

Hyperspectral Image Compression Using Implicit Neural Representations

Shima Rezasoltani and Faisal Z. Qureshi

Faculty of Science

University of Ontario Institute of Technology

Oshawa, ON L1H 0C5 Canada

{shima.rezasoltani,faisal.qureshi}@ontariotechu.ca

Abstract—Hyperspectral images, which record the electromagnetic spectrum for a pixel in the image of a scene, often store hundreds of channels per pixel and contain an order of magnitude more information than a similarly-sized RGB color image. Consequently, concomitant with the decreasing cost of capturing these images, there is a need to develop efficient techniques for storing, transmitting, and analyzing hyperspectral images. This paper develops a method for hyperspectral image compression using implicit neural representations where a multi-layer perceptron network f_{Θ} with sinusoidal activation functions “learns” to map pixel locations to pixel intensities for a given hyperspectral image I . f_{Θ} thus acts as a compressed encoding of this image, and the original image is reconstructed by evaluating f_{Θ} at each pixel location. We have evaluated our method on four benchmarks—Indian Pines, Jasper Ridge, Pavia University, and Cuprite—and we show that the proposed method achieves better compression than JPEG, JPEG2000, and PCA-DCT at low bitrates.

Index Terms—hyperspectral image compression, implicit neural representations

I. INTRODUCTION

Unlike a grayscale image that records a single intensity value per pixel, hyperspectral images capture electromagnetic spectrum per pixel [31], [21]. Therefore, each pixel in a hyperspectral image contains tens or hundreds of values, representing recorded reflectance at various frequency bands. As a result, hyperspectral images offer greater possibilities for object detection, material identification, and scene analysis than those provided by a typical color RGB image. The costs associated with capturing high-resolution (both spatial and spectral) hyperspectral images continue to decrease, and it is no surprise that hyperspectral images have found widespread use in areas such as remote sensing, biotechnology, crop analysis, environmental monitoring, food production, medical diagnosis, pharmaceutical industry, mining, and oil & gas exploration, etc. [20], [24], [2], [19], [33], [9], [3], [30], [47], [41], [17], [25], [18], [12]. Hyperspectral images requires two orders of magnitude more space than what is needed to store a similarly sized color RGB image. Consequently, there is a need to develop efficient schemes for capturing, storing, transmitting, and analyzing hyperspectral images. This work studies the problem of hyperspectral image compression, with a view that it serves an important role in the storage and transmission of these images.

Recently, there has been a surge in interest in learning-based compression schemes. For example, autoencoders [1] and rate-distortion autoencoders [4], [5] have been used to learn compact representations of the input signals. Here network weights together with the signal signature—latent representation in the case of autoencoders—serve as the compressed representation of the input signal. Other concurrent works are exploring the use of Implicit Neural Representations (INRs) for signal compression. INRs are particularly well-suited to describe data that lives on an underlying grid, and as such, these offer a new paradigm for signal representation. In INRs, the goal is to learn a mapping between a location, say an (x, y) pixel location, and the signal value at that location, e.g., the pixel intensity $I[x, y]$. This mapping is subsequently used to recreate the original signal. It is as simple as evaluating the INR at various locations. In the case of INRs, network parameters serve as the learned representation of the input signal.

We investigate the use of INRs for hyperspectral image compression and show that it is possible to achieve high rates of compression while maintaining acceptable Peak Signal-to-Noise Ratio (PSNR) values. Figure 1 provides an overview of the proposed compression and decompression pipeline. We evaluate the proposed approach on four benchmarks (Figure 2)—(1) Indian Pines, (2) Jasper Ridge, (3) Pavia University, and (4) Cuprite—and show that our method achieves better PSNR values than those posted by three popular hyperspectral image compression schemes—(1) JPEG [23], [44], (2) JPEG2000 [14], and (3) PCA-DCT [38]—at comparable bits-per-pixel-per-band (bpppb) values. The results confirm that our method achieves better PSNRs at low compression rates than those obtained by other methods.

The rest of the paper is organized as follows. We discuss the related work in the next section. Section III describes the proposed method along with evaluation metrics. Datasets, experimental setup, and compression results are discussed in the following section. Section V concludes the paper with a summary and possible directions for future work.

II. RELATED WORK

Hyperspectral images exhibit both spatial and spectral redundancies that can be exploited to achieve compression. Lossless compression schemes—e.g., those that use quantization or rely upon entropy-coding—where it is possible to recover the

