

Beyond the Visible: Jointly Attending to Spectral and Spatial Dimensions with HSI-Diffusion for the FINCH Spacecraft

Ian Vyse, Rishit Dagli, Dav Vrat Chadha, John P. Ma, Hector Chen, Isha Ruparelia, Prithvi Seran, Matthew Xie, Eesa Aamer, Aidan Armstrong, Naveen Black, Ben Borstein, Kevin Caldwell, Orrin Dahanaggamaarachchi, Joe Dai, Abeer Fatima, Stephanie Lu, Maxime Michet, Anoushka Paul, Carrie Ann Po, Shivesh Prakash, Noa Prosser, Riddhiman Roy, Mirai Shinjo, Iliya Shofman, Coby Silayan, Reid Sox-Harris, Shuhan Zheng, Khang Nguyen

University of Toronto Aerospace Team - Space Systems¹
 55 St. George Street, Myhal Centre, Room 618, Toronto, ON
 ss-outreach@utat.ca

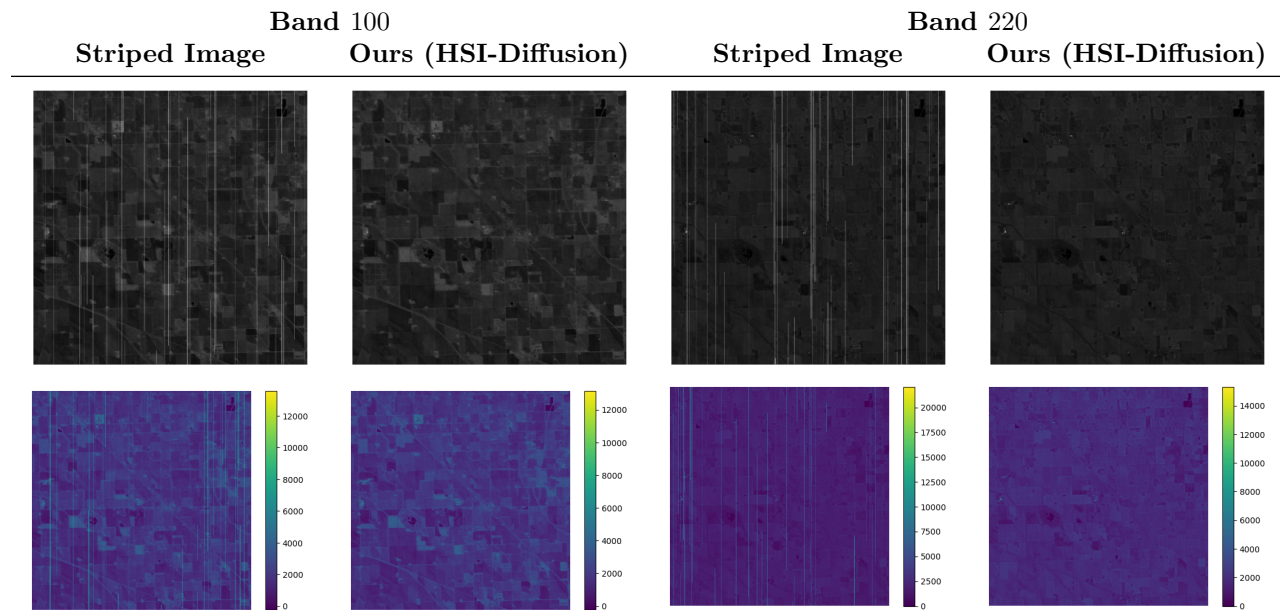


Figure 1: *Hyperpectral Destriping*. We present Hyperspectral Diffusion, a technique that can denoise or destripe satellite hyperspectral data cubes. We demonstrate results from real collected data from the EnMAP hyperspectral satellite mission^{1,2} which is analogous to images we expect to be captured from our FINCH satellite. Best viewed with color and zoom.

ABSTRACT

Satellite remote sensing missions have gained popularity over the past fifteen years due to their ability to cover large swaths of land at regular intervals, making them ideal for monitoring environmental trends. The FINCH mission, a 3U+ CubeSat equipped with a hyperspectral camera, aims to monitor crop residue cover in agricultural fields. Although hyperspectral imaging captures both spectral and spatial information, it is prone to various types of noise, including random noise, stripe noise, and dead pixels. Effective denoising of these images is crucial for downstream scientific tasks. Traditional methods, including hand-crafted techniques encoding strong priors, learned 2D image denoising methods applied across different hyperspectral bands, or diffusion generative models applied independently on bands, often struggle with varying noise strengths across spectral bands, leading to significant spectral distortion. This paper presents a novel approach to hyperspectral image denoising using latent diffusion models that integrate spatial and spectral information.

¹University of Toronto Aerospace Team is a student group affiliated with The University of Toronto, Toronto, Canada.

We particularly do so by building a 3D diffusion model and presenting a 3-stage training approach on real and synthetically crafted datasets. The proposed method preserves image structure while reducing noise. Evaluations on both popular hyperspectral denoising datasets and synthetically crafted datasets for the FINCH mission demonstrate the effectiveness of this approach.

Our code can be found at github.com/utat-ss/FINCH-destriping

1 INTRODUCTION

The FINCH spacecraft^{3,4} (Field Imaging Nanosatellite for Crop residue Hyperspectral mapping) represents a significant advancement in the field of agricultural monitoring via remote sensing technologies. Currently under development by the University of Toronto Aerospace Team’s Space Systems division, this 3U+ CubeSat is designed to precisely estimate the percentage of crop residue cover across Canadian agricultural fields employing advanced spectral unmixing techniques facilitated by its onboard hyperspectral remote sensing payload. The spacecraft integrates a pushbroom sensor^{5,6}, necessitating a robust data processing pipeline to transform raw spectral data into scientifically usable information.

Hyperspectral imaging has been pivotal in remote sensing, providing detailed spectral information across a plethora of contiguous bands, and enabling precise analysis of the Earth’s surface or, in the case of the FINCH mission, crop residue mapping. This technology continues to become more accessible as an “image-based acquisition tool for physically meaningful measurements”⁷. Hyperspectral denoising, critical to ensuring that hyperspectral data is scientifically useful, has undergone significant evolution over the years, moving from traditional image-processing techniques^{8,9} to advanced deep-learning approaches¹⁰⁻¹⁴. However, the deployment of hyperspectral sensors, such as the pushbroom sensor^{5,6} aboard FINCH, introduces inherent challenges, notably the presence of striping and random noise^{15,16}. Striping noise, which results in linear distortions along the sensor’s scanning direction, is akin to artifacts observed in other hyperspectral satellites like EO-1 Hyperion¹⁷. This type of noise, if unaddressed, can severely compromise the spatial and spectral integrity of the data^{8-10,18}, rendering it less reliable for scientific analysis. The removal of striping noise, caused by the operation of the push-broom sensor, is particularly relevant for FINCH, where stripes can compromise both the spatial and spectral integrity of the data collected.

