


Practical Transform Coding of Multispectral Imagery



The Stock Market/Lightscaapes

JOHN A. SAGHRI, ANDREW G. TESCHER,
and JOHN T. REAGAN

Future land remote sensing satellite systems will likely be constrained in terms of downlink communication bandwidth. To alleviate this limitation, the data must be compressed. Images obtained from satellite and airborne multispectral collection platforms exhibit a high degree of spatial and spectral correlations that must be properly exploited in any multispectral bandwidth compression scheme. While spatial correlation is readily exploited by various im-

age compression techniques [1, 2, 3, 4, 5], relatively little attention has been given to spectral correlation across bands. In addition, the compression technology must provide a range of fidelity options for the reconstructed data tailored to the end use. This use includes quick-look and browse, which exploits human visual perception deficiencies and offers the opportunity for high compression ratios; machine exploitation, where visually-lossless coding is appropriate; and pre-

cise radiometric calculations, which require certain bands to be coded at near-lossless quality.

Recent work in the area of multispectral bandwidth compression can be categorized into three groups: (1) Three-dimensional transform-based techniques [6, 7, 8, 9]; (2) Vector quantization-based schemes [10], and (3) predictive techniques [11, 12]. A comparative performance evaluation of various multispectral compression techniques is given in [3].

In general, transform based techniques outperform the classical predictive methods. Assessment of the effects of compression on multispectral imagery is discussed in references [14] and [15]. In general, evaluation of the compression degradation should be based on a combination of statistical, visual, and machine based descriptors. Experimental results indicate that compression at moderate levels, e.g., 10:1 CR, have little or no impact on the results of statistical, visual, and machine-based exploitation.

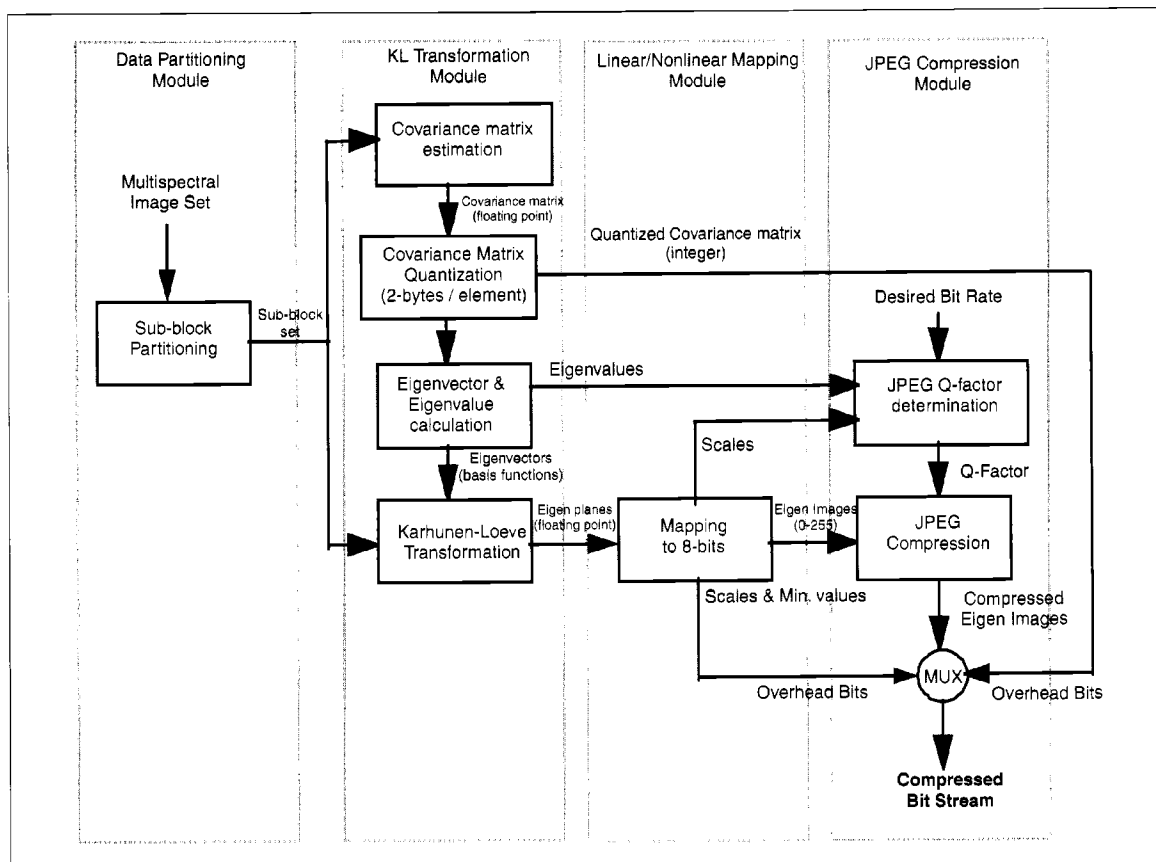
In this article, we present a robust implementable compression algorithm for multispectral imagery with a selectable quality level within the near-lossless to visually lossy range. The three-dimensional terrain-adaptive transform-based algorithm involves a one dimensional Karhunen-Loeve transform followed by two-dimensional discrete cosine transform. The images are spectrally decorrelated via the KLT to produce the eigen images. The resulting spectrally-decorrelated eigen

