

Real-time Hyperspectral Data Compression Using Principal Components Transformation

¹Suresh Subramanian,¹ Nahum Gat,² Alan Ratcliff,² Michael Eismann

1. INTRODUCTION

The advantages of Hyperspectral Imaging (HSI) over conventional remote sensing imaging has been well recognized (Schowengerdt, 1983). HSI provides the user with spatial and spectral information about scene objects and facilitates improved detection and identification over panchromatic and multispectral imagery. The predecessors of present day HSI systems were the broad band spaceborne multi-spectral imaging systems like MSS and the more recent TM, that provided data in 4 and 7 spectral bands respectively. Present airborne HSI systems like AVIRIS and HYDICE operate in the VIS-SWIR spectral region (0.4-2.5 μm). They acquire data at a higher spectral resolution (~ 10 nm) and produce data cubes in 224 and 210 bands respectively. The SEBASS sensor acquires data in the MWIR and LWIR regions in 100 and 128 spectral bands each.

The many spectral of bands of HSI sensors results in large sized data sets. For example, an AVIRIS image of size 614 x 512 spatial pixels, occupies about 140 Mbytes in a 2 byte binary storage format. Future generation ultraspectral sensors based on Imaging Fourier Transform spectrometers, will have an order of magnitude higher resolution than the present grating spectrometers, and place a correspondingly higher burden on data storage and transmission requirements. Data compression becomes increasingly important in (1) spaceborne HSI applications that have constraints placed on onboard storage, and data transmission bandwidths, and (2) real-time tactical UAV applications required to transmit imagery to ground stations. Increased onboard storage enhances the strategic and commercial significance of HSI by permitting larger area coverage. From an applications and tactical perspective, algorithms that can operate directly on compressed data, benefit from quicker turnaround times.

This work discusses issues related to real-time, lossy, hyperspectral data compression. The compression /decompression algorithm's performance is evaluated in terms of the reconstructed data's fidelity, using raw sensor data. It is important to develop suitable figures of merit (FOM), since the advantages of HSI in detecting small objects based on subtle spectral differences, can be lost if the reconstruction process does not reproduce the original data with sufficient fidelity. Several questions arise concerning definitions of data fidelity. For example, should this be defined in terms of pure "SNR" type quantities or in terms of applications involving data exploitation algorithms. Each position has its merits and demerits. These issues including real-time implementation are addressed in this paper. The well-known Principal Component Transformation (PCT) is used as a reference technique.

2. PCT BASED REAL-TIME COMPRESSION

2.1. Algorithm Description

PCT has been extensively utilized in remote sensing applications especially in the areas of image classification and compression (Richards et al, 1999, Wang et al, 1995, Fung et al, 1987). PCT compresses by removing redundant information due to spectral correlations that exist between bands in hyperspectral data. It is a linear transformation technique that rotates data into a coordinate system where the variables (spectral bands) are no longer correlated (Jolliffe, 1986, Richards and Jia, 1999, Therein, 1992, Parsons 1987). The rotation matrix that decorrelates the spectral bands comprises of the eigenvectors obtained by diagonalizing the covariance matrix of the data. Since the off diagonal elements of the covariance matrix measure the correlations between bands, the transformation finds the coordinate frame in which these terms are zero. The covariance matrix is symmetric in

¹ Suresh Subramanian (contact author), Nahum Gat: Opto-Knowledge Systems Inc., 4030 Spencer Street, Suite 108, Torrance, CA 90503-2442. Email: suresh@oksi.com, nahum@oksi.com

²Alan Ratcliff, Michael Eismann, Air Force Research Laboratory, Sensors Directorate, Building 622, 3109 P Street, Wright Patterson AFB, OH 45433-7700. Email: alan.ratcliff@sensors.wpafb.af.mil, michael.eismann@sensors.wpafb.af.mil

structure and can therefore be diagonalized using an orthonormal transformation (Golub et al, 1996, Horn *et-al*, 1985). The new axes (eigenvectors or principal components, PCs) are organized in the order of decreasing variance in the data. When this transformation is applied to the data by using only the first few principal axes or eigenvectors, we capture most of the energy and also the dominant spectral features in the signal in very few dimensions. This compact low dimensional representation of PCT transformed data is responsible for compression. The geometric significance of the reduction in the data dimensionality with minimal loss of information is that in order to identify L distinct classes of objects with spectra in K bands (where $K \gg L$), we require at the most L distinct coordinate axes, and often fewer than L suffice. The PCT transformed coefficients may themselves be subjected to further compression by ways thresholding, quantization, entropy coding, predictor based encoding, etc. However, we limit the current discussion to the PCT process only, and its effects on data fidelity.

The mean and covariance matrix of an image cube are computed using all the pixels. The spectrum of each pixel in the scene is represented by a K dimensional vector \mathbf{x} (for example, $K = 224$ for AVIRIS data). Let N be the total number of pixels in the image. The mean vector of this set, \mathbf{m} is denoted by

$$\mathbf{m} = E(\mathbf{x}) = \frac{1}{N} \sum_{i=1}^N \mathbf{x}_i \quad (1)$$

where $E(\mathbf{x})$ is the expectation of \mathbf{x} . The covariance matrix of the data, $\mathbf{\Sigma}_x$, is given by

$$\mathbf{\Sigma}_x = E\{(\mathbf{x} - \mathbf{m})(\mathbf{x} - \mathbf{m})^T\} = \frac{1}{N} \sum_{i=1}^N (\mathbf{x}_i - \mathbf{m})(\mathbf{x}_i - \mathbf{m})^T \quad (2)$$