Addressing the striping noise challenge is critical for performing any science on images from the FINCH mission. We present a 3D hyperspectral dif-

fusion model designed to effectively mitigate striping noise. We particularly do so by building a 3D diffusion model and perform a 3-stage training. We train our approach on real and synthetic data. We also compare our approach with existing denoising models, where we find that our approach not only demonstrates superior performance in destriping on satellite images but also shows potential for broader inpainting applications within remote sensing.

The implementation of the hyperspectral diffusion destriping model within the FINCH spacecraft’s operational schema marks a pivotal development in remote sensing technology. By effectively addressing the issue of striping noise in hyperspectral data, this model significantly improves the data quality, thereby enhancing the reliability of remote sensing analyses. The application of deep learning techniques in the noise reduction process not only showcases the model’s technical prowess but also expands the potential for these methodologies in refining data accuracy across various remote sensing applications. Moreover, the scalability of this solution presents opportunities for integration into additional satellite systems, potentially broadening its impact on the field.

Contributions. Our contributions can be summarized as follows,

- We present a hyperspectral diffusion destriping model to remove the various types of striping noise that can occur within FINCH data. Our approach can significantly enhance the quality of data by mitigating striping noise, which is critical for maintaining both the spatial and spectral integrity of hyperspectral remote sensing data.
- We perform comprehensive comparisons with existing denoising models, demonstrating that the proposed model performs favorably in removing various types of synthetic and real noise from hyperspectral data.

2 RELATED WORKS

Noise is a common issue in remote sensing imagery, often appearing in the form of stripes due

to factors such as poor sensor calibration¹⁹. These striping artifacts degrade data quality, cause poor interpretability, and may also cause failures in image processing and analysis systems^{20–29}. Therefore, it is highly valuable to develop algorithms for identifying and correcting or at least alleviating striping noise²⁸. Dead pixel correction is an additional challenge in hyperspectral images (HSI) as dead pixels can significantly affect the data quality of hyperspectral data. Methods that detect and correct these issues ensure the continuity of the spectral information across the image, thereby enhancing the overall quality of the data.

Many traditional approaches to denoising hyperspectral images are statistical-based methods that have been used for denoising typical triband images. Some of these approaches includes using wavelet transformations²⁶, low-rank decomposition³⁰, total variation³¹, and sparse representation³². Although these methods have shown some effectiveness, they are limited by their tendency to generalize to very specific and varying types of stripes that may occur in the model. This generalization can make it difficult to achieve proper restoration across a multitude of cases, leading to the creation of large, complex statistical models that lack interpretability and are difficult to reuse.

Machine learning approaches, particularly those involving convolutional neural networks (CNNs), such as HSI-DeNet³³, have recently been explored for restoring hyperspectral images. Within this context, various approaches have been explored to achieve the denoising task, such as decomposing the task into multiple types of stripes³⁴, as well as employing wavelet transforms³⁵. Another significant approach is the Multi-scale Adaptive Feature Network (MAFNet)³⁶. This method is designed to handle varying noise levels and structures by adaptively learning from the data to identify a complex mapping between noisy and clean hyperspectral imaging³⁷. GLCSA-Net³⁸, which employs a generative adversarial network (GAN) as the backbone, has been developed for describing and denoising based on deep prior models³⁹. It is important to note that hyperspectral images contain low-level information, which has been demonstrated to be present in both two-dimensional and three-dimensional convolutional networks. This information can serve as prior information in various other models⁴⁰.

As transformer-based architectures⁴¹ have been explored for various applications, it is intuitive that many have also considered using them for hyperspectral image (HSI) denoising. The Mixed Attention Network (MAN)⁴² proposes methods to dynamically

control the flow of information through a transformer-based model. One problem with hyperspectral images is information imbalance among spectral-spatial bands. By applying attention to assign importance to skip connections and spectral bands, the model learns to become extremely adept at interpreting complicated hyperspectral feature spaces. The Hyperspectral Denoising Transformer (HSDT)⁴³ builds upon the idea of attentive information flow by modifying the attention blocks to look directly at the full 3D spectral-spatial space. This allows the model to generalize much better than previous models. Several advanced deep learning models employed for tasks such as image super-resolution and reconstruction share various degrees of similarity with hyperspectral denoising. The Spectral Enhanced Rectangle Transformer (SERT)⁴⁴ uses transformers to handle complex noise patterns in HSIs.

Recently, diffusion has been used in various tasks, such as denoising^{45–49}, generating images from text^{50–53}, and inpainting⁵⁴. These methods have also spread to hyperspectral image denoising. Models such as SpectralDiff⁵⁵ have used diffusion as an approach to denoise hyperspectral images. However, it is essential to recognize that it uses 2D convolutions instead of 3D. With the three-dimensional nature of hyperspectral images, where multiple bands can provide cross-information of the noise in the image, it is possible to fully exploit the 3D data in service of the denoising task. This could involve developing models that can simultaneously leverage information across all bands, improving the overall destriping and denoising performance by utilizing the inherent spectral correlations present in hyperspectral images. In this work, we propose a 3D diffusion model to tackle this problem.

3 PRELIMINARY: DIFFUSION MODELS

Diffusion models^{45,56,57} are a class of generative models that iteratively transform noise into structured data through a sequence of denoising steps. Unlike traditional generative models that directly model the data distribution, diffusion models utilize a forward and backward diffusion process to generate data. Specifically, the forward diffusion process gradually adds noise to the input data over a series of time steps, creating increasingly noisy versions of the data. Given an initial data point \mathbf{x}_0 , the forward process produces \mathbf{x}_t by adding Gaussian noise in t discrete steps:

$$q(\mathbf{x}_t|\mathbf{x}_{t-1}) = \mathcal{N}\left(\mathbf{x}_t; \sqrt{1 - \beta_t}\mathbf{x}_{t-1}, \beta_t I\right), \quad (1)$$

