

IMAGE DEQUANTIZATION FOR HYPERSPECTRAL LOSSY COMPRESSION WITH CONVOLUTIONAL NEURAL NETWORKS

Diego Valsesia, Enrico Magli

Department of Electronics and Telecommunications – Politecnico di Torino, Italy

ABSTRACT

Significant work has been devoted to methods based on predictive coding for onboard compression of hyperspectral images. This is supported by the new CCSDS 123.0-B-2 recommendation for lossless and near-lossless compression. While lossless compression can achieve high throughput, it can only achieve limited compression ratios. The introduction of a quantizer and local decoder in the prediction loop allows to implement lossy compression with good rate-performance. However, the need to have a locally decoded version of a causal neighborhood of the current pixel under coding is a significant limiting factor in the throughput such encoder can achieve. In this work, we study the rate-distortion performance of a significantly simpler and faster onboard compressor based on prequantizing the pixels of the hyperspectral image and applying a lossless compressor (such as the lossless CCSDS 123.0-B-2) to the quantized pixels. While this is suboptimal in terms of rate-distortion performance compared to having an in-loop quantizer, we compensate the lower quality with an on-ground post-processor based on modeling the distortion residual with a convolutional neural network. The task of the neural network is to learn the statistics of the quantization error and apply a dequantization model to restore the image.

1. INTRODUCTION

Hyperspectral imaging from spaceborne instruments 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 peculiar 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], and predictive coding [6, 7, 8]. Predictive coding 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 [9, 10, 11, 12]. The CCSDS has been working on extending the CCSDS 123 recommendation for predictive lossless compression to include lossy compression modes based on the introduction of a quantizer and a local decoder inside the prediction loop, resulting in the new CCSDS 123.0-B-2 standard [13]. It is well-known [14] 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 increases the complexity of each prediction iteration and may hinder 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. CNNs proved to be able to achieve state-of-the-art performance on a wide variety of tasks including classification [15], segmentation [16], object detection [17] and inverse problems such as denoising [18] and superresolution [19, 20].

In this work, we study the rate-distortion performance of an onboard compressor based on prequantizing the pixels of the hyperspectral image and applying a lossless compressor (such as the lossless CCSDS 123.0-B-2) to the quantized pixels. While this is suboptimal in terms of rate-distortion performance compared to having an in-loop quantizer, we compensate the lower quality with an on-ground post-processor based on a convolutional neural network whose goal is to learn the statistics of the compression distortion. This solution can achieve a much higher throughput than the one with an in-loop quantizer, so it is worth studying whether it can be competitive in terms of rate-distortion performance against the revised CCSDS 123.0-B-2 recommendation.

2. BACKGROUND

Any kind of lossy compression introduces artifacts on the images, changing 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-

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.

posed inverse problem, as there are infinitely many solutions. However, a better estimate of the original can be computed by proper modelling of 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 saying that the energy of the gradients in a natural image should be small, and casts image recovery from a compressed image as the solution to an optimization problem.

Recently, convolutional neural networks (CNNs) have shown remarkable results in a variety of inverse problems, including denoising and superresolution. Their success lies in their capacity 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 noise or superresolution residuals). In particular, residual networks (see ResNet [21]) are an established solution in the literature on denoising [18], as they allow to solve a simpler task by using the input as a prediction of lower-frequency content and learning to compensate only the residual. This is used in the CNN architecture as a global input-output residual connection or across a block made of a few (typically 3-4) hidden layers.

3. PROPOSED METHOD

The proposed dequantization CNN restores compressed hyperspectral images to recover an estimate of the original. Its training objective is to minimize the mean squared error (MSE) between the restored image and the original. It is important to notice that the restoration 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 [18, 22]. In our case, we train a CNN to restore images produced a specific compression algorithm (e.g., near-lossless CCSDS 123) at a quality point fixed by the compression system design (e.g., a fixed quantizer step size for near-lossless compression).

Fig. 1 shows an overview of the network. The input to the network is a slice of a hyperspectral image of size $N_r \times N_c \times 8$. While the spatial dimensions can be arbitrary, the number of bands is fixed to 8 in the design we proposed. 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. 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 process a slice of an image with more bands at a time 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. Convolutional layers are interleaved with instance normalization layers and leaky ReLU nonlinearities [23]. Finally, the last layer allows to enforce some prior knowledge about the image

by clipping the values of the correction estimated by the neural network. 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 $|\mathbf{I}^Q - \mathbf{I}| \preceq \Delta$. If we call \mathbf{E}^{CLIP} the correction term estimated by the network, then it must obey $|\mathbf{E}^{\text{CLIP}}| \preceq \Delta$ since we know that the quantized pixel is never further than Δ from the original.