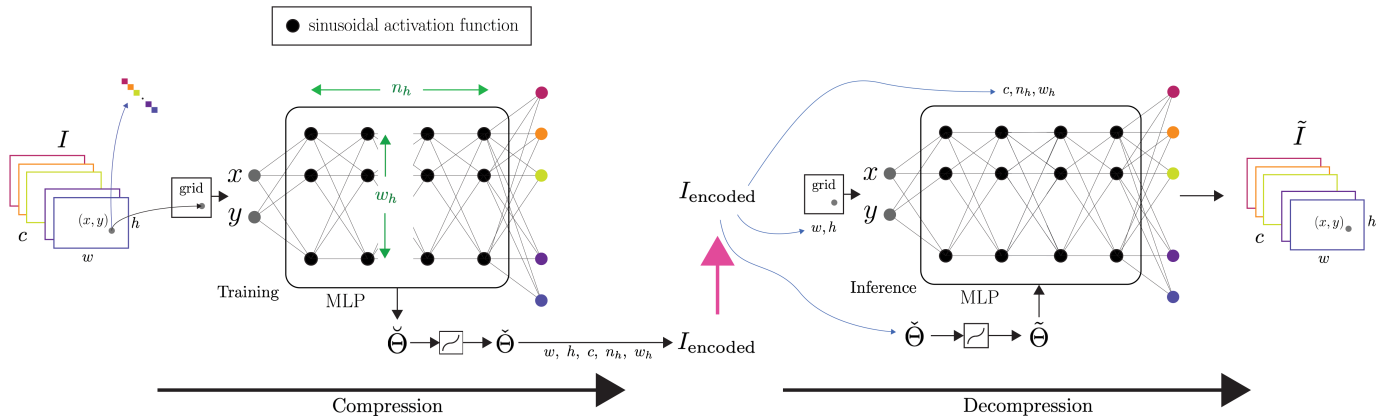


Fig. 1: **Compression and Decompression Pipeline.** (left) An MLP with a periodic activation function is *trained* to map pixel locations to the pixel’s spectral signature. (right) Once fitted, MLP is used to reconstruct the hyperspectral image by performing *inference* at various pixel locations.




Fig. 2: Datasets used in this study shown in pseudo-colors. (L2R) Indian Pines, Jasper Ridge, Pavia University, and Cuprite.

original signal precisely often do not yield large savings in terms of memory required to store or transmit a hyperspectral image [35], [40]. Lossy compression schemes, on the other hand, promise large savings while maintaining acceptable reconstruction quality. Inter-band compression techniques aim to eliminate spectral redundancy [10], while intra-band compression techniques aim to exploit spatial correlations. Intra-band compression techniques often follow the ideas developed for color image compression. [55] exploit the fact that groups of pixels that are around the same location in two adjacent bands are strongly correlated and propose schemes that perform both inter-band and intra-band compression. Principal Component Analysis (PCA) is a popular dimensionality reduction technique that has been used for hyperspectral image compression. PCA offers strong spectral decorrelation and it can be used to reduce the number of channels in a hyperspectral image. The remaining channels are subsequently compressed using Joint Picture Expert Group (JPEG) or JPEG2000 standard [7], [15], [52], [43].

Along similar lines, tensor decomposition methods have also been applied to the problem of hyperspectral image compression [54]. Tensor decomposition achieves dimensionality reduction while maintaining the spatial structure. Transform coding schemes that achieve image compression by reducing spatial correlation have also been used to compress hyperspectral data. Discrete Cosine Transform (DCT) has been used to perform intra-band compression; however, it ignores

inter-band (or spectral) redundancy. 3D-DCT that divides a 3-dimensional hyperspectral image into $8 \times 8 \times 8$ datacubes is proposed to achieve both inter-band and intra-band compression [45]. Similar to JPEG, which uses 8×8 blocks, 3D-DCT exhibits blocking effects in reconstructed hyperspectral images. The blocking effects can be removed to some extent by using wavelet transform instead [46], [22].

Video coding-based methods that treat each channel of a hyperspectral image as a frame in a video have also been used to perform hyperspectral image compression. These models rely upon inter-band spectral prediction to compress a hyperspectral image. This is similar to how inter-frame motion prediction is used for video encoding.

As mentioned earlier, recently, there has been a lot of interest in developing learning-based approaches for signal representation and compression. E.g., autoencoder-based techniques have been proposed to compress hyperspectral images [4], [5]. Hierarchical variational autoencoders have also been used for the purposes of hyperspectral image compression. Here the latent variables are discretized for entropy encoding purposes [6], [32], [37].

Work in the area of implicit neural representations has shown that it is possible to represent a signal by overfitting an appropriately designed neural network to it. Here the parameters (weights) of the neural network serve as the compact representation of the signal, and it is possible to reconstruct the original signal by sampling the neural network at various input locations [39], [42], [11], [49], [36], [48], [50]. [16] shows that implicit neural representations with periodic activation functions are able to represent signals, including images, with high-fidelity. This work serves as an inspiration for our work.

Similar to previous research on latent variable models [27], [28], [29], [34], numerous studies [8], [26], [53] make an effort to close the amortization gap [13] by combining the usage of amortized inference networks with iterative gradient-based optimization procedures. Using inference time per instance optimization, [53] also identifies and makes an effort to bridge




Fig. 3: **Model capacity.** PSNR vs. bpppb for (L2R) Indian Pines, Jasper Ridge, Pavia University, and Cuprite datasets. The trend of these plots confirms our intuition that PSNR values increase as bpppb numbers are increased. The plots are not monotonically non-decreasing. This has to do with the stochastic nature of MLP overfitting.

the discretization gap caused by quantizing the latent variables. The concept of per-instance model optimization is expanded upon in [51], which fine-tunes the decoder for each instance and transmits updates to the quantized decoder’s parameters together with the latent code to provide better rate-distortion performance.