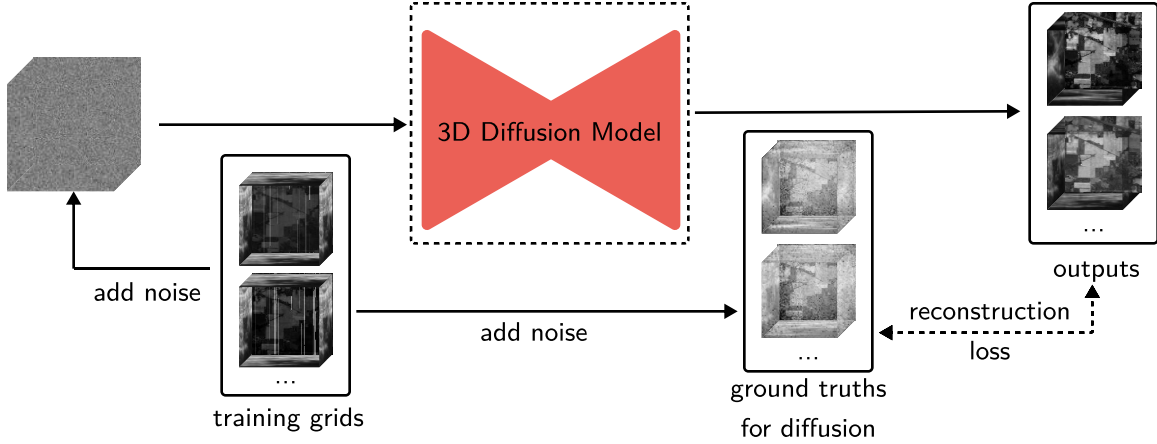


Figure 2: Our method trains a 3D diffusion model across 3 stages to perform destriping on hyperspectral satellite images like the ones that would be captured by FINCH.

where β_t is a variance schedule that controls the amount of noise added at each step. The backward diffusion process, also known as the denoising process, reverses this noising process. It aims to reconstruct the original data from the noisy versions. This is achieved by training a neural network ϵ_θ to predict the noise added at each step. The model is trained to minimize the difference between the predicted noise and the actual noise added in the forward process, often using a mean squared error loss:

$$\mathbb{E}_{\mathbf{x}_0, t, \epsilon \sim \mathcal{N}(0, I)} [\|\epsilon - \epsilon_\theta(\mathbf{x}_t, t)\|_2^2], \quad (2)$$

where \mathbf{x}_t is the noisy image at time step t , and ϵ is the added Gaussian noise. During inference, the model starts with a random noise \mathbf{x}_T and iteratively denoises it using the trained model ϵ_θ , producing a data sample \mathbf{x}_0 that resembles the training data.

4 METHODS

We next describe our approach to train our diffusion model (Section 4.1) and create the synthetic training data (Section 4.2).

4.1 Training the Diffusion Model

Following recent progress in generative models^{45, 51, 56, 57}, we build a 3D diffusion model for hyperspectral image destriping that operates on $32 \times 32 \times 32$ occupancy grids of hyperspectral images. To do so, we adopt a three-stage training approach, which we describe. Initially, we adapt a U-Net architecture, previously employed for the Pavia University⁵⁸ dataset’s 2D image segmentation, to support 3D

hyperspectral input. This adaptation includes an expansion of the model’s convolutional layers to process an additional spectral dimension, thus accommodating the intricate characteristics of hyperspectral images. Subsequently, the model undergoes a denoising training phase using the ICVL-HSI³² dataset, which reflects diverse real-world noise conditions, excluding satellite hyperspectral imagery. The final stage involves fine-tuning our model on a synthetic dataset derived from multiple hyperspectral scenes, ensuring robust performance across various spectral scenarios. We illustrate this in Figure 2.

Stage I. We start by first reconfiguring a U-Net trained on the Pavia University dataset⁵⁸ for image segmentation. We expand the underlying convolutional layers by adding a channel and modifying the weights denoted as $\mathbf{W}^{2D} \in \mathbb{R}^{m \times n \times c_{in} \times c_{out}}$, into weights for the new network with an additional channel denoted as, $\mathbf{W}^{3D} \in \mathbb{R}^{m \times n \times l \times c_{in} \times c_{out}}$. We do so by replicating 2D filters to expand into the spectral domain by averaging these filters across the additional spectral dimension,

$$\mathbf{W}^{3D}[:, :, k, :, :] = \frac{1}{l} \sum_{i=1}^l \mathbf{W}^{2D}_{\dots, i, \dots}, \quad (3)$$

$$\forall k \in 1, \dots, l,$$

where l represents the number of spectral layers. This transformation allows us to incorporate the essential spectral information for processing hyperspectral images.

¹We were unable to reproduce results from most of the methods and we take metrics directly from the papers.

Table 1: *Results on Real Data.* We report the PSNR, SSIM, SAM, and LPIPS metrics for different models¹. The **best**, **second best**, and **third best** results for each metric are color coded.

Modification	Model	PSNR \uparrow	SSIM \uparrow	SAM \downarrow	LPIPS \downarrow
-	GRN-HSI ⁵⁹	40.58	0.9840	0.0730	-
	Deep HS Prior 3D ⁴⁰	23.24	0.8520	9.9100	-
	MAFNet ³⁷	42.15	0.9630	0.0290	-
	SERT ⁶⁰	43.68	0.9969	1.9700	-
	HSDT ⁶¹	44.69	0.9940	0.0380	-
	MAN ⁴²	43.44	0.9890	0.0440	-
	Ours (HSI-Diffusion)		39.2274	0.8817	0.0423
Non-i.i.d Gaussian Noise w/ $\sigma \in [0, 15]$	BM4D ⁶²	44.39	0.9683	0.0692	-
	MTSNMF ⁶³	45.39	0.9592	0.0845	-
	LLRT ⁶⁴	45.74	0.9657	0.0832	-
	NGMeet ⁶⁵	39.63	0.8612	0.2144	-
	LRMR ⁶⁶	41.50	0.9356	0.1289	-
	FastHyDe ⁶⁷	48.08	0.9917	0.0404	-
	LRTF L ₀ ⁶⁸	43.41	0.9315	0.0570	-
	E-3DTV ⁶⁹	46.05	0.9811	0.0560	-
	T3SC ⁷⁰	49.68	0.9912	0.0486	-
	MAC-Net ⁷¹	48.21	0.9915	0.0387	-
	NSSNN ⁷²	49.83	0.9934	0.0302	-
	TRQ3D ⁷³	46.43	0.9878	0.0437	-
	SST ⁷⁴	50.87	0.9938	0.0298	-
	SSRT-UNET ⁷⁵	52.12	0.9950	0.0225	-
	SSUMamba ⁷⁶	51.34	0.9946	0.0256	-
	Ours (HSI-Diffusion)		38.3892	0.9864	0.0474

