

Reduction of Interband Correlation for Landsat Image Compression

Daniel G. Acevedo and Ana M. C. Ruedin
Departamento de Computación,
Facultad de Ciencias Exactas y Naturales,
Universidad de Buenos Aires
anita@dc.uba.ar

Abstract

We present a lossless compressor for multispectral images that exploits interband correlations. Each band is divided into blocks, to which a wavelet transform is applied. The wavelet coefficients are predicted by means of a linear combination of coefficients belonging to the same orientation and spatial location. The prediction errors are then encoded with an entropy-based coder. Our original contributions are i) the inclusion, among the candidates for prediction, of coefficients of the same location from other spectral bands, ii) the calculation of weights tuned to the landscape being processed, iii) a fast block classification and a different band-ordering for each landscape. Our compressor reduces the size of an image to about a fourth of its original size. Our method is equivalent to LOCO-I, on 3 of the images tested it was superior. It is superior to other lossless compressors: WinZip, JPEG2000 and PNG.

1. Introduction

The high cost of transmitting and storing large amounts of data is a major concern. It imposes a permanent challenge as to the design of efficient techniques for data compression. Satellite images present a problem, due to their size: e.g., an 8-band Landsat7 image, plus a panchromatic band, typically requires 600 MB to be stored. The storage of original satellite images imposes in addition one nontrivial condition: compression must be lossless.

Despite the existence of different excellent lossless image compressors, such as JPEG-LS or LOCO-I [9], SPIHT [12], CALIC [17], CREW [2], JPEG2000[14], and commercial products such as MrSID, we think that there is room for improvement. Landsat 7 images are eight-band multispectral images obtained from spectrally filtering radiation at visible (bands 1,2,3), near-infrared (band 4), short-wave infrared (bands 5,7), and thermal (bands 6.1,6.2) frequency bands from the sun-lit earth. In figure 1 we have a Landsat




Figure 1. Eight bands of a LANDSAT 7 image.

Band	Wavelength
1	0.45 – 0.52 μm
2	0.52 – 0.60 μm
3	0.63 – 0.69 μm
4	0.76 – 0.90 μm
5	1.55 – 1.75 μm
6.1,6.2	10.4 – 12.5 μm
7	2.08 – 2.35 μm

Table 1. Wavelength of each band.

7 image, and in table 1 are listed the wavelengths of each spectral band.

The bands of multispectral images are highly correlated. This fact should be taken into consideration when designing a compressor for such images. To our knowledge, none of the mentioned image coders – and no available satellite image coder– exploits the interband correlations of multispectral images in order to improve compression.

LOCO-I [9] is context-based predictive code. However, its otherwise excellent performance decreases in the case of noisy satellite SAR images, which have multiplicative speckle noise. It becomes necessary to find algorithms

based on the deepest structure of the image. For that reason, we will apply the multiresolution analysis that wavelets provide.

Traditional wavelet transforms [7] yield real coefficients, which is not advisable for lossless compression. Integer-to-integer wavelets [1] have been developed, which are used in the SPIHT compressor [12], and in the JPEG 2000 norm [14]. Yet they do not take advantage of the correlations between spectral bands of a satellite image.

In section 2 we introduce the integer-to-integer wavelet transform that we apply on blocks of the image. In section 3 we explain our different models for linear prediction of the wavelet coefficients, which is followed by entropy-based encoding of the prediction errors. For the sake of reducing the entropy of these errors, we take advantage of the correlations between bands in the multispectral images, making predictions with the coefficients belonging to other spectral bands that have already been encoded. Another original contribution consists of tuning the weights for the linear prediction according to the landscape being processed, which requires a fast block classification.

In section 4 we explain our fast block classification, in section 5 we have our numerical results and in section 6 our concluding remarks.

2 Wavelet Transforms

The wavelet transform has been successfully applied to solve different signal processing problems [15, 8]. It is invertible and efficient. When applied to an image, the image is split into its details at different scales, orientations and positions. The transformed image is decorrelated.

The smaller the entropy of the data, the better it can be compressed with entropy-based coders [3]. For a given variance, the differential entropy is maximized when the distribution of the data is Gaussian. The non-Gaussian distribution of wavelet coefficients indicates that the wavelet transform is an adequate tool for compression.

We work with multispectral images typically having 8 spectral bands; the wavelet transform is applied to each band separately.

The large size of each band (around 7000 rows and columns) makes it unpractical to transform it as a whole; instead, we calculate the wavelet transform on blocks of 256×256 . This partition will not have any side effect because compression is lossless.