images are then compressed using the JPEG algorithm. The key feature of this approach is that it incorporates the best methods available to fully exploit the spectral and spatial correlation in the data. KLT is theoretically the optimum method to spectrally decorrelate the data. The standard DCT-based JPEG image compression algorithm is considered to be the most viable practical technique available today.

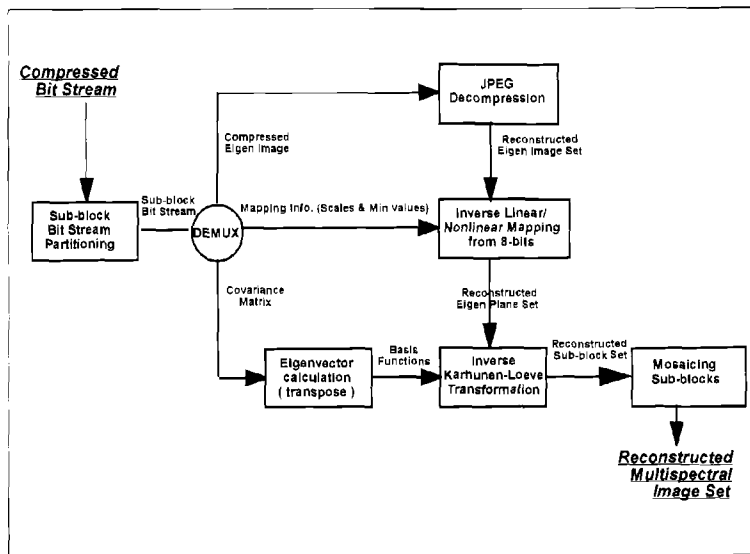
The novelty of this technique lies in its unique capability to adaptively vary the characteristics of the spectral decorrelation transformation based upon variations in the local terrain. The algorithm is conveniently parameterized to accommodate reconstructed image fidelities. For a set of 11 spectral bands, the reconstructed image fidelity ranges from near-lossless at about 5:1 CR to visually lossy beginning at around 40:1 CR. The compression ratio increases with the number of bands included in the multispectral set.

The spectral and spatial modularity of the algorithm architecture allows the replacement of the JPEG module by a different coder (e.g., DPCM). However, the significant practical advantage of the discussed approach is that it is leveraged on the standard and highly developed JPEG compression technology.

The required architecture can be implemented via recently developed multimedia chipsets [17, 18]. The obvious advantage of this system design is that it is highly leveraged on



1. Terrain-adaptive compression block diagram.



2. Decoder block diagram.

state-of-the-art VLSI components. As such, the hardware realization of this algorithm becomes viable, cost effective, and low-risk. In addition, the use of programmable chipsets makes the design flexible, allowing up-to-date refinements and modification to the algorithm to be made during and/or after the hardware implementation stage.

Compression System Overview

Figure 1 depicts the block diagram of the compression system. It consists of four modules: (1) Data partitioning; (2) Karhunen-Loeve transformation; (3) mapping eigen planes to 8-bit eigen images; and (4) JPEG compression of the eigen images. In the data partitioning module, the set of multispectral images are partitioned into sets of non-overlapping images: sub-block sets, which are sequentially fed to the KLT transformation module for spectral decorrelation. In the KLT transformation module the multispectral sub-block set is spectrally decorrelated to produce a set of eigen planes. The basis functions for the KLT transformation are the eigenvectors of the cross-covariance matrix associated with the multispectral sub-block set. The covariance matrix is estimated first and then quantized to two bytes per element.


