

High-throughput onboard hyperspectral image compression with ground-based CNN reconstruction

Original

High-throughput onboard hyperspectral image compression with ground-based CNN reconstruction / Valsesia, D.; Magli, E.. - In: IEEE TRANSACTIONS ON GEOSCIENCE AND REMOTE SENSING. - ISSN 1558-0644. - 57:12(2019), pp. 9544-9553. [10.1109/TGRS.2019.2927434]

Availability:

This version is available at: 11583/2781092 since: 2020-01-17T09:29:05Z

Publisher:

Institute of Electrical and Electronics Engineers Inc.

Published

DOI:10.1109/TGRS.2019.2927434

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2019 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

High-throughput Onboard Hyperspectral Image Compression with Ground-based CNN Reconstruction

Diego Valsesia, *Member, IEEE*, Enrico Magli, *Fellow, IEEE*

Abstract—Compression of hyperspectral images onboard of spacecrafts is a tradeoff between the limited computational resources and the ever-growing spatial and spectral resolution of the optical instruments. As such, it requires low-complexity algorithms with good rate-distortion performance and high throughput. In recent years, the Consultative Committee for Space Data Systems (CCSDS) has focused on lossless and near-lossless compression approaches based on predictive coding, resulting in the recently published CCSDS 123.0-B-2 recommended standard. While the in-loop reconstruction of quantized prediction residuals provides excellent rate-distortion performance for the near-lossless operating mode, it significantly constrains the achievable throughput due to data dependencies. In this paper, we study the performance of a faster method based on prequantization of the image followed by a lossless predictive compressor. While this is well known to be suboptimal, one can exploit powerful signal models to reconstruct the image at the ground segment, recovering part of the suboptimality. In particular, we show that convolutional neural networks can be used for this task and that they can recover the whole SNR drop incurred at a bitrate of 2 bits per pixel.

Keywords—Hyperspectral image compression, convolutional neural networks

I. INTRODUCTION

Hyperspectral imaging from spaceborne spectrometers enables a wide range of applications, including material identification, terrain analysis and military surveillance. The ever-increasing spectral and spatial resolution of such instruments allows to create higher and higher quality products for the final user but it poses challenges in handling such wealth of data. In particular, onboard compression is critical to overcome the limited downlink bandwidth. This area of research poses specific challenges due to the strict complexity limitations on the payload hardware. Several solutions based on different techniques have been proposed, such as low-complexity spatial [1] and spectral transforms [2], distributed source coding [3], compressed sensing [4], [5], [6], and predictive coding [7], [8], [9], [10]. Predictive coding has emerged as one of the most popular solutions, as it enables low-complexity, high-throughput solutions, excellent rate-distortion performance and flexibility in the definition of image quality policies [11],

[12], [13], [14]. The CCSDS has been working on extending the CCSDS 123.0-B-1 recommendation [10] for predictive lossless compression, resulting in the recent publication of the 123.0-B-2 recommendation [15]. The new standard extends the previous one in the lossless mode, and includes lossy compression modes based on the introduction of a quantizer and a local decoder inside the prediction loop. It is well known [16] that an in-loop quantizer provides better rate-distortion performance than quantization followed by lossless predictive coding. However, one must consider that the need for a local decoder to reconstruct pixel values in the prediction neighborhood creates data dependencies which prevent parallelization and, consequently, high-throughput operations.

Meanwhile, recent years have seen the rise of neural networks as data-driven methods to solve problems previously tackled with hand-crafted models. In particular, imaging problems have been revolutionized by convolutional neural networks (CNNs). CNNs are able to capture very complex models about natural images because the convolution operation exploits powerful image priors such as shift invariance, and compositionality, where a complex global model is constructed from nonlinear hierarchies of local features. Ultimately, CNNs have proved to be able to achieve state-of-the-art performance on a wide variety of tasks including classification [17], segmentation [18], object detection [19] and regularization of inverse problems such as denoising [20] and superresolution [21], [22], [23].

In this paper, we propose to combine a low-complexity onboard compressor of hyperspectral images with a CNN-based reconstruction algorithm working at the ground segment. The main objective is to study its the rate-distortion performance with respect to the latest CCSDS standard. It is known that midpoint reconstruction from quantized data is not always optimal for image reconstruction, and e.g., using uniform-threshold quantization and a Laplacian assumption on the residuals is better. CNNs do not require an a priori model of the residuals, but are able to learn this model from training data. We show that the CNN learns to exploit the spatial and spectral correlation patterns of natural images to regularize the inverse reconstruction problem, and can be very effective at improving the quality of the image. Armed with such a powerful tool that runs at the ground segment where computational resources are abundant, one may wonder how much complexity is really needed onboard where resources are scarce. Preliminary FPGA implementations of the CCSDS 123.0-B-2 standard (using the Golomb entropy encoder) show

The authors are with Politecnico di Torino – Department of Electronics and Telecommunications, Italy. email: {name.surname}@polito.it. The research leading to this publication has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 776311.

that the lossless algorithm can achieve throughputs in excess of 100 Msamples/s [24], [25], while its lossy counterpart is limited to 20 Msamples/s [26] due to the aforementioned data dependencies. The new standard addresses this issue with a coding mode dedicated to high-throughput scenarios by removing some data dependencies, at a cost in terms of rate-distortion performance. In this paper, we propose to replace the lossy standard compressor with a different scheme based on prequantization of the raw pixels followed by the lossless CCSDS 123.0-B-2 encoder and a CNN reconstructor at the ground segment. The throughput of this compressor is essentially limited by the lossless predictor which is fast due to the lack of data dependencies. We show that the suboptimality due to moving the quantizer outside the prediction loop can be fully recovered by the CNN reconstruction and the same rate-distortion performance as lossy CCSDS 123.0-B-2 (without the CNN) is achieved, while potentially achieving the same throughput of the lossless version of the recommendation.