Each step of the wavelet transform is carried out by filtering the data – with a lowpass filter and a highpass filter – followed by downsampling. In the case of traditional wavelets [7], the filters are made of real numbers. The values of the original image are integers in the range $[0 - 255]$, and after filtering we obtain real coefficients: the transformed image takes more place than the original one, which is not advis-

able for compression. Different integer-to-integer wavelets [1, 13] have been developed, which are used in the SPIHT compressor [12], and in the JPEG 2000 [14]. These transforms are reversible when the values of the original image are integers, as is the case.

The most simple integer to integer transform is a variant of the Haar wavelet, known as the S transform [11]. Given the original signal $s_n, n = 1 \dots N - 1$, we define one step of the S transform as the sequences c_n (approximation or low-pass coefficients) and d_n (detail or high-pass coefficients):

$$c_n = \left\lfloor \frac{s_{2n} + s_{2n+1}}{2} \right\rfloor, \quad (1)$$

$$d_n = s_{2n} - s_{2n+1}. \quad (2)$$

To obtain more steps of the transform, the formula is to be applied successively to the approximation coefficients. The S+P transform was introduced by Said and Pearlman [11]. This consists in applying the S transform and then operating a differential prediction on the detail coefficients, rounding the result and subtracting it from d_n to give d'_n .

Let $\{c_n, d_n\}$ be the S transform of the signal s_n . Calculate

$$\hat{d}_n = \frac{1}{4} \Delta c_n + \frac{3}{8} \Delta c_{n+1} - \frac{1}{4} d_{n+1}, \quad (3)$$

where

$$\Delta c_n = c_{n-1} - c_n, \quad (4)$$

and obtain the new detail coefficients

$$d'_n = d_n - \left\lfloor \hat{d}_n + \frac{1}{2} \right\rfloor. \quad (5)$$

Then $\{c_n, d'_n\}$ is the S+P transform of s_n (one step). To obtain the transform of an image, the above formulae are to be applied to the rows of the image and to the resulting columns.

In order to evaluate the performance of S+P compared to different integer-to-integer wavelet transforms: (2,2), (6,2) and (2+2,2) [1], the entropy of the transformed file was computed on a Landsat 7 image – see figure 2. Clearly the performance of the S+P transform was superior to the others. The same behaviour was observed when the transforms were tested on other images. So the S+P was the transform chosen for our compressor.

The S+P is not an orthogonal transform; errors in the transformed coefficients – such as quantization errors – are magnified in the reconstructed image. This, however, is not an issue: we aim at lossless compression and the wavelets coefficients are not be quantized.




Figure 2. Entropies of different integer to integer wavelet transforms.



Figure 3. A block of an image.

3 Prediction of Wavelet Coefficients

Wavelet transforms decorrelate images, eliminating spatial redundancy, i.e. the entropy of the transformed coefficients is much lower than that of the original image. Even so, similarities can be observed between the absolute values of the coefficients corresponding to the same spatial location at different scales – see the block of an original image in figure 3, and its wavelet transform in figure 4, where we have the 9 detail subbands and 1 approximation subband of 3 steps of the S+P transform. Inside white circles and squares are observable similarities.

Figure 5 shows the order in which the wavelet coefficients are encoded or sent. We take advantage of the mentioned similarities, confirmed by numerous tests [5], in order to improve compression. Instead of encoding the




Figure 4. S+P transformed block. In white, coefficients above a threshold. Circled and squared: similarities between subbands.

wavelet coefficients directly, each wavelet coefficient is predicted with a linear combination of coefficients that have already been encoded: we encode the difference between the predicted coefficient and its real value. If prediction is accurate, prediction errors are small and this favours compression. Afterwards, these prediction errors are encoded by means of an entropy-based coder.

In figure 6 we have, in black, the actual wavelet coefficient to be encoded, in light grey, the coefficients that have been encoded at this stage – these coefficients are available both to the encoder and to the decoder–, and in dark grey, the coefficients to be encoded.

We follow Buccigrossi and Simoncelli [5] in the prediction of wavelet coefficients: their work was applied to the lossy compression of photographic images. The prediction of each wavelet coefficient is performed with a linear combination of i) its neighbours at the same scale, and ii) coefficients corresponding to coarser scales in the same spatial location. This implies a model for the statistical dependencies between wavelet coefficients at different scales: this model states that the dependencies are fixed inside each subband, and do not depend on the coefficients' actual value. Candidates for prediction are shown in figure 7; they must have been encoded before the actual coefficient, in order to be able to make the same prediction at the decoder.

To include all candidates would be a heavy computational overload. For the training phase, offline we ran a greedy algorithm to determine a few of the best candidates for each subband. This algorithm added candidates, one by one. The weights for the linear combination were calculated by least squares, over many transformed blocks of images, and averaged. We chose the candidates that minimized the




Figure 5. Ordering of the wavelet subbands.

entropy of the prediction errors. After adding 3 candidates, there was no significant improvement if we added more. The weights and the position of the best predictors (candidates) for each subband were then incorporated into the compressor.