Stage II. We now train the model from Stage I for denoising on the ICVL-HSI dataset³². The ICVL-HSI dataset³² provides a variety of real-world scenarios reflective of common noise patterns in hyperspectral imaging. However, this dataset does not include any satellite hyperspectral images. We introduce this step to perform large-scale pre-training since our synthetic dataset is very small in size and is sampled from only 7 hyperspectral datacubes. During this stage, we train the network as a 3D denoising diffusion model with the objective,

$$\mathcal{L}_{\text{denoise}}(\theta) = \mathbb{E}_{\mathbf{X}, \mathbf{N}} \left[|\epsilon_{\theta}(\mathbf{X}_{\text{noisy}}, t) - \mathbf{N}|_2^2 \right], \quad (4)$$

where $\mathbf{X}_{\text{noisy}} = \mathbf{X}_{\text{clean}} + \mathbf{N}$ and $\mathbf{N} \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$.

Stage III. We now fine-tune the model from Stage II using a synthetic dataset derived from the Hyperspectral Remote Sensing Scenes database⁵⁸ comprised of the Indian Pines⁷⁷, Salinas⁵⁸, Pavia Centre⁵⁸, Pavia university⁵⁸, Cuprite⁵⁸, KSC⁵⁸, and Botswana⁵⁸ datasets. We describe the creation of

this synthetic dataset in Section 4.2. We fine-tune the 3D denoising diffusion model from stage II using the following objective,

$$\mathcal{L}_{\text{fine-tune}}(\theta) = \mathbb{E}_{\mathbf{X}, \mathbf{N}} \left[|\epsilon_{\theta}(\mathbf{X}_{\text{noisy}}, t) - \mathbf{N}|_2^2 \right]. \quad (5)$$

4.2 Creating Synthetic Data

Given the absence of real striped hyperspectral satellite data, we devise a method to synthesize such data for model training. This approach is essential for our denoising models, which rely on accurately replicating and subsequently removing common artifacts like stripes from hyperspectral images. Our methodology begins with ground-truth hyperspectral cubes which are free from any artifacts, providing a clean baseline from which noise can be methodically introduced.

To emulate realistic stripe noise, we introduce Gaussian noise stripes to the clean hyperspectral

Table 2: Results on Synthetic Data. We report the PSNR, SSIM, SAM, and LPIPS metrics for various methods tested on synthetic datasets for the FINCH mission. The **best**, **second best**, and **third best** results for each metric are color coded.

Method	Pavia				WDC			
	PSNR	SSIM	SAM	LPIPS	PSNR	SSIM	SAM	LPIPS
BM4D ⁶²	25.0946	0.6944	0.2325	-	24.5261	0.5675	0.2198	-
LRMR ⁶⁶	31.3789	0.8760	0.1746	-	31.3062	0.8015	0.1255	-
LRTV ⁷⁸	33.5149	0.9078	0.1566	-	32.7237	0.8561	0.0939	-
LRTF-DFR ⁷⁹	33.3251	0.9286	0.1062	-	33.5054	0.9059	0.0668	-
SSLR-SSTV ²³	35.9781	0.9674	0.0860	-	34.6840	0.9494	0.0806	-
Ours (HSI-Diffusion)	38.9210	0.9789	0.0923	0.1531	36.2172	0.9473	0.0928	0.2213

data. The stripe noise is generated as follows,

$$\mathbf{S}_i(x, y, z) = \mathbf{G}(0, \sigma^2) \cdot \text{intensity_factor}(i), \quad (6)$$

where $\mathbf{G}(0, \sigma^2)$ represents a Gaussian noise matrix with mean 0 and variance σ^2 , and $\text{intensity_factor}(i)$ denotes the stripe intensity for the i -th band. The stripe intensity is calculated as a percentage of the dynamic range of each spectral band, which varies from 0.1% to 5%. Specifically, $\text{intensity_factor}(i)$ is determined by,

$$\text{intensity_factor}(i) = \max(0.001, \text{stripe_intensity} - 0.05 + \mathcal{U}(0, 0.1)), \quad (7)$$

where stripe_intensity is sampled from a uniform distribution $\mathcal{U}(0.01, 0.3)$ initially, and $\mathcal{U}(a, b)$ represents a uniform random variable between a and b . Based on the stripe frequency and with uniform probability, columns are randomly selected for striping. These stripes are fragmented across the column, with the number of fragments following a uniform distribution. In addition, the stripes can vary in size and may not extend throughout the column, and the size is also uniformly distributed. Each band will have stripes generated with this method.

Following the introduction of synthetic stripe noise, we sample $32 \times 32 \times 32$ cubes from the larger hyperspectral images. These cubes are randomly sampled across spatial and spectral dimensions to ensure variability in training data, thus preparing the model to handle a wide range of real-world scenarios.

To generate additional data, we also use augmentation techniques such as CutMix⁸⁰ and Mixup⁸¹, which we apply specifically in the frequency bands of the hyperspectral cubes. For each sampled hyperspectral cube, we apply the frequency-specific

augmentation as follows:

$$\mathbf{X}_{\text{aug}}^{(f)} = \lambda \mathbf{X}_1^{(f)} + (1 - \lambda) \mathbf{X}_2^{(f)}, \quad (8)$$

where f indicates the frequency band, $\mathbf{X}_1^{(f)}$ and $\mathbf{X}_2^{(f)}$ are two randomly chosen hyperspectral cubes, and λ is sampled from a Beta distribution, $\text{Beta}(\alpha, \alpha)$ for CutMix and a uniform distribution $\mathcal{U}(0, 1)$ for Mixup. Here, α is a hyperparameter to adjust the amount of mixing.

5 EXPERIMENTS

We also perform experiments to assess the performance of our model. We perform experiments on the ICVL-HSI³² dataset where we evaluate our method on performing denoising. Many modern approaches evaluate on the ICVL dataset³² by adding a synthetic Gaussian noise and we evaluate our model on this benchmark too to compare our approach with other models. While we present these results for completeness, performing well on the ICVL dataset³² is not a goal of this work, as our method aims to solely perform well on satellite hyperspectral images. We also evaluate our method on our synthetic dataset created with the approach outlined in Section 4.2 on remote sensing hyperspectral images. Furthermore, we also evaluate our method on the EnMAP datacube in the region that our FINCH mission would image that is thus comparable to our satellite’s destriping needs.

