

# Effect of Hyperspectral Data Compression on Data Pre-processing: Analysis of Reconstruction Error Propagation

Jannick Kuester<sup>1,3</sup>, Wolfgang Gross<sup>1</sup>, Simon Schreiner<sup>1</sup>, Andreas Michel<sup>1</sup>, Jannik Sheikh<sup>1</sup>, Joshua Dare-Cullen<sup>2</sup>, Michael Heizmann<sup>3</sup>

<sup>1</sup> Fraunhofer IOSB, Image Analysis Group, Ettlingen, Germany

<sup>2</sup> Department of Mathematics and Statistics, University of Exeter, Exeter, United Kingdom

<sup>3</sup> Institute of Industrial Information Technology, Karlsruhe Institute of Technology, Karlsruhe, Germany

**Keywords:** Hyperspectral Data Compression, Lossy Data Compression, Spectral Compression, Reconstruction Error Propagation, Remote Sensing, Deep Learning, Autoencoder.

## Abstract

Hyperspectral imaging platforms such as UAVs and small satellites face strict constraints in data transmission and onboard storage. Lossy compression applied directly to raw sensor measurements offers substantial benefits for bandwidth efficiency. However, most existing studies evaluate compression only after radiometric and atmospheric correction, leaving the propagation of compression-induced errors through the complete pre-processing chain poorly understood. This study addresses this gap through a quantitative analysis of how reconstruction errors evolve from compressed raw data to georeferenced surface reflectance.

A representative set of state-of-the-art learning-based compression methods, including A1D-CAE, NLPCA, HyCoT, SSCNet and 3D-CAE, was evaluated. All models were trained on UAV-borne HySpex data and tested at a fixed compression rate of  $c_R = 4$ . This setting was chosen as a controlled and practically relevant operating point for a consistent comparison across all investigated methods. Original and reconstructed raw test data were processed with identical metadata through radiometric calibration, georeferencing and atmospheric correction. Reconstruction fidelity was assessed at four pre-processing stages using complementary spectral and spatial metrics.

The results show that spectral models retain high reconstruction accuracy throughout the workflow, with minimal error accumulation during pre-processing. Spatial and spatio-spectral architectures introduce spectral distortions that persist after atmospheric correction. These findings indicate that, for the investigated UAV-based HySpex dataset and the fixed compression setting of  $c_R = 4$ , spectral compression of raw hyperspectral data can preserve high-quality georeferenced reflectance products throughout the considered pre-processing workflow.

## 1. Introduction

The demand for hyperspectral data in remote sensing applications, including disaster management (Krekeler et al., 2023), land cover classification (Vali et al., 2020), and surveillance (Gross et al., 2022), has significantly increased. Hyperspectral sensors capture extensive spectral information across hundreds of bands, offering unique advantages for material discrimination and environmental monitoring. However, the resulting data volumes pose major challenges for transmission and storage, particularly on resource-limited platforms such as drones and satellites (Villafranca et al., 2012, Zandler et al., 2015, Melián et al., 2021). Lossy hyperspectral data compression enables compression rates  $\geq 4$ , thereby reducing storage requirements and supporting real-time transmission while striving to preserve reconstruction accuracy (Christophe, 2011).

Existing compression methods are typically evaluated on benchmark datasets (Baumgardner et al., 2015, Prasad et al., 2020, Fuchs and Demir, 2023) or satellite mission data such as PRISMA (Loizzo et al., 2018), EnMAP (Guanter et al., 2015), and DESIS (Müller et al., 2016). As these datasets are only provided after conversion to radiance or reflectance, the propagation of lossy compression errors through subsequent pre-processing steps remains largely unexplored. Compressing raw data would optimise bandwidth utilisation and reduce onboard storage demands, but requires that compression-induced errors remain stable throughout the processing chain.

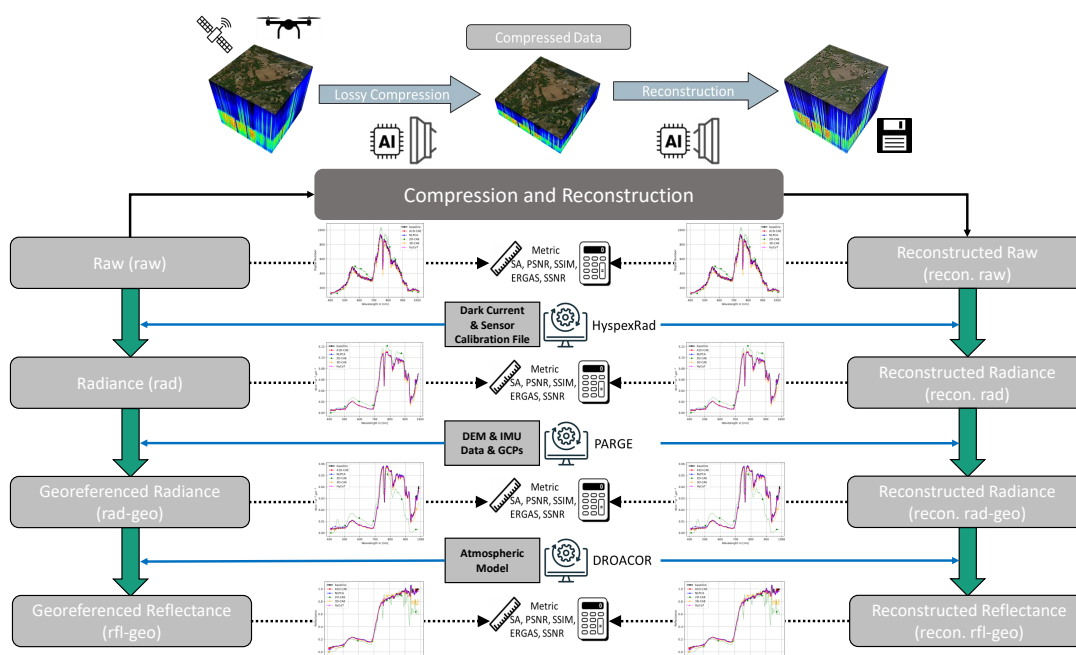
Unlike previous studies, this work systematically analyses how