III. IMAGE COMPRESSION USING INRS

Let us consider a w -by- h grayscale image. We can represent this image as a function

$$I_{\text{grayscale}} : U \mapsto [0, 1],$$

where

$$U = \{(1, 1), \dots, (w, 1), \dots, (1, h), \dots, (w, h)\}.$$

This notation captures the intuition that an image is a function over a 2d grid that defines the pixel locations. The intensity at each pixel is then $I_{\text{grayscale}}(x, y)$. It is straightforward to extend this notation to hyperspectral image I as follows

$$I = \begin{pmatrix} I_1 : U \mapsto [0, 1] \\ \vdots \\ I_c : U \mapsto [0, 1] \end{pmatrix}. \quad (1)$$

Here for the sake of simplicity, we assume a w -by- h hyperspectral image comprising c channels. Using this notation, we can find the spectral signature of the pixel at location $(x, y) \in U$ as follows:

$$\{I_1(x, y), \dots, I_c(x, y)\}.$$

Given this setup, it is possible to imagine a neural network that models the functions I_1, \dots, I_c . Specifically, work on implicit neural representations suggests using a multilayer perceptron (MLP) with periodic activation functions to represent functions of the form shown in Equation 1. Consider an MLP f_{Θ} with parameters Θ that maps locations in U to pixel spectral signatures:

$$f_{\Theta} : U \mapsto \{I_1, \dots, I_c\}.$$

Under this regime, training can be defined as

$$\check{\Theta} = \arg \min_{\Theta} \mathcal{L}(I, f_{\Theta}),$$

	w	h	c	n_h	w_h	q	Θ
#bits	16	16	16	8	8	1	$\text{bpp} \times n_{\Theta}$

TABLE I: Disk layout for I_{encoded} . Here q denotes if parameters Θ were quantized at compression time. bpp (or #bits-per-parameter) is either 32 or 16.

Where \mathcal{L} is a loss function that is differentiable and that captures the error between the original hyperspectral image and the decompressed hyperspectral image. We use the mean-squared error to compute this loss. Given $\check{\Theta}$, it is possible to reconstruct the original image I by sampling $f_{\check{\Theta}}$ at the relevant locations. Parameters $\check{\Theta}$, along with w, h, c , plus the structure of the MLP, represent an encoding I_{encoded} of the hyperspectral image I that was used to train the MLP f_{Θ} . It is expected that the memory required to store I_{encoded} is an order of magnitude less than the memory needed to store the hyperspectral image. The decompression process requires constructing the sampling locations U , setting up MLP $f_{\check{\Theta}}$ and initializing its weights to $\check{\Theta}$, and evaluating $f_{\check{\Theta}}$ at locations in U .

A. Metrics for Measuring Compression Quality

Similar to previous studies, we use Peak Signal-to-Noise Ratio (PSNR) to compare the proposed method with JPEG, JPEG2000, and PCA-DCT approaches for hyperspectral image compression. PSNR, measured in decibels, is a frequently used metric in image compression. It measures the difference in “quality” between the original image and its compressed version. Higher PSNR values suggest that the compressed image is more similar to the original image, i.e., the compressed image preserves more of the information present in the original image and that it has higher quality. In addition, we also compare the compressed image using Mean Squared Error (MSE), which computes the cumulative error between the original image and its compressed version. Lower values of MSE mean better reconstruction quality.

MSE is computed as follows

$$\text{MSE} = \sum_i \frac{|I[i] - \tilde{I}[i]|^2}{i}, \quad (2)$$




Fig. 4: **Architecture search.** Exploring MLP structure that achieves the best PSNR for different datasets (for a fixed bpppb budget). For our purposes, the MLP structure is defined by the number of hidden layers and the width of these layers. Together the number and width of the hidden layers define network capacity.

where \tilde{I} denotes the compressed image and i indices over the pixels. MSE is used to calculate PSNR

$$\text{PSNR} = 10 \log_{10} \left(\frac{R^2}{\text{MSE}} \right), \quad (3)$$

where R is the largest variation in the input image in the previous equation. For instance, R is 1 if the input image is of the double-precision floating-point data. R is 255, for instance, if the data is an 8-bit unsigned integer.

In addition, the number of bits-per-pixel-per-band (bpppb) captures the level of compression achieved by a model. Lower values of bpppb indicate higher compression rates. The bpppb value for an uncompressed hyperspectral image is either 8 or 32 bits, depending upon how the pixels are stored. It is common to store hyperspectral pixels value (for each channel) as a 32-bit floating point. The parameter bpppb is calculated as follows:

$$\text{bpppb} = \frac{\#\text{parameters} \times (\text{bits per parameter})}{(\text{pixels per band}) \times \#\text{bands}}. \quad (4)$$

Figure 3 plots PSNR vs. bpppb for the four datasets that we are using in this work. The plots confirm our intuition that higher bpppb leads to better compression quality as measured by PSNR values. For our method, bpppb calculations do not include the storage required to keep network structures.

B. Compression Pipeline

The proposed compression method consists of two steps. Step 1 performs an architecture search. The goal here is to find an MLP that achieves the highest reconstruction accuracy for a given bpppb budget. Architecture search is performed by overfitting multiple MLPs having different numbers of hidden layers and hidden layers' widths to the hyperspectral image. Architecture search, however, means longer compression times. Step 2 involves quantizing and storing the parameters of the overfitted MLP to disk. The caveat here is that this further reduces the quality of the reconstructed image.

1) *Overfitting a SIREN network:* The compression procedure comprises overfitting a SIREN network f_{Θ} to a hyperspectral image I [16]. The width w and height h of the hyperspectral image are used to set up an input location grid on $[-1, +1] \times [-1, +1]$, and the MLP is trained to reconstruct a pixel's spectral signature given its location. The parameters Θ of this overfitted MLP are quantized $\hat{\Theta}$. MLP structure that

contains the number of hidden layers n_h , widths of these layers w_h , and the width w , height h , and the number of channels c of the original hyperspectral image I along with $\hat{\Theta}$ serve as a compressed encoding I_{encoded} of the hyperspectral image I (Table I). Parameters $\hat{\Theta}$ are either stored as 32-bit floats or as 16-bit floats. Training and inference require 32-bit floats, and quantization/dequantization is performed to move between 32 and 16 bits representations. We have yet to try an 8-bit, fixed-point representation for parameters.