The eigen planes are formed by matrix multiplication of the sub-block set and the basis functions [6]. The eigen planes are in floating point format and assume both positive and negative values. In the next module, the eigen plane set is converted into the 8-bit eigen images set via linear/nonlinear mapping of each plane into the 0-255 range. The spectrally-decorrelated eigen images are then compressed in the next module using the JPEG algorithm. The quantized covari-

ance matrix and mapping information are transmitted along with the compressed bit stream as overhead information. This three-dimensional transform-based compression algorithm efficiently exploits the spectral and spatial correlations in the data. Another significant and practical advantage is that it is leveraged on the highly developed JPEG compression technology. The algorithm adopts well to the local terrain variation since the covariance matrix, from which the transformation basis functions are derived, is updated very frequently over each small sub-block set of multispectral data. The bit requirement for the sub-block covariance matrix and mapping information is negligible and, as such, there is no need to resort to a stored covariance look-up-table to minimize the overhead bit information.

The block diagram for the decoder is shown in Figure 2. The mapping information and covariance matrix are extracted from the received bit stream. The eigen images are reconstructed from their compressed bit stream using the JPEG decoder. The mapping information is used to span the 8-bit dynamic range of the decoded eigen images back to their original range. The inverse KLT transformation basis functions are obtained from the extracted covariance matrix via transposing its eigenvectors. An inverse KLT is performed to produce the reconstructed multispectral sub-block set. The sub-block sets are mosaiced together to form the reconstruction of the multispectral image set.

Spectral Decorrelation

Removing the inherent spectral correlation in the data results in a significant compaction of data to be coded. This can be




3. Removing spectral correlation via a Karhunen-Loeve transformation.

done via the optimum Karhunen-Loeve (KLT) transformation. Figure 3 illustrates the KLT or orthogonal transformation. For the N spectral images to be decorrelated, we form an N -by- N KLT transformation matrix composed of the N ordered eigenvectors of the associated cross-covariance matrix [6]. Each vector of the spectrally-correlated components from identical locations in each band is multiplied by the KLT transformation matrix to form an output vector of spectrally-decorrelated components. The output vectors are placed adjacent to one another, in the same order as the input vectors, to form the stack of the spectrally decorrelated eigen planes. The decorrelation property of the transform images, in this sense, has a diagonal covariance matrix only.

Test Data Sets for Spectral Decorrelation Experiments

Figure 4 shows band 10 of two 512 x 512 test image sets selected for spectral decorrelation experiments. The sets are from the M7 sensor platform. The M7 sensor is a multispectral scanner developed by the Environmental Research Institute of Michigan (ERIM) for use on an airborne platform. It covers the visible through infrared region in sixteen unequal bands. The ASAS (Advanced Solid-state Array Spectroradiometer) imagery consists of six spectral bands, a 512 by 1024 image having 12 bit dynamic resolution. The typical ground sampling distance (GSD) is 4-6 meters.

In this experiment, sixteen bands in the spectral range of 0.36 to 12.11 micron are included in each set. These test images are selected because they contain a diverse range of natural and urban terrain and as such are very challenging for spectral decorrelation and subsequent coding experiments. The correlation coefficient is a convenient and useful method to measure the inherent spectral correlation. The correlation coefficient matrix is defined as the normalized covariance matrix. That is, the coefficient for each pair of bands is equal




4. Sample test images.


to their covariance value divided by the square root of the product of their individual variances. Figure 5 shows the correlation coefficient matrix (only the non-redundant half shown) and its three dimensional surface rendition for the Airfield test image.

Spectral Decorrelation Efficiency

As discussed earlier, the KLT is theoretically the optimum method to spectrally decorrelate a set of multispectral images. Figure 6 shows the first nine spectrally-decorrelated eigen planes associated with the 16-band Airfield test image set. For display purposes, the images have been linearly mapped from their original dynamic range to the 0-255 (8 bits) range. The



5. Correlation coefficient matrix.



6. Eigen images of the test image set (first 9 of the total 16).


KLT performed in this experiment is non-adaptive. That is, the parameters of the transformation, the KLT basis functions, are fixed for the entire image set. The compaction of the data as the result of KLT operation is clearly evident as more than 80 percent of the energy, or information content of the data set resides in the first two or three eigen planes. The remaining eigen planes have very little information content and as such require substantially fewer bits to be coded. A convenient approach to measure the amount of data compaction is obtained via the assessment of the eigenvalues of the cross covariance matrix. The eigenvalues of the covariance matrix correspond to the variances of the eigen planes. The variance of an image reflect its busyness or, excluding noise, information content.