The mean vector, \mathbf{m} is also a K dimensional vector and the covariance matrix $\mathbf{\Sigma}_x$ is a $K \times K$ symmetric matrix. Diagonalizing this covariance matrix $\mathbf{\Sigma}_x$ gives the required transformation matrix, G . An important property of $\mathbf{\Sigma}_x$ is that it is not only symmetric but also positive, semi-definite. In other words, its eigenvalues, λ , are always ≥ 0 . The eigenvectors (columns of G) corresponding to the largest few eigenvalues define the axes of maximal variance in the data. The data show very little variance along the remaining principal directions (eigenvectors) and are rejected. The initial principal components may be viewed as the bands where maximum energy is confined. The transformed data are computed by retaining only the eigenvectors corresponding to the largest few (L) eigenvalues and zeroing the rest. As a result, the transformation matrix G becomes rectangular with dimensions $N \times L$. The transformed vectors (pixels) are obtained by using the equation

$$\mathbf{y} = G^T (\mathbf{x} - \mathbf{m}) \quad (3)$$

where L is the number of eigenvectors retained in the transformation. Hence the dimension of a transformed pixel, \mathbf{y} , is L and the data compression ratio (CR) = $NK/(NL + 2(K+1)L)$. This expression for CR includes the burden of the transformation matrix and the mean pixel value that must also be transmitted as floating point numbers, together with the transformed data. When $N \gg 2(K+1)$, the expression for CR $\sim K/L$. The inverse transform used to recover the original data is a single step process given by:

$$\tilde{\mathbf{x}} = G \mathbf{y} + \mathbf{m} \quad (4)$$

2.2. Computational Complexity

The computational complexity of the PCT algorithm can be derived in terms of the previously defined image parameters like N , K , and L . The computational requirements (MFLOPS) can be calculated by assuming N to be a pixel rate instead of the number of pixels. The pixel rate is itself defined by the sensor data rate. There are 5 steps in the PCT compression calculation: (i) mean, (ii) mean removal (iii) covariance matrix, (iv) diagonalization of covariance matrix, and (v) data transformation/compression. Decompression is a single step process. The algorithm's computational complexity is summarized in Table I following Eqs. (1) to (4), in terms of atomic processor operations like read, write, add, and multiply (R, W, A, M). The formulas for the complexity are scaled according to the number of cycles required. For the Analog Devices SHARC DSP used in this work, R, A, and M are single cycle operations whereas W requires 2 cycles. Hence the mean removal step requires NK writes operations but a total of $2NK$ CPU cycles to compute. Further, operations involving R, A, and M can be done concurrently (in a single cycle) if the data are pipelined. Any special computations like division (D), etc., must be accounted for appropriately by determining the number of CPU cycles they require. Note that scaling or division by a constant α can be treated as a multiplication by $1/\alpha$, and subtractions are equivalent to additions. The diagonalization procedure requires special attention. It has a complexity $\sim O(K^3)$ but it is somewhat complicated to express in terms of atomic operations. Instead, a numerical factor of 5 derived from simulations conducted on various hyperspectral data sets, has been used

to express its complexity (see table I). The exact time required here is non-deterministic as computation proceeds until a convergence criterion is satisfied. The last column indicates the inherent parallelism or the parallelizability of the various steps in the algorithm.

Table I: Computational complexity of PCT algorithm in terms of atomic CPU cycles. The numbers in parentheses alongside the operations indicate the number of CPU cycles required.

No.	Operation	R (1)	A(1)	M(1)	W(2)	Parallelism	
1	Mean Calculation	NK	NK		K	Fine grain	
2	Mean removal	NK	NK		2NK	Fine “	
3	Covariance	NK^2	NK^2	NK^2, K^2	K^2	Fine “	
4	Diagonalization	$O(\sim 5K^3)$					Coarse “
5	Transform Data/ Compression	NKL	NKL	NKL	2NKL	Fine “	
6	Decompression	NKL	NKL	NKL	NK^2	Fine “	

The algorithm’s complexity allows us to develop computational strategies and design the multiprocessor real-time compression system. The quadratic and cubic dependence of the algorithm’s complexity in the covariance matrix calculations and diagonalization steps consume the bulk of the computation time. Spectral partitioning of data therefore reduces the burden imposed by these stages drastically. In our implementation, we partition the data into segments containing ~ 32 bands each. The resulting multiprocessor design is over 90% efficient in terms of processor utilization. The effects of data partitioning on data fidelity are discussed later.

The PCT algorithm belongs to a class of parallelizable algorithms where a single instruction/program can be simultaneously executed on multiple data streams (SIMD/SPMD). This is an important distinction from other classes comprising of sequential single instruction, single data problems, or those belonging to the multiple instruction, multiple data (MIMD) type. Algorithms belonging to various classes are most efficiently implemented only if the multiprocessor architecture and the processor attributes satisfy the class paradigm. The PCT algorithm supports fine grain and coarse grain parallelism in different steps of the computation. We note from Table I that all stages of the algorithm but one, support fine grain parallelism. This means that the data can be partitioned into very small sets, and each set can be independently executed on a separate processor for parallelization and speedup. Intermediate results computed on the various processors can be combined to get the overall result. This is possible only because of the absence of data dependencies within a particular stage of the computation.

