
FRN: Fractal-Based Recursive Spectral Reconstruction Network

Ge Meng¹, Zhongnan Cai¹, Ruizhe Chen¹, Jingyan Tu¹, Yingying Wang¹,
Yue Huang¹, Xinghao Ding^{1,*}

¹Key Laboratory of Multimedia Trusted Perception and Efficient Computing,
Ministry of Education of China, Xiamen University, Xiamen, Fujian, China
mengge0001@gmail.com

Abstract

Generating hyperspectral images (HSIs) from RGB images through spectral reconstruction can significantly reduce the cost of HSI acquisition. In this paper, we propose a Fractal-Based Recursive Spectral Reconstruction Network (FRN), which differs from existing paradigms that attempt to directly integrate the full-spectrum information from the R, G, and B channels in a one-shot manner. Instead, it treats spectral reconstruction as a progressive process, predicting from broad to narrow bands or employing a coarse-to-fine approach for predicting the next wavelength. Inspired by fractals in mathematics, FRN establishes a novel spectral reconstruction paradigm by recursively invoking an atomic reconstruction module. In each invocation, only the spectral information from neighboring bands is used to provide clues for the generation of the image at the next wavelength, which follows the low-rank property of spectral data. Moreover, we design a band-aware state space model that employs a pixel-differentiated scanning strategy at different stages of the generation process, further suppressing interference from low-correlation regions caused by reflectance differences. Through extensive experimentation across different datasets, FRN achieves superior reconstruction performance compared to state-of-the-art methods. Code is available at <https://github.com/mongko007/frn>.

1 Introduction

Hyperspectral images (HSIs) contain more spectral bands (channels) than RGB images, enabling them to capture richer emission information that more accurately reflects the properties of objects. As a result, HSIs are commonly used in applications such as medical imaging [33, 35, 3], remote sensing [55, 46, 34, 19], material classification [25, 26], and object tracking [48, 31], etc.

Conventional hyperspectral imaging systems typically employ a single 1D or 2D sensor to scan the scene along the spatial or spectral dimension, capturing hyperspectral information through prolonged, repeated exposures. However, this approach is not well-suited for dynamic scenes. The coded aperture snapshot spectral imaging (CASSI) system takes advantage of the sparsity of spectral data, acquiring




Figure 1: PSNR-Parameters comparisons of FRN and SOTA methods. FRN achieves outstanding HSI reconstruction performance with only a minimal number of parameters.

*Corresponding author.




Figure 2: Overview of the Fractal-Based Progressive Spectral Reconstruction Paradigm: (a) illustrates how FRN reconstructs images at specific wavelengths in a coarse-to-fine manner, transitioning from wide spectral bands to narrow ones across multiple levels. (b) demonstrates the structural self-similarity of the modules within each level.

compressed 2D measurements by modulating spectral signals at different wavelengths. The original HSIs are then reconstructed from these measurements using reconstruction algorithms [7, 9, 23, 24, 8, 36]. Despite its effectiveness, the high cost of CASSI devices has led researchers to explore more affordable alternatives. Given the widespread availability of RGB cameras, spectral reconstruction (SR) algorithms have been developed to recover HSIs from RGB input [10, 28, 20, 16, 52, 15].

The SR task is inherently ill-posed [15, 52]. Due to the variability in camera response functions, a single RGB image may correspond to multiple HSIs, complicating the accurate retrieval of hyperspectral information from limited RGB data. Traditional methods rely on statistical principles or hand-crafted sparse priors [2, 44]. However, the scarcity of paired RGB-HSI data limits the ability to fully capture and explore these prior assumptions. Convolutional neural networks (CNNs) have demonstrated strong performance in addressing various ill-posed problems, with some studies exploring their application to the SR task [15, 45, 29, 20]. Nevertheless, CNNs struggle to capture non-local spatial dependencies and model correlations between spectral bands. In contrast, transformers, owing to their multi-head self-attention mechanism, are better suited for handling long-range spatial dependencies and inter-band relationships, and have thus been widely applied to the SR task [10, 52, 11, 9, 8, 27]. However, such methods typically integrate spectral information from RGB images in a brute-force manner, leading to significant computational overhead and increased training complexity for neural networks.

Fractals are common patterns observed in neural networks [30]. Numerous studies have demonstrated the effectiveness of fractal or scale-invariant small-world network structures in the brain and its functional networks [4, 47, 6]. This inspires the idea that a larger network can be recursively constructed from smaller atomic modules. In addition to neural networks, fractal or near-fractal patterns are also commonly observed in natural data [30]. HSI data consists of sub-images from different bands, where sub-bands with slow variations in object emission characteristics can be treated as atomic cubes. Therefore, the SR task can be decomposed into the reconstruction of these sub-bands.

Based on the above analysis, we attempt to introduce the concept of fractals into the SR task. Fig. 1 shows the superior performance of our method. By employing the recursive principle, our method achieves superior HSI reconstruction performance while requiring only a minimal number of parameters. Fig. 2 (a) illustrates the progressive spectral reconstruction from wide to narrow bands. During the generation of a specific wavelength image, only spectral information from neighboring channels is used as cues for reconstruction, which aligns with the low-rank properties of HSI. Unlike integrating all hyperspectral information from an RGB image in a single step, this recursive generation approach mitigates the ill-posed problem by increasing the input and reducing the output at each

level. Fig. 2 (b) depicts our recursive generation framework, which exhibits self-similarity across levels by recursively invoking atomic reconstruction modules within atomic reconstruction modules.