All of our final code was optimized for a 1 x A100-80 GB GPU. For our Diffusion model, we trained our model for 30K steps per stage with a batch size of 32 on the $32 \times 32 \times 32$ resolution. We use the Adam optimizer kingma2017adam with $\beta_1 = 0.9$, $\beta_2 = 0.999$, and $\epsilon = 10^{-8}$ with an initial learning rate of 10^{-4} . We use cosine decay for the learning rate. We use the same optimization setup throughout all of our stage, however, we scale the learning rate according




Figure 3: *Qualitative Results.* We present striped (Row 1) and the corresponding destriped bands produced from our methods (Row 2) from the EnMAP^{1,2} image of FINCH’s imaging site, with stripes added as described in Section 4.2. From left to right, the images correspond to bands 100, 124, 148, 172, 196, and 220. Best viewed with zoom.

to our dataset size for each of the stages.

5.1 Quantitative Results

We demonstrate our results on the ICVL-HSI dataset³² in Table 1 and our results on the synthetically crafted satellite hyperspectral images in Table 2. We evaluate our method and other approaches based on the PSNR, SSIM, spectral angle mapper (SAM)⁸³, and LPIPS⁸⁴ metrics which are commonly used to evaluate denoising. While our method is unable to perform well on the ICVL-HSI dataset³² we report these for completeness and not as the focus of our results. We demonstrate that our approach beats state-of-the-art approaches on the synthetic dataset.

5.2 Qualitative Results

We present qualitative results on the EnMAP datacube^{1,2} from the region that would be imaged by our FINCH mission in Figure 3 and Figure 4. These images appear consistent with the original images and maintain scene data such as field boundaries, which are important for the FINCH mission. Of critical importance to the FINCH mission is that the destriped images visually appear to maintain similar intensity values to their striped counterparts (besides the striping artifacts). This suggests that our model preserves spectral information which is critical for achieving FINCH’s mission of crop residue mapping.

Of note for the FINCH mission is the PSNR that our model achieved on the task of destriping remote sensing data. The PSNRs of 36 and above are well above the FINCH mission’s requirement of a SNR of 30 for performing scientific retrieval³.

6 DISCUSSION AND CONCLUSION

We introduce a 3D diffusion model specifically tailored for hyperspectral remote sensing data from the FINCH spacecraft. We train the model on widely available hyperspectral data and fine-tune the training with remote sensing data synthetically striped to prepare the model for the task of destriping FINCH data. We observe favorable results in the task of destriping hyperspectral remote sensing data.

Limitations. The model has not been trained or tested on actual FINCH data, so we cannot be certain of how it will perform on said data. The model will likely need fine-tuning once FINCH data becomes available for several reasons: we expect FINCH data to have a much lower spatial resolution and larger swath width than existing datasets, and we expect that FINCH’s spectral range and spectral resolution may be different from many of the publicly available datasets. Regarding the spatial resolution and swath width, we can expect that FINCH data may present different features or the same features at different scales than existing hyperspectral datasets. This may require a degree of retraining in our model. Similarly, regarding the spectral range and resolution, since many of the publicly available remote sensing datasets may cover spectral ranges that only partly overlap with FINCH’s spectral range, we may see spectral features in FINCH data that were not seen during training or said features may appear at a resolution that was not seen during training. This last limitation is particularly relevant for the FINCH

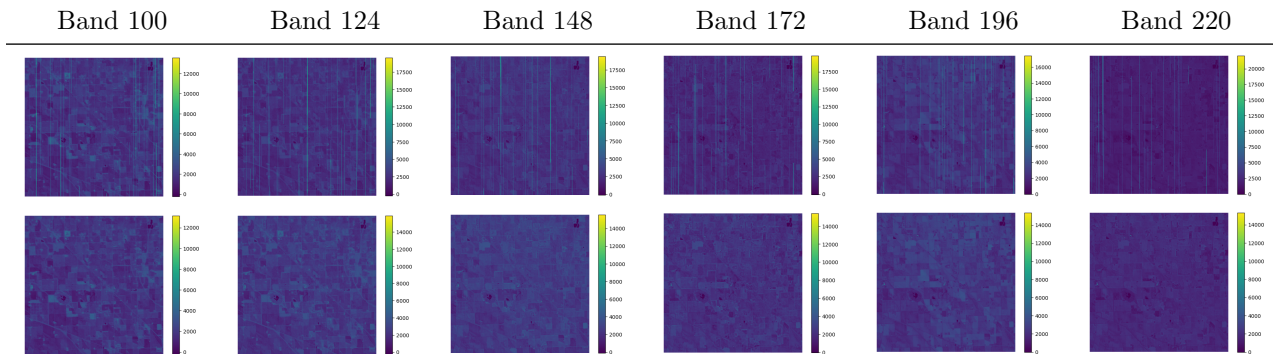


Figure 4: *Qualitative Results.* We present striped (Row 1) and the corresponding destriped (Row 2) bands from the EnMAP^{1,2} image of FINCH’s imaging site, with stripes added as described in Section 4.2. From top left to bottom right, the images correspond to bands 100, 124, 148, 172, 196, and 220. We present them in the Viridis⁸² colour scheme in this figure, with intensity bars to aid the reader in understanding the image. Best viewed in color and in conjunction with Figure 3.

mission, where accurate spectral information is critical for achieving the mission’s scientific objective, as mentioned earlier.

ety³ for their substantial support through grants and funding.

Future Applications. Within the context of the FINCH mission, there are several opportunities for future work with this model. The small form factor and focal length of the FINCH spacecraft and payload motivate the requirement for a super-resolution technique to improve the spatial resolution of FINCH data for tasks such as georeferencing and imaging site identification. Additionally, different types of noise, such as dark current noise, may require more sophisticated machine learning-based denoising approaches, depending on the precision of thermal sensors employed within the final space-ready version of FINCH. Beyond our FINCH mission, our approach would be very useful for denoising hyperspectral images in general. Our hyperspectral diffusion model may lend itself well to being applied to both of these tasks in the future. Finally, future work could analyze and fine-tune this model’s performance on actual FINCH data, after FINCH has been launched and has started operating.