A preliminary version of this work appeared in [27]. With respect to the conference version, the method and its analysis are more thoroughly explained, we expand the treatment by also considering a relative error objective, present new experiments on a larger test set, and discuss transfer learning to different sensors. The paper is organized as follows. Sec. II provides some background material on the CCSDS 123.0-B-2 recommendation for lossy compression. Sec. III details the CNN used for image reconstruction. Sec. IV outlines the two approaches to onboard compression analyzed in the paper, i.e., lossy CCSDS 123.0-B-2 and prequantization followed by lossless CCSDS 123.0-B-2, for two quality objectives, namely bounded absolute or relative error. Sec. V discusses the experimental results. Finally, Sec. VI draws some conclusions.

II. BACKGROUND ON CCSDS 123.0-B-2

The CCSDS issued the Blue book for the 123.0-B-1 recommendation in May 2012 [10] and an Issue 2 in February 2019 [15]. The original recommendation focused on defining a method for lossless compression of hyperspectral images based on predictive coding. In particular, it is based on the fast lossless [28] predictor, which uses an adaptive filter to estimate a pixel value from information in a causal neighborhood. The prediction residual is then entropy coded by means of Golomb power of 2 (GPO2) codes [29]. This recommendation has been recently subject to a revision in order to extend it to lossy compression, resulting in the CCSDS 123.0-B-2 standard [15]. This extension is essentially based on the near-lossless coding principle, whereby a prediction residual, i.e., the difference between the predicted and the original pixel values, is quantized and locally decoded in order to update the weights of the prediction filter with the sign algorithm [30]. The extended recommendation also introduces a new prediction mode, namely *narrow local sums*, which essentially avoids using the pixel immediately on the left and in the same band of the pixel being coded. This mode is motivated by reasons of implementation efficiency: due to the local decoder in the prediction loop, the current pixel cannot be predicted unless every pixel in the causal neighborhood under consideration

has already been coded and decoded. The pixel on the left is especially important because it is coded immediately before the current one in the popular BSQ and BIL orderings and it is the main bottleneck in hardware implementations.

More in detail, the algorithm computes a local sum $\sigma_{x,y,z}$ which is defined as

$$\sigma_{x,y,z} = \begin{cases} s_{x-1,y,z}^R + s_{x-1,y-1,z}^R + s_{x,y-1,z}^R + s_{x+1,y-1,z}^R, & y > 0, 0 < x < N_x - 1 \\ 4s_{x-1,y,z}^R, & y = 0, x > 0 \\ 2(s_{x,y-1,z}^R + s_{x+1,y-1,z}^R), & y > 0, x = 0 \\ s_{x-1,y,z}^R + s_{x-1,y-1,z}^R + 2s_{x,y-1,z}^R, & y > 0, x = N_x - 1 \end{cases}$$

for the wide, neighbor-oriented mode and as

$$\sigma_{x,y,z} = \begin{cases} s_{x-1,y-1,z}^R + 2s_{x,y-1,z}^R + s_{x+1,y-1,z}^R, & y > 0, 0 < x < N_x - 1 \\ 4s_{x-1,y,z-1}^R, & y = 0, x > 0 \\ 2(s_{x,y-1,z}^R + s_{x+1,y-1,z}^R), & y > 0, x = 0 \\ 2(s_{x-1,y-1,z}^R + s_{x,y-1,z}^R), & y > 0, x = N_x - 1 \\ 4s_{mid}, & y = 0, x > 0, z = 0 \end{cases}$$

for the narrow, neighbor-oriented mode, being $s_{x,y,z}^R$ the reconstructed pixel at position (x, y, z) . Column-oriented modes also exist but will not be considered in this paper, as they are mostly intended for images with striping artifacts. The reduced prediction mode only uses the central local difference $d_{x,y,z} = 4s_{x,y,z}^R - \sigma_{x,y,z}$ while the full prediction mode also uses directional local differences $d_{x,y,z}^N, d_{x,y,z}^W, d_{x,y,z}^{NW}$ (we refer the reader to [15] for more details on the definitions). The predicted central difference $\hat{d}_{x,y,z}$ is obtained by multiplying the adaptive filter weights with the vector of differences, i.e.,

$$\hat{d}_{x,y,z} = \mathbf{W}_{x,y,z} \begin{bmatrix} d_{x,y,z}^N \\ d_{x,y,z}^W \\ d_{x,y,z}^{NW} \\ d_{x,y,z-1} \\ d_{x,y,z-2} \\ \vdots \\ d_{x,y,z-P} \end{bmatrix}$$

for full mode and

$$\hat{d}_{x,y,z} = \mathbf{W}_{x,y,z} \begin{bmatrix} d_{x,y,z-1} \\ d_{x,y,z-2} \\ \vdots \\ d_{x,y,z-P} \end{bmatrix}$$

for reduced mode. The predicted central difference is then transformed to obtain the predicted pixel value $\hat{s}_{x,y,z}$.

Finally, the recommendation also provides new tools such as sample representatives and new hybrid entropy coder able to reach rates lower than 1 bit per pixel (bpp), overcoming the limit of the original GPO2 encoder. Two main objectives can also be specified to drive the in-loop quantizer: bounded absolute error or bounded relative error.




Fig. 1: Reconstruction CNN. C: 2D convolution, R: leaky ReLU, IN: 2D instance normalization, CLIP: residual clipping. Input and output sizes are $N_l \times N_c \times 8$.

III. RECONSTRUCTION USING CONVOLUTIONAL NEURAL NETWORKS

This section presents the proposed approach to recover part of the image information lost during the lossy compression process. Any kind of lossy compression introduces artifacts which change the distribution of pixel values with respect to the one exhibited by natural uncompressed images. Recovering the original image from its distorted version is an ill-posed inverse problem, as there are infinitely many solutions. However, it is possible to compute a better estimate of the original image by properly modelling what constitutes a natural image.