A major challenge of recursively invoking the model is the substantial increase in the number of parameters and FLOPs. Visual Mamba (VMamba) models sequential dependencies by dividing an image into sequential blocks. A key advantage of VMamba is its linear computational complexity [57, 13, 42, 41, 43], which is particularly critical for handling high-dimensional hyperspectral features. Therefore, we design a state space model (SSM) with an adaptive band-aware mask (BAMamba) that filters out pixels with low spatial correlation before cross-scanning. This approach reduces computational cost while enabling the network to learn spatial sparsity effectively. In summary, our contributions are listed as follow:

- 1) We propose a fractal-based recursive spectral reconstruction network (FRN), to the best of our knowledge, it is the first attempt to introduce fractals to the SR task. FRN decomposes SR task into the reconstruction of sub-band images, establishing a new paradigm for spectral reconstruction.
- 2) FRN effectively leverages the low-rank characteristics of HSI while reducing the complexity of solving the ill-posed problem. Moreover, it exhibits self-similarity between atomic reconstruction modules across different levels.
- 3) We design a SSM with an adaptive band-aware mask (BAMamba), which reduces the computational cost by filtering out pixels with lower spatial correlation and forces the network to learn more compact pixel-wise inductive biases.
- 4) Through experiments conducted on different datasets, FRN demonstrates outstanding performance in both qualitative and quantitative metrics.

2 Related Work

2.1 Hyperspectral Image Reconstruction

Conventional hyperspectral imaging systems typically use spectrometers to scan scenes along the spatial or spectral dimensions. These scanners—such as pushbroom and whiskbroom types—have been widely applied in remote sensing, medical imaging, and environmental monitoring [5, 40]. However, they require long exposure times during scanning, rendering them unsuitable for dynamic scenes, and the imaging devices are too large to be easily portable. To overcome these limitations, snapshot compressive imaging (SCI) systems have been developed [12, 18, 32, 49]. These systems compress 3D HSI data into 2D measurements. The original HSIs are then reconstructed from these measurements using reconstruction algorithms. A representative example is the CASSI system. Despite their advantages, SCI systems are costly. As a more accessible alternative, many researchers have explored spectral reconstruction algorithms that aim to recover hyperspectral data from conventional RGB images [56, 53, 52, 10, 11], leveraging the widespread availability of RGB cameras.

2.2 Model-Based Methods

Model-based SR methods generally introduce prior knowledge to help the model reduce the difficulty of the ill-posed problem. Given the intrinsic low-rank nature of HSIs, sparse representation has become one of the most typically techniques for incorporating prior knowledge [38, 2, 44]. Some methods assume that the camera response function is known and learn a mapping from RGB images to hyperspectral reflectance [37, 27, 53]. However, this assumption is overly restrictive in practice. In addition, mathematical techniques such as singular value decomposition (SVD) [21] and Gaussian processes [1] are also employed to reconstruct HSIs. These methods place excessive reliance on priors, which constrains the model's representational ability and generalization capability.

2.3 Deep-Learning-Based Methods

Thanks to the powerful nonlinear fitting capabilities of CNNs, many researchers have applied them on SR. Convolutional blocks are used for spectral upsampling [53] or stacked layer by layer to enable the network to model more complex functions [45]. Considering the significant emission differences of objects across different spectral bands, efficient spatial-spectral attention mechanisms have been




Figure 3: The details of BAMamba. BAMamba is a U-Net style network built with state space models equipped with band-aware masks (BSSM). BSSM introduces a band-aware spatial mask that adaptively perceives the reflectance of objects at specific wavelengths, suppressing interference from pixels with lower correlation.

designed to achieve more realistic reconstruction results [22, 28, 39]. To address the limitations of CNNs in capturing long-range spatial dependencies, a variety of Transformer-based SR methods have since emerged [10, 11, 52]. These methods learn inter-channel relationships by designing effective spectral-wise self-attention mechanisms. However, the above methods follow a one-shot reconstruction paradigm for HSI, overlooking the potential of recursive generation strategy to reduce the complexity of SR.