Acknowledgements

We gratefully acknowledge the University of Toronto Student Union and the student body for their generous funding of UTAT-Space Systems through the UTAT-Innovation Fund². Additionally, we also thank the University of Toronto Engineering Soci-

²<https://www.utsu.ca/u-of-t-service-groups/>

³<https://skule.ca/>

References

- [1] Tobias Storch, Hans-Peter Honold, Sabine Chabrilat, Martin Habermeyer, Paul Tucker, Maximilian Brell, Andreas Ohndorf, Katrin Wirth, Matthias Betz, Michael Kuchler, Helmut Mühle, Emiliano Carmona, Simon Baur, Martin Mücke, Sebastian Löw, Daniel Schulze, Steffen Zimmermann, Christoph Lenzen, Sebastian Wiesner, Saika Aida, Ralph Kahle, Peter Willburger, Sebastian Hartung, Daniele Dietrich, Nicolae Plesia, Mirco Tegler, Katharina Schork, Kevin Alonso, David Marshall, Birgit Gerasch, Peter Schwind, Miguel Pato, Mathias Schneider, Raquel de los Reyes, Maximilian Langheinrich, Julian Wenzel, Martin Bachmann, Stefanie Holzwarth, Nicole Pinnel, Luis Guanter, Karl Segl, Daniel Scheffler, Saskia Foerster, Niklas Bohn, Astrid Bracher, Mariana A. Soppa, Ferran Gascon, Rob Green, Raymond Kokaly, Jose Moreno, Cindy Ong, Manuela Sornig, Ricarda Wernitz, Klaus Bagschik, Detlef Reintsema, Laura La Porta, Anke Schickling, and Sebastian Fischer. The enmap imaging spectroscopy mission towards operations. *Remote Sensing of Environment*, 294:113632, 2023.
- [2] Luis Guanter, Hermann Kaufmann, Karl Segl, Saskia Foerster, Christian Rogass, Sabine Chabrilat, Theres Kuester, André Hollstein, Godela Rossner, Christian Chlebek, Christoph Straif, Sebastian Fischer, Stefanie Schrader, Tobias Storch, Uta Heiden, Andreas Mueller, Martin Bachmann, Helmut Mühle, Rupert Müller, Martin Habermeyer, Andreas Ohndorf, Joachim Hill, Henning Buddenbaum, Patrick Hostert, Sebastian Van der Linden, Pedro J. Leitão, Andreas Rabe, Roland Doerffer, Hajo Krasemann, Hongyan Xi, Wolfram Mauser, Tobias Hank, Matthias Locherer, Michael Rast, Karl Staenz, and Bernhard Sang. The enmap spaceborne imaging spectroscopy mission for earth observation. *Remote Sensing*, 7(7):8830–8857, 2015.
- [3] University of Toronto Aerospace Team. Finch: A blueprint for accessible and scientifically valuable remote sensing satellite missions. In *36th Annual Small Satellite Conference*, 2022.
- [4] University of Toronto Aerospace Team. Hot & cold: The finch thermal simulator. In *37th Annual Small Satellite Conference*, 2023.
- [5] R. Gupta and R.I. Hartley. Linear pushbroom cameras. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(9):963–975, 1997.
- [6] James E. Fowler. Compressive pushbroom and whiskbroom sensing for hyperspectral remote-sensing imaging. In *2014 IEEE International Conference on Image Processing (ICIP)*, pages 684–688, 2014.
- [7] M. H. Kim. The three-dimensional evolution of hyperspectral imaging. *Springer EBooks*, pages 64, 84, 2015.
- [8] Qiangqiang Yuan, Liangpei Zhang, and Huanfeng Shen. Hyperspectral image denoising employing a spectral–spatial adaptive total variation model. *IEEE Transactions on Geoscience and Remote Sensing*, 50(10):3660–3677, 2012.
- [9] Wei He, Quanming Yao, Chao Li, Naoto Yokoya, and Qibin Zhao. Non-local meets global: An integrated paradigm for hyperspectral denoising. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [10] Alessandro Maffei, Juan M. Haut, Mercedes Eugenia Paoletti, Javier Plaza, Lorenzo Bruzzone, and Antonio Plaza. A single model cnn for hyperspectral image denoising. *IEEE Transactions on Geoscience and Remote Sensing*, 58(4):2516–2529, 2020.
- [11] Qian Shi, Xiaopei Tang, Taoru Yang, Rong Liu, and Liangpei Zhang. Hyperspectral image denoising using a 3-d attention denoising network. *IEEE Transactions on Geoscience and Remote Sensing*, 59(12):10348–10363, 2021.
- [12] Yong-Qiang Zhao and Jingxiang Yang. Hyperspectral image denoising via sparse representation and low-rank constraint. *IEEE Transactions on Geoscience and Remote Sensing*, 53(1):296–308, 2015.
- [13] Weisheng Dong, Huan Wang, Fangfang Wu, Guangming Shi, and Xin Li. Deep spatial–spectral representation learning for hyperspectral image denoising. *IEEE Transactions on Computational Imaging*, 5(4):635–648, 2019.
- [14] Qiangqiang Yuan, Qiang Zhang, Jie Li, Huanfeng Shen, and Liangpei Zhang. Hyperspectral image denoising employing a spatial–spectral deep residual convolutional neural network. *IEEE Transactions on Geoscience and Remote Sensing*, 57(2):1205–1218, 2019.