In order to reduce interband correlations for Landsat images, the model proposed in [5] was slightly modified. We develop here what was sketched in our earlier work [4]. Coefficients are now predicted by taking into account: iii) the coded coefficients in the same position belonging to other spectral bands, as well as i) and ii). As a consequence, although each block is processed independently, the information of the 8 blocks having the same position in different bands has to be available in memory. In figure 8 are shown all the candidates for prediction of one wavelet coefficient. The actual coefficient is colored in black, the candidates in the same spectral band in white, and the candidates from other bands in light grey. As a result of the high interband correlation, the candidates from other spectral bands were frequently chosen when our compressor was running. Since the candidates have to be available to the decoder, in order to recover the image from the compressed code, they have to belong to already processed bands. This induces an order for processing the spectral bands of each block, where highly correlated bands are consecutive [16].

At this stage our results were not too good. In the training phase we noticed great variations both in the positions of the best candidates for linear prediction and in the weights; these variations were noticeable when the landscape changed.




Figure 6. Actual coefficient (black) and encoded coefficients (light grey).

4 Block classification

We subsequently modified the model proposed in section 3 for the statistical dependencies of wavelet coefficients. In this new model, the dependencies are different according to the landscape of the block. That is, we consider that inside a block belonging to a certain landscape, each wavelet coefficient is statistically dependent of i) its neighbours at the same scale in the same subband, ii) the coefficients corresponding to coarser scales in the same spatial location from other subbands, and iii) the coefficients in the same position belonging to other spectral bands.

A quick block classification was needed, in order to: a) run an algorithm offline that calculated the positions of the best predictors and their weights for each landscape, these were included into the compressor; b) when the compressor was running, to determine to which class the block belonged, so as to operate with the predictors and weights of that class.

If the reflectance of the earth is plotted in terms of the wavelength, different typical curves are obtained according to each landscape. These reflectance curves are called the spectral signature of the landscape [6, 10]. The reflectance curves for water, snow, vegetation and bare soil are shown in figure 9. For our block classification, first each pixel is classified according to its spectral signature. We developed short ad hoc classification algorithms based on the shape of these reflectance curves.

For example, in the case of vegetation: notice (figure 9) that the reflectance curve for vegetation has a sharp increase at around wavelength $0.7 \mu m$. Now a Landsat image consists of 8 samples at the wavelengths given in table




Figure 7. Candidates (white) for the prediction of an actual coefficient (black).




Figure 8. Candidates from other bands (in grey).

1 for each position (pixel). As a consequence of this increase we have $Band3 < Band4$. Our algorithm classifies a pixel as vegetation if the well-known vegetation index $Band4 - Band3$ is positive. Similarly, a pixel is classified as water if $Band4 - Band3 < 0$ and the values for band 4 and band 7 are small. If there are many classes inside a block, the class having the most pixels determines the classification of the whole block. For the block of figure 10, the pixels classified as water are shown in figure 11, the ones classified as vegetation are shown in figure 12, and the ones classified as city are shown in figure 13. If a block contains pixels of different classes, the calculated predictors and weights are not optimal for that block. Since we were pursuing lossless compression – and not an exact class segmentation – the only effect is that the prediction is less accurate and our compression rate is slightly less for blocks which have a mixture of classes, but there is no loss




Figure 9. Reflectance curves for different types of landscapes.



Figure 10. A block to be classified.

of information. The same holds if a pixel is missclassified : there will be less compression but no loss of information.

5 Results

Our compressor was tested on 4 Landsat 7 images from Argentina: Buenos Aires, Santa Cruz, San Luis and Mendoza. In subsequent graphs we show how much we gained by the algorithms applied by our compressor for the image of San Luis. In figure 14 we have the entropy of each band of the original image (dotted) and the entropy of the wavelet S+P coefficients (solid); the gap between both curves indicates how much we have gained by applying the wavelet transform.

In figure 15 we have the bpp (bits per pixel) for each band of the S+P wavelet coefficients (dotted), versus our method based on the prediction of wavelet coefficients. The gap indicates how much we have gained by prediction.

In figure 16 we have the performance of our method, in bpp, compared to the results of other lossless compressors : Winzip, LOCO-I, JPEG2000 and PNG, applied to the image



Figure 11. Pixels classified as water (in white).




Figure 13. Pixels classified as city (in white).



Figure 12. Pixels classified as vegetation (white).