The overfitted MLP contains

$$(w_h \times 2) + (w_h \times w_h)^{(n_h-1)} + (c \times w_h) \quad (5)$$

parameters.

2) *Decompressing I_{encoded} :* The hyperspectral image is reconstructed from its compressed encoding I_{encoded} as follows: 1) use n_h , w_h , and c to reconstruct f_{Θ} , 2) dequantize $\hat{\Theta}$ to Θ and use it to initialize the parameters of f_{Θ} , 3) use the width w and height h to set up the input grid between $[-1, +1] \times [-1, +1]$, and 4) evaluate $f_{\tilde{\Theta}}$ at each location in the input grid to reconstruct the image \tilde{I} .

IV. EXPERIMENTS

We have used four datasets to evaluate our approach:

- **Indian Pines:** This is a $145 \times 145 \times 220$ hyperspectral image. It was collected using the AVIRIS sensor in 1992, and it spans a region over NW Indiana. The hyperspectral image contains a mix of farms and wooded areas. Additionally, the image contains low-density built-up regions, houses, a number of secondary roads, a rail line, and two dual-lane motorways.
- **Jasper Ridge:** This is a $100 \times 100 \times 224$ hyperspectral image. It was also captured using the AVIRIS sensor.
- **Pavia University:** It is a $610 \times 340 \times 103$ hyperspectral image. This dataset was captured by the ROSIS-03 aerial instrument that was flown by the German Aerospace Centre as part of the HySens project.
- **Cuprite:** The Cuprite dataset contains one $614 \times 512 \times 224$ hyperspectral image.

We have selected these datasets since others have used these previously to study hyperspectral image compression. These datasets have been collected using NASA's Airborne Visible/Infrared Imaging Spectrometer (AVIRIS) sensor.

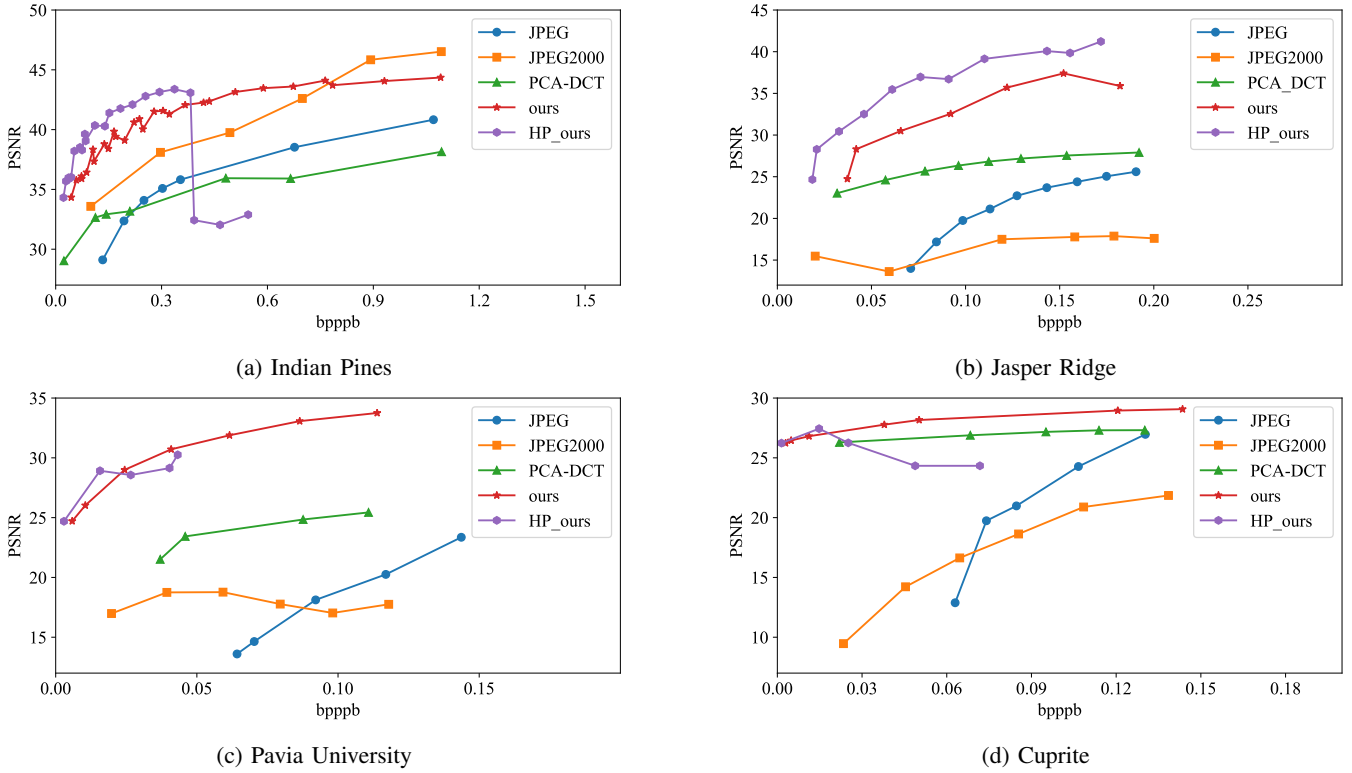


Fig. 5: **Compression results.** PSNR values achieved at various bpppb for our method, along with those obtained by JPEG, JPEG2000, and PCA-DCT schemes. Here “ours” refer to our method where parameters are stored at 32-bit precision, and “HP_ours” refer to results when parameters are stored at 16-bit precision.