3 Method

3.1 Problem Formulation

The sensor of an RGB camera transfers the incident light to the R, G, and B channels through filters. This process can be regarded as an interaction between the camera response function (CRF) and the hyperspectral image

$$X(h, w, \lambda) = \int_{\lambda_{min}}^{\lambda_{max}} \phi(\lambda_i) \cdot Y(h, w, \lambda_i) d\lambda, \quad (1)$$

where $X \in \mathbb{R}^{H \times W \times 3}$ represents the RGB image and $Y \in \mathbb{R}^{H \times W \times L}$ is the corresponding HSI. $Y(h, w, \lambda_i)$ denotes the spectral reflectance in the location (h, w) at the wavelength of λ_i . $\phi(\lambda_i)$ represents the spectral response of the sensors at the wavelength of λ_i , and $[\lambda_{min}, \lambda_{max}]$ is the band range of $X(h, w, \lambda)$. The spectral information within the specified band range is integrated and stored in the channels of the RGB image. Eq. (1) can be rewritten in a discrete form

$$X(h, w, c) = \sum_{i=1}^K \phi(\lambda_i) \cdot Y(h, w, \lambda_i), \quad (2)$$

where $c \in [R, G, B]$ and K is the number of channels within the corresponding band. Furthermore, Eq. (2) can be simplified as a matrix form

$$\mathbf{X} = \mathbf{Y}\Phi. \quad (3)$$

where $\mathbf{X} \in \mathbb{R}^{HW \times 3}$ denotes the vectorial representation of RGB image, and $\mathbf{Y} \in \mathbb{R}^{HW \times L}$ is the corresponding HSI with L channels, $\Phi \in \mathbb{R}^{L \times 3}$ represents the CRF.

3.2 Spectral Reconstruction via Fractal Generator

Due to the self-similarity observed among local bands in HSI, we propose to construct a structurally self-similar SR network by following a recursive principle. We define a fractal generator g_i as an

atomic module that generates next-level data x_{i+1} from the previous-level result x_i : $x_{i+1} = g_i(x_i)$, as illustrated in Fig. 2 (b). Since the generator at each level can produce multiple outputs from a small amount of input, the fractal framework enables exponential growth of generated outputs with only a linear number of recursive levels [30], as shown in Fig. 2 (a). This property makes it particularly suitable for modeling high-dimensional HSI data using only a limited number of recursive levels. Specifically, we design an SSM with an adaptive band-aware mask (BAMamba) as the atomic generator, which will be described in Section 3.3.

The neural network learns the recursive principle from inter-band spectral correlations. For SR task, the objective is to learn the joint distribution of images at all wavelengths $p(y_{\lambda_1}, y_{\lambda_2}, \dots, y_{\lambda_N})$. However, it is difficult to model the joint distribution in a single step. To address this, we adopt a progressive strategy, which can be viewed as a divide-and-conquer approach, which model the conditional distribution $p(y | x)$ at different generation levels. Assume that each atomic generator produces a data sequence of length n , and the number of channels in the HSI is K , let $K = n^m$, where $m = \log_n K$ is the number of recursive levels. The atomic generator at the first level divides the joint distribution $p(y_1, y_2, \dots, y_K)$ into n subsets, each containing n^{m-1} variables. The joint distribution is decomposed

$$p(y_1, y_2, \dots, y_K) = \prod_{i=1}^n p(y_{(i-1) \cdot n^{m-1} + 1}, \dots, y_{i \cdot n^{m-1}} | y_1, \dots, y_{(i-1) \cdot n^{m-1}}). \quad (4)$$

Each conditional distribution is modeled by the atomic generator at the corresponding level. Fig. 2 illustrates the overall process. Through this typical divide-and-conquer strategy, FRN models the joint distribution over K variables by employing m levels of generators. The self-similarity of HSI along the spectral (channel) dimension, also known as its low-rank property, is effectively captured via the recursive principle, enabling a progressive approximation of the CRF Φ in Eq. (3).

3.3 Architecture of BAMamba

Visual Mamba (VMamba) exhibits linear computational complexity, enabling it especially well-suited for processing HSI data, which has a dimensionality much higher than that of RGB images. To address the computational burden introduced by recursive calls, we propose a VMamba-based sub-band generator that balances efficiency with performance, as shown in Fig. 3 (a).

The SSM employs a system of linear ordinary differential equations to connect inputs and outputs via intermediate hidden state representation. For a system with input signal $x(t) \in \mathbb{R}^L$, hidden state $h(t) \in \mathbb{C}^N$ and output response $y(t) \in \mathbb{R}^L$, the model can be formulated as

$$\begin{aligned} h'(t) &= \mathbf{A}h(t) + \mathbf{B}x(t), \\ y(t) &= \mathbf{C}h(t) + \mathbf{D}x(t), \end{aligned} \quad (5)$$

where $\mathbf{A} \in \mathbb{C}^{N \times N}$, $\mathbf{B}, \mathbf{C} \in \mathbb{C}^N$ and $\mathbf{D} \in \mathbb{C}^1$ are weighting parameters. Eq. (5) is typically discretized using a zero-order keeper (ZOH)

$$\begin{aligned} \bar{\mathbf{A}} &= \exp(\Delta \mathbf{A}), \\ \bar{\mathbf{B}} &= (\Delta \mathbf{A})^{-1}(\exp(\Delta \mathbf{A}) - \mathbf{I}) \cdot \Delta \mathbf{B}, \end{aligned} \quad (6)$$

where Δ is a time scale parameter used to transform the continuous parameters $\mathbf{A}, \mathbf{B}$ into discrete parameters $\bar{\mathbf{A}}, \bar{\mathbf{B}}$. The discretized Eq. (5) can be written as

$$\begin{aligned} h_t &= \bar{\mathbf{A}}h_{t-1} + \bar{\mathbf{B}}x_t, \\ y_t &= \mathbf{C}h_t + \mathbf{D}x_t, \end{aligned} \quad (7)$$

Different objects may exhibit substantial differences in emissivity at the same wavelength, which implies that the energy intensity in HSI can undergo significant spatial variations. Furthermore, as illustrated in Fig. 4, objects exhibit varying degrees of distinction across different spectral bands.