reconstruction errors from lossy compression propagate from compressed raw data to radiance, georeferenced radiance, and atmospherically corrected reflectance. To test this, we evaluate a representative set of state-of-the-art machine learning models, namely A1D-CAE (Kuester et al., 2023), NLPCA (Licciardi and Chanussot, 2018), HyCoT (Fuchs et al., 2024), SSCNet (Grassa et al., 2022), and 3D-CAE (Chong et al., 2021). Each model is trained on raw UAV-borne hyperspectral data acquired during the HyperThun'22 campaign (Vögtli et al., 2023), and reconstruction quality is assessed at multiple stages of the pre-processing pipeline using an independent test set.

To the best of our knowledge, this is the first study to quantitatively evaluate the propagation of compression-induced reconstruction errors through the hyperspectral data pre-processing chain. We hypothesise that suitable lossy compression approaches do not significantly amplify reconstruction errors through successive processing steps within the investigated processing chain, thereby supporting their potential suitability for deployment on airborne and spaceborne platforms.

## 2. Hyperspectral Data Pre-processing

Pre-processing is a crucial step for transforming raw hyperspectral sensor data into georeferenced reflectance products that can be reliably used in scientific analysis and applications. Each step in the processing chain introduces transformations, corrections, and potential error sources that may interact with the reconstruction errors caused by lossy compression. Figure 1 provides an



**Figure 1.** Overview of the evaluated processing chain. Original and reconstructed hyperspectral raw data are propagated through identical pre-processing steps, including radiometric calibration, georeferencing, and atmospheric correction. Reconstruction fidelity is assessed at the raw, radiance, georeferenced radiance, and georeferenced reflectance levels.

overview of the evaluated processing chain, including the original and reconstructed data streams from raw digital numbers to georeferenced reflectance, as well as the auxiliary data required at each processing stage. In the following, the main pre-processing steps are described in more detail.

## 2.1 Radiometric Calibration

The first stage is the conversion of raw digital numbers into physically interpretable radiance values. This involves several sub-steps:

- **Dark current correction:** The intrinsic sensor noise due to thermal electrons is measured by recording dark frames with closed shutter or masked pixels. This signal is subtracted from the raw data to improve signal-to-noise ratio (SNR).
- **Radiometric calibration:** Sensor-specific calibration coefficients, usually provided by the manufacturer, are applied to convert digital numbers into at-sensor radiance values ( $W m^{-2} sr^{-1} nm^{-1}$ ). This includes band-dependent gain and offset parameters.
- **Sensor-related corrections:** Pixel-wise sensitivity variations, sensor non-linearities, and spectral smile (wavelength-dependent pixel shifts across the detector array) are corrected using calibration files obtained from laboratory characterisation.
- **Spectral calibration:** Wavelength alignment is performed to map detector channels to central wavelengths. This ensures spectral comparability between acquisitions and with reference libraries.
- **Data binning:** Many hyperspectral sensors apply binning in spectral or spatial domain to reduce noise and data size. While this improves SNR, it may lead to information loss and interacts with compression-induced errors.

Radiometric calibration is critical because it defines the physical basis of hyperspectral analysis. Errors at this stage propagate into all subsequent steps and may distort spectral indices or material classification.

## 2.2 Geometric Correction and Orthorectification

Hyperspectral imaging from airborne or UAV platforms is prone to geometric distortions due to platform motion, sensor perspective, and terrain relief. Geometric correction aims to remove these effects and assign each pixel a correct ground coordinate. The procedure typically involves:

- **Platform trajectory estimation:** Using GPS and Inertial Measurement Unit (IMU) data, the position and orientation of the sensor are reconstructed with high temporal resolution.
- **Projection onto a map grid:** Each recorded scan line is projected onto a standard cartographic coordinate system such as WGS84/UTM, which requires accurate sensor model parameters including focal length, field of view, and detector geometry.
- **Orthorectification with DEM:** A Digital Elevation Model (DEM) is used to correct parallax errors caused by topography. Pixels are shifted according to the terrain height to achieve true orthographic projection.
- **Resampling:** The irregularly distributed projected pixels are interpolated onto a regular grid. Nearest-neighbour resampling is often preferred to avoid mixing spectral information, although bilinear or cubic methods may yield smoother images at the cost of spectral fidelity.

Geometric correction ensures spatial comparability of hyperspectral data with other geospatial datasets. However, resampling may amplify reconstruction artefacts, e.g. if compression errors vary locally across scan lines.

### 2.3 Atmospheric Correction

The final step is the conversion of at-sensor radiance into surface reflectance. Since hyperspectral sensors record radiation after it has been modified by atmospheric absorption and scattering, correction of these effects is essential. Typical approaches include:

- **Radiative transfer modelling:** Physically based models, such as ATCOR (Richter and Schläpfer, 2002), simulate atmospheric interactions including gas absorption and aerosol scattering. Required inputs include aerosol optical thickness, water vapour content, ozone concentration, and solar geometry.
- **Topographic illumination correction:** DEMs are used to compensate for relief-induced illumination differences, which can strongly influence reflectance values in mountainous terrain, dense vegetation, or urban environments with significant height variation.
- **Empirical line method:** Reflectance calibration can be supported by *in-situ* reference targets of known reflectance that are spectrally characterised in the field (e.g., calibrated tarps), providing anchor points for model adjustment.
- **Normalisation procedures:** Additional corrections, such as BRDF normalisation, mitigate anisotropy caused by variations in viewing and illumination geometry.