Figure 7 shows the ordered variances of the eigen planes of the Airfield test image set. The variance drops almost exponentially with the order of the eigen planes. The steeper the drop in the variance in going from low to high order eigen planes is, the more efficient compaction, spectral decorrelation, is achieved.


KLT vs. DCT for Spectral Decorrelation

As an alternative, the one-dimensional DCT can substitute for the optimum KLT for spectral decorrelation. The advantage

of using DCT is that the basis functions are fixed regardless of the characteristics of the data. Unlike the KLT, the DCT does not require covariance and eigenvectors calculations and, as such, is much simpler to implement. The drawback is that the spectral decorrelation efficiency will be significantly lower. Figure 8 shows the comparison of the spectral decorrelation efficiency of the DCT and KLT for the Airfield test image. The variance of each of the 16 DCT coefficients is larger than the variance of its corresponding eigen image by several orders of



7. Ordered variances of the test eigen images.



8. KLT versus DCT for spectral decorrelation.

magnitude. Since the bit rate required to code an image increases with the value of its variance, the KLT approach will result in a substantially lower bit rate requirement.

Decorrelation Efficiency vs. Number of Bands

The spectral decorrelation efficiency (compactness) increases with the number of bands in the data set. A greater compaction of data results in a lower bit rate requirement. Figure 9 shows the plots of the variances of the eigen images associated with Airfield test image sets consisting of 16, 12, 8, and 4 bands, respectively. The area under the variance curve reflects the amount of information to be coded. This area is substantially greater for the 4-band set than the 16-band set.

As the number of bands in the set decreases, the performance of the algorithm tends to be more influenced by the power of the standard DCT-based JPEG algorithm.

Quantization of the Spectrally-Decorrelated Eigen Planes

The spectrally decorrelated eigen planes are represented in floating point real numbers. They must be quantized to 8-bit, 10-bit, or 12-bit, to be subsequently coded by the DCT-based JPEG compression algorithm.

Linear Quantization for 8-bit Multispectral Imagery

For 8-bit multispectral imagery, all eigen planes are linearly quantized into 8-bit images: they are the eigen images. Alternatively, the high-dynamic range, low-order, eigen planes may be quantized into 10 or 12-bit images. However, the experimental results do not indicate any appreciable improvement in results using the dual quantizer. For the 10/8-bit quantizer, only the first two eigen planes whose dynamic ranges are greater than 256 are quantized with a 10-bit quantizer. In order to isolate the quan-

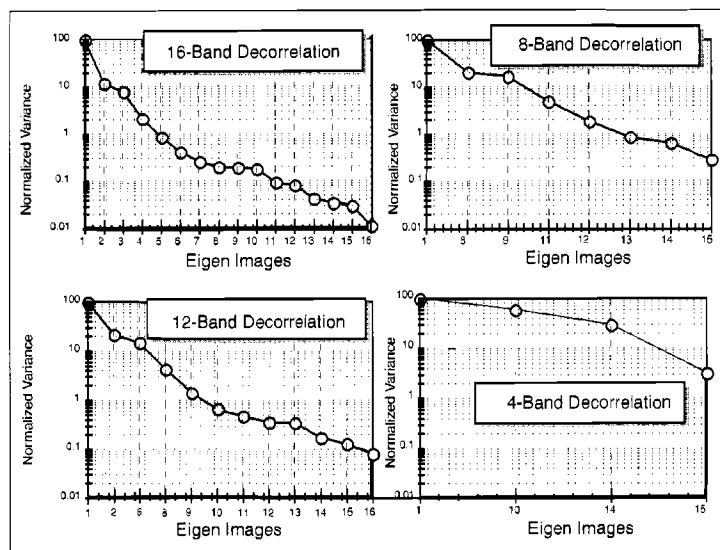
tization round-off error, the resulting quantized eigen images were losslessly coded. For the 10/8-bit quantization scheme, the average mean squared error (MSE) decreases from 0.36 to 0.28. However, the maximum induced error remains at one count in 256 for either of the two quantization schemes. Compared to the nominal 8-bit quantization of the eigen images, the dual 10/8-bit quantizer results in approximately 5 percent increase in bit rate.