2.3. Multiprocessor System Design

The SPMD paradigm supported by the PCT algorithm has been implemented using the Analog Devices Super Harvard Architecture (SHARC, model 21060) processor. The SHARC processor has a 40 MHz CPU and an independent I/O processor. Each SHARC has 6 link ports operating at 40 Mbytes/s for multiprocessor communications support. The independent I/O processor’s permits concurrent processing with I/O (ADSP User Manual, 1997). Clusters of SHARCs are used to process large amounts of data, and each SHARC’s effective link port bandwidth of 240 Mbytes/s permits efficient data flow across large networks. In recent years, real-time systems involving over 100 SHARC processors have been developed for real-time radar and sonar signal processing (Huxtable *et-al*, 1998, Mansfield *et-al*, 1998).

Fig. 1 shows a schematic layout of the multiprocessing real-time compression system implemented using board-level, commercial off the shelf components. It consists of a number of different industry standard SHARCpac that comprise the input, processing, and output modules. The input SHARCpac contains a single SHARC processor and a field programmable gate array (FPGA) device, operating under the SHARC’s control. The FPGA connects to the hyperspectral sensor using a custom interface. Data are double buffered on two embedded SHARCs (ES0 and ES1) and alternately sent to a pipeline of 12 compression SHARCs after spectral partitioning. The compressed data are double buffered on a pair of output SHARCs (ES2 and ES3), and finally written to a SCSI disk that is also under SHARC control.

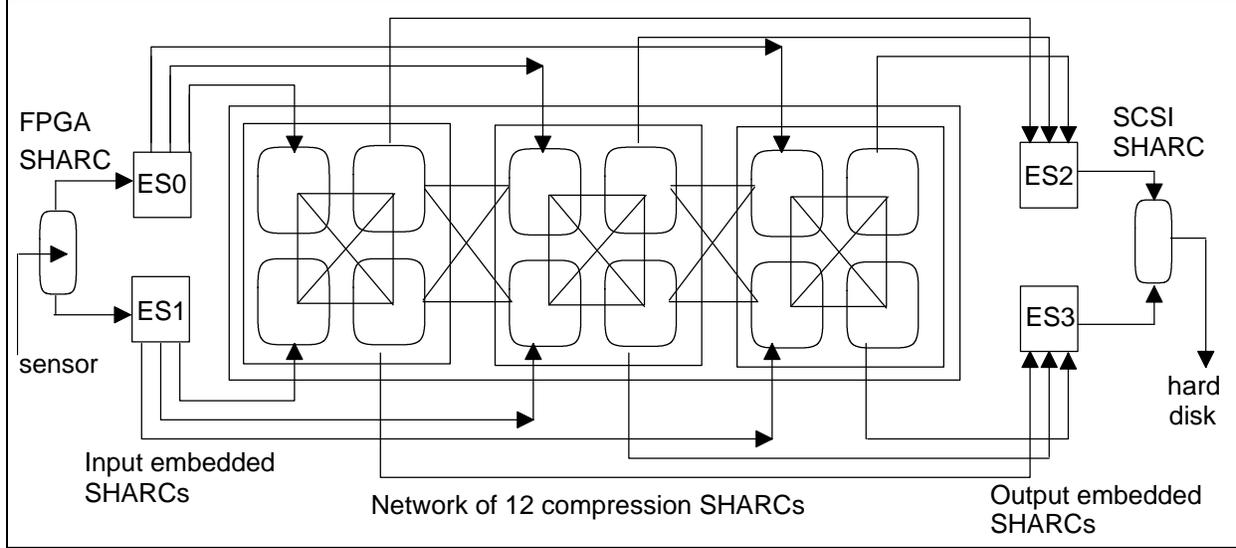


Figure 1: Schematic of multiprocessor architecture, interprocessor communication, and dataflow.

Interprocessor communications are handled using a combination of low level direct memory access (DMA) transfers as well as high level communication libraries. Fig. 1 shows the data flow and link port connections in the network. The input and output double buffers (on embedded SHARCs) involve direct processor to processor data transfers. They are implemented using highly efficient low level DMA transfers that achieve close to the theoretical maximum of 40 Mbytes/s, for even very small data packets (< 64 bytes). Data movement from the input double buffers to the processing network of 12 SHARCs, and from there to the output double buffers, is a complex process occurring over multiple link ports, and sometimes requires routing through intermediate SHARCs. High level communication libraries are used here as they offer programming ease and flexibility. The penalty for using these abstract high level functions is speed and performance. The maximum data transfer rate achievable for a 8K size packet is about 32 Mbytes/s. All the communications functions used are non-blocking and therefore support concurrent I/O and data processing. The present system is designed to handle sensor data rates up to 100 Mbits/s and can be scaled to handle higher data rates if required, by adding more SHARCs.

3. DATA FIDELITY FIGURES OF MERIT

3.1. Error metrics

We develop several error metrics to quantify reconstructed data fidelity after PCT compression. The various metrics bear relevance to different applications. For example, global error metrics like SNR curves and error distributions may be useful indicators of overall data quality for scientific applications. In addition to the global metrics, pixel by pixel local image metrics based on signal to clutter ratio type of quantities, and rms error maps, will likely be useful in tactical applications involving CFAR type detection algorithms. The objective here is to develop a number of application independent metrics that will nevertheless provide useful indicators of the likely performance in a particular application.