The AVIRIS sensor gathers geometrically coherent spectro-radiometric data that can be used to characterize the Earth’s surface. The captured data has found use in areas ranging from oceanography, snow hydrology, geology and volcanology, and limnology to environmental studies, atmospheric and aerosol studies, agriculture, and land management and use.

A. Setup

We compare our work with three hyperspectral image compression methods: 1) JPEG [23], [44]; 2) PCA-DCT [38], and 3) JPEG2000 [14]. JPEG method for hyperspectral image compression uses JPEG standard to encode each channel (band) separately. JPEG2000, instead, uses the JPEG2000 standard for encoding the hyperspectral image. It, too, treats each channel separately. PCA-DCT uses PCA-based analysis to reduce the number of channels, followed by a DCT-based method for encoding these channels. PCA-DCT method posts low signal-to-noise ratios; however, this can be fixed somewhat by keeping more of the original channels. We have chosen these hyperspectral image compression techniques since they are widely used for reducing the size of hyperspectral data in hyperspectral analysis pipelines.

B. Architecture Search

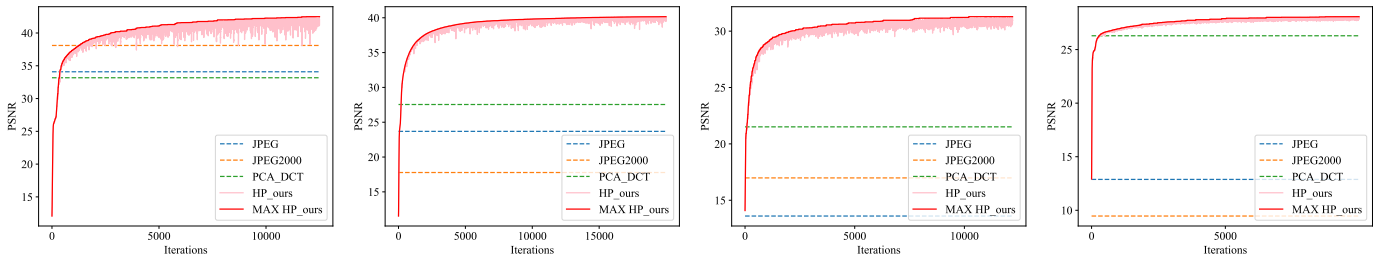
Given an image and our (MLP) parameter budget, which is measured in bits per pixel per band, or bpppb for short), the

first goal is to select the MLP structure—i.e., the number of hidden layers and their widths—that is able to represent this image with an acceptable PSNR value. Figure 4 shows PSNR values achieved for different architectures for the four datasets having a fixed bpppb budget. This suggests that network structure, in addition to network capacity, affects how well a network represents the hyperspectral image.

MLP structure is chosen via hyperparameter search, which involves training feasible designs containing the right number of hidden layers having the correct width on a given hyperspectral image. The result of this process is a single MLP that is able to reconstruct the hyperspectral image with the desired PSNR value. The parameters of the final MLP are then quantized to 16-bit precision, which leads to further savings in terms of the storage needed to represent the hyperspectral image. Our experiments suggest that reducing the MLP parameters from 32-bit to 16-bit precision did not increase distortion and that it had little effect on the signal-to-noise ratio.

C. Comparison with other methods

Figure 5 shows PSNR values at various compression rates for different methods. Specifically, we compare our approach, labeled as *ours* and *hp_ours*, with JPEG, JPEG2000, and PCA-DCT methods. Here, *ours* method stores MLP weights as 32-bit floating point values, whereas *hp_ours* stores MLP



(a) Model training on Indian Pines dataset at 0.2 bpppb. (b) Model training on Jasper Ridge dataset at 0.15 bpppb. (c) Model training on Pavia University dataset at 0.025 bpppb. (d) Model training on Cuprite dataset at 0.02 bpppb.

Fig. 6: Encoding procedure. Model training on (counter-clockwise from top-left) Indian Pines, Jasper Ridge, Pavia University, and Cuprite datasets. At around 2000 iteration mark, our method is already achieving better PSNR values than those for JPEG, JPEG2000, and PCA-DCT. Furthermore, the PSNR value for our methods continues to improve with more iterations (up to a point).

weights at half-precision as 16-bit floating point values that are constructed by quantizing the MLP weights. These plots illustrate that our methods achieve higher compression quality, i.e., better PSNR, for a given value of bpppb. This is especially true for lower bpppb values.

For the Indian Pines dataset, *ours* method achieves better PSNR up to around 0.7 bpppb, at which point JPEG2000 obtains better PSNR. What is curious is that the PSNR for *hp_ours* drops drastically at around 0.4 bppb. This merits further investigation. For Jasper Ridge, *hp_ours* performs better than *ours*. However, both *ours* and *hp_ours* achieve higher PSNR values than other methods. For Pavia University and Cuprite datasets, *our* method obtains better PSNR values than other methods.

We draw the following conclusions from these results: 1) the proposed method, both *ours* and *hp_ours*, perform high-quality compression at high compression rates; 2) it is beneficial to perform architecture search plus examine the effects of quantization at compression time since on some datasets *hp_ours* outperforms *ours*; and 3) the compression quality obtained by the proposed method compares favorably with the three commonly used compression methods for hyperspectral images.

D. Encoding Considerations

Our method belongs to the class of “slow-encoding-fast-decoding” compression methods. The method needs to train, actually *overfit*, multiple MLPs at encoding (compression) time. This is needed to find the MLP structure that best represents the hyperspectral image given a particular storage budget. Decoding, however, only requires evaluating this MLP at various pixel locations. Decoding is fast. It can be made even faster by exploiting the parallelism inherent to this procedure. The “slow-encoding-fast-decoding” nature of this method makes it particularly suitable for applications where the hyperspectral image is compressed once only, say at capture time.

We show an example of the overfitting procedure in Figure 6. These plots show the encoding procedure on the four

datasets: (1) Indian Pines at 0.2 bpppb; (2) Jasper Ridge at 0.15 bpppb; (3) Pavia University at 0.025 bpppb; and (4) Cuprite at 0.02 bpppb. JPEG, JPEG2000, and PCA-DCT methods do not require iterations. Consequently, their respective PSNR values are denoted with the horizontal dashed lines. The method proposed in this paper is iterative. Note that PSNR values for *HP_ours* continue to increase with the number of iterations (up to a point). Improvement in PSNR values saturates at around 10,000, 15,000, 10,000, and 5,000 iterations for Indian Pines, Jasper Ridge, Pavia University, and Cuprite datasets, respectively. This hints at the upper bound on encoding, or compression, time for our method. Note also that at around 2,000 iteration mark *HP_ours* method starts to obtain better PSNR values than the other three methods. As stated earlier, our method involves model fitting, which is inherently stochastic. Therefore, throughout the iterative process, we store the model parameters that obtained the highest value for PSNR thus far. In these plots, *MAX HP_ours* denote these PSNR scores. This guarantees that the model does not get worse over time.

E. Model Fitting

The number of inputs for all our models was 2, and the number of outputs was equal to the number of channels (or bands) of the hyperspectral image. The activation functions for hidden layers were sinusoidal. We initialized the MLP using the guidelines provided in [48]. Adam optimizer was used during training, and the learning rate was set to $2e - 4$. All experiments were conducted on an Intel i7 desktop with Nvidia RTX 2080 GPU.

F. Compression Results

Table II lists compression results obtained by *ours*, *HP_ours*, JPEG, JPEG2000, and PCA-DCT methods on the four datasets. The table also shows the size of the original, uncompressed hyperspectral images. For these results, we fix the bpppb for each method, and we measure the performance of each method using PSNR. Notice that the proposed method achieves higher PSNR values than those achieved by JPEG, JPEG2000, and PCA-DCT methods.




Fig. 7: **Reconstructed images** shown in pseudo-color. (L2R) Indian Pines, Jasper Ridge, Pavia University, and Cuprite. The image on the left in each pair is the original hyperspectral image, whereas the image on the right is the reconstructed, i.e., decompressed, hyperspectral image. The zoomed-in portions show that the structure is preserved in the reconstructed image. Images are shown in pseudo-color.

Method	Dataset	Size	PSNR	bpppb	n_h, w_h	Dataset	Size	PSNR	bpppb	n_h, w_h
-	Indian Pines	9251 KB	∞	16	-,-	Jasper Ridge	4800 KB	∞	16	-,-
JPEG		115.6 KB	34.085	0.2	-,-		30 KB	21.130	0.1	-,-
JPEG2000		115.6 KB	38.098	0.2	-,-		30 KB	17.494	0.1	-,-
PCA-DCT		115.6 KB	33.173	0.2	-,-		30 KB	26.821	0.1	-,-
ours		115.6 KB	40.61	0.2	15,40		30 KB	35.696	0.1	10,20
hp_ours		57.5 KB	40.35	0.1	15,40		15 KB	35.467	0.06	10,20
-	Pavia University	42724 KB	∞	16	-,-	Cuprite	140836 KB	∞	16	-,-
JPEG		267 KB	20.253	0.1	-,-		880.2 KB	24.274	0.1	-,-
JPEG2000		267 KB	17.752	0.1	-,-		880.2 KB	20.889	0.1	-,-
PCA-DCT		267 KB	25.436	0.1	-,-		880.2 KB	27.302	0.1	-,-
ours		267 KB	33.749	0.1	20,60		880.2 KB	28.954	0.1	25,100
hp_ours		133.5 KB	20.886	0.05	20,60		440.1 KB	24.334	0.06	25,100

TABLE II: Compression results

V. CONCLUSION

In this work, we employ implicit neural representations to compress hyperspectral images. Multi-layer perceptron neural networks with sinusoidal activation layers are overfitted to a hyperspectral image. The network is trained to map pixel locations to pixels' spectral signatures. The parameters of the network, along with its structure, represent a compressed encoding of the original hyperspectral image. We have tested our approach on four datasets, and the proposed method achieves better PSNR than those achieved by JPEG, JPEG2000, and PCA-DCT methods, especially at lower bitrates.

In the future, we plan to experiment with fixed-point representations to achieve smaller sizes for the compressed encodings. Additionally, we plan to explore techniques for lowering the compression times.

REFERENCES

- [1] Autoencoders, minimum description length, and helmholtz free energy. In *Advances in neural information processing systems*, pages 3–10, 1994.
- [2] Elhadi Adam, Onesimo Mutanga, and Denis Rugege. Multispectral and hyperspectral remote sensing for identification and mapping of wetland vegetation: a review. *Wetlands Ecology and Management*, 18(3):281–296, 2010.
- [3] Martin A Fromowitz, James B Callis, David M Heimbach, Larry A DeSoto, and Mary K Norton. Multispectral imaging of burn wounds. In *Medical Imaging II*, volume 914, pages 500–504. SPIE, 1988.
- [4] A. Alemi, B. Poole, I. Fischer, J. Dillon, R.A. Saurous, and K. Murphy. Fixing a broken elbow. In *Proc. 35th International Conference on Machine Learning*, volume 80, pages 159–168, 2018.
- [5] J. Ballé, V. Laparra, and E.P. Simoncelli. End-to-end optimized image compression. In *Proc. International Conference on Learning Representations (ICLR)*, June 2017.
- [6] Johannes Ballé, David Minnen, Saurabh Singh, Sung Jin Hwang, and Nick Johnston. Variational image compression with a scale hyperprior. *arXiv preprint arXiv:1802.01436*, 2018.
- [7] Ian Blanes and Joan Serra-Sagristà. Cost and scalability improvements to the karhunen–loève transform for remote-sensing image coding. *IEEE Transactions on Geoscience and Remote Sensing*, 48(7):2854–2863, 2010.
- [8] Joaquim Campos, Simon Meierhans, Abdelaziz Djelouah, and Christopher Schroers. Content adaptive optimization for neural image compression. *arXiv preprint arXiv:1906.01223*, 2019.
- [9] Oscar Carrasco, Richard B Gomez, Arun Chainani, and William E Roper. Hyperspectral imaging applied to medical diagnoses and food safety. In *Geo-Spatial and Temporal Image and Data Exploitation III*, volume 5097, pages 215–221. SPIE, 2003.
- [10] C. I. Chang. *Hyperspectral Data Processing: Algorithm Design and Analysis*. Wiley, 2013.
- [11] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5939–5948, 2019.
- [12] Roger N Clark and Gregg A Swayze. Mapping minerals, amorphous materials, environmental materials, vegetation, water, ice and snow, and other materials: the usgs tricorder algorithm. In *JPL, Summaries of the Fifth Annual JPL Airborne Earth Science Workshop. Volume 1: AVIRIS Workshop*, 1995.
- [13] Chris Cremer, Xuechen Li, and David Duvenaud. Inference suboptimality in variational autoencoders. In *International Conference on Machine Learning*, pages 1078–1086. PMLR, 2018.
- [14] Qian Du and James E Fowler. Hyperspectral image compression using jpeg2000 and principal component analysis. *IEEE Geoscience and Remote sensing letters*, 4(2):201–205, 2007.
- [15] Qian Du and James E Fowler. Hyperspectral image compression using jpeg2000 and principal component analysis. *IEEE Geoscience and Remote sensing letters*, 4(2):201–205, 2007.
- [16] Emilien Dupont, Adam Goliński, Milad Alizadeh, Yee Whye Teh, and Arnaud Doucet. Coin: Compression with implicit neural representations. *arXiv preprint arXiv:2103.03123*, 2021.
- [17] Gerda J Edelman, Edurne Gaston, Ton G Van Leeuwen, PJ Cullen, and Maurice CG Aalders. Hyperspectral imaging for non-contact analysis of forensic traces. *Forensic science international*, 223(1-3):28–39, 2012.
- [18] Yao-Ze Feng and Da-Wen Sun. Application of hyperspectral imaging