Traditional techniques relied on hand-crafted image priors to model image data. For instance, a popular technique is total variation minimization, which amounts to requiring that the energy of the gradients in a natural image should be small. Image recovery from a compressed image I^Q is cast as the solution to the following minimization problem:

$$I^{DQ} = \arg \min_I \left[\|I - I^Q\|_2^2 + \lambda \sum_{x,y,z} (|I_{x+1,y,z} - I_{x,y,z}| + |I_{x,y+1,z} - I_{x,y,z}| + |I_{x,y,z+1} - I_{x,y,z}|) \right]. \quad (1)$$

Recently, convolutional neural networks (CNNs) have shown remarkable results in a variety of inverse problems, including denoising and superresolution. Their success lies in their ability to create more sophisticated models of complex image data as well as being able to handle perturbations with non-trivial statistics (e.g., non-Gaussian).

A. Proposed CNN

The proposed CNN reconstructs a better estimate of the original image from decoded hyperspectral images after lossy compression. Its training objective is to minimize the mean squared error (MSE) between the reconstructed image and the original. It is important to notice that the reconstruction depends on the specific algorithm used for compression and also the chosen quality level. This is similar to the denoising problem where several algorithms are based on knowing the noise variance [20], [31]. In our case, we train a CNN to invert a specific compression algorithm (e.g., near-lossless CCSDS 123.0-B-2) at a specific quality point which is known from the compression system design (e.g., a fixed quantizer step size for bounded absolute error near-lossless compression). We also argue that the trained model is optimal for new images acquired by the same sensor, as the network learns to exploit the peculiar spatial and spectral correlation patterns produced by that sensor. Nevertheless, the CNN has some generalization

capability to unseen sensors as some feature extraction steps are common for all sensors, thus only requiring fine-tuning with a smaller amount of data. Concerning the MSE training loss, some works have addressed image restoration using adversarial losses [32], [33], i.e., a game between two networks, one restoring the image, the other discriminating whether its input is an original or restored image. We will not consider this kind of loss because it tends to hallucinate image details which might be visually pleasing [34], but not really part of the original image and, in fact, such objective typically yields higher MSE values.

Fig. 1 shows an overview of the network. The input to the network is a slice of a hyperspectral image of size $N_l \times N_c \times 8$, where N_l and N_c are the number of lines and columns, respectively. While the spatial dimensions can be arbitrary, the number of bands is fixed to 8 in our proposed design. The main reason for this choice is the use of two-dimensional convolutional layers instead of three-dimensional ones. The first convolutional layer of the network has 64 filters of size $3 \times 3 \times 8$, thus merging the information from the 8 bands without sliding the kernel in the spectral dimension. A three-dimensional convolutional layer would have had a sliding kernel over all the three dimension and would have allowed an arbitrary number of spectral channels in the input. However, we found two main issues with this approach: *i)* the large size of hyperspectral images calls for careful memory usage and 3D convolutions require a very large amount of memory; *ii)* after reducing the use of memory to an acceptable value we found training to be highly unstable and providing results worse than those of the architecture with 2D convolutions. This is also an important design point in order to deal efficiently with images of large size. Notice that having a fixed number of input bands does not mean that only images with 8 bands can be processed. In fact, it is sufficient to slide a window over the spectral dimension of an image with more bands to process each slice and then merge the results. If partially overlapping slices are processed then the results are averaged by weighing each band by the number of times it has gone through the network.

The global input-output residual connection in the architecture means that the network learns to estimate the perturbation of the input image. This is an established solution in the literature on denoising [20], as it allows solving a simpler task by removing low-frequency content predicted by the input image. The inner layers of network show two main residual blocks composed of alternating convolutions, instance normalization layers and leaky ReLU nonlinearities [35]. The use of residual

blocks was introduced by the ResNet architecture [36] for image classification and has multiple benefits such as reducing the vanishing gradient problem thanks to one of the addends skipping several layers and improved learning capability due to the need to only learn the residual of an identity mapping instead of the full mapping. Instance normalization [37] normalizes activations to be approximately zero mean and unit standard deviation but, contrary to batch normalization [38], has different normalization factors for each image in the batch. Intuitively, this acts as a “contrast normalization” across the batch and helps dealing with perturbations that have more complex statistics than Gaussian noise, such as the case for reconstruction of compressed images.

Finally, the last layer allows to enforce consistent reconstruction, i.e., it ensures that the reconstructed pixel values fall in the same quantization bins as the original pixels by clipping the values of the correction estimated by the neural network. This is a design point that is specific to the reconstruction problem presented in this paper and also depends on the choice of the quantizer in the compression algorithm. In order to understand this, let us study a simple example. Suppose that the compression algorithm consists of simple uniform scalar quantization of the integer pixel values, i.e. $\mathbf{I}^Q = Q \lfloor \frac{\mathbf{I}}{Q} + \frac{1}{2} \rfloor$, with $Q = 2\Delta + 1$ for some integer Δ . Then, we know that the error is bounded as $|\mathbf{I}^Q - \mathbf{I}| \leq \Delta$. If we call \mathbf{E}^{CLIP} the correction term estimated by the network, then it must obey $|\mathbf{E}^{\text{CLIP}}| \leq \Delta$ since we know that the quantized pixel is never further than Δ from the original. Also notice that the bound on maximum error on the reconstructed image is, inevitably, twice the original bound.

We want to emphasize that proposing an entirely novel CNN architecture is outside the scope of this paper. Instead, we are interested in assessing how a baseline design inspired by recent results in the literature can already show that the proposed approach is competitive. Further optimization is certainly possible, e.g., by exploiting non-local features [39], [40]. However, this further strengthens the main point of this paper, which is about showing that coupling a simpler on-board compressor with a CNN at the ground segment allows higher throughput and has competitive rate-distortion performance with respect to the lossy CCSDS 123.0-B-2 standard.