Figure 4: Residual maps across the R, G, and B channels from a CAVE dataset sample.

In SSMs, hidden states capture long-range dependencies by propagating historical information along the sequence. By accumulating and carrying previous data from earlier time steps, they enable

the model to retain past context and effectively establish a global receptive field. However, variations in the spatial distribution of spectral features across different bands can negatively impact the network's generalization ability. As shown in Fig. 4, significant differences may exist between the same (or different) objects across different (or the same) spectral bands. To mitigate the influence of accumulated band-specific information in hidden states, we attempt to suppress the interference from low-correlation regions by generating band-aware masks. According to the Eq. (6), $\Delta \mathbf{A}$ controls the impact of the current input sequence x_t on the hidden states, with a positive correlation. Based on the value of the coefficients in $\Delta \mathbf{A}$, a band-aware mask $\mathbf{M}$ is generated during each sub-band generation

$$\mathbf{M} = \begin{cases} 1, & \Delta \mathbf{A} \geq \epsilon \\ 0, & \Delta \mathbf{A} < \epsilon \end{cases}, \quad (8)$$

$$y_t = (\mathbf{C} \odot \mathbf{M}) h_t, \quad (9)$$

where $\epsilon \in [0, \alpha]$, and α is a hyperparameter. According to Eq. (8) and Eq. (9), features in the hidden states associated with coefficient $\Delta \mathbf{A}$ that fall below a defined threshold are suppressed. As illustrated in Fig. 3 (b), some pixels with low spatial correlation (shaded areas) are filtered out. The overall working mechanism of the Band-aware SSM block (BSSM) is as follows

$$feat_{k+1} = \text{CS}(\text{LN}(feat_k)) \odot \text{SiLU}(\text{LN}(feat_k)) + feat_k. \quad (10)$$

where $feat$ denotes the backbone features and $\text{CS}(\cdot)$ represents the band-aware scanning operation which employs the following operation sequence: $DWConv \rightarrow SiLU \rightarrow SSM \rightarrow LN$.

3.4 Loss Function

In this paper, we use $\mathcal{L}_1$ loss to optimize the reconstructed HSI at the pixel level

$$\mathcal{L}_{rec} = \frac{1}{H \times W \times C} \sum_{i=0} \left| \hat{Y}(i) - Y(i) \right|. \quad (11)$$

where C is the number of channels in the reconstructed HSI, $\hat{Y}(i)$ represents the predicted value for pixel i , and $Y(i)$ is its corresponding ground truth.

4 Empirical Results

4.1 Experimental Settings

Dataset. To validate the effectiveness of the proposed network, we conducted experiments on two datasets. The first dataset is the CAVE dataset [54] provided by Columbia University, which contains 32 HSIs. Each HSI consists of 31 spectral bands with a spectral interval of 10 nm, covering the spectral range from 400 nm to 700 nm. We randomly selected 20 HSIs for training, 6 HSIs for validation, and 6 HSIs for testing. The second dataset is the Harvard dataset [14] provided by Harvard University. It contains 50 HSIs covering both indoor and outdoor scenes. Each HSI consists of 31 spectral bands with a 10 nm interval, covering the spectral range from 420 nm to 720 nm. We randomly selected 30 HSIs for training, 10 HSIs for validation, and 10 HSIs for testing.

Implementation Details. We implemented our network on the PC with a single NVIDIA RTX 4090 GPU and built it in the PyTorch framework. In the training phase, the Adam optimizer [17] was used to optimize the model parameters. The initial learning rate was set to 4×10^{-4} , and the learning rate was decayed using a cosine annealing schedule with a minimum value of 1×10^{-6} . The batch size was set to 32. We cropped 64×64 patches from 3D cubes and input them into the network. We set the number of recursive levels to 5, where each atomic generation module reconstructs a two-channel image. The threshold parameter α in Eq. (8) is empirically set to 0.5.

4.2 Baseline Methods

We compared our FRN with seven SOTA spectral reconstruction methods: HSACS [29], SSRNet [16], HSRNet [20], AWAN [28], MST++ [10], LTRN [15], and MSFN [52].

4.3 Metrics

To evaluate the reconstruction quality of HSIs, we adopt four widely used IQA metrics: peak signal-to-noise ratio (PSNR), structural similarity index (SSIM) [51], root mean square error (RMSE), and universal image quality index (UIQI) [50]. PSNR quantifies reconstruction quality by computing the ratio of signal variance to noise, providing a measure sensitive to pixel-level errors. SSIM assesses the perceptual similarity between the reconstructed image and the ground truth by jointly considering luminance, contrast, and structural information. UIQI evaluates the consistency of pixel distributions by comparing the means and variances of the reconstruction and ground truth. Moreover, we compared the number of parameters of FRN with those methods.

Table 1: Performance comparison of different methods on CAVE and Harvard datasets. The best values are bolded. The up or down arrow indicates a higher or lower metric, corresponding to better performance.