Atmospheric correction is particularly sensitive to compression-induced spectral distortions, as subtle absorption features are directly linked to material properties. Even small deviations may bias parameter retrieval. Although atmospheric effects are generally less pronounced for UAV-based acquisitions due to shorter atmospheric path lengths, robust correction remains crucial, and becomes even more critical for airborne and satellite missions. Therefore, quantifying how compression-induced errors propagate through this step is a central aspect of this study.

## 3. Compression Methods

This section summarises the core characteristics of the compression methods introduced in Section 1. It provides a detailed description of the lossy hyperspectral compression methods evaluated in this work. The selection covers a broad spectrum of state-of-the-art (SOTA) learning-based approaches, ranging from shallow autoencoders and convolutional architectures to modern transformer-based designs. The models differ in whether they exploit spectral, spatial, or joint spatio-spectral redundancies, and were chosen to represent complementary paradigms. The study adopts a controlled comparison of representative learning-based paradigms that differ in the type of redundancy they exploit. Learning-based approaches are particularly relevant in current hyperspectral compression research because they enable data-driven modelling of complex and nonlinear spectral structures without relying on predefined transformations. This focus also allows us to analyse how architectural design affects the propagation of reconstruction errors through subsequent pre-processing steps. A classical signal-processing baseline was not included, which should be considered when interpreting the scope of the comparison. For comparability, all methods are evaluated at a fixed compression rate of  $c_R = 4$ .

### 3.1 A1D-CAE

The Adaptive 1D Convolutional Autoencoder (A1D-CAE) (Kuester et al., 2023) is a spectral compression model designed to achieve high reconstruction accuracy while maintaining transferability across sensors with different spectral resolutions. The model operates on individual spectral signatures  $\mathbf{x} \in \mathbb{R}^b$ , with  $b$  denoting the number of spectral bands. A key innovation is the use of *adaptive reflection padding*, which ensures divisibility of the spectrum by the desired compression rate  $c_R$ . This enables consistent encoding and decoding across sensors while preventing artificial discontinuities at the spectrum edges. The number of required padding bands is given by

$$\Delta_{pad} = \left\lceil \left\lfloor \frac{b}{c_R} \right\rfloor \cdot c_R \right\rceil - b. \quad (1)$$

The encoder consists of multiple convolutional modules for feature extraction, downsampling, and bottleneck compression. The number of downsampling blocks corresponds to  $\log_2(c_R)$ , ensuring consistency with the compression factor. The number of channels increases progressively up to 256, enabling the model to capture complex spectral dependencies, while kernel sizes decrease from 11 to 5 to adapt from long-range to local correlations. The decoder mirrors the encoder structure, employing transposed convolutions for upsampling, complemented by a cropping layer to remove padding and a final Hardtanh activation to scale the output to  $[0, 1]$ .

The compressed representation is extracted at the bottleneck, where spectral features are encoded into a compact latent space. The A1D-CAE demonstrates high spectral fidelity and sensor transferability, but does not exploit spatial redundancies.

### 3.2 Nonlinear PCA (NLPCA)

Nonlinear Principal Component Analysis (NLPCA) (Licciardi and Chanussot, 2018) is a shallow autoencoder architecture that applies three fully connected hidden layers with sigmoid activations. Its objective is to project the spectral signature into a compact nonlinear subspace, defined by the number of bottleneck neurons, and to reconstruct the original spectrum from this reduced representation.

The advantages of NLPCA are its conceptual simplicity, low computational cost, and ability to capture limited nonlinear dependencies in spectral data. However, its shallow structure restricts the capacity to represent complex spectral features, limiting reconstruction accuracy for high-variability datasets. In addition, NLPCA lacks transferability across sensors and ignores spatial redundancies, making it less effective for large-scale or multi-sensor applications.

### 3.3 SSCNet

The Spectral Signals Compressor Network (SSCNet) (Grassa et al., 2022) is a fully convolutional 2D autoencoder designed to jointly capture spectral and spatial features. The encoder comprises three convolutional layers with Parametric ReLU (PReLU) activations, each followed by  $2 \times 2$  max pooling for spatial downsampling. The compressed latent representation is obtained by flattening the feature maps. The decoder mirrors the encoder, employing transposed convolutions with PReLU activations for upsampling, and a final sigmoid activation for normalization to  $[0, 1]$ .

By incorporating 2D convolutions, SSCNet is able to exploit local spatial structures, which purely spectral methods neglect. The learnable PReLU further enhances nonlinear modeling. However, SSCNet does not explicitly model spectral dependencies and thus exhibits degraded spectral fidelity. Moreover, the architecture is sensor-specific and requires retraining for each spectral configuration.

### 3.4 3D-CAE

The 3D Convolutional Autoencoder (3D-CAE) (Chong et al., 2021) jointly compresses spectral and spatial information using 3D convolutions. The encoder comprises 3D convolutional layers with batch normalization and LeakyReLU activations, augmented by residual blocks to enhance feature extraction. Compression is applied to volumetric tiles (e.g.,  $32 \times 32 \times b$ ). The decoder reconstructs the data using 3D transposed convolutions and residual connections. Quantization is performed via stochastic rounding. Since the 3D-CAE training objective differs substantially from the single-metric losses used by the other models (see Table 1), its composite loss is detailed here for clarity:

$$\mathcal{L} = 1000 \times \text{SAM} + (45 - \text{PSNR}) + 1000 \times (1 - \text{SSIM}), \quad (2)$$

which jointly penalises spectral angle deviation, low peak signal-to-noise ratio, and structural dissimilarity.