IV. ONBOARD COMPRESSION APPROACHES

This section discusses two approaches to lossy onboard compression of hyperspectral images, namely the new CCSDS 123.0-B-2 recommendation and a simpler algorithm based on scalar quantization of the pixel values followed by a lossless predictive coding scheme, which we choose to be the lossless mode of CCSDS 123.0-B-2. We will refer to this method as “prequantization”. Fig. 2 visually depicts the two methods. We study the performance of both algorithms for two quality objectives: bounded absolute error and bounded relative error. We also study the performance impact of an on-ground reconstruction stage using the CNN presented in the previous section.

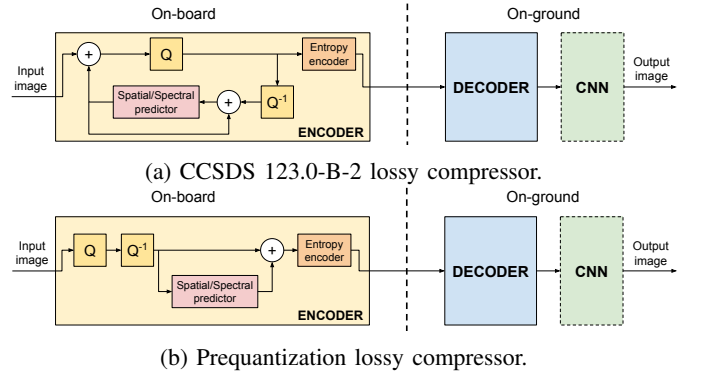


Fig. 2: Two predictive compression approaches. CCSDS 123.0-B-2 uses a quantizer inside the prediction loop. Prequantization quantizes raw pixel data and then applies a lossless predictor.

A. Complexity and data dependencies

The main reason to compare the two methods is to assess the most efficient way to employ the revised recommendation for lossy hyperspectral image compression. Scenarios requiring high-throughput implementations are particularly interesting, whereby the in-loop quantizer significantly limits the CCSDS algorithm. Recalling the notation of Sec. II, let us consider the wide, neighbor-oriented coding mode of lossy CCSDS 123.0-B-2 under band interleaved by line (BIL) coding order. The computation of the current local sum $\sigma_{x,y,z}$ requires knowing the value of $s_{x-1,y,z}^R$, i.e., the reconstructed pixel value on the left of the current pixel in the same band. In the BIL order, the $(x-1, y, z)$ pixel is coded immediately before the (x, y, z) pixel, which implies that all computations for $(x-1, y, z)$ must be terminated before starting coding (x, y, z) . This prevents building efficient parallel pipelines where the computation of the local sum can be started for several pixels ahead of the one being coded. The lossless version of CCSDS 123.0-B-2 does not suffer from such dependency as it only requires the original pixel values, not the reconstructed ones. In fact, space-grade FPGA implementations [24], [41] of the lossless algorithm achieved a throughput in excess of 100 Msamples/s while a comparable FPGA implementation of the lossy standard [26] was only limited to 20 Msamples/s due to this dependency issue.

The prequantization approach removes the quantizer from the prediction loop and therefore does not suffer from the same bottleneck. The prediction loop is lossless and can therefore achieve very high throughput, while the prequantization operation of the input data has negligible complexity compared to the predictor. Therefore, the prequantization method essentially shifts part of the complexity from the on-board encoder to the CNN needed after the decoder at the ground segment in order to recover the sub-optimal rate-distortion performance compared to the in-loop quantizer. The ground segment has fewer complexity issues and the main limitation is the memory usage of the GPU while reconstructing the image. This is limited by the design in Sec. III-A which uses 2D convolutions instead of more expensive 3D convolutions. The memory required by each 2D convolutional layer is $N_x N_y F$ floating point values




Fig. 3: Relative error quantizer for prequantization method. Dashed lines show the $\pm 10\%$ error bound.

instead of $N_x N_y N_z F$ floating point values required by 3D convolutions, being F the number of layer filters (64 in our design) and $N_x \times N_y$ the spatial dimensions of the input image.

B. Bounded absolute error

A guarantee bounding the absolute error is achieved by both the CCSDS and the prequantization methods by using a uniform scalar quantizer. In the former case, the quantizer operates on the prediction residuals, while in the latter case it is directly applied to the pixel values.

C. Bounded Relative error

A method to compress hyperspectral images using the CCSDS 123.0-B-2 standard with a target on relative error, rather than absolute error has been first proposed by Conoscenti et al. [14] and it is included in the revised recommendation. The main idea is to use an in-loop uniform scalar quantizer whose quantization step size changes at every pixel as it depends on the predicted pixel value to approximate the desired relative error. In particular, the following formula is used:

$$Q = 2 \lfloor R |\hat{s}_{x,y,z}| \rfloor + 1,$$

being R the target relative error and $\hat{s}_{x,y,z}$ the predicted pixel value. Notice that the predicted pixel value is used rather than the original pixel value in order to maintain causal decodability. This does not provide a hard bound on the relative error, but the use of a safety margin in the formula to compute the desired quantization step size showed good performance, with rare instances of error beyond the chosen limit.

It is obvious that the prequantization method can achieve a bounded relative error guarantee by designing a non-uniform scalar quantizer, where large pixel values are more coarsely quantized according to the desired relative error. Fig. 3 shows a sample design [42] of such quantizer, obtained by successive greedy extension of each quantization interval to match the relative error constraint.

V. EXPERIMENTS

This section presents an experimental assessment of the performance of the proposed CNN reconstruction when combined with the two compression approaches presented in Sec. IV. For both approaches we set the CCSDS predictor in its full prediction mode with wide neighbor-oriented local sums. Their rate-distortion performance is measured against a number of baseline methods. A first baseline is a transform-coding approach to onboard hyperspectral image compressor where the CCSDS 122 recommendation [1] for spatial compression using wavelets is combined with the Pairwise Orthogonal Transform (POT) to remove spectral correlation [2]. Another comparison is drawn with the CCSDS lossy compressor set in reduced prediction mode with narrow neighbor-oriented local sums. This is the recommended mode of the CCSDS standard to achieve high throughput at the expense of some compression performance.