We consider two approaches to lossy onboard compression of hyperspectral images, namely the revised 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 lossless CCSDS 123.0-B-2. We will refer to this method as “pre-quantization”. Fig. 2 visually depicts the two methods. We study the performance of the two methods with and without an onground post-processing stage using the CNN for image dequantization. In the study we also consider a prediction mode offered by the new CCSDS 123.0-B-2 which excludes the pixel on the left of the pixel under coding from the prediction neighbor. This mode is intended to offer the highest throughput as in-loop decoding of the pixel on the left is the main speed bottleneck.

4. RESULTS

This section presents an experimental assessment of the performance of the proposed CNN for image dequantization when combined with the compression approaches presented in the previous section. For both approaches we set the CCSDS predictor in its full prediction mode with wide neighbor-oriented local sums. The rate-distortion performance is also measured against 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 123.0-B-2 lossy compressor set in reduced prediction mode with narrow neighbor-oriented local sums, which is the high-throughput mode provided by the recommendation.

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. Patches have been extracted from the decoded images. For the experiments, 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. We used the Adam optimization algorithm [24] with a learning rate equal to 10^{-8} for a total number of iterations corresponding to 1000 epochs. The convolutional layers have a fixed number of filters equal to 64.

The testing dataset is composed of the *sc0* scene from the AVIRIS Yellowstone images. For testing purposes the input




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




Fig. 2: Two predictive compression approaches. CCSDS 123 uses a quantizer inside the prediction loop. Prequantization quantizes raw pixel data and then applies a lossless predictor.

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 provided to the network by a sliding window selecting and results are obtained through a weighted average over the overlapping parts.

Quality is measured by the SNR computed as

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

Fig. 3 shows the rate-SNR curve. 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 suboptimality 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 high-throughput mode of the recommendation even without the CNN.

5. CONCLUSIONS

We proposed a CNN to dequantize hyperspectral images compressed by onboard compression methods. We also studied the performance of the revised CCSDS recommendation for onboard predictive compression and how it is a significantly simpler and faster method that eliminates data-dependency issues by moving quantization outside the prediction loop




Fig. 3: Rate-SNR performance of various compression methods with and without onground CNN for the *sc0* test image. 123-NL: lossy CCSDS 123.0-B-2 (full, wide, neighbor-oriented); Q+123-LS: prequantization followed by lossless CCSDS 123.0-B-2 (full, wide, neighbor-oriented); 123-NL-RED-NARROW: lossy CCSDS 123.0-B-2 (reduced, narrow, neighbor-oriented); 122-POT: CCSDS 122 and POT; CNN: CNN as post-processing.

can outperform it with the help of the onground CNN. We also discovered that high-throughput coding mode provided by CCSDS 123.0-B-2 is outperformed by the prequantization method, even without the CNN.

6. REFERENCES

- [1] Consultative Committee for Space Data Systems (CCSDS), “Image Data Compression,” *Blue Book*, November 2005.
- [2] Ian Blanes and Joan 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] Andrea Abrardo, Mauro Barni, Enrico Magli, and Filippo Nencini, “Error-resilient and low-complexity on-

- board 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] D. Valsesia and P. T. Boufounos, “Multispectral image compression using universal vector quantization,” in *2016 IEEE Information Theory Workshop (ITW)*, Sep. 2016, pp. 151–155.
- [6] 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.
- [7] Aaron B. Kiely and Matthew 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.
- [8] Consultative Committee for Space Data Systems (CCSDS), “Lossless Multispectral and Hyperspectral Image Compression,” *Silver Book*, , no. 1, May 2012.
- [9] 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.
- [10] D. Valsesia and E. Magli, “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.
- [11] D. Valsesia and E. Magli, “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.
- [12] Marco Conoscenti, Riccardo Coppola, and Enrico Magli, “Constant-SNR, rate control and entropy coding for predictive lossy hyperspectral image compression,” *IEEE Transactions on Geoscience and Remote Sensing*, (accepted for publication).
- [13] Consultative Committee for Space Data Systems (CCSDS), “Low-Complexity Lossless and Near-Lossless Multispectral and Hyperspectral Image Compression,” *Blue Book*, , no. 1, February 2019.
- [14] Nugehally S Jayant and Peter Noll, “Digital coding of waveforms: principles and applications to speech and video,” *Englewood Cliffs, NJ*, pp. 115–251, 1984.
- [15] 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.
- [16] 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.
- [17] 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.
- [18] Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei 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.
- [19] 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.
- [20] 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.
- [21] 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.
- [22] Jaakko Lehtinen, Jacob Munkberg, Jon Hasselgren, Samuli Laine, Tero Karras, Miika Aittala, and Timo Aila, “Noise2Noise: learning image restoration without clean data,” in *International Conference on Machine Learning (ICML)*, 2018.
- [23] Bing Xu, Naiyan Wang, Tianqi Chen, and Mu Li, “Empirical evaluation of rectified activations in convolutional network,” *arXiv preprint arXiv:1505.00853*, 2015.
- [24] Diederik P Kingma and Jimmy Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.