Compression of Multispectral Images

M. Klimesh
Communications Systems and Research Section

Compression of multispectral images (obtained from sensors that sample in both the spatial and the spectral domains) is important for reducing the transmission and storage requirements of such data. This article describes versatile software developed to simulate a family of compression algorithms. Two algorithms are selected as the most suitable for implementation. The first is a moderately high complexity algorithm consisting of the Karhunen–Loève transform (KLT) in the spectral dimension, the discrete cosine transform (DCT) for spatial decorrelation of the resulting bands, and the DCT on the residual. The other is a medium-complexity algorithm that uses predictive coding for spectral decorrelation and the DCT for spatial decorrelation. Performance results are given for these algorithms. A low-complexity algorithm is also discussed.

I. Introduction

A multispectral image is a data set with two spatial dimensions and a spectral dimension. These images are obtained from a multispectral scanner (or imaging spectrometer) on board an aircraft or spacecraft. Due to the spectral and spatial detail of multispectral images, they have a wide variety of Earth remote-sensing uses in studies of the oceans (including sea ice), the atmosphere (including cloud cover), and land. Specific applications include global environmental monitoring, mapping, natural resource management, and land use planning. Fine spectral resolution allows for applications such as estimation of crop health and yields, evaluation of vegetation stresses, and selective mapping of various aerosols over clear water [4]. Multispectral scanners also are used in deep-space missions to map the surface spatial distribution of the mineral and chemical features of various targets such as planets and their satellites. For example, the Cassini visible and infrared mapping spectrometer subsystem (VIMS) will gather multispectral images of Saturn, its rings, and its satellites.

Existing and proposed multispectral scanners for Earth remote sensing produce multispectral image data at rates ranging from about 4 Mb/s to more than 4 Gb/s. Multispectral scanners for use in deep-space missions produce data at a somewhat lower rate but still produce enormous data sets. In order to make transmission and storage of these data practical and less expensive, it is advantageous to compress these data. In particular, a spacecraft with a multispectral scanner may be unable to transmit all of the image data it produces in uncompressed form due to the rate limitation of its communication link; thus, onboard compression is needed to make full use of the scanner.

In this article, we investigate multispectral image compression techniques. All of our tests are done on an airborne visible and infrared imaging spectrometer (AVIRIS) data set. This data set contains 224 spectral bands, two of which are shown in Fig. 1.

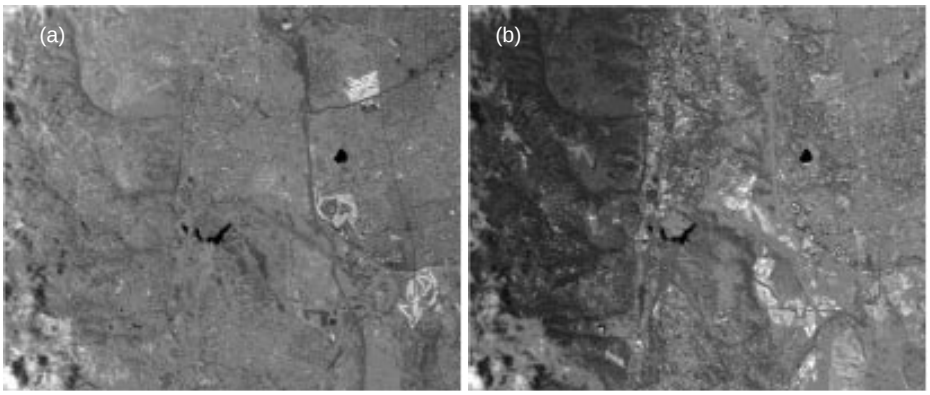

Fig. 1. Two spectral bands from an AVIRIS image: (a) band 60 and (b) band 140.

We focus on lossy compression techniques that have the potential for being suitable for onboard compression. In a later article, we plan to give a discussion of the considerations used in selecting the algorithms, a detailed literature survey, a comprehensive discussion of the algorithms implemented, and several possible improvements.

II. The Compression Software

In order to test a large number of algorithms and variations, we developed C software to determine the performance of a variety of algorithms. The individual components of this general algorithm have been used for multispectral image compression in other studies. For example, the Karhunen–Loève transform (KLT) has been used in a similar way in [5,9,10], predictive coding has been used in a similar way in [1], and the use of the discrete cosine transform (DCT) for spatial decorrelation is very common [8]. The selective combination of these steps results in several algorithms, some of which are novel.