The symbol \tilde{x} in Eq. (4) indicates that the image pixel recovered after PCT compression and decompression is not identical to the original pixel x because all eigenvectors were not used in the compression/decompression process. However, the error or difference between x and \tilde{x} can be quantified in a mean square sense using the eigenvalues of the transformation matrix G . It can be shown that the mean square difference between x and \tilde{x} is given by:

$$R = \sum_{j=1}^K \lambda_j - \sum_{j=1}^L \lambda_j = \sum_{j=L+1}^K \lambda_j \quad (5)$$

Thus the mean square error of the PCT algorithm is given by the sum of the residual eigenvalues whose eigenvectors were not employed for the transformation. If all the eigenvectors are used in the transformation matrix G then $L=K$ and the error becomes zero. This condition is of course useless for data compression, as the transformed

vector \mathbf{y} is also K dimensional (like the input vector \mathbf{x}). Eq. (5) can be used to maintain a constant error level by varying the compression ratio (i.e., by varying the quantity L or the number of eigenvectors retained in G the transformation). The eigenvalues usually decrease rapidly (and monotonically). Hence very few of the corresponding eigenvectors are required to render the above error small, thus achieving a reasonable compression factor at an acceptable/user defined error level. The operating compression ratio for a given sensor may be determined from the data fidelity requirements of the application, and the sensor's SNR specifications.

A basic global metric is the SNR of the PCT processed data. We derive the SNR from the *true rms distortion error* caused by the PCT. The rms error e_{rms} , is a K dimensional vector defined as

$$e_{rms} = \sqrt{\left(\frac{1}{N-1}\right) \sum_{i=1}^N \frac{(x_i - \tilde{x}_i)^2}{x_i^2}} \quad (6)$$

where \mathbf{x} and $\tilde{\mathbf{x}}$ are the raw and PCT processed images respectively, and N is the number of scene pixels. The above summation is implied for each band of the image. The SNR metric is then defined as $SNR = 1/e_{rms}$, and it is also a K dimensional vector. The so-called noise (of our SNR metric) is really the error distortion introduced in the data by the compression/decompression algorithm. We note that the SNR obtained here will be lower than the well known PSNR which normalizes the rms error using the peak signal (for example, the denominator in eqn. (6) would read 255 instead of x_i , for 8 bit data). In addition to SNR curves at different operating points or compression ratios (CR), other quantities like min/max error, average error, and error standard deviation, all calculated band by band, are useful global data fidelity metrics. Measures like SNR may be used to set operating points on ROC curves, in target detection applications.

Local pixel metrics in the form of error maps, like band averaged error and signal to clutter ratio (SCR) type metrics, allow direct error visualization by comparing these maps to the original image. SCR is a measure of a pixels detectability and is a relevant metric for target detection applications. It is a scalar quantity defined as

$$scr = \sqrt{\mathbf{y}^T \Sigma^{-1} \mathbf{x}} \quad (7)$$

where Σ^{-1} is the image background covariance matrix, \mathbf{x} is a scene pixel, and \mathbf{y} is a required target pixel selected from a library. The geometric significance of this expression is one of a dot product of \mathbf{x} and \mathbf{y} after projecting the vectors into the subspace of Σ^{-1} (Theerein, 1992, Stocker et al, 1990). In order to use this metric meaningfully, we calculate the ratio of SCR (ρ) after and before PCT processing.

$$\rho = \frac{\sqrt{\tilde{\mathbf{x}}^T \Sigma^{-1} \tilde{\mathbf{x}}}}{\sqrt{\mathbf{x}^T \Sigma^{-1} \mathbf{x}}} \quad (8)$$

The \sim variables in the numerator correspond to the PCT processed image and the ones in the denominator refer to the original image. Since there is no a priori target pixel, we compare a pixel to itself in its surrounding defined by Σ^{-1} . The ratio ρ should remain as close to one as possible, thus indicating that a given pixel has not changed relative to its background. If $\rho > 1$, it indicates the possibility of increased false alarms, and when $\rho < 1$, we have a loss of detection sensitivity. Another local metric is a pixel error map that provides a band averaged map of errors. This is a useful visual aid for characterizing the algorithm's performance, by determining the pixels where the greatest errors occur.

3.2. Data Analysis

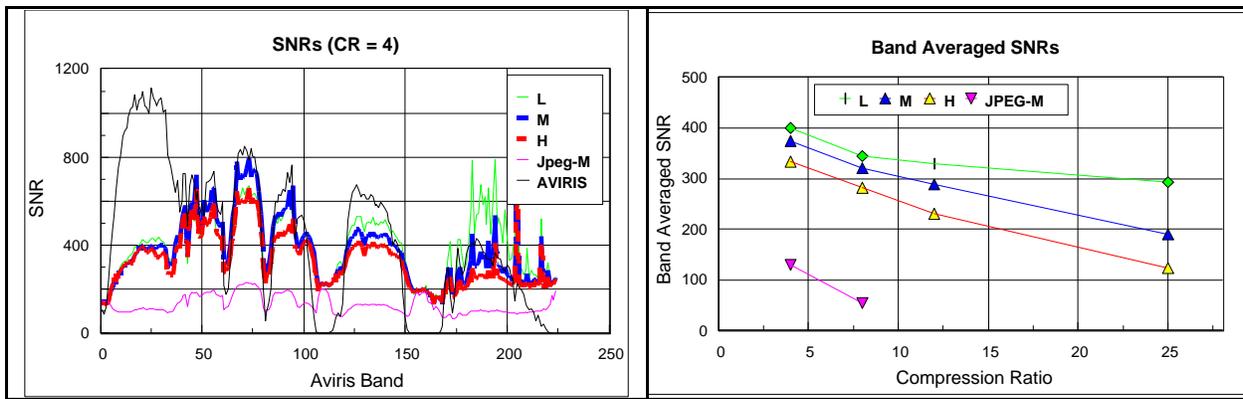
It is essential to note that PCT is an optimal data compression technique based on second order Gaussian statistics. It compresses by decorrelating data to remove redundant information, and organizing the results along axes of decreasing variance. Hence very few axes are required to capture most of the signal's energy. These "few" axes determine the natural minimum dimensionality of the data. Clearly, the fewer the number of axes, the lower the minimum dimensionality, and the higher the compression that can be attained. The eigenvalues of the PCT transformation are used to determine this dimensionality of the compressed data. Hence, the issue of data dimensionality becomes central to PCT compression. Data dimensionality is influenced by a scene's content; the larger the number of "distinct features or variability", the higher the dimensionality. In general, we expect spatially large scenes to have a higher dimensionality than smaller ones. Urban scenes will normally have higher dimensionality than corresponding size rural ones. In addition to scene related dimensionality, the sensor's noise characteristics also influence data dimensionality. High sensor noise will produce scatter in the data distributions and