A. CNN training and testing details

The CNN described in Sec. III-A is trained from scratch with patches from scenes acquired by the target sensor. The number of patches should be large enough to represent the variability in the acquired scenes. Patches, instead of full scenes, can be used since the CNN is learning the distortion introduced by the compression process, which is local in nature. Once trained, the CNN can be used to restore any new scene acquired by that sensor without further fine-tuning. In a real operating scenario, one may not have realistic training data to begin with, e.g., just after the launch of the satellite. This can be easily solved by downloading a few scenes with lossless compression as one of the first tasks after deployment, and train the neural network using those (their compressed versions at different quality points can be easily produced by running the compression algorithm directly on the ground).

In our experiments, the CNN has been trained using 70000 patches of size $32 \times 32 \times 8$ randomly extracted from AVIRIS images from the Cuprite, Jasper and Moffett scenes. Notice that these are older scenes and have some artifacts with respect to newer scenes, showing that the proposed CNN is also robust to perturbations and that the overall performance could be further improved with a higher quality training set. Nevertheless, we used them as they are well-known and readily available to create a training set with sufficiently varied scenes. Patches have been extracted from the decoded images. Concerning the experiments on bounded absolute error, the following quantization step sizes have been chosen: $Q \in \{3, 7, 11, 15, 21, 31, 41, 61, 101\}$ for both the CCSDS and prequantization compressors to let the networks operate at roughly the same quality point. On the other hand, the following maximum absolute relative errors (defined as $R = \max_{x,y,z} \frac{|I_{x,y,z}^Q - I_{x,y,z}|}{I_{x,y,z}}$), have been chosen for the experiments on bounded relative error: $R \in \{0.01, 0.001, 0.0075, 0.005, 0.0025, 0.0005\}$. An independent model has been trained for each value of Q and R and each compression method. The clipping layer in the CNN




Fig. 4: Rate-SNR performance of various compression methods with and without onground CNN. 123-NL: lossy CCSDS 123.0-B-2 (full, wide, neighbor-oriented mode); Q+123-LS: prequantization followed by lossless CCSDS 123.0-B-2 (full, wide, neighbor-oriented mode); 123-NL-RED-NARROW: lossy CCSDS 123.0-B-2 (reduced, narrow, neighbor-oriented mode); 122-POT: CCSDS 122 and POT; CNN: CNN reconstruction.

implements the following operation

$$E_{x,y,z}^{\text{CLIP}} = \begin{cases} -\Delta & \text{if } E_{x,y,z} \leq -\Delta \\ \Delta & \text{if } E_{x,y,z} \geq \Delta \\ E_{x,y,z} & \text{otherwise} \end{cases},$$



for the bounded absolute value experiments, and the following

$$E_{x,y,z}^{\text{CLIP}} = \begin{cases} -RI_{x,y,z}^Q & \text{if } E_{x,y,z} \leq -RI_{x,y,z}^Q \\ RI_{x,y,z}^Q & \text{if } E_{x,y,z} \geq RI_{x,y,z}^Q \\ E_{x,y,z} & \text{otherwise} \end{cases},$$

for the bounded relative error experiments. As a remark, one might wonder why using an additive residual also for the reconstruction problem with bounded relative error, instead of a multiplicative residual: we found that a multiplicative residual caused instability in the training process. We used the Adam optimization algorithm [43] with a learning rate equal to 10^{-8} for a total number of iterations corresponding to 1000 epochs. It was noticed that models for small values of Q and R especially benefited from the low learning rate. The convolutional layers have a fixed number of filters equal to 64. The CCSDS predictor has been set to use 3 prediction bands

TABLE I: SNR (dB) for test set

	123-NL	123-NL + CNN	Q + 123-LS	Q + 123-LS + CNN	123-NL-RED-NARROW	122-POT
1.5 bpp	52.23 ± 0.49	53.38 ± 0.50	49.58 ± 0.69	51.10 ± 0.73	50.85 ± 0.64	53.13 ± 0.54
2.0 bpp	57.60 ± 0.34	58.19 ± 0.36	57.15 ± 0.33	57.65 ± 0.36	56.36 ± 0.33	55.53 ± 0.31
3.0 bpp	64.88 ± 0.34	64.93 ± 0.33	64.80 ± 0.34	64.92 ± 0.35	63.81 ± 0.32	60.09 ± 0.28
4.0 bpp	71.57 ± 0.36	71.55 ± 0.36	71.53 ± 0.37	71.56 ± 0.36	70.49 ± 0.34	65.33 ± 0.36

Fig. 5: Error distribution for $sc0$ for $Q = 61$.Fig. 6: CNN reconstruction residual $\mathbf{I}^{DQ} - \mathbf{I}^Q$ for $Q = 31$. $sc0$ image, rows 150-300, all columns, band 47.

for both the lossy compressor and the lossless prediction after prequantization.

The testing dataset is strictly disjoint from the training data and it is composed of the $sc0$, $sc3$, $sc10$, $sc11$, $sc18$ scenes from the AVIRIS Yellowstone images. We remark that these images have not been used during the training phase. For testing purposes the input to the network is a slice of the image with 8 bands and full spatial resolution ($512 \times 680 \times 8$). All the possible slices of 8 bands out of the available 224 bands are fed to the network by moving the window selecting the bands by one band at a time and finally merging the resulting images with a weighted average of the overlapped parts. Reconstructing one full image of size $512 \times 680 \times 224$ takes 64 seconds on an Nvidia GTX 1080 Ti with a peak GPU memory utilization of 4096 MB. A C-language reference implementation of the CCSDS standard has

been used to generate compression results while the CNN has been implemented with the PyTorch library. Code and pretrained models are available online¹.