Method	Params(M)	CAVE				Harvard			
		PSNR $\uparrow$	RMSE $\downarrow$	UIQI $\uparrow$	SSIM $\uparrow$	PSNR $\uparrow$	RMSE $\downarrow$	UIQI $\uparrow$	SSIM $\uparrow$
HSACS	19.74	37.9112	5.4099	0.8389	0.9765	42.1360	3.3203	0.8586	0.9762
SSRNet	0.39	38.6807	4.9824	0.8573	0.9794	42.1070	3.5370	0.8608	0.9760
HSRNet	0.77	38.4459	4.6511	0.8527	0.9801	41.6952	3.5459	0.8571	0.9747
MST++	1.62	38.5511	4.6304	0.8731	0.9832	42.4756	3.2552	0.8623	0.9773
AWAN	21.36	39.4262	4.8245	0.8597	0.9798	42.2312	3.5034	0.8616	0.9769
LTRN	0.67	39.7349	4.3095	0.8702	0.9832	42.4953	3.3112	0.8632	0.9770
MSFN	2.48	39.8430	4.0372	0.8877	0.9860	42.6455	2.9916	0.8715	0.9771
Ours	0.30	41.0522	3.6243	0.9010	0.9900	42.8762	2.8933	0.8791	0.9774




Figure 5: Comparison of the reconstruction results of different methods on one scene from the CAVE dataset, including seven SOTA methods and our FRN. We select three bands (20, 25, and 31) for visualization.

4.4 Performance Evaluation

Numerical Results. The quantitative results of different methods on the CAVE and Harvard datasets are presented in Tab. 1. Our method consistently achieves superior performance across all evaluation metrics. On the CAVE dataset, the average PSNR and SSIM of our method reach 41.05 dB and 0.99, respectively, outperforming the second-best results by 1.2 dB and 0.004. On the Harvard dataset, our method achieves a PSNR of 42.87 dB, surpassing the second-best results by 0.23 dB. Additionally, it can be observed that, compared to other methods, FRN achieves higher reconstruction quality while requiring the fewest model parameters.




Figure 6: Comparison of the reconstruction results of different methods on one scene from the Harvard dataset, including seven SOTA methods and our FRN. We select three bands (10, 20, and 31) for visualization.




Figure 7: Comparison of the spectral curves among different methods on CAVE (left) and Harvard (right) datasets.

Visual Results. Fig. 5 and Fig. 6 show the reconstruction results of different methods on the CAVE and Harvard datasets, respectively. We select three channels from each scene for qualitative comparison. To provide a more intuitive comparison, we also present the residual maps between the predicted results and the ground truth. By zooming into local regions, it can be observed that our method reconstructs results that are closer to the ground truth in terms of spatial details and contrast. The corresponding spectral curve in Fig. 7 indicates FRN achieves higher spectral accuracy.

Table 2: Ablation studies of suppression threshold α on CAVE dataset. ‘w/o’ refers to the setting where the band-aware mask is not used (only with vanilla VMamba).

Config	PSNR $\uparrow$	RMSE $\downarrow$	UIQI $\uparrow$	SSIM $\uparrow$
w/o	39.7482	4.2998	0.8758	0.9840
$\alpha = 0.2$	39.9022	4.1562	0.8820	0.9852
$\alpha = 0.3$	40.3285	3.8683	0.8911	0.9866
$\alpha = 0.5$	41.0522	3.6243	0.9010	0.9900
$\alpha = 0.7$	40.1822	4.1022	0.8854	0.9860
$\alpha = 0.8$	37.4822	5.8448	0.7822	0.9573

Table 3: Ablation studies of the number of recursive levels on CAVE dataset. ‘w/o’ means reconstruct the HSI from the RGB image in one step.

Config	PSNR $\uparrow$	RMSE $\downarrow$	UIQI $\uparrow$	SSIM $\uparrow$
w/o	39.8644	3.8683	0.8898	0.9883
$M=2$	40.2210	3.7996	0.9001	0.9884
$M=3$	40.6382	3.7057	0.9004	0.9886
$M=5$	41.0522	3.6243	0.9010	0.9900

Table 4: Ablation studies of the number of reference spectral on CAVE dataset. ‘w/o RGB’ refers to excluding the RGB image from each reconstruction level.

Config	PSNR $\uparrow$	RMSE $\downarrow$	UIQI $\uparrow$	SSIM $\uparrow$
w/o RGB	38.6244	5.1320	0.8448	0.9773
$S=2$	40.3981	3.8862	0.8962	0.9884
$S=3$	40.8286	3.7900	0.8989	0.9886
$S=4$	41.0522	3.6243	0.9010	0.9900
$S=5$	40.8744	3.7458	0.8993	0.9888

4.5 Ablation Study

Band-aware Mask. We investigated the impact of the suppression threshold α . Tab. 2 shows that setting α either too high or too low negatively affects the network performance. When the threshold is too low, the suppression effect is weakened, and low-correlation regions introduce redundant interference to feature learning. Conversely, an excessively high α causes the SSM to suppress informative features, resulting in information loss.