- [15] Luis Gómez-Chova, Luis Alonso, Luis Guanter, Gustavo Camps-Valls, Javier Calpe, and José Moreno. Correction of systematic spatial noise in push-broom hyperspectral sensors: application to chris/proba images. *Appl. Opt.*, 47(28):F46–F60, Oct 2008.
- [16] Christian Rogass, Christian Mielke, Daniel Scheffler, Nina K Boesche, Angela Lausch, Christin Lubitz, Maximilian Brell, Daniel Spengler, Andreas Eisele, Karl Segl, et al. Reduction of uncorrelated striping noise—applications for hyperspectral pushbroom acquisitions. *Remote Sensing*, 6(11):11082–11106, 2014.
- [17] T. McVicar, T. Van Niel, D. Jupp, and J. Pearlman. Preprocessing eo-1 hyperion hyperspectral data to support the application of agricultural indexes. *Geoscience and Remote Sensing, IEEE Transactions on*, 41:1246–1259, 07 2003.
- [18] Tao Zhang, Ying Fu, and Cheng Li. Hyperspectral image denoising with realistic data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2248–2257, October 2021.
- [19] P. Rakwatin, W. Takeuchi, and Y. Yasuoka. Stripe noise reduction in modis data by combining histogram matching with facet filter. *IEEE Transactions on Geoscience and Remote Sensing*, 45(6):1844–1856, 2007.
- [20] Jianxin Jia, Xiaorou Zheng, Shanxin Guo, Yueming Wang, and Jinsong Chen. Removing stripe noise based on improved statistics for hyperspectral images. *IEEE Geoscience and Remote Sensing Letters*, 19:1–5, 2022.
- [21] Pablo Meza, Jorge E. Pezoa, and Sergio N. Torres. Multidimensional striping noise compensation in hyperspectral imaging: Exploiting hypercubes’ spatial, spectral, and temporal redundancy. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 9(9):4428–4441, 2016.
- [22] Yi Chang, Luxin Yan, Tao Wu, and Sheng Zhong. Remote sensing image stripe noise removal: From image decomposition perspective. *IEEE Transactions on Geoscience and Remote Sensing*, 54(12):7018–7031, 2016.
- [23] Fang Yang, Xin Chen, and Li Chai. Hyperspectral image destriping and denoising using stripe and spectral low-rank matrix recovery and global spatial-spectral total variation. *Remote Sensing*, 13(4), 2021.
- [24] Bing Zhang, Yashinov Aziz, Zhicheng Wang, Lina Zhuang, Michael K. Ng, and Lianru Gao. Hyperspectral image stripe detection and correction using gabor filters and subspace representation. *IEEE Geoscience and Remote Sensing Letters*, 19:1–5, 2022.
- [25] Chengjun Wang, Miao Zhong Xu, Yonghua Jiang, Guo Zhang, Hao Cui, Litao Li, and Da Li. Translution-snet: A semisupervised hyperspectral image stripe noise removal based on transformer and cnn. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–14, 2022.
- [26] Roshan Pande-Chhetri and Amr Abd-Elrahman. De-striping hyperspectral imagery using wavelet transform and adaptive frequency domain filtering. *ISPRS Journal of Photogrammetry and Remote Sensing*, 66(5):620–636, 2011.
- [27] Qiang Zhang, Qiangqiang Yuan, Jie Li, Xinxin Liu, Huanfeng Shen, and Liangpei Zhang. Hybrid noise removal in hyperspectral imagery with a spatial-spectral gradient network. *IEEE Transactions on Geoscience and Remote Sensing*, 57(10):7317–7329, 2019.
- [28] F. Tsai and Chen W.W. Striping noise detection and correction of remote sensing images. *IEEE Transactions on Geoscience and Remote Sensing*, 46(12):4122–4131, 2008.
- [29] Lina Zhuang and Michael K. Ng. Fasthymix: Fast and parameter-free hyperspectral image mixed noise removal. *IEEE Transactions on Neural Networks and Learning Systems*, 34(8):4702–4716, 2023.
- [30] Y. Chen, T Huang, and X.L. Zhao. Destriping of multispectral remote sensing image using low-rank tensor decomposition. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 11(12):4950–4967, 2018.
- [31] Y. Chang, L. Yan, H. Fang, and C. Luo. Anisotropic spectral-spatial total variation model for multispectral remote sensing image destriping. *IEEE Transactions on Image Processing*, 24(6):1852–1866, 2015.
- [32] B. Arad and O. Ben-Shahar. Sparse recovery of hyperspectral signal from natural rgb images. In B. Leibe, J. Matas, N. Sebe, and M. Welling, editors, *Computer Vision – ECCV 2016*, pages 19–34, Cham, 2016. Springer International Publishing.

- [33] Y. Chang, L. Yan, H. Fang, S. Zhong, and W. Liao. Hsi-denet: Hyperspectral image restoration via convolutional neural network. *IEEE Transactions on Geoscience and Remote Sensing*, 57(2):667–682, 2019.
- [34] E. Pan, Y. Ma, X. Mei, J. Huang, Q. Chen, and J. Ma. Hyperspectral image destriping and denoising from a task decomposition view. *Pattern Recognition*, 144:109832, 2023.
- [35] E. Pan, Y. Ma, X. Mei, F. Fan, J. Huang, and J. Ma. Progressive hyperspectral image destriping with an adaptive frequencial focus. *IEEE Transactions on Geoscience and Remote Sensing*, 61:1–12, 2023.
- [36] Ronghua Shang, Jiyu Zhang, Licheng Jiao, Yangyang Li, Naresh Marturi, and Rustam Stolkin. Multi-scale adaptive feature fusion network for semantic segmentation in remote sensing images. *Remote Sensing*, 12(5), 2020.
- [37] H. Pan, F. Gao, J. Dong, and Q. Du. Multiscale adaptive fusion network for hyperspectral image denoising. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 16:3045–3059, 2023.
- [38] H. Chen, J. Li, J. Zhang, Y. Fu, C. Yan, and D. Zeng. Glcsa-net: global–local constraints-based spectral adaptive network for hyperspectral image inpainting. *The Visual Computer*, 40(5):3331–3346, 2024.
- [39] H. Chen, J. Li, J. Zhang, Y. Fu, C. Yan, and D. Zeng. Glcsa-net: global–local constraints-based spectral adaptive network for hyperspectral image inpainting. *The Visual Computer*, 40(5):3331–3346, 2024.
- [40] O. Sidorov and J.Y. Hardeberg. Deep hyperspectral prior: Denoising, inpainting, super-resolution. *CoRR*, abs/1902.00301, 2019.
- [41] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need, 2023.
- [42] Z. Lai and Y. Fu. Mixed attention network for hyperspectral image denoising. *arXiv preprint arXiv:2301.11525*, 2023.
- [43] Z. Lai, Y. Chenggang, and Y. Fu. Hybrid spectral denoising transformer with guided attention. In *Proceedings of the IEEE International Conference on Computer Vision*, 2023.
- [44] Amin Shoari Nejad, Rocío Alaiz-Rodríguez, Gerard D. McCarthy, Brian Kelleher, Anthony Grey, and Andrew Parnell. Sert: A transformer based model for spatio-temporal sensor data with missing values for environmental monitoring, 2023.
- [45] J. Ho, A. Jain, and P. Abbeel. Denoising diffusion probabilistic models, 2020.
- [46] Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising diffusion restoration models. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 23593–23606. Curran Associates, Inc., 2022.
- [47] Yuanzhi Zhu, Kai Zhang, Jingyun Liang, Jiezhong Cao, Bihan Wen, Radu Timofte, and Luc Van Gool. Denoising diffusion models for plug-and-play image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 1219–1229, June 2023.
- [48] Vladimir Kulikov, Shahar Yadin, Matan Kleiner, and Tomer Michaeli. SinDDM: A single image denoising diffusion model. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett, editors, *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 17920–17930. PMLR, 23–29 Jul 2023.
- [49] Rishit Dagli. Diffuseraw: End-to-end generative raw image processing for low-light images, 2023.
- [50] Bowen Li, Xiaojuan Qi, Thomas Lukasiewicz, and Philip Torr. Controllable text-to-image generation. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [51] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10684–10695, June 2022.
- [52] L. Zhang, A. Rao, and M. Agrawala. Adding conditional control to text-to-image diffusion models, 2023.