B. Bounded absolute error

The first experiment regards the rate-distortion performance of the two compressors and the relative gain provided by the CNN for the bounded absolute error scenario. Quality is measured by the SNR computed as

$$\text{SNR} = 10 \log_{10} \frac{\sum_{i=1}^{N_{\text{pixel}}} s_i^2}{\sum_{i=1}^{N_{\text{pixel}}} (s_i - s_i^R)^2}.$$

Other metrics such as the maximum spectral angle and the average spectral angle have been studied in the literature [44], but we omit them as they follow the same trends observed for SNR. Fig. 4 shows the rate-SNR curves for four test scenes. Table I reports the average SNR over the test set achieved by the various methods at four fixed rates (SNR values are linearly interpolated from the two closest available rate-distortion points). First, it can be noticed that the CNN provides more than 1 dB of improvement at 1.5 bpp, around 0.5 dB at 2.0 bpp and very small gains at high rates. Then, it is very interesting to notice that the sub-optimality of the prequantized method is quite limited and can be fully recovered by the CNN at all rates above or equal to 2.0 bpp. We also notice that the prequantized method is always better than lossy CCSDS 123.0-B-2 in reduced mode with narrow, neighbor-oriented local sums, which enables higher-throughput implementations, even without the help of the CNN.

Fig. 5 shows the distribution of the error between the original $sc0$ image, the compressed version and the reconstructed version using the CNN for $Q = 61$, for both compression techniques. It can be noticed that the CNN is able to reduce the average error amplitude, explaining the excess distribution around zero. We can also notice the longer tail of the error for the reconstructed image which is due to the ability to only guarantee twice the original bound after the reconstruction process, as explained in Sec. III-A. Fig. 6 visually shows the residual correction, i.e., $\mathbf{E}^{\text{CLIP}} = \mathbf{I}^{DQ} - \mathbf{I}^Q$, estimated by the network to restore the image. We can notice that the action of the CNN is particularly significant around edges.

Finally, we remark that we also tested total variation regularization as defined in Eq. (1) but the gain was limited to 0.1 dB at 1.5 bpp, 0.05 dB at 2 bpp and no gain was observed at higher rates, for both compression techniques. This confirms that CNNs are able to exploit much more complex models to regularize the reconstruction problem.

¹<https://github.com/diegovalsesia/hyperspectral-dequantization>




Fig. 7: Rate-MARE performance of various compression methods with and without onground CNN. 123-NL: lossy CCSDS 123.0-B-2 (full, wide, neighbor-oriented mode); Q+123-LS: prequantization followed by lossless CCSDS 123.0-B-2 (full, wide, neighbor-oriented mode); 123-NL-RED-NARROW: lossy CCSDS 123.0-B-2 (reduced, narrow, neighbor-oriented mode); CNN: CNN reconstruction.

C. Bounded relative error

In the experiments on bounded relative error we measure image quality in terms of mean absolute relative error (MARE) defined as:


$$\text{MARE} = \frac{1}{N_{\text{pixel}}} \sum_{i=1}^{N_{\text{pixel}}} \frac{|s_i - s_i^R|}{s_i}.$$

Fig. 7 shows the MARE as function of the rate for some test scenes. Table II also reports the achieved MARE for the different methods at fixed rate points. It can be noticed that CCSDS 123.0-B-2 in full, wide, neighbor-oriented mode followed by the CNN is confirmed as the best method.

However, the gain provided by the CNN is quite limited with respect to the absolute error case. This may be due to the more challenging error statistics, being dependent on the signal in a multiplicative way. The prequantization method followed by the CNN is competitive with the CCSDS 123.0-B-2 full, wide, neighbor-oriented baseline, and can outperform the fast CCSDS 123.0-B-2 reduced, narrow, neighbor-oriented method. Fig. 8 reports the relative error distribution with and without the CNN, again showing an excess around zero thanks to the CNN and a tail extending to twice the original maximum error target.

TABLE II: Percentage mean absolute relative error for test set

	123-NL	123-NL + CNN	Q + 123-LS	Q + 123-LS + CNN	123-NL-RED-NARROW
1.5 bpp	$(0.258 \pm 0.023)\%$	$(0.255 \pm 0.020)\%$	$(0.348 \pm 0.033)\%$	$(0.284 \pm 0.027)\%$	$(0.288 \pm 0.024)\%$
2.0 bpp	$(0.138 \pm 0.011)\%$	$(0.136 \pm 0.012)\%$	$(0.153 \pm 0.012)\%$	$(0.145 \pm 0.010)\%$	$(0.160 \pm 0.013)\%$
3.0 bpp	$(0.060 \pm 0.003)\%$	$(0.060 \pm 0.003)\%$	$(0.066 \pm 0.004)\%$	$(0.065 \pm 0.004)\%$	$(0.067 \pm 0.004)\%$
4.0 bpp	$(0.026 \pm 0.002)\%$	$(0.026 \pm 0.002)\%$	$(0.032 \pm 0.003)\%$	$(0.032 \pm 0.003)\%$	$(0.030 \pm 0.002)\%$

Fig. 8: Relative error distribution for $sc0$ for $R = 0.01$.

D. Transfer learning experiment

The optimal reconstruction results from the CNN can be obtained when the network is trained on images generated by the same sensor, so that the specific spatial and spectral correlation patterns or artifacts generated by that instrument can be exploited. However, the CNN works as a feature extractor and some of the features may generalize to different sensors. Table III reports the results obtained by using the same CNNs trained from the AVIRIS images on the *gran9* scene from the AIRS ultraspectral instrument, for the bounded absolute error mode. The size of this scene is equal to $135 \times 90 \times 1501$, thus having lower spatial resolution but higher spectral resolution with respect to the AVIRIS scenes. The results show that the CNNs perform well even if not trained specifically for the AIRS instrument.