In the following, we describe our general compression algorithm. As we will discuss shortly, by eliminating certain steps and varying some of the algorithm parameters, we obtain different algorithms with various performance and complexity characteristics.

The major steps performed by our general algorithm are as follows:

(1) Scale the data: So that different spectral bands can be encoded at different fidelities, each band is independently linearly scaled. In our tests, the scaling factors were all greater than or equal to one and resulted in most bands having root-mean-squared (RMS) sensor noise levels that were in the range of 50 to 100 (a few were higher). Most scaled bands had a range of about 30,000. The same scaling was used for all of our tests.

(2) Perform the KLT in the spectral dimension: This transform is adaptive in that the transform is recomputed for each $x_{\mathrm{KLT}} \times y_{\mathrm{KLT}}$ (spatial) block of the image. This computation may be performed using the statistics from the original data (in which case the transform parameters must be transmitted as overhead) or from the final reconstructed image based on the $x_{\mathrm{KLT}} \times y_{\mathrm{KLT}}$ block to the left of the block to be transformed (in which case no overhead is required except for the first block in each row).

(3) Compress the KLT bands: The n_{KLT} most-significant (highest-energy) KLT-transformed bands are compressed using the DCT (on $n_{\mathrm{DCT}} \times n_{\mathrm{DCT}}$ blocks), uniform quantization of the resulting coefficients (step size q_{KLT}), and entropy coding of the quantized coefficients. The remaining KLT-transformed bands are discarded.

(4) Compute the residual: Invert Steps (3) and (2) to obtain an approximation of the original multispectral image block. Then form a difference image by subtracting this approximation from the original image. We refer to this multispectral difference image as difference image 1 (DI1). If no KLT transform is performed (by setting $n_{\mathrm{KLT}} = 0$), then DI1 is the original image.

(5) Compress the residual (DI1) bands: One $x_{\mathrm{pred}} \times y_{\mathrm{pred}}$ (spatial) block is compressed at a time. Within each of these blocks, compress the bands consecutively as follows:

   (a) (Optional.) Estimate the DI1 band from the previous (reconstructed) DI1 bands. A first-, second-, or third-order affine predictor is used (the order indicates how many previous bands are used to estimate a value). The predictor parameters may be computed using the actual DI1 (in which case these parameters must be transmitted as overhead) or they may be computed from the final reconstructed DI1 of the $x_{\mathrm{pred}} \times y_{\mathrm{pred}}$ block to the left of the current one (in which case no overhead is required except for the first block of each row).

   (b) Subtract the prediction from DI1, forming difference image 2 (DI2). If no prediction was performed, then DI2 is identical to DI1.

   (c) (Optional.) Encode DI2 using the DCT (on $n_{\mathrm{DCT}} \times n_{\mathrm{DCT}}$ blocks) with a fixed uniform quantizer (step size q_{DCT}) and entropy coding of the resulting quantized coefficients.

   (d) Reconstruct DI2 by inverting Step (5c) and subtract this reconstructed DI2 from the actual DI2, forming difference image 3 (DI3).

   (e) (Optional.) Encode DI3 using a fixed uniform quantizer (step size q_{final}) and entropy coding of the resulting coefficients.

Figure 2 contains a block diagram of some of the steps performed by this algorithm. Steps (2) through (4) can be effectively skipped by choosing $n_{\mathrm{KLT}} = 0$; this eliminates the KLT. Although the KLT is the best linear transform for performing spectral decorrelation, in some cases the rate-distortion performance improvement resulting from its use does not justify its relatively high computational complexity. Steps (5a) and (5b) may be skipped by choosing no prediction. If the spectral decorrelation achieved by the KLT transform is adequate, these predictive coding steps may not result in a performance improvement. Steps (5c) and (5e) are optional. If DI2 does not contain significant spatial correlation, then it may be appropriate to use the straight quantization of Step (5e) rather than the DCT coding of Step (5c). Even if the DCT step is included, it may be reasonable to also use straight quantization to reduce the maximum difference between the original and reconstructed images.

We require that x_{KLT} must be a multiple of x_{pred}, which must be a multiple of n_{DCT}. Similarly, y_{KLT} must be a multiple of y_{pred}, which must be a multiple of n_{DCT}. The meanings of these parameters are illustrated in Fig. 3. For our experiments, we always use $n_{\mathrm{DCT}} = 8$. A prediction block is a spatial region where the estimator of the pixel values of a given band is a fixed function of the pixel values in the adjacent spectral bands (but varies with the band being estimated), and a KLT block is a spatial region where, for each spatial location, the spectral samples are transformed by the same orthogonal transform.