We may conclude that it suffices to use an 8-bit quantizer for eigen images of 8-bit data. The dual 10/8-bit quantizer does not offer better performance. The reasons are (1) the maximum induced error is only 1 count in 256; (2) the maximum induced error still remains at 1 count in 256 using the dual 10/8-bit quantizer; (3) the quantization round-off error is far smaller than the error that is incurred in the subsequent JPEG coding of the eigen images; (4) the advantage of the lower induced MSE for the dual 10/8-bit quantizer is offset by the 5 percent increase in the bit rate requirement and the additional implementation complexity.

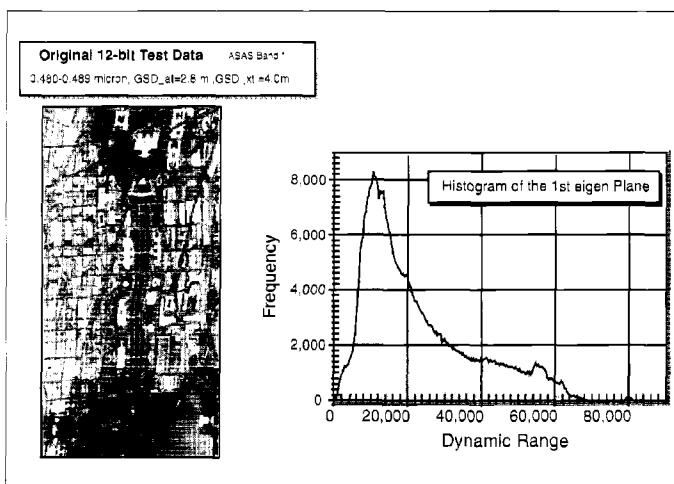
Nonlinear/Linear Quantization for High Dynamic Range

For high dynamic range (12-bit) imagery, an optimized nonlinear 8-bit or 12-bit quantizer may be used to quantize the first one or two high-dynamic range eigen planes to 8-bit eigen images. The remaining lower-dynamic range eigen images may be quantized using a linear 8-bit quantizer. The high-dynamic range of the first few eigen planes typically have an uneven (unequalized) histogram with long narrow tails. A nonlinear quantizer achieves a substantially lower overall quantization error by allocating finer quantizer steps to the large central portion of the dynamic range. Figure 10 shows band 1 of a 12-bit ASAS multispectral test data set, and the histogram of its first eigen plane.


A parametric nonlinear 8-bit quantizer that attempts to



9. Decorrelation efficiency versus the number of bands in the data set.



10. Skewed histogram of the first eigen plane of 12-bit test data.



11. Nonlinear quantization of eigen planes.

achieve the optimum mapping of the dynamic range to the 0-255 range was selected. This nonlinear quantizer is defined as,

$$Y = 255 / (MAX^{exp}) X^{exp} \quad (1)$$

where: X is the eigen plane value (floating point), Y is the eigen image value (an integer from 0 to 255), exp is the nonlinearity index (0-1 range); and MAX is the maximum value of the eigen plane.

Figure 11 shows the plot of this nonlinear quantizer for different values of the nonlinearity index, exp . This process results in a varying (nonuniform) quantization step width that is optimized, for minimum mean-square-quantization error, for each segment of the input image's dynamic range.

The above nonlinear quantizer was applied to the first two eigen planes of the 12-bit ASAS test image set. The selected values for the nonlinearity index, exp , were determined experimentally: 0.7 and 0.8 for the first and second eigen

images, respectively. The remaining eigen planes were linearly quantized as before. The maximum and the mean square quantization errors were measured to be 6, and 1.4 counts, in 4096. The relatively small quantization error indicates that an optimized nonlinear/linear 8-bit quantizer suffices to quantize the eigen planes associated with high-dynamic range data. The advantage of a 10 or 12-bit quantizer, in terms of achieving a lower quantization error, is offset by the higher bit rate requirement to code the resulting 10, or 12-bit eigen images.

Coding the Eigen Images

Two issues are related to coding the eigen images: (1) modifying the JPEG parameters to suit the characteristics of the eigen images; and (2) arriving at an optimum coding bit allocation among the eigen images.

Modifying JPEG for Coding Eigen Images

The standard JPEG compression algorithm includes various parameters which may be modified to suit the specific characteristics of the image. Two such parameters are the quantization scale table for the DCT coefficients (Q table), and the internal Huffman table used to entropy code the quantized DCT coefficients. In this phase of the study, the Q table has been modified to match the spatial characteristics of the eigen images.

Bit Rate Assignment For Eigen Images

Compression performance is strongly influenced by the selected bit assignment scheme for the spectrally-decorrelated eigen images. Since the variances of the ordered eigen planes decrease almost exponentially, it may be suggested to code the eigen images at rates proportional to the logarithm or square roots of their variances. Although this bit assignment strategy may result in the lowest MSE, it is not suitable for multispectral images. In multispectral imagery, the subtle variations in the spectral signatures of certain terrains manifest themselves in the lower eigen images. Therefore, in order to preserve the spectral fidelity of the data, it is imperative to code the lower eigen images with the same level of accuracy as the others, irrespective of their variances or the resulting overall mean square coding error. In general, the criterion for bit assignment should be to induce a uniform overall coding error on all the reconstructed eigen planes.

The overall coding error on the reconstructed eigen planes originate from two sources: (1) the quantization error incurred in linear mapping of the eigen planes into 8-bit eigen images;

and (2) the error induced by the JPEG compression of the 8-bit eigen images. Thus, the bit rates should be adjusted so that the resulting mean square error in the reconstructed eigen planes are approximately the same. The selected approach should be a compromise between performance and implementation complexity. Again, these results are derived from an extensive set of experimentation, and cannot be derived from theoretical principles.

To implement the optimum bit assignment scheme, an arbitrary MSE for the first eigen image is selected. The effective MSE for the remaining eigen images are calculated as:

$$MSE(n) = L(1) \times B(n)/B(1) \quad (2)$$

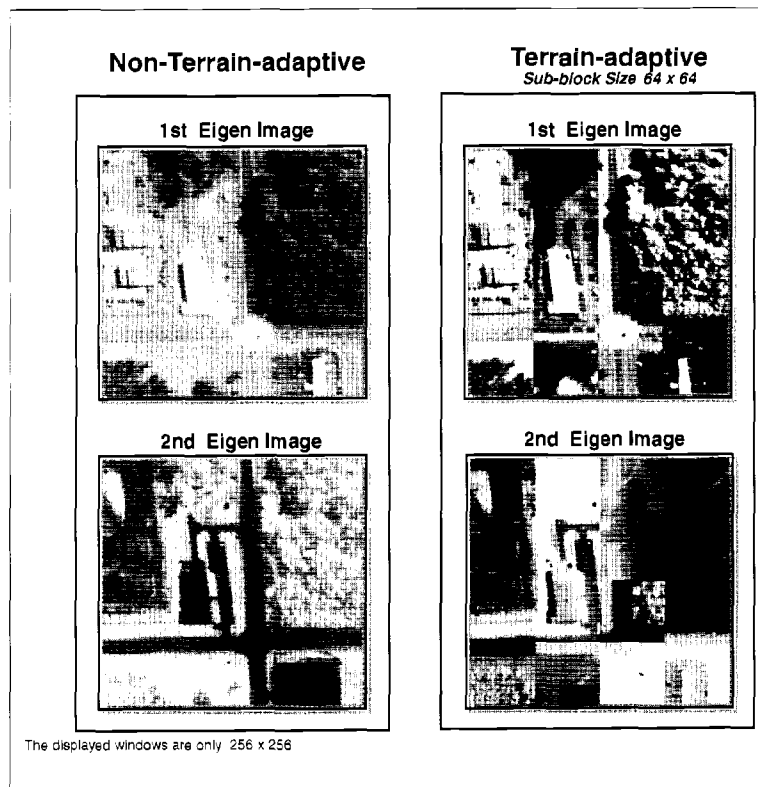
where n (2, 3, ...) is the order of the eigen image, $L(n)$ = MSE for eigen image n , $B(n)$ is the bin width to quantize eigen plane n .

The JPEG coding bit rate is changed via the parameter Q , the quality factor, which affects the quantization of the DCT coefficients, and where Q ranges from 1 to 100. For each eigen image, we select the lowest JPEG quality factor that results in an effective MSE less than or equal to that found from the equation above. Because the dynamic range of the first few eigen planes typically exceeds $2^8 = 256$, the appropriate quantization steps are less than 1. As such, according to the above scheme, the first few eigen images are coded at a relatively higher

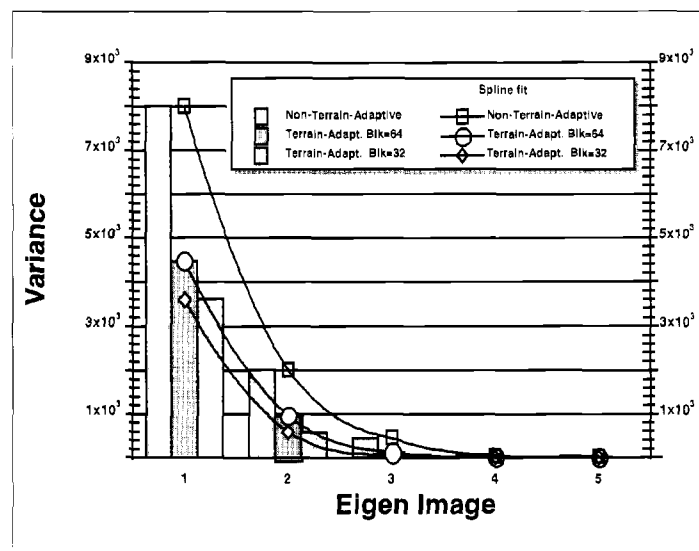
JPEG Q factor to compensate for the higher quantization error they incurred. The drawback with this bit assignment scheme is the complexity of implementation. Since it is not known in advance what MSE results for a particular eigen image at a given JPEG Q factor, it is necessary to construct a look-up-table of MSE versus Q for each eigen image. Thus, all eigen images must be coded at all possible Q factors. This approach is prohibitive for practical implementation.

Other possibilities can be utilized for eigen image bit allocation. A second method is similar to that described above in the "optimum scheme" except that some of the low-variance eigen images are replaced by their mean value. The advantages of this technique are (1) good performance in terms of MSE; and (2) considerable savings in computation and power requirements, since fewer number of eigen images need to be coded.

A third method uses the same JPEG quality factor Q to code all eigen images. This scheme does not adjust the coding bit rate to compensate for the error induced by the initial quantization. The advantage of this technique is that it is simple to implement. However, performance is sacrificed in terms of MSE and maximum error.



12. Eigen images of the terrain-adaptive and non-terrain-adaptive schemes.



13. Comparing the spectral decorrelation efficiency of the terrain-adaptive versus the non-terrain adaptive approaches.

A fourth method is a simplified and practical version of the optimum scheme. It is an empirical technique which is a compromise between complexity and performance. Here the JPEG Q factor is varied for each eigen image based upon its variance and the initial quantization error it incurred.

Terrain Adaptive Approach

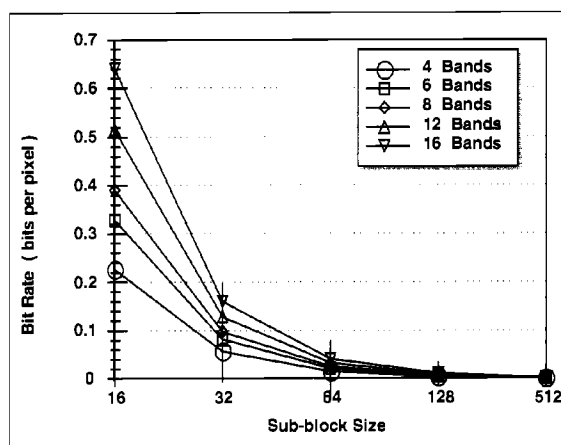
A typical multispectral image set obtained from either satellite or airborne collection platforms exhibits a number of different terrains, such as water, forest, cloud, ice, and desert. Each terrain has a unique spectral signature. Thus, to achieve the highest compactness of spectral information, the spectral transformation parameters must adopt to the local terrain characteristics. In the terrain-adaptive approach, the covariance matrix, from which the spectral transformation basis functions are derived, is updated frequently. The smaller the block-size over which the covariance is updated is, the more efficient is the spectral decorrelation process. The drawback with the selection of a small sub-block size is the resulting increase in the overhead bit rate due to an increase in the number of sub-blocks. Let's discuss the tradeoff between this "window" size and overhead bit rate.