thus increase their intrinsic dimensionality. We note here that since PCT makes an implicit assumption of Gaussian statistics, the higher order moments in data distribution should be identically zero. Violation of this condition will result in poor data quality after reconstruction. It is well known that natural scenes rarely exhibit Gaussian data distributions, and this may be a natural limitation of the PCT compression technique (Chen et al, 1987). The situation may be alleviated by selecting a large enough number of pixels (typically 10-100 x number of bands) so that the data distributions in each band may approach Gaussian. Despite the above limitation, our results from AVIRIS indicate an excellent compression performance using PCT.

The PCT algorithm's performance was benchmarked using raw data from the AVIRIS sensor, prior to radiometric calibration or geo-rectification. Since PCT removes spectral correlations to achieve compression, its real-time performance can be truly characterized only by testing on raw data. Most HSI instruments employ multiple spectrometers and detectors to span the sensor's spectral range (for example, AVIRIS has 4 spectrometers and 4 linear detector arrays). The raw data from these instruments will exhibit different correlation properties from their calibrated counterparts for reasons such as different -detector QEs, pixel nonuniformities and dark current, amplifier gains, and optical throughput of the different spectrometer assemblies. All these factors affect the data's correlation properties and hence impact the PCT algorithm's performance.



(a) Low variability (L) (b) Medium variability (M) (c) High variability (H)
Figure 2: AVIRIS scenes from low altitude 1998 flight of San Joaquin Valley



(a) (b)
Figure 3: (a) SNRs for 3 scenes at fixed CR = 4, and (b) band averaged SNR vs. CR. Bottom curve in both graphs shows results from 12 bit JPEG for scene M.

Fig. 2 shows three test images of size 250 x 150 spatial pixels x 224 spectral bands (AVIRIS Data, 1998). These images qualitatively represent scenes of low, medium, and high variability and are labeled as L, M, and H respectively. The terms low, medium and high “variability” are used for these scenes in the context of their data dimensionality, determined by their eigenvalues, as discussed earlier. Fig. 2(a) comprises of 2 or 3 homogenous fields and uncultivated land. Fig. 2(b) is a mix of several distinct cultivated and uncultivated fields and Fig. 2(c) is a mixture of urban and rural areas.

Fig. 3(a) shows the SNR curves for the 3 images at a fixed compression ratio (CR) of 4. The SNRs in the visible bands are comparable for the 3 images but the SNR values for low variability image L are higher in the SWIR bands. The bottom-most curve shows a 12 bit JPEG compression for the medium variability scene M, for comparison. It is clear from the poor SNR curve for JPEG that it is not a very suitable compression technique for

hyperspectral data. The SNR curve for the AVIRIS sensor is also shown alongside for comparison. The AVIRIS SNR values are ~ 1000 in the VNIR bands whereas the PCT SNR are ~ 400 . The results from PCT are however comparable to AVIRIS SNR in the SWIR bands. We therefore have an apparent loss of data fidelity in the VIS-NIR bands after PCT compression. The reason for this is that the raw signal intensity in these bands is of the order of 500-800 DN (mean pixel spectrum in Figs. 4(a) and 4(b), for example). Hence an error of 1-2 DN introduced during the PCT process would give rise to an SNR of about 400. Included in this is the numerical round off error of the order of 1 DN (two truncation errors of $\pm 1/2$ of floating point results to integer during compression and decompression). This leaves room of a computational error of about 1 DN during PCT. In reality, the error magnitudes resulting from PCT compression are of the order of 1-2 DN (see Fig. 4(b)) which is numerically near lossless. Fig. 3(b) shows the band averaged SNRs as a function of CR. As we increase the CR the effect of intrinsic data dimensionality becomes clear. While the SNRs are close together at CR = 4, they are quite different when CR = 25. Here, the image with the least spatial variability (L) has the highest SNR whereas the one with highest variability (H) suffers maximum loss.