Figure 14. Entropy of the S+P transform (solid line) vs original image (dotted line).

of San Luis. The different shades of the bars indicate the volume occupied by each band in the compressed image. We have the final results of the 5 mentioned lossless compressors applied to the mentioned 4 images of Argentina, in table 2. Landsat 8-band images have 8 bpp. For the image of San Luis, each band has 7044 rows and 7476 columns. For San Luis, our method gives 2.07 bpp, that is, a saving of 5.93 bpp, and a compression rate of 3.86:1. In table 3 are the execution times taken for a whole image –on a Pentium IV running at 2.4 GHz– by different compressors and ours. Although our compressor is slower than several other methods, our aim was to construct a lossless compressor that would have high compression rates for storing images at the base. In the balance between compression and speed, we chose to favour compression.

6 Conclusions

We have developed an image compressor that exploits the high correlation between bands of the Landsat7 images. The model for dependencies of wavelet coefficients proposed by Simoncelli and Buccigrossi was modified so as to include information from other spectral bands for an efficient prediction of the wavelet coefficients. This model was further modified so as to tune the weights and the positions of the predicting coefficients to the landscape of the block being processed. This required a fast block classification, and a different band ordering according to the landscape.

Our results are highly satisfactory, and on all the images tested the compression rates of our compressor are superior to those of PNG, JPEG2000 and WinZip. The performance of our method is comparable to that of LOCO-I; in several




Figure 15. S+P transform(dotted) vs our method (solid) in bpp.




Figure 16. Results in accumulated bpp.

cases our method outperforms LOCO-I.

For future work we plan to perform a nonlinear prediction of wavelet coefficients, and integrate this prediction into our compressor.

Acknowledgements: we thank CONAE (Comisión Nacional de Actividades Espaciales), Argentina, for gently providing the test images, and Haydée Karszenbaum (Instituto de Astronomía y Física del Espacio) for her generous help.

References

[1] W. S. A. R. Calderbank, I. Daubechies and B. Yeo. Wavelet transforms that map integer to integers. *Applied and Computational Harmonics Analysis*, 5(3):332–369, 1998.

[2] E. S. A. Zandi, J. Allen and M. Boliek. Crew: Compression with reversible embedded wavelets. *Proceedings. DCC 1995 Data Compression Conference*, pages 212–221, 1995.

[3] N. Abramson. *Teoría de la Información y Codificación*. Paraninfo, 1986.

	PRED	WINZIP	LOCO	JPG2000	PNG
Buenos As	104.76	161.05	106.88	113.53	121.67
Santa Cruz	90.37	139.94	92.43	98.54	105.61
San Luis	103.98	159.24	105.33	111.3	117.92
Mendoza	116.72	179.8	114.83	125.1	136.65

Table 2. Results in MB.

PRED	402
WINZIP	124
LOCO	44
JPG2000	242
PNG	427

Table 3. Execution time (in seconds).

[4] D. Acevedo and A. Ruedin. Lossless compression of landsat images. *Argentine Symposium of Technology, 33 IAIO*, 2004.

[5] R. Buccigrossi and E. Simoncelli. Image compression via joint statistical characterization in the wavelet domain. *IEEE Trans Signal Proc.*, 8:1688–1701, 1999.

[6] E. Chuvieco. *Fundamentos de teledetección espacial*. RI-ALP, 1996.

[7] I. Daubechies. *Ten lectures on wavelets*. Society for Industrial and Appl Mathematics, 1992.

[8] S. Mallat. *A Wavelet Tour of Signal Processing*. Academic Press, 1999.

[9] G. S. Marcelo Weinberger and G. Sapiro. The loco-i lossless image compression algorithm: principles and standardization into jpeg-ls. *IEEE Transactions on Image Processing*, 9:1309–1324, 2000.

[10] NASA. *Landsat 7 Science Data Users Handbook*. <http://ltpwww.gsfc.nasa.gov/IAS/handbook>.

[11] A. Said and W. Pearlman. An image multiresolution representation for lossless and lossy compression. *IEEE Trans Image Proc.*, 5(9), 1996.

[12] A. Said and W. Pearlman. A new fast and efficient image codec based on set partitioning in hierarchical trees. *IEEE Transactions on Circuits and Systems for Video Technology*, 6:243–250, 1996.

[13] F. Sheng, A. Bilgin, P. Sementilli, and M. Marcellin. Lossy and lossless image compression using reversible integer wavelet transforms. *ICIP*, 1998.

[14] A. Skodras, C. Christopoulos, and T. Ebrahimi. The jpeg 2000 still image compression standard. *IEEE Signal Processing Magazine*, 18(5):36–58, September 2001.

[15] G. Strang and T. Nguyen. *Wavelets and Filter Banks*. Wellesley Cambridge Press, 1996.

[16] S. R. Tate. Band ordering in lossless compression of multispectral images. *IEEE Transactions on Computers*, 46(4):477–483, 1997.

[17] N. M. X. Wu. Context-based, adaptive, lossless image coding. *IEEE Transactions on Communications*, 45(4):437–444, 1997.