The 3D-CAE preserves spectral–spatial consistency as encouraged by its composite training objective (see Equation 2) and was originally reported to achieve high reconstruction accuracy on benchmark datasets. However, its architecture is tied to a fixed band configuration, preventing application to other sensors without modification. Moreover, computational costs are higher compared to purely spectral models, limiting suitability for resource-constrained platforms.

### 3.5 HyCoT

HyCoT (Fuchs et al., 2024) employs transformer-based self-attention to capture long-range spectral relationships. The encoder is based on the SpectralFormer principle and groups spectral bands into tokens. A dedicated compression token collects the spectral information and then maps it to a low-dimensional latent space with a lightweight MLP. The decoder reconstructs the spectrum from this latent representation. The compression rate is adjusted by modifying the latent dimensionality, so the backbone remains unchanged as  $c_R$  varies, which stabilises computational cost.

HyCoT performs spectral compression only. Although training and inference operate on tiles for batching, the compression acts solely along the spectral dimension and does not exploit spatial redundancies. The model does not provide sensor generalization. Padding is used to meet divisibility constraints for token grouping or the selected compression rate, but it does not enable operation on data with a different number or placement of spectral bands. As a result, a new model must be trained for each sensor configuration.

### 3.6 Comparison of Methods

The investigated methods represent complementary paradigms of hyperspectral compression, covering spectral, spatial, and joint spatio-spectral modelling. Their diversity ensures that both local spectral correlations and broader spatial structures

are considered, providing a comprehensive basis for evaluating reconstruction accuracy and error propagation in subsequent experiments.

## 4. Experimental Setup

This section details the datasets, pre-processing workflow, data partitioning, training protocol, and evaluation metrics used to assess (i) the reconstruction accuracy at the raw level and (ii) the propagation of reconstruction errors through the complete pre-processing chain.

### 4.1 Datasets and Sampling Strategy

All data were acquired during the HyperThun'22 campaign near Thun, Switzerland, using a drone-mounted HySpex Mjolnir V-1240 sensor (Vögtli et al., 2023). The at-sensor raw line-scan data cover the 405 nm to 1015 nm range with  $b=488$  spectral bands and a native ground sampling distance of 21.6 mm.

A representative flight strip of size  $11,028 \times 1,392$  pixels yields about 30.7 million spectra. These data form the basis for training the learning-based compression models and are split into 80% training and 20% validation using a fixed random seed for reproducibility. Pixel-based models (A1D-CAE, NLPCA) draw individual spectra uniformly along scan lines, whereas tile-based models (HyCoT, SSCNet, 3D-CAE) extract non-overlapping  $32 \times 32$  tiles from the same strips (approximately 29,584 training tiles).

The test set corresponds to a second acquisition on the following day, recorded under different illumination and partly different scene content. This ensures independent evaluation and enables assessment of out-of-distribution behaviour.

### 4.2 Pre-processing Workflow

To quantify error propagation from raw to surface reflectance, both the original and reconstructed test data are processed with identical settings and metadata following the manufacturer's recommended workflow (see Figure 1) using proprietary tools:

1. **Radiometric calibration:** Raw digital numbers are converted to radiance using HySpexRad (HySpex, 2026). As part of this step, spectral binning is applied according to the manufacturer's specifications, reducing the number of bands from  $b=488$  to  $b=200$  in the VNIR range.
2. **Georeferencing/orthorectification:** Radiance is projected onto a standard cartographic coordinate system such as WGS84/UTM using PARGE (HySpex, 2026) with GPS/IMU data, a DEM, and ground control points. Zero-valued spectra introduced at boundaries during resampling are removed prior to analysis in the reflectance domain.
3. **Atmospheric correction:** Georeferenced radiance is converted to surface reflectance with DROACOR (Schläpfer et al., 2021).

Reflectance outputs are scaled to  $[0, 1]$  and stored as `float32`, with identical processing parameters applied to both streams.

### 4.3 Compression Protocol

After training, each model encodes the raw test dataset at a fixed compression rate of  $c_R=4$  and reconstructs it in the raw domain. The compression rate was fixed at  $c_R=4$  to define a controlled comparison point across all investigated models while still achieving substantial data reduction. More aggressive compression settings were not considered in this study, as the primary objective was to isolate the propagation of compression-induced errors through the subsequent pre-processing chain under one uniform operating condition. Both the baseline (original raw) and reconstructed raw test data are then propagated through the identical pre-processing chain (Section 4.2). This yields four paired levels for evaluation: Raw, Radiance, Georeferenced Radiance, and Georeferenced Reflectance, enabling step-wise attribution of error changes to specific pre-processing stages. These four evaluation levels correspond to the stages illustrated in Figure 1 as raw, rad, rad-geo, and rfl-geo for the original and reconstructed processing streams.

### 4.4 Training Configuration

The hyperparameters and training settings for all models are summarised in Table 1. Pixel-based models (A1D-CAE, NLPKA) operate on individual spectral signatures, whereas tile-based models (SSCNet, 3D-CAE, HyCoT) process  $32 \times 32$  pixel patches.

Training was performed on the Day 1 dataset as described in Section 4.1, using normalised inputs without data augmentation. All experiments were executed on a single node with two NVIDIA A100 GPUs (80 GB each) using multi-GPU training. Random seeds were fixed across libraries for reproducibility, and model weights as well as training logs were archived for later inspection.

For HyCoT, the original publication does not report a batch size or learning-rate scheduler, and training was carried out for 2000 epochs. This substantial budget increases the risk of overfitting and reduces experimental comparability. Due to the lack of detailed information and to harmonise the overall setup, reasonable default settings were adopted for batch size and scheduler, and the number of epochs was reduced to 150 across all models.

### 4.5 Evaluation Metrics

Reconstruction quality is evaluated between baseline and reconstructed data at each pre-processing level using complementary spectral and spatial metrics:

- **Spectral Angle (SA, °)** measures the angle between two spectra. Lower values indicate better agreement.
- **Signal-to-Spectral-Noise Ratio (SSNR, dB)** measures the ratio between signal power and spectral reconstruction error. Higher values indicate better reconstruction quality.

- **Structural Similarity (SSIM, [0, 1])** measures how well spatial structures are preserved. Higher values indicate better similarity.
- **Spectral Information Divergence (SID)** measures the difference between two spectral signatures. Lower values indicate better agreement.
- **ERGAS (Relative Global Dimensional Error)** measures the overall normalised reconstruction error. Lower values indicate better reconstruction quality.

Computing metrics across Raw, Radiance, Georeferenced Radiance, and Georeferenced Reflectance quantifies how radiometric calibration, resampling, and atmospheric corrections affect error propagation.

### 4.6 Implementation Notes and Quality Control

All model inferences for the test set are performed *before* any pre-processing is applied to ensure that differences arise solely from compression. Manufacturer-recommended spectral binning to  $b=200$  bands is applied during radiometric calibration and therefore impacts both the original and reconstructed streams identically. All outputs are validated for value range, band ordering, and geodetic metadata consistency prior to metric computation.

**Summary:** This setup provides (i) a like-for-like comparison at a fixed  $c_R=4$  with identical pre-processing, (ii) an out-of-distribution assessment via day-shifted acquisitions, and (iii) a step-wise analysis of how compression-induced errors interact with each pre-processing stage from raw to reflectance.

## 5. Results & Discussion

The quantitative evaluation in Table 2 reveals clear differences in the reconstruction capabilities of the investigated compression methods. Across all domains, the A1D-CAE achieves the most accurate results, with consistently low Spectral Angle (SA  $< 1.1^\circ$ ), low Spectral Information Divergence (SID  $< 0.0005$ ), and high SSNR ( $> 34$  dB) and SSIM ( $\approx 1.0$ ). These values indicate that the A1D-CAE achieves reconstructions with minimal spectral and structural deviations from the original hyperspectral data. Its strong performance is attributed to the convolutional design that is optimised for spectral correlations while remaining flexible across different sensor configurations due to adaptive padding. The method therefore provides a robust compromise between compression efficiency and spectral fidelity.

The transformer-based HyCoT and the shallow autoencoder NLPKA also demonstrate high reconstruction accuracy, achieving SA values between  $1.2^\circ$  and  $1.4^\circ$ , which according to (Licciardi and Chanussot, 2018) indicate reconstructions with low

**Table 1.** Overview of training hyperparameters for all investigated models. Compression rates are fixed to  $c_R = 4$ .

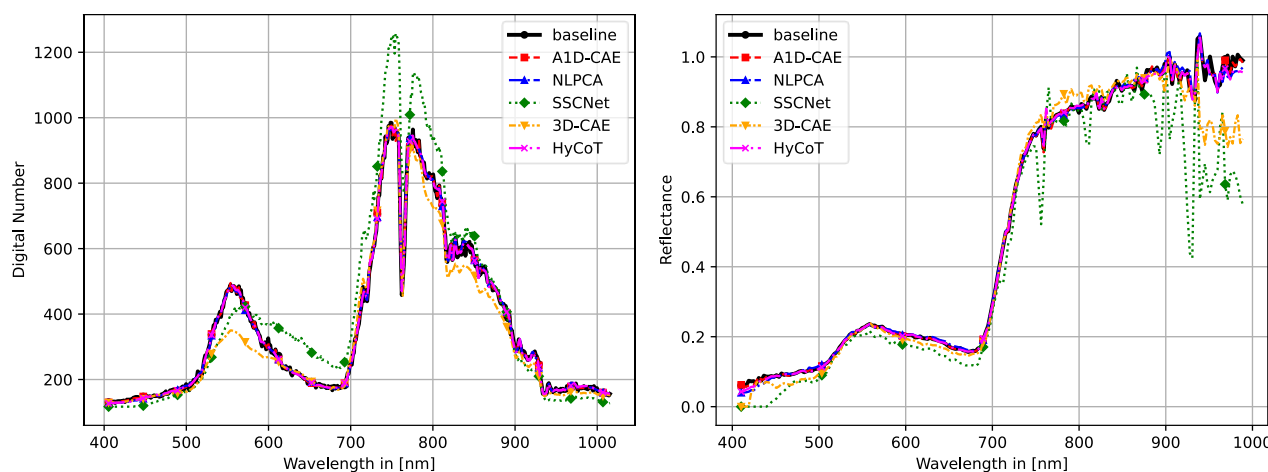
Model	Type	Optimizer	Learning rate	Scheduler	Batch size	Epochs	Loss
A1D-CAE	Pixel	Adam	$5 \times 10^{-4}$	ReduceLROnPlateau	1024	150	MSE
NLPKA	Pixel	Adam	$1 \times 10^{-4}$	—	128	150	MSE
3D-CAE	Tile	Adam	$1 \times 10^{-3}$	StepLR	16	150	Composite
SSCNet	Tile	Adam	$1 \times 10^{-3}$	MultiStepLR	16	150	BCE
HyCoT	Tile	Adam	$1 \times 10^{-4}$	CosineAnnealing	32	150	MSE

**Table 2.** Evaluation of the reconstruction accuracy for the raw, radiance, georeferenced radiance, and georeferenced reflectance data based on SA, SSNR, SSIM, SID, and ERGAS at  $c_R = 4$ . Arrows indicate whether smaller or larger values are favourable.

Model	Raw data					Radiance data				
	SA ↓	SSNR ↑	SSIM ↑	SID ↓	ERGAS ↓	SA ↓	SSNR ↑	SSIM ↑	SID ↓	ERGAS ↓
A1D-CAE	<b>1.04</b>	<b>34.77</b>	<b>0.997</b>	<b>0.0004</b>	<b>2.18</b>	<b>0.70</b>	<b>38.34</b>	<b>1.000</b>	<b>0.0004</b>	<b>1.66</b>
NLPCA	1.41	32.15	0.995	0.0007	3.04	1.16	34.19	0.996	0.0011	3.19
HyCoT	1.21	33.50	0.995	0.0005	2.55	0.91	36.37	0.997	0.0008	2.71
3D-CAE	3.00	24.17	0.995	0.0102	88.4	4.52	20.81	0.994	0.0377	35.8
SSCNet	7.43	16.97	0.823	0.0239	18.9	12.1	13.42	0.929	0.3628	29.2

Model	Radiance georeferenced data					Reflectance georeferenced data				
	SA ↓	SSNR ↑	SSIM ↑	SID ↓	ERGAS ↓	SA ↓	SSNR ↑	SSIM ↑	SID ↓	ERGAS ↓
A1D-CAE	<b>0.70</b>	<b>38.36</b>	<b>1.000</b>	<b>0.0004</b>	<b>1.69</b>	<b>0.90</b>	<b>36.29</b>	<b>0.994</b>	<b>0.0004</b>	<b>1.95</b>
NLPCA	1.16	34.20	0.996	0.0011	3.21	1.34	32.96	0.988	0.0010	3.05
HyCoT	0.91	36.38	0.997	0.0007	2.73	1.20	33.95	0.991	0.0008	2.61
3D-CAE	4.51	20.81	0.994	0.0373	35.6	5.70	26.02	0.918	0.4369	36.9
SSCNet	12.1	13.45	0.929	0.3580	29.1	15.7	11.50	0.641	0.0301	15.4



**Figure 2.** Left: Comparison of a reconstructed vegetation-spectrum of the raw data with the baseline for each tested method. Right: Comparison of a reconstructed vegetation-spectrum of the reflectance data with the baseline for each tested method.


spectral distortion. HyCoT benefits from its ability to capture long-range dependencies, whereas NLPCA remains competitive despite its simplicity, underlining that nonlinear dimensionality reduction can provide a compact yet informative spectral representation. Nevertheless, both models show slightly reduced robustness compared to A1D-CAE, particularly under varying illumination conditions.

In contrast, spatially oriented architectures such as SSCNet and 3D-CAE exhibit pronounced weaknesses. SSCNet, which relies exclusively on 2D convolutions, shows strong spectral distortions with SA values exceeding  $7^\circ$  and markedly reduced SSNR ( $< 17$  dB). These results confirm that neglecting spectral correlations leads to significant information loss. The 3D-CAE, while explicitly combining spectral and spatial context, still performs poorly in spectral metrics (SA  $3^\circ$ , SID  $> 0.01$ ). The limited robustness can be attributed to the architectural complexity and reliance on 3D convolutions, which may not capture fine-grained spectral correlations as effectively as spectral-focused models. Figure 2 and Figure 4 support these findings. SSCNet and 3D-CAE differ clearly from the baseline, while A1D-CAE, HyCoT,

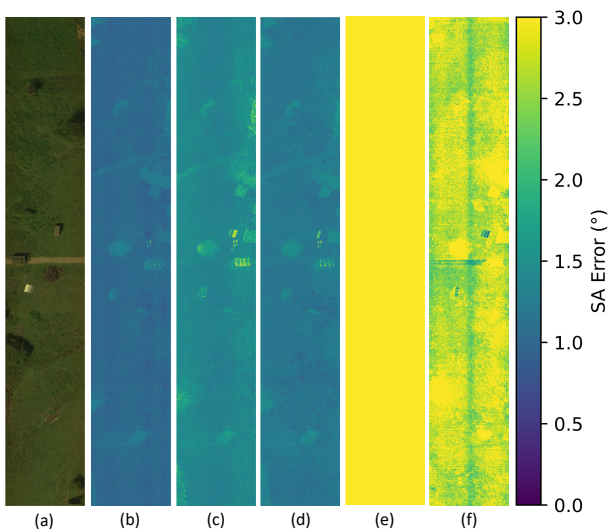
and NLPCA remain close to the reference spectra.

As illustrated in Figure 2, the spectral signatures undergo expected shape transformations between the raw digital numbers (left) and the surface reflectance domain (right) due to radiometric and atmospheric corrections. Importantly, these transformations affect both the baseline and the reconstructed data equally. Consequently, the relative ranking of the compression models remains consistent across domains. The spectral models A1D-CAE, HyCoT, and NLPCA remain close to the baseline. In contrast, SSCNet and 3D-CAE continue to show clear deviations. This supports our finding that pre-processing amplifies existing reconstruction errors rather than introducing new discrepancies between models.

The analysis of pre-processing stages provides further insight into error propagation. The conversion from raw data to radiance, which involves spectral binning by the manufacturer’s software, has a differential effect across methods. For spectral models (A1D-CAE, HyCoT, NLPCA), binning reduces noise and leads to slight improvements in SA and SSNR. In contrast, spatial models



**Figure 3.** RGB representation of the test dataset subset (a) and corresponding SA error maps of georeferenced reflectance data for AID-CAE (b), NLPCA (c), HyCoT (d), SSCNet (e), and 3D-CAE (f). All SA maps use the same colour scale, clipped at  $3^\circ$ , where lower values indicate higher spectral agreement with the baseline.



**Figure 4.** RGB representation of the dataset subset (a) and corresponding SA error maps of raw data for AID-CAE (b), NLPCA (c), HyCoT (d), SSCNet (e), and 3D-CAE (f). All SA maps use the same colour scale, clipped at  $3^\circ$ , where lower values indicate higher spectral agreement with the baseline.

(SSCNet, 3D-CAE) are adversely affected, as their already higher reconstruction errors distort spectral signatures in a way that is further amplified by the binning step. This highlights how compression-induced distortions can interact with sensor-specific pre-processing procedures.

Georeferencing introduces only marginal deviations, as expected, since resampling essentially duplicates or omits spectra without fundamentally altering spectral content. More pronounced differences occur in the reflectance domain, particularly in shaded regions (Figure 3), where lower signal intensity amplifies the relative effect of reconstruction errors. In Figure 3, panel (a) shows the RGB reference, while panels (b) to (f) display SA error maps in which dark blue and purple indicate low error and yellow indicates stronger deviations from the baseline. Here, SSCNet and 3D-CAE exhibit strong local distortions, while AID-CAE, HyCoT, and NLPCA remain stable. These observations highlight that high spectral fidelity is especially critical in low intensity

areas, where downstream tasks such as material classification are most sensitive to errors.

In this context, no explicit bad-band masking was applied, as the aim was to analyse error propagation through the complete manufacturer-recommended preprocessing workflow. For the investigated VNIR dataset, the representative spectra in Figure 2 do not indicate obvious absorption-dominated intervals or severe band failures. Additional application-specific band selection was not considered here.

In addition to illumination-related effects, rare spectral signatures such as man-made materials or unique object classes pose a particular challenge. These spectra typically occur only sparsely in the scene and are underrepresented during model training. As a result, compression errors tend to be more pronounced for such samples, as illustrated in Figure 2 and Figure 4. Since precisely those rare spectral signatures often carry the highest semantic value for downstream applications (e.g., material discrimination, anomaly or object detection), maintaining spectral fidelity in the distribution tails is critical for practical deployment of lossy compression methods.

Overall, the results confirm three central findings. First, spectral models outperform spatial and spatio-spectral models for hyperspectral compression when low spectral distortion is required. Second, for the investigated dataset and compression setting, error propagation through the pre-processing chain remains limited, suggesting that raw-domain lossy compression can be feasible without substantial degradation of reflectance products under these conditions. Third, interactions between compression and pre-processing (e.g., spectral binning) must be carefully considered, as they can either attenuate or amplify existing reconstruction errors depending on their magnitude and structure.

## 6. Conclusion

This study systematically evaluated state-of-the-art learning-based hyperspectral compression methods with a focus on reconstruction accuracy and error propagation across the complete pre-processing workflow, from raw sensor data to atmospherically corrected reflectance. Using UAV-based HySpex data, we

demonstrated that, for the investigated sensor, scene configuration, and fixed compression setting of  $c_R = 4$ , lossy compression can be applied at the raw stage while preserving high reconstruction fidelity in the final reflectance products.

Among the tested models, A1D-CAE achieved the highest spectral fidelity, consistently maintaining SA values below  $1.1^\circ$  and near-perfect SSIM. Transformer-based HyCoT and shallow NLPCA also performed well, remaining below the  $3^\circ$  low-distortion threshold (Licciardi and Chanussot, 2018). In contrast, spatially focused architectures such as SSCNet and 3D-CAE exhibited severe distortions that persisted throughout the pre-processing chain, rendering them less suitable for applications requiring high spectral accuracy.

Importantly, our results show that compression-induced errors do not substantially accumulate during pre-processing. Radiometric calibration with spectral binning attenuates noise for spectral models but amplifies distortions for spatial architectures, while georeferencing introduces only negligible variations. The largest deviations occur in the reflectance domain for weaker models. In addition, rare spectral signatures such as man-made materials remain particularly challenging, since they are sparsely represented in the data yet carry high semantic relevance for downstream applications (e.g., anomaly or object detection). These findings emphasise that compression architecture must be aligned with both sensor design and pre-processing specifics, and that reconstruction fidelity in the distribution tails is critical for practical deployment.

In practical terms, the results highlight that spectral compression methods such as A1D-CAE are well-suited for resource-constrained platforms like UAVs and satellites, enabling efficient transmission while maintaining high spectral fidelity in the evaluated processing chain.

Future work should extend this analysis to multi-sensor datasets, investigate higher compression ratios, and evaluate the impact on end-user tasks such as classification, unmixing, and biophysical parameter retrieval. In addition, future studies should include classical signal-processing baselines to place learning-based methods into a broader methodological context. Extending the analysis to SWIR data will be particularly relevant, as stronger atmospheric absorption features and low-quality bands may interact differently with compression-induced distortions. This will support a more comprehensive assessment of the operational viability of lossy hyperspectral compression for Earth observation missions.

## References

- Baumgardner, M., Biehl, L., Landgrebe, D., 2015. 220 band aviris hyperspectral image data set: June 12, 1992 indian pine test site 3. Purdue University Research Repository.
- Chong, Y., Chen, L., Pan, S., 2021. End-to-end joint spectral-spatial compression and reconstruction of hyperspectral images using a 3D convolutional autoencoder. *Journal of Electronic Imaging*, 30(4), 041403–041403.
- Christophe, E., 2011. Hyperspectral data compression tradeoff. *Optical Remote Sensing: Advances in Signal Processing and Exploitation Techniques*, Springer, 9–29.
- Fuchs, M. H. P., Demir, B., 2023. Hyspecnet-11k: a large-scale hyperspectral dataset for benchmarking learning-based hyperspectral image compression methods. *IGARSS 2023 - 2023 IEEE International Geoscience and Remote Sensing Symposium*, IEEE.
- Fuchs, M. H. P., Rasti, B., Demir, B., 2024. Hycot: A transformer-based autoencoder for hyperspectral image compression. *2024 14th Workshop on Hyperspectral Imaging and Signal Processing: Evolution in Remote Sensing (WHISPERS)*, IEEE, 1–5.
- Grassa, R. L., Re, C., Cremonese, G., Gallo, I., 2022. Hyperspectral Data Compression Using Fully Convolutional Autoencoder. *Remote Sensing*, 14(10), 2472.
- Gross, W., Queck, F., Schreiner, S., Vöggtli, M., Kuester, J., Mispelhorn, J., Kneubühler, M., Middelmann, W., 2022. A multi-temporal hyperspectral camouflage detection and transparency experiment. *Target and Background Signatures VIII SPIE*, 12270, 11–19.
- Guanter, L., Kaufmann, H., Segl, K., Foerster, S., Rogass, C., Chabrilat, S., Kuester, T., Hollstein, A., Rossner, G., Chlebek, C. et al., 2015. The EnMAP Spaceborne Imaging Spectroscopy Mission for Earth Observation. *Remote Sensing*, 7(7), 8830–8857.
- HySpex, 2026. Hyspex processing software. Accessed March 24, 2026.
- Krekeler, M. P. S., Burke, M., Allen, S., Sather, B., Chappell, C., McLeod, C. L., Loertscher, C., Loertscher, S., Dawson, C., Brum, J. et al., 2023. A novel hyperspectral remote sensing tool for detecting and analyzing human materials in the environment: a geoenvironmental approach to aid in emergency response. *Environmental Earth Sciences*, 82(4), 1–7.
- Kuester, J., Gross, W., Schreiner, S., Middelmann, W., Heizmann, M., 2023. Adaptive two-stage multisensor convolutional autoencoder model for lossy compression of hyperspectral data. *IEEE Transactions on Geoscience and Remote Sensing*, 61, 1–22.
- Licciardi, G., Chanussot, J., 2018. Spectral transformation based on nonlinear principal component analysis for dimensionality reduction of hyperspectral images. *European Journal of Remote Sensing*, 51(1), 375–390.
- Loizzo, R., Guarini, R., Longo, F., Scopa, T., Formaro, R., Facchinetti, C., Varacalli, G., 2018. Prisma: The italian hyperspectral mission. *IGARSS 2018 - 2018 IEEE International Geoscience and Remote Sensing Symposium*, 175–178.
- Melián, J. M., Jiménez, A., Díaz, M., Morales, A., Horstrand, P., Guerra, R., López, S., López, J. F., 2021. Real-Time Hyperspectral Data Transmission for UAV-Based Acquisition Platforms. *Remote Sensing*, 13(5), 850.
- Müller, R., Avbelj, J., Carmona, E., Gerasch, B., Graham, L., Günther, B., Heiden, U., Kerr, G., Knodt, U., Krutz, D. et al., 2016. The New Hyperspectral Sensor DESIS on the Multi-Payload Platform MUSES Installed on the ISS. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 41, 461–467.
- Prasad, S., Saux, B. L., Yokoya, N., Hansch, R., 2020. 2018 IEEE GRSS data fusion challenge – fusion of multispectral lidar and hyperspectral data. IEEE Dataport.
- Richter, R., Schläpfer, D., 2002. Geo-atmospheric processing of airborne imaging spectrometry data. Part 2: Atmospheric/topographic correction. *International Journal of Remote Sensing*, 23(13), 2631–2649.

Schläpfer, D., Richter, R., Popp, C., Nygren, P., 2021. DROACOR® REFLECTANCE RETRIEVAL FOR HYPER-SPECTRAL MINERAL EXPLORATION USING A GROUND-BASED ROTATING PLATFORM. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLIII-B3-2021, 209–214. <https://isprs-archives.copernicus.org/articles/XLIII-B3-2021/209/2021/>.

Vali, A., Comai, S., Matteucci, M., 2020. Deep learning for land use and land cover classification based on hyperspectral and multispectral earth observation data: A review. *Remote Sensing*, 12(15), 2495.

Villafranca, A. G., Corbera, J., Martín, F., Marchán, J. F., 2012. Limitations of Hyperspectral Earth Observation on Small Satellites. *Journal of Small Satellites*, 1(1), 19–29.

Vögtli, M., Sierro, L., Kneubühler, M., Schreiner, S., Gross, W., Queck, F., Kuester, J., Mispelhorn, J., Middelman, W., 2023. Hyperthun'22: A multi-sensor multi-temporal camouflage detection campaign. *IGARSS 2023-2023 IEEE International Geoscience and Remote Sensing Symposium*, 2153–2156.

Zandler, H., Brenning, A., Samimi, C., 2015. Potential of Space-Borne Hyperspectral Data for Biomass Quantification in an Arid Environment: Advantages and Limitations. *Remote Sensing*, 7(4), 4565–4580.