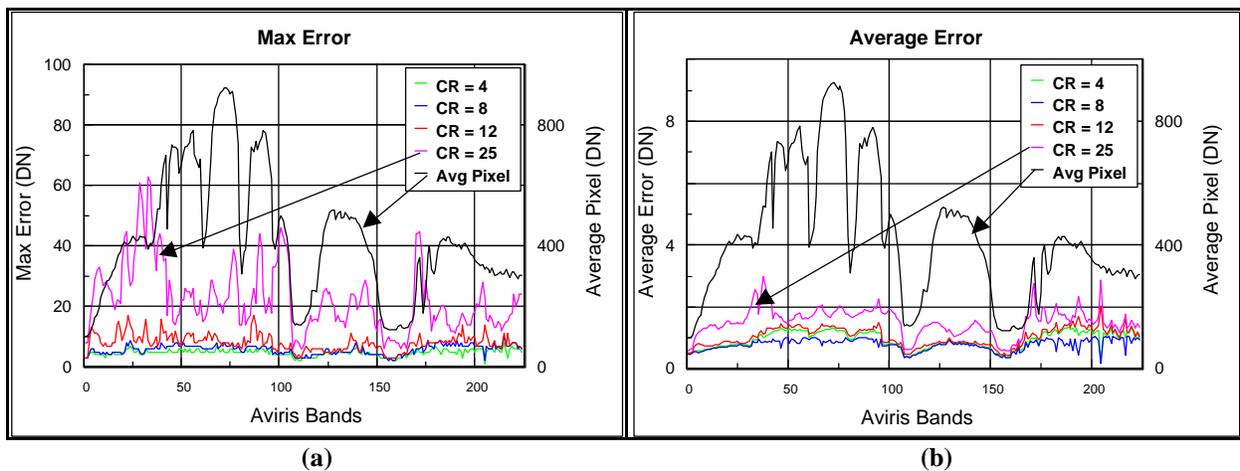


Figure 4: Band by band (a) maximum error and (b) average error, for medium variability image M as a function of compression ratio. Top curves shows the mean spectrum as a reference (second y-axis on right).

Fig. 4 shows the maximum and average errors that occur in each band as a function of CR, in the medium variability image M. Note that the mean error is less than 2 counts in nearly all bands even at the highest CR. In comparison the mean pixel spectrum is several 100's of counts. The error standard deviation curves (not shown) have magnitudes approximately equal to the mean errors. All these metrics together furnish a global picture of the kind of data fidelity that can be expected by applying PCT on AVIRIS data. The mean image spectrum in Figs. 4(a) and 4(b) shows sharp absorption features from atmospheric constituents like O_2 , H_2O and CO_2 . After radiometric calibration, data from these absorption bands are used to retrieve the atmospheric parameters, and to subsequently convert the calibrated radiance data cube into a reflectance cube. It is worth noting that the PCT errors in these bands are not necessarily higher than in others. This observation is significant for atmospheric retrieval.

Global metrics involving the SCR ratio ρ and error maps are illustrated in Fig. 5 using image L as an example. Fig. 5(a)-(c) show image L, the normalized error map, and the SCR ratio, all at CR=12. Fig. 5(a) also shows a boxed area indicating outlier pixels. The normalized error (ϵ) for a pixel is obtained by dividing the band averaged absolute pixel error, by the band averaged signal (Eq. (9)). The bright portions of Figs. 5(b) and 5(c) indicate higher values. Note that the boxed outlier region appears as bright spots in both Figs. 5(b) and 5(c). Fig. 5(b) also manifests the roads as regions of higher error whereas they do not appear in the SCR ratio map.

$$\epsilon = \frac{\sum_{i=1}^K |(x_i - \tilde{x}_i)|}{\sum_{i=1}^K x_i} \quad (9)$$

Fig. 5(d) shows the mean spectrum of image L together with an outlier pixel from the boxed area. The same pixel spectrum reconstructed after PCT (at CR=12) lies on top of the outlier, and the two are visually indistinguishable. The figure also shows the difference between the two spectra plotted against a second y-axis. The error magnitude is < 2 counts while the signal values are in the 150-200 DN range. Figs. 5(e) and 5(f) show horizontal profiles of the normalized error map (Fig. 5(b)) and the SCR ratio map (Fig. 5(c)). These profiles run

through the outlier pixel area and appear as spikes. The dual spikes in the Fig. 5(e) correspond to the outliers in the boxed region and the road that appears nearby. The same outlier appears as a spike in the SCR ratio profile of Fig. 5(f) indicating that this pixel is now more prominent as an outlier after PCT. In general, the values of ρ remain very close to 1 thereby indicating that the pixel properties have changed very little relative to their background (average $\rho = 0.999990$, and standard deviation = 0.001049 for image L at CR=12). The robustness of the SCR ratio even at high CR may indicate the viability of PCT based compression in applications involving subpixel analysis.

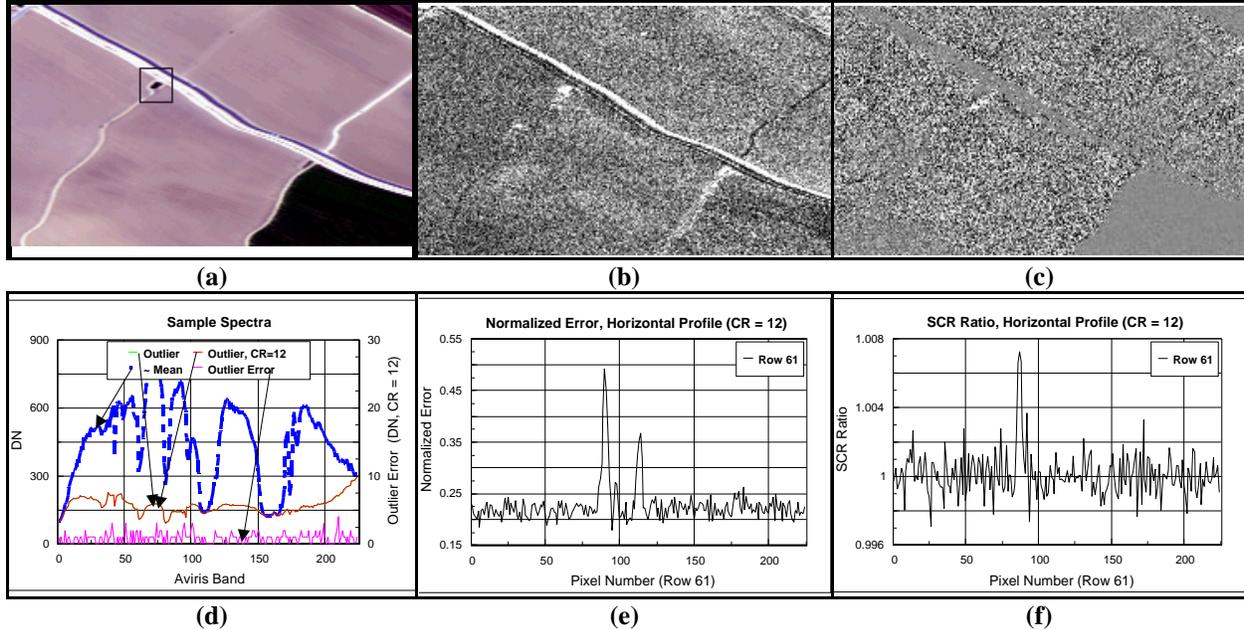


Figure 5: (a) Image L with box enclosure showing outlier pixels, (b) Normalized error map for CR=12, where bright regions indicate larger errors, (c) SCR ratio ρ for CR = 12, outlier pixels appear bright, (d) Mean spectrum for image L together with a single outlier pixel spectrum from the boxed region. The PCT processed spectrum for the same outlier lies over the original spectrum, and is within the figure’s visual resolution. The error between them is plotted on the bottom against a 2nd Y-axis. (e) Horizontal profile of normalized error map in 5(b). The two spikes correspond to an outlier in the boxed area shown in 5(a) and the road nearby, (f) Horizontal profile of SCR ratio in 5(c). The spike corresponds to the outlier.

3.3. Effects of Spectral Partitioning

Spectral partitioning of data is desirable from a computational standpoint. From Table I, we note the $O(K^2)$ and $O(K^3)$ complexities of the covariance matrix calculation and its subsequent diagonalization. Hence, reducing the number of bands (K) via partitioning can therefore reduce the computational burden significantly. The best method for selecting breakpoints to spectrally partition data is by examine the correlation matrix of the unpartitioned image. The natural partition points occur when the correlation matrix values change abruptly. This is shown in Fig. 6 using the correlation matrices from images H and L. The bright regions have mean correlation values > 0.9 and the dark regions have values that are near zero or negative.

We use Fig. 6(a) to examine the effects of spectral partitioning on the SNR and compression ratio. Data from scenes H and M are spectrally partitioned into two sets comprising of 112 bands each, as shown by the white squares in Fig. 6(a). Recall that M was considered as having moderate variability, and H that contains a mixture of rural and urban settings, was considered a high variability scene. Fig. 7(a) shows the band averaged SNRs plotted vs. CR for both spectral partitions of the above images. The first partition corresponds to the VNIR bands, and the second to the SWIR bands. Fig. 7(b) shows the normalized eigenvalues after PCT in image H. The eigenvalues for the SWIR bands reduce in magnitude faster than the VNIR bands, indicating a lower dimensionality of the SWIR data. The results for M are similar.

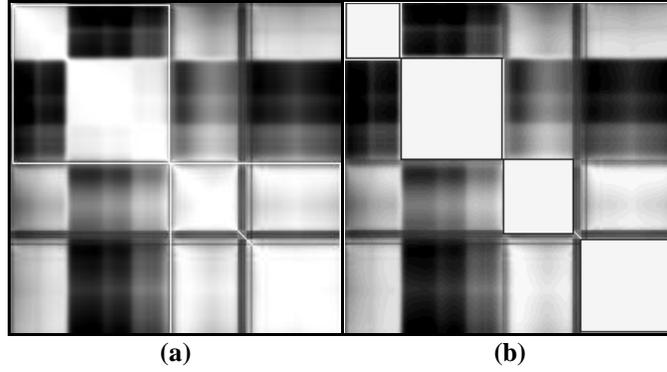


Figure 6: Spectral partitioning based on correlation matrices. Bright regions indicate correlation values > 0.9 , (a) Image H, white squares indicate 2×112 block partitions, and (b) Image L, black squares indicate 4 variable size partitions of $\sim 40, 70, 50,$ and 60 bands each.

At low CR, the VNIR bands exhibit a higher SNR than the SWIR bands. This is because the natural SNR of the sensor data are high in the VNIR bands due to the greater solar radiation present in these bands, as compared with the SWIR bands. The data dimensionality problem does not manifest at low CR because the large number of PCs used here sufficiently captures all the signal. However, the SNR curves for H are always below those of M because H has greater variability. As CR increases, we notice a crossover between the VNIR and SWIR snr values, and the latter are now larger. This is because the VNIR and SWIR bands themselves have considerably different dimensionality or scene variability. The SWIR spectrum is dominated by the tapering solar blackbody radiation curve strongly modulated by H_2O and CO_2 absorption bands. Most natural materials also have very slowly varying reflectance spectra in these bands. All these factors together contribute to making the dimensionality of the SWIR bands very low. Hence the SWIR bands can withstand higher CRs than the higher dimension VNIR bands.

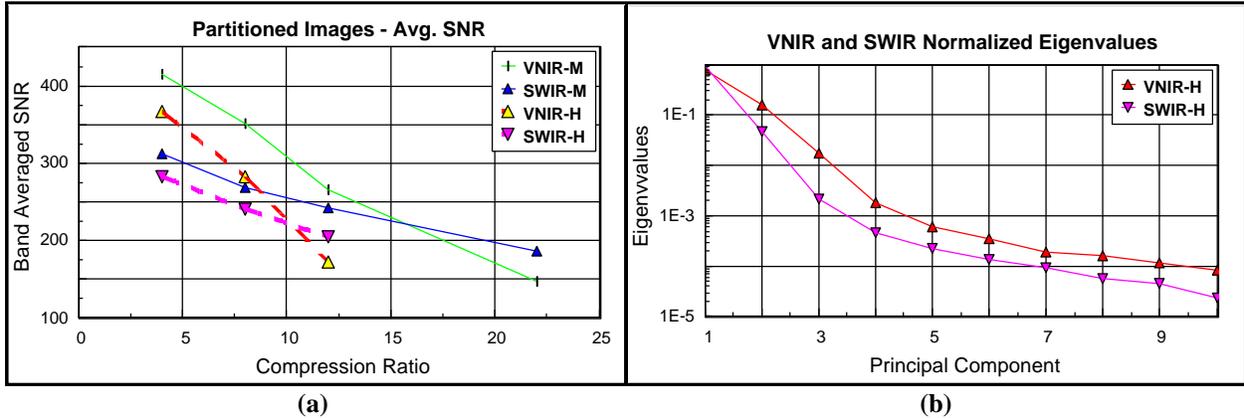


Figure 7: (a) Average SNRs vs. CR for scenes M and H, and (b) Normalized eigenvalues for VNIR and SWIR bands of image H.

4. SUMMARY AND CONCLUSIONS

This paper analyzed properties of raw AVIRIS hyperspectral data from a data compression perspective, using the PCT as a reference method. The algorithm's SPMD nature was effectively exploited to implement a real-time multiprocessor data compression system, using off the shelf components. A number of benchmark criteria error visualization methods were developed to analyze data fidelity. It was shown that data dimensionality that is related to scene content/variability played a critical role in PCT based compression. Spectral partitioning revealed that higher SNRs were obtainable in the SWIR than the VIS-NIR bands, at high compression ratios. Even under high CRs (>10), the SCRs appear to remain robust throughout the image. Subtle spectral features of distant outlier pixels were also seen to be well preserved at high CRs. The implication for spectral library-based detection algorithms is that PCT could be a reliable hyperspectral compression technique, in subpixel analysis type applications. This bears further investigation.

Overall, our studies indicate that PCT holds excellent promise as a hyperspectral data compression technique. The success of PCT using raw AVIRIS data, is in part due to the high quality of the sensor data. SNR values of several 100s were obtained for compression ratios in the range of 4-25. These numbers can be further increased by incorporating standard entropy coding techniques such as Huffman or Arithmetic coding, as a part of the compression algorithm. We note that the above SNRs values correspond to rms errors in the range of 0.1% - 0.5%. Ultimately the extent of “acceptable” error which in turn determines how much compression can be achieved, depends on the application. However, given that the accuracy of radiometric calibration is of the order of 1%-2%, lossy hyperspectral data compression using a technique such as PCT may be viable even in scientific data collections.

ACKNOWLEDGMENTS

This work was supported under the DOD SBIR Phase-II contract #F33615-99-C-1408. We thank Dr. Robert Green of NASA-JPL and the AVIRIS Data Distribution Center for providing raw AVIRIS data.

REFERENCES

- ADSP-2106X SHARC Users Manual, 2nd Ed., 1997, Analog Devices.
- AVIRIS raw data low altitude flight from 1998 over San Joaquin Valley, 1998, CA (flight number f981006t01p01_r02). Actual scenes are subsets from image cubes SC-03, SC-04, and SC-06 of this flight.
- Chen, J.Y., Reed, I.S., 1987, “A detection algorithm for Optical Targets in Clutter”, IEEE Trans. Aerospace Electronic Systems, Vol. AES-23, No. 1.
- Fung, T., LeDrew, E., 1987, “Applications of Principal Components Transformation to Change Detection”, Photogrammetry Rem. Sensing, Vol. 53, pp 1649, 1987.
- Golub, G.H., Van Loan, 1996, C.F., Matrix Computations, 3rd ed., Johns Hopkins University Press.
- Horn, R, Johnson, C.R., 1985, Matrix Analysis, Cambridge University Press.
- Huxtable, B.D., Jackson, C.R., Skaron, S.A., 1998, “SAR Processing using SHARC Signal Processing Systems”, SPIE Vol. 3371, pp. 226.
- Jia, X., Richards, J.A., 1999, “Segmented Principal Components Transformation for Efficient Hyperspectral Remote Sensing Image Display and Classification”, IEEE Trans Geosci. Rem. Sensing, Vol. 37, pp 538.
- Jolliffe, I.T., 1986, Principal Component Analysis, Springer Verlag New York.
- Mansfield, A.W., Rogers, G.W., Rais, H., 1998, “Real-time SAR processing for Search and Rescue”, SPIE Vol., 3371, pp. 233.
- Parsons, T., 1987, Voice and Speech Processing, McGraw-Hill.
- Richards, J.A., Jia, X., 1999, Remote Sensing and Digital Analysis, Springer Verlag, 3rd ed.
- Schowengerdt, R.A., 1983, Techniques for Image Proc. and Classification in Remote Sensing, Academic Press.
- Stocker, A., Reed, I.S., Yu, J., 1990, “Multi-dimensional Signal Processing for Electro-optical Target Detection”, SPIE Vol. 1305, pp 218.
- Therrein, C.W., Discrete Time Signals and Statistical Signal Processing, Prentice Hall, (1992).
- Wang, J., Zhang, K., 1995, “Spectral and Spatial Decorrelation of Landsat-TM Data for Lossless Compression”, Vol. 33, pp 1277.