A. Observations

The general algorithm compresses sections of (usually y_{KLT}) rows of the image (including all bands) separately. This property could be utilized to confine the effects of a channel error to one such section if appropriate synchronization markers are inserted between the sections.

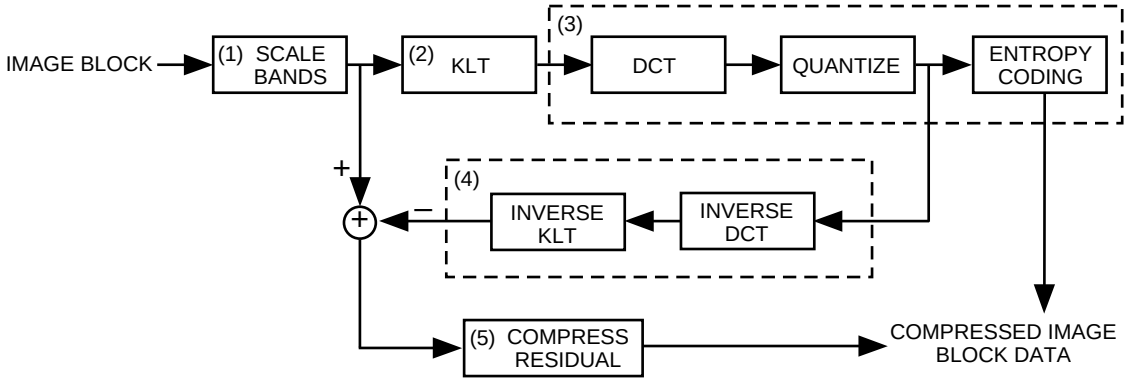

Fig. 2. The steps the general algorithm performs on an x_{KLT} by y_{KLT} block of the image, without the details of Step (5).

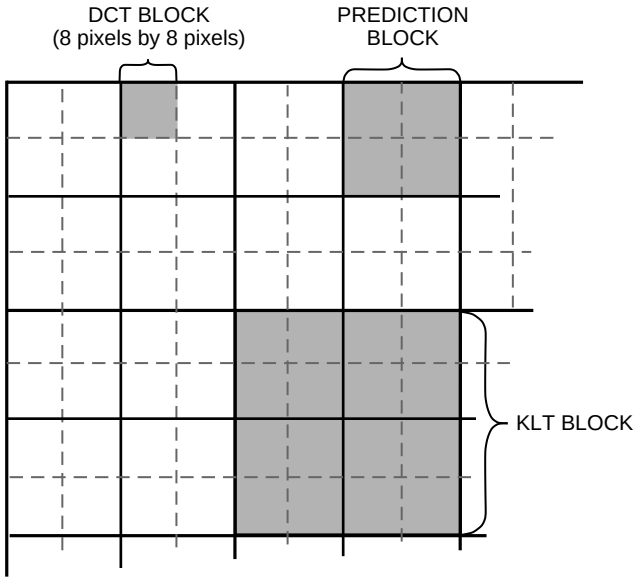

Fig. 3. The various blocks used by the general algorithm, with $n_{DCT} = 8$, $x_{pred} = y_{pred} = 16$, and $x_{KLT} = y_{KLT} = 32$.

In our tests, Step (5e) appears never to reduce the mean-squared error (MSE) of the reconstructed image. However, this step does limit the maximum error of a sample. Note that we use uniform quantization in this step (and the other quantization steps), which does not yield optimal performance [3].

Perhaps not surprisingly, for the cases we tested, the interband prediction steps [Steps (5a) and (5b)] produced no performance improvement when the KLT was used.

B. The KLT Algorithm

The recommended high-performance algorithm, which we refer to as the KLT algorithm, consists of the KLT, the DCT on the transformed bands, and the DCT on the residual. That is, Steps (1) through (4) and Step (5c) are included from our general algorithm. Although the DCT on the residual generally adds rate without significantly improving the performance in terms of MSE, it is important for accurate reconstruction of exceptional features of the images. Also, this step is needed if overhead is to be saved by computing the KLT parameters from the reconstructed image blocks, which may result in a net reduction in rate.

We use $x_{\mathrm{KLT}} = y_{\mathrm{KLT}} = 32$, and we vary q_{KLT}, q_{DCT}, and n_{KLT} to try to achieve the best performance. For most of our tests, we compute the KLT parameters from previously reconstructed horizontally adjacent image blocks (when possible) to minimize overhead.