VI. CONCLUSIONS

We proposed a method to compress hyperspectral images composed of an onboard predictive compressor and a ground-based CNN to reconstruct the decoded images and analyzed how it relates to the new CCSDS-123.0-B-2 recommendation. We showed that an onboard component based on prequantization followed by the lossless mode of CCSDS-123.0-B-2 can be significantly faster than the lossy mode of the standard and that, when coupled with the onground CNN, the same rate-distortion performance of the most efficient mode of lossy CCSDS-123.0-B-2 is achieved.

REFERENCES

- [1] Consultative Committee for Space Data Systems (CCSDS), "Image Data Compression," *Blue Book*, November 2005. [Online]. Available: <https://public.ccsds.org/Pubs/122x0b1c3s.pdf>
- [2] I. Blanes and J. Serra-Sagristà, "Pairwise orthogonal transform for spectral image coding," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 49, no. 3, pp. 961–972, 2011.
- [3] A. Abrardo, M. Barni, E. Magli, and F. Nencini, "Error-resilient and low-complexity onboard lossless compression of hyperspectral images by means of distributed source coding," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 48, no. 4, pp. 1892–1904, 2010.
- [4] D. Valsesia and P. T. Boufounos, "Universal encoding of multispectral images," in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing*, March 2016, pp. 4453–4457.
- [5] —, "Multispectral image compression using universal vector quantization," in *2016 IEEE Information Theory Workshop (ITW)*, Sep. 2016, pp. 151–155.
- [6] A. Barducci, D. Guzzi, C. Lastrì, V. Nardino, I. Pippi, and V. Raimondi, "Compressive sensing for hyperspectral Earth observation from space," *International Conference on Space Optics*, vol. 7, p. 10, 2014.
- [7] B. Aiazzi, P. Alba, L. Alparone, and S. Baronti, "Lossless compression of multi/hyper-spectral imagery based on a 3-d fuzzy prediction," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 37, no. 5, pp. 2287–2294, 1999.
- [8] E. Magli, G. Olmo, and E. Quacchio, "Optimized onboard lossless and near-lossless compression of hyperspectral data using CALIC," *IEEE Geoscience and Remote Sensing Letters*, vol. 1, no. 1, pp. 21–25, Jan 2004.
- [9] A. B. Kiely and M. A. Klimesh, "Exploiting calibration-induced artifacts in lossless compression of hyperspectral imagery," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 47, no. 8, pp. 2672–2678, 2009.
- [10] Consultative Committee for Space Data Systems (CCSDS), "Lossless Multispectral and Hyperspectral Image Compression," *Silver Book*, no. 1, May 2012. [Online]. Available: <https://public.ccsds.org/Pubs/123x0b1ec1s.pdf>
- [11] D. Valsesia and E. Magli, "A novel rate control algorithm for onboard predictive coding of multispectral and hyperspectral images," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 52, no. 10, pp. 6341–6355, Oct 2014.
- [12] —, "A hardware-friendly architecture for onboard rate-controlled predictive coding of hyperspectral and multispectral images," in *2014 IEEE International Conference on Image Processing*, Oct 2014, pp. 5142–5146.
- [13] —, "Fast and lightweight rate control for onboard predictive coding of hyperspectral images," *IEEE Geoscience and Remote Sensing Letters*, vol. 14, no. 3, pp. 394–398, March 2017.
- [14] M. Conoscenti, R. Coppola, and E. Magli, "Constant snr, rate control, and entropy coding for predictive lossy hyperspectral image compression," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 54, no. 12, pp. 7431–7441, Dec 2016.
- [15] Consultative Committee for Space Data Systems (CCSDS), "Low-Complexity Lossless and Near-Lossless Multispectral and Hyperspectral Image Compression," *Blue Book*, no. 1, February 2019. [Online]. Available: <https://public.ccsds.org/Pubs/123x0b2.pdf>
- [16] N. S. Jayant and P. Noll, "Digital coding of waveforms: principles and applications to speech and video," *Englewood Cliffs, NJ*, pp. 115–251, 1984.

TABLE III: Transfer learning on AIRS sensor

	Q	123-NL	123-NL + CNN	Q + 123-LS	Q + 123-LS + CNN
3	SNR (dB)	68.73	68.83	68.73	68.82
	Rate (bpp)		2.81		2.87
7	SNR (dB)	60.96	61.25	60.95	61.45
	Rate (bpp)		1.87		1.99
11	SNR (dB)	57.07	58.15	56.97	58.03
	Rate (bpp)		1.54		1.68
15	SNR (dB)	54.53	56.26	54.26	55.53
	Rate (bpp)		1.39		1.54
21	SNR (dB)	51.97	53.91	51.32	53.49
	Rate (bpp)		1.25		1.43
31	SNR (dB)	49.21	51.59	47.94	52.01
	Rate (bpp)		1.14		1.33
41	SNR (dB)	47.16	49.51	45.51	48.88
	Rate (bpp)		1.10		1.28
61	SNR (dB)	44.11	46.94	42.05	46.85
	Rate (bpp)		1.06		1.23
101	SNR (dB)	40.09	41.42	37.66	43.00
	Rate (bpp)		1.04		1.17