- in food safety inspection and control: a review. *Critical reviews in food science and nutrition*, 52(11):1039–1058, 2012.
- [19] Christian Fischer and Ioanna Kakoulli. Multispectral and hyperspectral imaging technologies in conservation: current research and potential applications. *Studies in Conservation*, 51(sup1):3–16, 2006.
- [20] Pedram Ghamisi, Naoto Yokoya, Jun Li, Wenzhi Liao, Sicong Liu, Javier Plaza, Behnood Rasti, and Antonio Plaza. Advances in hyperspectral image and signal processing: A comprehensive overview of the state of the art. *IEEE Geoscience and Remote Sensing Magazine*, 5(4):37–78, 2017.
- [21] Alexander FH Goetz, Gregg Vane, Jerry E Solomon, and Barrett N Rock. Imaging spectrometry for earth remote sensing. *science*, 228(4704):1147–1153, 1985.
- [22] Jorge González-Conejero, Joan Bartrina-Rapesta, and Joan Serra-Sagrasta. Jpeg2000 encoding of remote sensing multispectral images with no-data regions. *IEEE Geoscience and Remote Sensing Letters*, 7(2):251–255, 2009.
- [23] Walter F Good, Glenn S Maitz, and David Gur. Joint photographic experts group (jpeg) compatible data compression of mammograms. *Journal of Digital Imaging*, 7(3):123–132, 1994.
- [24] Megandhren Govender, Kershani Chetty, and Hartley Bulcock. A review of hyperspectral remote sensing and its application in vegetation and water resource studies. *Water Sa*, 33(2):145–151, 2007.
- [25] Aoife A Gowen, Colm P O’Donnell, Patrick J Cullen, Gérard Downey, and Jesus M Frias. Hyperspectral imaging—an emerging process analytical tool for food quality and safety control. *Trends in food science & technology*, 18(12):590–598, 2007.
- [26] Tiansheng Guo, Jing Wang, Ze Cui, Yihui Feng, Yuning Ge, and Bo Bai. Variable rate image compression with content adaptive optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 122–123, 2020.
- [27] Devon Hjelm, Russ R Salakhutdinov, Kyunghyun Cho, Nebojsa Jojic, Vince Calhoun, and Junyoung Chung. Iterative refinement of the approximate posterior for directed belief networks. *Advances in neural information processing systems*, 29, 2016.
- [28] Yoon Kim, Sam Wiseman, Andrew Miller, David Sontag, and Alexander Rush. Semi-amortized variational autoencoders. In *International Conference on Machine Learning*, pages 2678–2687. PMLR, 2018.
- [29] Rahul Krishnan, Dawen Liang, and Matthew Hoffman. On the challenges of learning with inference networks on sparse, high-dimensional data. In *International conference on artificial intelligence and statistics*, pages 143–151. PMLR, 2018.
- [30] Jaana Kuula, Ilkka Pölonen, Hannu-Heikki Puupponen, Tuomas Selander, Tapani Reinikainen, Tapani Kalenius, and Heikki Saari. Using vis/nir and ir spectral cameras for detecting and separating crime scene details. In *Sensors, and Command, Control, Communications, and Intelligence (C3I) Technologies for Homeland Security and Homeland Defense XI*, volume 8359, page 83590P. International Society for Optics and Photonics, 2012.
- [31] MCJJ L EH. Lightness and retinex theory. *J. Opt. Soc. Am.*, 61(1):1–11, 1971.
- [32] Jooyoung Lee, Seunghyun Cho, and Seung-Kwon Beack. Context-adaptive entropy model for end-to-end optimized image compression. *arXiv preprint arXiv:1809.10452*, 2018.
- [33] Haida Liang. Advances in multispectral and hyperspectral imaging for archaeology and art conservation. *Applied Physics A*, 106(2):309–323, 2012.
- [34] Joe Marino, Yisong Yue, and Stephan Mandt. Iterative amortized inference. In *International Conference on Machine Learning*, pages 3403–3412. PMLR, 2018.
- [35] J. Mielikainen and B. Huang. Lossless compression of hyperspectral images using clustered linear prediction with adaptive prediction length. *IEEE Geosci. Remote Sens. Lett.*, (9):1118–1121, 2012.
- [36] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *European conference on computer vision*, pages 405–421. Springer, 2020.
- [37] David Minnen, Johannes Ballé, and George D Toderici. Joint autoregressive and hierarchical priors for learned image compression. *Advances in neural information processing systems*, 31, 2018.
- [38] Yongjian Nian, Yu Liu, and Zhen Ye. Pairwise klt-based compression for multispectral images. *Sensing and Imaging*, 17(1):1–15, 2016.
- [39] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Occupancy flow: 4d reconstruction by learning particle dynamics. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5379–5389, 2019.
- [40] N. R. Mat Noor and T. Vladimirova. Investigation into lossless hyperspectral image compression for satellite remote sensing. *International Journal of Remote Sensing*, 34:5072–5104, 2013.
- [41] R Padoan, TA Steemers, M Klein, B Aalderink, and G De Bruin. Quantitative hyperspectral imaging of historical documents: technique and applications. *Art Proceedings*, pages 25–30, 2008.
- [42] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 165–174, 2019.
- [43] Barbara Penna, Tammam Tillo, Enrico Magli, and Gabriella Olmo. Transform coding techniques for lossy hyperspectral data compression. *IEEE Transactions on Geoscience and Remote Sensing*, 45(5):1408–1421, 2007.
- [44] Tong Qiao, Jinchang Ren, Meijun Sun, Jiangbin Zheng, and Stephen Marshall. Effective compression of hyperspectral imagery using an improved 3d dct approach for land-cover analysis in remote-sensing applications. *International Journal of Remote Sensing*, 35(20):7316–7337, 2014.
- [45] Tong Qiao, Jinchang Ren, Meijun Sun, Jiangbin Zheng, and Stephen Marshall. Effective compression of hyperspectral imagery using an improved 3d dct approach for land-cover analysis in remote-sensing applications. *International Journal of Remote Sensing*, 35(20):7316–7337, 2014.
- [46] Behnood Rasti, Johannes R Sveinsson, Magnus O Ulfarsson, and Jon Atli Benediktsson. Hyperspectral image denoising using 3d wavelets. In *2012 IEEE International Geoscience and Remote Sensing Symposium*, pages 1349–1352. IEEE, 2012.
- [47] Rebecca L Schuler, Paul E Kish, and Cara A Plese. Preliminary observations on the ability of hyperspectral imaging to provide detection and visualization of bloodstain patterns on black fabrics. *Journal of forensic sciences*, 57(6):1562–1569, 2012.
- [48] Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. *Advances in Neural Information Processing Systems*, 33:7462–7473, 2020.
- [49] Kenneth O Stanley. Compositional pattern producing networks: A novel abstraction of development. *Genetic programming and evolvable machines*, 8(2):131–162, 2007.
- [50] Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in Neural Information Processing Systems*, 33:7537–7547, 2020.
- [51] Ties van Rozendaal, Iris AM Huijben, and Taco S Cohen. Overfitting for fun and profit: Instance-adaptive data compression. *arXiv preprint arXiv:2101.08687*, 2021.
- [52] Lei Wang, Jiayi Wu, Licheng Jiao, and Guangming Shi. Lossy-to-lossless hyperspectral image compression based on multiplierless reversible integer tdl/klt. *IEEE Geoscience and remote sensing letters*, 6(3):587–591, 2009.
- [53] Yibo Yang, Robert Bamler, and Stephan Mandt. Improving inference for neural image compression. *Advances in Neural Information Processing Systems*, 33:573–584, 2020.
- [54] L. Zhang, L. Zhang, D. Tao, X. Huang, and B. Du. Compression of hyperspectral remote sensing images by tensor approach. *Neurocomputing*, (147):358–363, 2015.
- [55] Dongyu Zhao, Shiping Zhu, and Fengchao Wang. Lossy hyperspectral image compression based on intra-band prediction and inter-band fractal encoding. *Computers & Electrical Engineering*, 54:494–505, 2016.