C. The Interband Prediction Algorithm

For reduced complexity, we eliminate the KLT computation. We refer to this moderate-complexity algorithm as the interband prediction algorithm; it consists of the interband prediction in conjunction with the DCT. Only Step (1) and Steps (5a) through (5c) are included from our general algorithm.

In our tests, we use $x_{\mathrm{pred}} = y_{\mathrm{pred}} = 16$ and vary the predictor order. The value of q_{DCT} is varied to obtain the desired root-mean-squared error (RMSE). We minimize overhead by estimating predictor parameters from horizontally adjacent reconstructed values.

III. Performance

The KLT algorithm and the interband prediction algorithm were tested on a 608×512 AVIRIS image with 224 spectral bands. The performance will vary for images obtained from other multispectral scanners (due to differences in the type and number of spectral bands, the spatial resolution, the sensor noise, and other factors) and, to a lesser extent, for different images from the same scanner; however, the same trends will be observed. The original data were gathered with 10-bit samples. Some of our tests were performed on a subset of the bands. Tables 1 and 2 contain some representative results. The bit rates are estimated from an entropy measurement of the quantized coefficients; the rates are slightly pessimistic because some bits are left uncoded and correlations between the coefficients, which likely would be exploited in practice, are ignored. In these tables, the total bit rate includes all of the overhead necessary to decompress the image. Rates are expressed in bits/pixel/band (b/p/band).

Table 1. Sample compression performance results.

Method	Spectral bands	RMSE	Total rate, b/p/band
Prediction, first order	20	100	0.94
Prediction, second order	20	100	0.72
KLT algorithm	20	100	0.37
KLT algorithm	224	85	0.20

Compression results in Table 2 show the bit rates required to obtain an RMSE of 50. The algorithm and parameters used are indicated. Table 2 also lists the number of reconstruction errors that have a magnitude above a certain threshold (chosen to illustrate the differences between the algorithms); a small value suggests more accurate reconstruction of exceptional features of the image. It is instructive to note that the KLT algorithm has very little advantage over the interband prediction algorithm at this RMSE, due to the noise in the original image. Figure 4 further compares the performance of the KLT algorithm to that of the interband prediction algorithm.

The RMSE value of 100 represents very high fidelity compression, since we have multiplied most bands by constants that give RMS sensor noise levels of around 50 to 100. For this RMSE, typical bands have a signal variance-to-quantization (compression) noise ratio of about 27 dB. Because of the large dynamic range of the sensor (the high end of which is seldom reached), the traditional peak signal-to-quantization noise ratio (PSNR) is not very meaningful. However, it could be estimated at 49 dB (for an RMSE of 100).

Table 2. Compression performance for compression of all 224 bands with an overall RMSE of 50 (the KLT algorithms use $n_{KLT} = 10$).

Method	Parameters	Number of large errors[a] (band 60)	Rate, b/p/band	
			Partial[b]	Total
Prediction, second order	$q_{DCT} = 205$	24	0.826	0.829
KLT/overhead[c]	$q_{DCT} = 243$, $\frac{q_{KLT}}{q_{DCT}} = \frac{2}{3}$	63	0.604	0.722
KLT	$q_{DCT} = 232$, $\frac{q_{KLT}}{q_{DCT}} = \frac{2}{3}$	72	0.727	0.733
KLT	$q_{DCT} = 230$, $\frac{q_{KLT}}{q_{DCT}} = 1$	77	0.716	0.722
KLT	$q_{DCT} = 226$, $\frac{q_{KLT}}{q_{DCT}} = \frac{3}{2}$	70	0.722	0.728

[a] This column shows the number of errors with magnitude greater than 200.

[b] The partial rate is the bit rate without the overhead of the KLT or prediction parameters.

[c] The KLT/overhead method uses the original data to compute the KLT, using overhead for the KLT parameters.

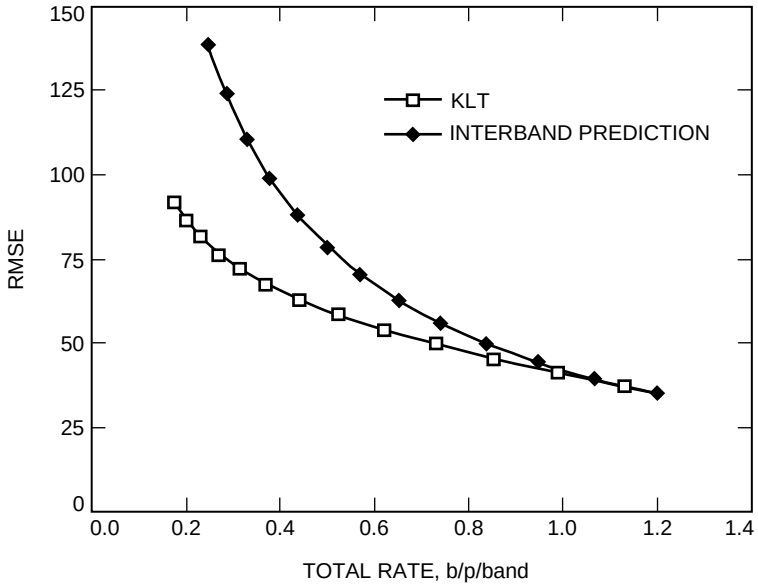

Fig. 4. Rate distortion performance of the KLT and interband prediction algorithms, where $n_{KLT} = 10$ and $q_{DCT} = q_{KLT}$ for the KLT algorithm, and second-order prediction was used in the interband prediction algorithm.

At an RMSE of 100, the reconstructed bands are completely indistinguishable from the originals, even when magnified, when displayed on a monitor that can display 256 shades of gray.

The interband prediction algorithm also was tested with a third-order predictor, but the improvement over the second-order predictor was insignificant.

IV. Low-Complexity Compression

If extremely high fidelity image reconstructions are desired, or if the interband prediction algorithm is too complex to implement, then it may be useful to consider a low-complexity algorithm. Although we did not test such an algorithm, it is reasonable to believe that the performance cost of using a simple algorithm diminishes as the quantization noise becomes much smaller than the sensor noise, because (conceptually) much of the bit rate is allocated to the compression of sensor noise, which is difficult to compress. One low-complexity algorithm is the use of some form of predictive coding (in all three dimensions) with very fine quantization steps, possibly in conjunction with some form of trellis-coded quantization [2,6,7]. This type of compression is not covered by our general algorithm.

V. Conclusion

We have described a general program for simulating a family of multispectral image compression algorithms. Two algorithms from this family have been discussed in detail, and performance results were presented. Very high fidelity compression is possible at bit rates as low as 0.2 b/p/band. We also have suggested a possible algorithm for use when complexity is a limiting constraint. Careful consideration of compression requirements is necessary for selecting an algorithm for a particular mission; however, we have provided a basis for evaluating the complexity/performance trade-offs involved.

References

[1] G. P. Abousleman, M. W. Marcellin, and B. R. Hunt, "Compression of Hyperspectral Imagery Using the 3-D DCT and Hybrid DPCM/DCT," *IEEE Trans. on Geoscience and Remote Sensing*, vol. 33, no. 1, pp. 26–34, January 1995.

[2] T. R. Fischer and M. Wang, "Entropy-Constrained Trellis-Coded Quantization," *IEEE Transactions on Information Theory*, vol. 38, no. 2, pp. 415–426, March 1992.

[3] A. Gersho and R. M. Gray, *Vector Quantization and Signal Compression*, Boston, Massachusetts: Kluwer, 1991.

[4] A. F. H. Goetz, G. Vane, J. E. Solomon, and B. N. Rock, "Imaging Spectrometry for Earth Remote Sensing," *Science*, vol. 228, pp. 1147–1153, June 1985.

[5] R. N. Hoffman and D. W. Johnson, "Application of EOF's to Multispectral Imagery: Data Compression and Noise Detection for AVIRIS," *IEEE Trans. on Geoscience and Remote Sensing*, vol. 32, no. 1, pp. 25–34, January 1994.

[6] M. W. Marcellin, "On Entropy-Constrained Trellis Coded Quantization," *IEEE Transactions on Communications*, vol. 42, no. 1, pp. 14–16, January 1994.

[7] M. W. Marcellin and T. R. Fischer, "Trellis Coded Quantization of Memoryless and Gauss–Markov Sources," *IEEE Transactions on Communications*, vol. 38, no. 1, pp. 82–93, January 1990.

[8] M. Rabbani and P. Jones, *Digital Image Compression*, Bellingham, Washington: SPIE Publications, 1991.

[9] J. A. Saghri, A. G. Tescher, and J. T. Reagan, "Practical Transform Coding of Multispectral Imagery," *IEEE Signal Processing Magazine*, vol. 12, no. 1, pp. 32–43, January 1995.

[10] V. D. Vaughn and T. S. Wilkinson, "System Considerations for Multispectral Image Compression Designs," *IEEE Signal Processing Magazine*, vol. 12, no. 1, pp. 19–31, January 1995.