- [53] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 8821–8831. PMLR, 18–24 Jul 2021.
- [54] A. Lugmayr, M. Danelljan, A. Romero, F. Yu, R. Timofte, and Gool L. Van. Repaint: Inpainting using denoising diffusion probabilistic models, 2022.
- [55] N. Chen, J. Yue, L. Fang, and S. Xia. Spectraldiff: A generative framework for hyperspectral image classification with diffusion models. *IEEE Transactions on Geoscience and Remote Sensing*, 61:1–16, 2023.
- [56] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pages 2256–2265. PMLR, 2015.
- [57] Y. Song and S. Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32, 2019.
- [58] M. Graña, M. Veganzons, and B. Ayerdi. Hyperspectral remote sensing scenes [database] from grupo de inteligencia computacional. 2021.
- [59] Xiangyong Cao, Xueyang Fu, Chen Xu, and Deyu Meng. Deep spatial-spectral global reasoning network for hyperspectral image denoising. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–14, 2022.
- [60] Miaoyu Li, Ji Liu, Ying Fu, Yulun Zhang, and Dejing Dou. Spectral enhanced rectangle transformer for hyperspectral image denoising. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5805–5814, 2023.
- [61] Zeqiang Lai, Chenggang Yan, and Ying Fu. Hybrid spectral denoising transformer with guided attention. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13065–13075, 2023.
- [62] M. Maggioni, V. Katkovnik, K. Egiazarian, and A. Foi. Nonlocal transform-domain filter for volumetric data denoising and reconstruction. *IEEE Transactions on Image Processing*, 22(1):119–133, 2013.
- [63] M. Ye, Y. Qian, and J. Zhou. Multitask sparse nonnegative matrix factorization for joint spectral-spatial hyperspectral imagery denoising. *IEEE Transactions on Geoscience and Remote Sensing*, 53(5):2621–2639, 2015.
- [64] Y. Chang, L. Yan, and S. Zhong. Hyperlaplacian regularized unidirectional low-rank tensor recovery for multispectral image denoising. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [65] W. He, Q. Yao, C. Li, N. Yokoya, Q. Zhao, H. Zhang, and L. Zhang. Non-local meets global: An iterative paradigm for hyperspectral image restoration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(4):2089–2107, 2022.
- [66] H. Zhang, W. He, L. Zhang, H. Shen, and Q. Yuan. Hyperspectral image restoration using low-rank matrix recovery. *IEEE Transactions on Geoscience and Remote Sensing*, 52(8):4729–4743, 2014.
- [67] L. Zhuang and J.M. Bioucas-Dias. Fast hyperspectral image denoising and inpainting based on low-rank and sparse representations. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 11(3):730–742, 2018.
- [68] F. Xiong, J. Zhou, and Y. Qian. Hyperspectral restoration via l_0 gradient regularized low-rank tensor factorization. *IEEE Transactions on Geoscience and Remote Sensing*, 57(12):10410–10425, 2019.
- [69] J. Peng, Q. Xie, Q. Zhao, Y. Wang, L. Yee, and D. Meng. Enhanced 3d tv regularization and its applications on hsi denoising and compressed sensing. *IEEE Transactions on Image Processing*, 29:7889–7903, 2020.
- [70] T. Bodrito, A. Zouaoui, J. Chanussot, and J. Mairal. A trainable spectral-spatial sparse coding model for hyperspectral image restoration. *Advances in Neural Information Processing Systems*, 34:5430–5442, 2021.

- [71] F. Xiong, J. Zhou, Q. Zhao, J. Lu, and Y. Qian. Mac-net: Model-aided nonlocal neural network for hyperspectral image denoising. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–14, 2022.
- [72] G. Fu, F. Xiong, J. Lu, J. Zhou, and Y. Qian. Nonlocal spatial-spectral neural network for hyperspectral image denoising. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–16, 2022.
- [73] L. Pang, W. Gu, and X. Cao. Trq3dnet: A 3d quasi-recurrent and transformer based network for hyperspectral image denoising. *Remote Sensing*, 14(18):4598, 2022.
- [74] M. Li, Y. Fu, and Y. Zhang. Spatial-spectral transformer for hyperspectral image denoising. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 1368–1376, 2023.
- [75] G. Fu, F. Xiong, J. Lu, J. Zhou, and Y. Qian. Hyperspectral image denoising via spatial-spectral recurrent transformer. *IEEE Transactions on Geoscience and Remote Sensing*, 2024.
- [76] G. Fu, F. Xiong, J. Lu, J. Zhou, and Y. Qian. Ssumamba: Spatial-spectral selective state space model for hyperspectral image denoising. *arXiv preprint arXiv:2405.01726*, 2024.
- [77] M. F. Baumgardner, L. L. Biehl, and D. A. Landgrebe. 220 band aviris hyperspectral image data set: June 12, 1992 indian pine test site 3, Sep 2015.
- [78] Wei He, Hongyan Zhang, Liangpei Zhang, and Huanfeng Shen. Total-variation-regularized low-rank matrix factorization for hyperspectral image restoration. *IEEE Transactions on Geoscience and Remote Sensing*, 54(1):178–188, 2016.
- [79] Yu-Bang Zheng, Ting-Zhu Huang, Xi-Le Zhao, Yong Chen, and Wei He. Double-factor-regularized low-rank tensor factorization for mixed noise removal in hyperspectral image. *IEEE Transactions on Geoscience and Remote Sensing*, 58(12):8450–8464, 2020.
- [80] S. Yun, D. Han, S.J. Oh, S. Chun, J. Choe, and Y. Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.
- [81] H. Zhang, M. Cisse, Y.N. Dauphin, and D. Lopez-Paz. mixup: Beyond empirical risk minimization, 2018.
- [82] Garnier, Simon, Ross, Noam, Rudis, Robert, Carmargo, Antônio Pedro, Sciaini, Marco, Scherer, and Cédric. *viridis(Lite) - Colorblind-Friendly Color Maps for R*, 2024. viridis package version 0.6.5.
- [83] Roberta H Yuhas, Alexander FH Goetz, and Joe W Boardman. Discrimination among semi-arid landscape endmembers using the spectral angle mapper (sam) algorithm. In *JPL, Summaries of the Third Annual JPL Airborne Geoscience Workshop. Volume 1: AVIRIS Workshop*, 1992.
- [84] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.