- [17] J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, June 2018, pp. 7132–7141.
- [18] P. Kaiser, J. D. Wegner, A. Lucchi, M. Jaggi, T. Hofmann, and K. Schindler, "Learning aerial image segmentation from online maps," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 55, no. 11, pp. 6054–6068, Nov 2017.
- [19] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016, pp. 779–788.
- [20] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang, "Beyond a Gaussian denoiser: residual learning of deep CNN for image denoising," *IEEE Transactions on Image Processing*, vol. 26, no. 7, pp. 3142–3155, 2017.
- [21] C. Dong, C. C. Loy, K. He, and X. Tang, "Image super-resolution using deep convolutional networks," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 38, no. 2, pp. 295–307, Feb 2016.
- [22] S. Lei, Z. Shi, and Z. Zou, "Super-resolution for remote sensing images via localglobal combined network," *IEEE Geoscience and Remote Sensing Letters*, vol. 14, no. 8, pp. 1243–1247, Aug 2017.
- [23] A. Bordone Molini, D. Valsesia, G. Fracastoro, and E. Magli, "Deep learning for super-resolution of unregistered multi-temporal satellite images," in *2019 10th Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing (WHISPERS)*, 2019.
- [24] L. Santos, L. Berrojo, J. Moreno, J. F. Lopez, and R. Sarmiento, "Multispectral and Hyperspectral Lossless Compressor for Space Applications (HyLoC): A Low-Complexity FPGA Implementation of the CCSDS 123 Standard," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 9, no. 2, pp. 757–770, Feb 2016.
- [25] University of Las Palmas de Gran Canaria, "SHyLoC IP Core," 2017. [Online]. Available: http://www.esa.int/Our_Activities/Space_Engineering_Technology/Microelectronics/SHyLoC_IP_Core
- [26] M. D. Nino, M. Romano, G. Capuano, and E. Magli, "Lossy multi/hyperspectral compression hw implementation at high data rate," in *Proceedings of International Astronautical Congress*, 2014.
- [27] D. Valsesia and E. Magli, "Image dequantization for hyperspectral lossy compression with convolutional neural networks," in *European Workshop on On-Board Data Processing (OBDP2019)*, 2019.
- [28] M. A. Klimesh, "Low-complexity lossless compression of hyperspectral imagery via adaptive filtering," 2005.
- [29] S. Golomb, "Run-length encodings," *IEEE Transactions on Information Theory*, vol. 12, no. 3, pp. 399–401, July 1966.
- [30] S. H. Cho and V. J. Mathews, "Tracking analysis of the sign algorithm in nonstationary environments," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 38, no. 12, pp. 2046–2057, 1990.
- [31] J. Lehtinen, J. Munkberg, J. Hasselgren, S. Laine, T. Karras, M. Aittala, and T. Aila, "Noise2Noise: learning image restoration without clean data," in *International Conference on Machine Learning (ICML)*, 2018.
- [32] N. Divakar and R. V. Babu, "Image denoising via CNNs: an adversarial approach," in *New Trends in Image Restoration and Enhancement, CVPR*, 2017.
- [33] J. Chen, J. Chen, H. Chao, and M. Yang, "Image blind denoising with generative adversarial network based noise modeling," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 3155–3164.
- [34] C. Ledig, L. Theis, F. Huszar, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, and W. Shi, "Photo-realistic single image super-resolution using a generative adversarial network," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017, pp. 105–114.
- [35] B. Xu, N. Wang, T. Chen, and M. Li, "Empirical evaluation of rectified activations in convolutional network," *arXiv preprint arXiv:1505.00853*, 2015.
- [36] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016, pp. 770–778.
- [37] D. Ulyanov, A. Vedaldi, and V. Lempitsky, "Instance normalization: The missing ingredient for fast stylization," *arXiv preprint arXiv:1607.08022*, 2016.
- [38] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *Proceedings of the 32Nd International Conference on Machine Learning - Volume 37*,

- ser. ICML'15. JMLR.org, 2015, pp. 448–456. [Online]. Available: <http://dl.acm.org/citation.cfm?id=3045118.3045167>
- [39] T. Plötz and S. Roth, "Neural nearest neighbors networks," in *Advances in Neural Information Processing Systems 31*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, Eds. Curran Associates, Inc., 2018, pp. 1087–1098. [Online]. Available: <http://papers.nips.cc/paper/7386-neural-nearest-neighbors-networks.pdf>
 - [40] D. Valsesia, G. Fracastoro, and E. Magli, "Image denoising with graph-convolutional neural networks," in *2019 26th IEEE International Conference on Image Processing (ICIP)*, 2019.
 - [41] J. Fjeldtvedt, M. Orlandic, and T. Arne Johansen, "An Efficient Real-Time FPGA Implementation of the CCSDS-123 Compression Standard for Hyperspectral Images," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. PP, pp. 1–12, 09 2018.
 - [42] A. Kiely, "Compression to Achieve a Relative Error Bound," *CCSDS MHDC WG meeting*, April 2017.
 - [43] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
 - [44] E. Christophe, D. Leger, and C. Mailhes, "Quality criteria benchmark for hyperspectral imagery," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 43, no. 9, pp. 2103–2114, Sep. 2005.



Diego Valsesia (S'13-M'17) received the Ph.D. degree in Electronic and Communication Engineering from the Politecnico di Torino, Turin, Italy, in 2016 and the M.Sc. degrees in Telecommunications Engineering from Politecnico di Torino and Electrical and Computer Engineering from the University of Illinois at Chicago, Chicago, IL, in 2012 and 2013 respectively. He is currently an Assistant Professor at the Department of Electronics and Telecommunications (DET), Politecnico di Torino. His main research interests include compression of remote

sensing images, compressed sensing, and deep learning.



Enrico Magli (S'97-M'01-SM'07-F'17) received the M.Sc. and Ph.D. degrees from the Politecnico di Torino, Torino, Italy, in 1997 and 2001, respectively. He is currently a Full Professor with Politecnico di Torino, Torino, Italy. His research interests include compressive sensing, image and video coding, and vision. He is an Associate Editor of the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY and the EURASIP Journal on Image and Video Processing, and a former Associate Editor of the IEEE TRANSACTIONS

ON MULTIMEDIA. He is a Fellow of the IEEE, and has been an IEEE Distinguished Lecturer from 2015 to 2016. He was the recipient of the IEEE Geoscience and Remote Sensing Society 2011 Transactions Prize Paper Award, the IEEE ICIP 2015 Best Student Paper Award (as senior author), and the 2010 and 2014 Best Associate Editor Award of the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY.