhizheng Zhang, Zhipeng Huang, Yan Lu, Cuiling Lan, Peng Chu, Quanzeng You, Jiang
523 Wang, Zicheng Liu, Amey Parulkar, et al. Deep frequency filtering for domain generalization.
524 In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp.
525 11797–11807, 2023.
- 526 Jianjun Liu, Zebin Wu, Liang Xiao, and Xiao-Jun Wu. Model inspired autoencoder for unsupervised
527 hyperspectral image super-resolution. *IEEE Transactions on Geoscience and Remote Sensing*, 60:
528 1–12, 2022.
- 529 Y. Liu, J. Liu, R. Dian, and S. Li. A selective re-learning mechanism for hyperspectral fusion imag-
530 ing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition
(CVPR)*, pp. 7437–7446, June 2025.
- 531 Y. Long, S. Ye, L. Wang, W. Wang, X. Liao, and S. Jia. Scale pyramid graph network for hyper-
532 spectral individual tree segmentation. *IEEE Transactions on Geoscience and Remote Sensing*, 62:
533 1–14, 2024.

- 540 A. L. Maas, A. Y. Hannun, and A. Y Ng. Rectifier nonlinearities improve neural network acoustic
 541 models. In *International Conference on Machine Learning (ICML)*, volume 30, pp. 3, 2013.
- 542
- 543 J. Qu, J. He, W. Dong, and J. Zhao. S2cyclediff: Spatial-spectral-bilateral cycle-diffusion frame-
 544 work for hyperspectral image super-resolution. In *Proceedings of the 38th AAAI Conference on*
 545 *Artificial Intelligence*, pp. 4623–4631, Mar. 2024.
- 546 J. Qu, X. Wu, W. Dong, J. Cui, and Y. Li. *IRArF*: Toward deep interpretable arbitrary resolution
 547 fusion of unregistered hyperspectral and multispectral images. *IEEE Transactions on Image Pro-
 548 cessing*, 34:1934–1949, 2025.
- 549
- 550 X. Su, X. Shen, M. Wan, J. Nie, L. Chen, H. Liu, and X. Zhou. *EigenSR*: Eigenimage-bridged pre-
 551 trained rgb learners for single hyperspectral image super-resolution. In *Proceedings of the 39th*
 552 *Annual AAAI Conference on Artificial Intelligence (AAAI)*, 2025.
- 553 H. Wang, X. Wu, Z. Huang, and E. P. Xing. High-frequency component helps explain the gener-
 554 alization of convolutional neural networks. In *IEEE/CVF Conference on Computer Vision and*
 555 *Pattern Recognition (CVPR)*, pp. 8681–8691, 2020.
- 556 H. Wang, Y. Xu, Z. Wu, and Z. Wei. Unsupervised hyperspectral and multispectral image blind
 557 fusion based on deep tucker decomposition network with spatial–spectral manifold learning. *IEEE*
 558 *Transactions on Neural Networks and Learning Systems*, 36(7):12721–12735, 2025.
- 559
- 560 Zhaoyang Wang, Dongyang Li, Mingyang Zhang, Hao Luo, and Maoguo Gong. Enhancing hy-
 561 perspectral images via diffusion model and group-autoencoder super-resolution network. In *Pro-
 562 ceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*, pp. 5794–5804, Mar.
 563 2024.
- 564 Q. Xie, M. Zhou, Q. Zhao, D. Meng, W. Zuo, and Z. Xu. Multispectral and hyperspectral image
 565 fusion by ms/hs fusion net. In *2019 IEEE/CVF Conference on Computer Vision and Pattern*
 566 *Recognition (CVPR)*, pp. 1585–1594, 2019.
- 567
- 568 J. Yang, L. Xiao, Y.-Q. Zhao, and J. C.-W. Chan. Unsupervised deep tensor network for hyperspec-
 569 tral–multispectral image fusion. *IEEE Transactions on Neural Networks and Learning Systems*,
 570 35(9):13017–13031, 2024a.
- 571 X. Yang, J. Chen, and Z. Yang. Hyperspectral image reconstruction via combinatorial embedding
 572 of cross-channel spatio-spectral clues. In *Proceedings of the 38th AAAI Conference on Artificial*
 573 *Intelligence (AAAI)*, pp. 6567–6575, Mar. 2024b.
- 574 J. Yao, D. Hong, J. Chanussot, D. Meng, X. Zhu, and Z. Xu. Cross-attention in coupled unmix-
 575 ing nets for unsupervised hyperspectral super-resolution. In *European Conference on Computer*
 576 *Vision (ECCV)*, pp. 208–224, 2020.
- 577
- 578 F. Yasuma, T. Mitsunaga, D. Iso, and S. K Nayar. Generalized assorted pixel camera: postcapture
 579 control of resolution, dynamic range, and spectrum. *IEEE transactions on image processing*, 19
 580 (9):2241–2253, 2010.
- 581 Haijin Zeng, Benteng Sun, Yongyong Chen, Jingyong Su, and Yong Xu. Spectral compressive imag-
 582 ing via unmixing-driven subspace diffusion refinement. In *Thirteenth International Conference*
 583 *on Learning Representations (ICCV)*, 2025.
- 584
- 585 L. Zhang, J. Nie, W. Wei, Y. Zhang, S. Liao, and L. Shao. Unsupervised adaptation learning for
 586 hyperspectral imagery super-resolution. In *2020 IEEE/CVF Conference on Computer Vision and*
 587 *Pattern Recognition (CVPR)*, pp. 3070–3079, 2020.
- 588 L. Zhang, J. Nie, Wei Wei, and Yanning Zhang. Unsupervised test-time adaptation learning for
 589 effective hyperspectral image super-resolution with unknown degeneration. *IEEE Transactions*
 590 *on Pattern Analysis and Machine Intelligence*, 46(7):5008–5025, 2024.
- 591 J. Zhu, H. Wang, Y. Xu, Z. Wu, and Z. Wei. Self-learning hyperspectral and multispectral image
 592 fusion via adaptive residual guided subspace diffusion model. In *Proceedings of the IEEE/CVF*
 593 *Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 17862–17871, Jun. 2025.