Recursive Levels. We evaluated the impact of the number of recursive levels M on the reconstruction performance on CAVE dataset. By default, we set $M = 5$, which means that each BAMamba generated 2 new spectral channels based on the input from the previous level ($2^5 = 32$). In addition, we also attempted to reconstruct the HSI from the RGB image in one step. Tab. 3 shows that the reconstruction quality improves as the number of recursive levels increases. This is because each atomic generation step deals with fewer unknowns, thereby reducing the difficulty of solving the ill-posed problem.

Spectral Cues. We conducted ablation studies on the number of reference channels (wavelengths), denoted as S , fed into each atomic generator. By default, we input RGB images into each atomic generator to provide spectral priors. In Tab. 4, we found that setting $S = 4$ yields a relatively optimal performance. A smaller value of S limits the amount of reference information available to the network, while a larger value may introduce noisy features. Another notable finding is that excluding the RGB image from each reconstruction level leads to a substantial decline in reconstruction quality. This degradation is attributed to the loss of structural and contrast priors inherently embedded in the RGB image, which are essential for accurate spectral learning.

5 Conclusion

In this paper, we propose a fractal-based recursive spectral reconstruction network (FRN). FRN establishes a new paradigm for spectral reconstruction by recursively invoking an atomic reconstruction module to progressively predict spectra from wide to narrow bands. By introducing the concept of fractals, FRN aligns with the low-rank nature of HSI data and exhibits structural self-similarity across different levels of the network. Furthermore, to alleviate the computational burden caused by the recursive design, we develop BAMamba, an atomic generation module based on SSM. Extensive experiments on multiple datasets demonstrate that FRN achieves outstanding performance in HSI reconstruction. Nevertheless, the recursive calling mechanism introduces significant computational overhead, which represents an important direction for future optimization.

6 Acknowledgments

This work was supported in part by the Dreams Foundation of Jianghuai Advance Technology Center project under Grant 2023-ZM01D002; ; in part by the National Natural Science Foundation of China under Grant 82172073, and Grant 62271430; and in part by the Open Fund of the National Key Laboratory of Infrared Detection Technologies.

References

- [1] N. Akhtar and A. Mian. Hyperspectral recovery from rgb images using gaussian processes. *IEEE transactions on pattern analysis and machine intelligence*, 42(1):100–113, 2018.
- [2] B. Arad and O. Ben-Shahar. Sparse recovery of hyperspectral signal from natural rgb images. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, the Netherlands, October 11–14, 2016, Proceedings, Part VII 14*, pages 19–34. Springer, 2016.
- [3] V. Backman, M. B. Wallace, L. Perelman, J. Arendt, R. Gurjar, M. Müller, Q. Zhang, G. Zonios, E. Kline, T. McGillican, et al. Detection of preinvasive cancer cells. *nature*, 406(6791):35–36, 2000.
- [4] D. S. Bassett, A. Meyer-Lindenberg, S. Achard, T. Duke, and E. Bullmore. Adaptive reconfiguration of fractal small-world human brain functional networks. *Proceedings of the National Academy of Sciences*, 103(51):19518–19523, 2006.
- [5] M. Breuer and J. Albertz. Geometric correction of airborne whiskbroom scanner imagery using hybrid auxiliary data. *International Archives of Photogrammetry and Remote Sensing*, 33(B3/1; PART 3):93–100, 2000.
- [6] E. Bullmore and O. Sporns. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature reviews neuroscience*, 10(3):186–198, 2009.
- [7] Y. Cai, X. Hu, H. Wang, Y. Zhang, H. Pfister, and D. Wei. Learning to generate realistic noisy images via pixel-level noise-aware adversarial training. *Advances in Neural Information Processing Systems*, 34:3259–3270, 2021.
- [8] Y. Cai, J. Lin, X. Hu, H. Wang, X. Yuan, Y. Zhang, R. Timofte, and L. Van Gool. Coarse-to-fine sparse transformer for hyperspectral image reconstruction. In *European conference on computer vision*, pages 686–704. Springer, 2022.
- [9] Y. Cai, J. Lin, X. Hu, H. Wang, X. Yuan, Y. Zhang, R. Timofte, and L. Van Gool. Mask-guided spectral-wise transformer for efficient hyperspectral image reconstruction. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 17502–17511, 2022.
- [10] Y. Cai, J. Lin, Z. Lin, H. Wang, Y. Zhang, H. Pfister, R. Timofte, and L. Van Gool. Mst++: Multi-stage spectral-wise transformer for efficient spectral reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 745–755, 2022.
- [11] Y. Cai, J. Lin, H. Wang, X. Yuan, H. Ding, Y. Zhang, R. Timofte, and L. V. Gool. Degradation-aware unfolding half-shuffle transformer for spectral compressive imaging. *Advances in Neural Information Processing Systems*, 35:37749–37761, 2022.
- [12] X. Cao, T. Yue, X. Lin, S. Lin, X. Yuan, Q. Dai, L. Carin, and D. J. Brady. Computational snapshot multispectral cameras: Toward dynamic capture of the spectral world. *IEEE Signal Processing Magazine*, 33(5):95–108, 2016.
- [13] Y. Cao, C. Liu, Z. Wu, W. Yao, L. Xiong, J. Chen, and Z. Huang. Remote sensing image segmentation using vision mamba and multi-scale multi-frequency feature fusion. *arXiv preprint arXiv:2410.05624*, 2024.
- [14] A. Chakrabarti and T. Zickler. Statistics of real-world hyperspectral images. In *CVPR 2011*, pages 193–200. IEEE, 2011.
- [15] R. Dian, Y. Liu, and S. Li. Spectral super-resolution via deep low-rank tensor representation. *IEEE Transactions on Neural Networks and Learning Systems*, 2024.
- [16] R. Dian, T. Shan, W. He, and H. Liu. Spectral super-resolution via model-guided cross-fusion network. *IEEE Transactions on Neural Networks and Learning Systems*, 2023.

- [17] P. K. Diederik. Adam: A method for stochastic optimization. (*No Title*), 2014.
- [18] H. Du, X. Tong, X. Cao, and S. Lin. A prism-based system for multispectral video acquisition. In *2009 IEEE 12th International Conference on Computer Vision*, pages 175–182. IEEE, 2009.
- [19] A. F. Goetz, G. Vane, J. E. Solomon, and B. N. Rock. Imaging spectrometry for earth remote sensing. *science*, 228(4704):1147–1153, 1985.
- [20] J. He, J. Li, Q. Yuan, H. Shen, and L. Zhang. Spectral response function-guided deep optimization-driven network for spectral super-resolution. *IEEE Transactions on Neural Networks and Learning Systems*, 33(9):4213–4227, 2021.
- [21] W. He, N. Yokoya, and X. Yuan. Fast hyperspectral image recovery of dual-camera compressive hyperspectral imaging via non-iterative subspace-based fusion. *IEEE Transactions on Image Processing*, 30:7170–7183, 2021.
- [22] J. Hu, L. Shen, and G. Sun. Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7132–7141, 2018.
- [23] X. Hu, Y. Cai, J. Lin, H. Wang, X. Yuan, Y. Zhang, R. Timofte, and L. Van Gool. Hdnet: High-resolution dual-domain learning for spectral compressive imaging. in 2022 ieee. In *CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 17521–17530, 2022.
- [24] T. Huang, W. Dong, X. Yuan, J. Wu, and G. Shi. Deep gaussian scale mixture prior for spectral compressive imaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16216–16225, 2021.
- [25] N. Keshava. Distance metrics and band selection in hyperspectral processing with applications to material identification and spectral libraries. *IEEE Transactions on Geoscience and remote sensing*, 42(7):1552–1565, 2004.
- [26] M. J. Khan, H. S. Khan, A. Yousaf, K. Khurshid, and A. Abbas. Modern trends in hyperspectral image analysis: A review. *Ieee Access*, 6:14118–14129, 2018.
- [27] J. Li, Y. Leng, R. Song, W. Liu, Y. Li, and Q. Du. Mformer: Taming masked transformer for unsupervised spectral reconstruction. *IEEE Transactions on Geoscience and Remote Sensing*, 61:1–12, 2023.
- [28] J. Li, C. Wu, R. Song, Y. Li, and F. Liu. Adaptive weighted attention network with camera spectral sensitivity prior for spectral reconstruction from rgb images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 462–463, 2020.
- [29] J. Li, C. Wu, R. Song, Y. Li, W. Xie, L. He, and X. Gao. Deep hybrid 2-d–3-d cnn based on dual second-order attention with camera spectral sensitivity prior for spectral super-resolution. *IEEE Transactions on Neural Networks and Learning Systems*, 34(2):623–634, 2021.
- [30] T. Li, Q. Sun, L. Fan, and K. He. Fractal generative models. *arXiv preprint arXiv:2502.17437*, 2025.
- [31] Y. Li, Y. Shi, K. Wang, B. Xi, J. Li, and P. Gamba. Target detection with unconstrained linear mixture model and hierarchical denoising autoencoder in hyperspectral imagery. *IEEE Transactions on Image Processing*, 31:1418–1432, 2022.
- [32] P. Llull, X. Liao, X. Yuan, J. Yang, D. Kittle, L. Carin, G. Sapiro, and D. J. Brady. Coded aperture compressive temporal imaging. *Optics express*, 21(9):10526–10545, 2013.
- [33] G. Lu and B. Fei. Medical hyperspectral imaging: a review. *Journal of biomedical optics*, 19(1):010901–010901, 2014.
- [34] F. Melgani and L. Bruzzone. Classification of hyperspectral remote sensing images with support vector machines. *IEEE Transactions on geoscience and remote sensing*, 42(8):1778–1790, 2004.
- [35] Z. Meng, M. Qiao, J. Ma, Z. Yu, K. Xu, and X. Yuan. Snapshot multispectral endomicroscopy. *Optics Letters*, 45(14):3897–3900, 2020.
- [36] Z. Meng, Z. Yu, K. Xu, and X. Yuan. Self-supervised neural networks for spectral snapshot compressive imaging. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 2622–2631, 2021.
- [37] R. M. Nguyen, D. K. Prasad, and M. S. Brown. Training-based spectral reconstruction from a single rgb image. In *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part VII 13*, pages 186–201. Springer, 2014.

- [38] M. Parmar, S. Lancel, and B. A. Wandell. Spatio-spectral reconstruction of the multispectral datacube using sparse recovery. In *2008 15th IEEE International Conference on Image Processing*, pages 473–476. IEEE, 2008.
- [39] H. Peng, X. Chen, and J. Zhao. Residual pixel attention network for spectral reconstruction from rgb images. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 486–487, 2020.
- [40] D. Poli and T. Toutin. Review of developments in geometric modelling for high resolution satellite pushbroom sensors. *The Photogrammetric Record*, 27(137):58–73, 2012.
- [41] J. Qiao, J. Liao, W. Li, Y. Zhang, Y. Guo, Y. Wen, Z. Qiu, J. Xie, J. Hu, and S. Lin. Hi-mamba: Hierarchical mamba for efficient image super-resolution. *arXiv preprint arXiv:2410.10140*, 2024.
- [42] K. Ren, X. Wu, L. Xu, and L. Wang. Remotedet-mamba: A hybrid mamba-cnn network for multi-modal object detection in remote sensing images. *arXiv preprint arXiv:2410.13532*, 2024.
- [43] Y. Ren, X. Li, M. Guo, B. Li, S. Zhao, and Z. Chen. Mambacr: Dual-interleaved scanning for compressed image super-resolution with ssms. *arXiv preprint arXiv:2408.11758*, 2024.
- [44] A. Robles-Kelly. Single image spectral reconstruction for multimedia applications. In *Proceedings of the 23rd ACM international conference on Multimedia*, pages 251–260, 2015.
- [45] Z. Shi, C. Chen, Z. Xiong, D. Liu, and F. Wu. Hscnn+: Advanced cnn-based hyperspectral recovery from rgb images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 939–947, 2018.
- [46] M. Shimoni, R. Haelterman, and C. Perneel. Hypersectral imaging for military and security applications: Combining myriad processing and sensing techniques. *IEEE Geoscience and Remote Sensing Magazine*, 7(2):101–117, 2019.
- [47] O. Sporns. Small-world connectivity, motif composition, and complexity of fractal neuronal connections. *Biosystems*, 85(1):55–64, 2006.
- [48] B. Uzktent, A. Rangnekar, and M. Hoffman. Aerial vehicle tracking by adaptive fusion of hyperspectral likelihood maps. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 39–48, 2017.
- [49] A. A. Wagadarikar, N. P. Pitsianis, X. Sun, and D. J. Brady. Video rate spectral imaging using a coded aperture snapshot spectral imager. *Optics express*, 17(8):6368–6388, 2009.
- [50] Z. Wang and A. C. Bovik. A universal image quality index. *IEEE signal processing letters*, 9(3):81–84, 2002.
- [51] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.
- [52] Y. Wu, R. Dian, and S. Li. Multistage spatial–spectral fusion network for spectral super-resolution. *IEEE Transactions on Neural Networks and Learning Systems*, 2024.
- [53] Z. Xiong, Z. Shi, H. Li, L. Wang, D. Liu, and F. Wu. Hscnn: Cnn-based hyperspectral image recovery from spectrally undersampled projections. In *Proceedings of the IEEE international conference on computer vision workshops*, pages 518–525, 2017.
- [54] F. Yasuma, T. Mitsunaga, D. Iso, and S. K. Nayar. Generalized assorted pixel camera: postcapture control of resolution, dynamic range, and spectrum. *IEEE transactions on image processing*, 19(9):2241–2253, 2010.
- [55] Y. Yuan, X. Zheng, and X. Lu. Hyperspectral image superresolution by transfer learning. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 10(5):1963–1974, 2017.
- [56] L. Zhang, Z. Lang, P. Wang, W. Wei, S. Liao, L. Shao, and Y. Zhang. Pixel-aware deep function-mixture network for spectral super-resolution. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 12821–12828, 2020.
- [57] Q. Zhu, Y. Cai, Y. Fang, Y. Yang, C. Chen, L. Fan, and A. Nguyen. Samba: Semantic segmentation of remotely sensed images with state space model. *Heliyon*, 10(19), 2024.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The abstract and introduction clearly articulate the main claims and contributions of the paper, accurately reflecting its scope and content.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: The conclusion section explicitly discusses the limitations of the proposed method.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: All assumptions are clearly stated or referenced in the statement of any theorem.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.

- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: The paper provides a detailed description of the experimental setup and releases the source code for reproducibility.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: The paper provides a detailed description of the experimental setup and releases the source code for reproducibility.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.

- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [\[Yes\]](#)

Justification: The paper provides a detailed description of the experimental setup.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [\[Yes\]](#)

Justification: The paper provides a detailed description and analysis of all experimental results.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [\[Yes\]](#)

Justification: The paper provides a detailed description of the experimental setup.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.

- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: The research conducted in the paper conforms, in every respect, with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: There is no societal impact of the work performed.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The paper poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.

- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [NA]

Justification: The paper does not use existing assets.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: The paper does not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: The core method development in this research does not involve LLMs as any important, original, or non-standard components.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.