Figure 12 shows the first and the second eigen images of the test image corresponding with the nonterrain-adaptive and the terrain-adaptive approaches for covariance update window sizes of 64 x 64 and 32 x 32, respectively. For the terrain-adaptive approach, the eigen images appear to have discontinuities over the edges of the selected covariance update window. This blocking effect clearly indicates the adaptation of the spectral transformation process to the characteristics of the terrain within the selected window. Notice the absence of blocking in the wooded section of the image where the texture is uniform. Figure 13 shows the relative sizes of the variance of each eigen image for the non-terrain-adaptive and the two terrain-adaptive approaches. The significant decrease in the variance as a result of terrain adaptation indicates a much greater compaction of the data and substantially lower performance.


Overhead Information

The terrain-adaptive approach entails transmission and processing some overhead information. The overhead information includes the approximated covariance matrix and the eigen plane quantization parameters for each non-overlapping sub-block of data over which the covariance matrix is updated. The covariance matrix is used to generate the KLT spectral transformation basis functions, the eigenvectors. Since the covariance matrix is diagonally symmetric, only the non-redundant half is transmitted. For each sub-block, the eigen plane quantization parameters are the scale factor (proportional to the inverse of the quantization step) and the minimum eigen plane value. The drawback with the selection of a small sub-block size, the covariance update window, is the resulting increase in the overhead bit rate due to an increase in the number of sub-blocks.

To reduce the computation burden for the terrain-adaptive approach, the calculation of the covariance matrix is not based on all the data points within the selected window size. Experimental results indicate that it suffices to use one data point per 8 x 8 section of the selected covariance update window to estimate the covariance matrix. A minor gain in performance



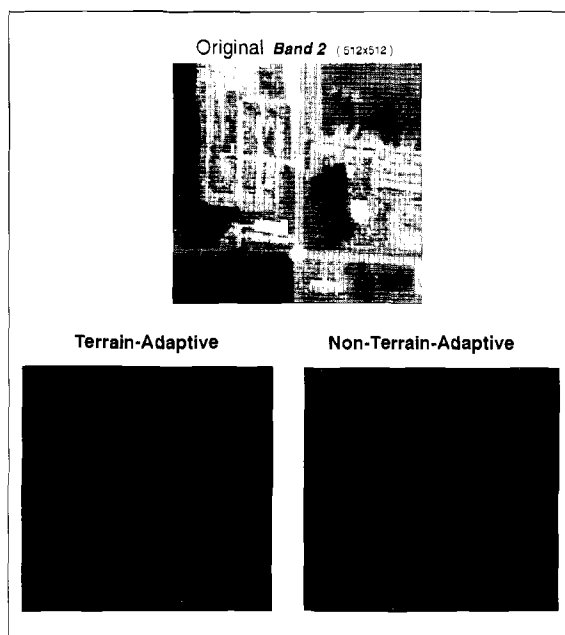
14. Overhead bit rate versus the number of bands and update window size.



15. Illustrative compression performance (military test image, 8 bpp, 11 bands, 5 m GSD).

results from the exact calculation of the covariance matrix based on all data points in the window. Further, to reduce the overhead bit rate, the floating-point covariance matrix is quantized to two bytes per element. The minimum eigen plane value and the scale factor used are quantized to 16 bits (2 bytes) and 8 bits (one byte), respectively.

Assume that a total of N bands are selected for simultaneous spectral decorrelation. Further, assume that the selected window size for covariance update is $B \times B$. The size of the covariance matrix is $N \times N$ and the number of non-redundant elements is $[(N^2/2) + N/2]$. The total bit requirement is $8(N^2 + N)$, where the number of minimum eigen plane values is N . There are 16 bits (two bytes) per minimum value. There are approximately two scale factors less than 1. Allocating 8 bits (one byte) per scale factor, the total bit requirement for eigen plane quantization adds up to $16(N + 1)$. Note that the




16. The original and difference images at 0.8 bpp for terrain-adaptive and non-terrain-adaptive schemes. Note: while the difference images indicate that the residual error is uniformly distributed (bright values represent error), the terrain-adaptive scheme yields a darker difference image, suggesting less degradation.

scales of unity, i.e., "1", are not transmitted. The overhead bit rate B_{oh} in bpp is thus approximately equal to:

$$B_{oh} = 8[(N+1)(N+2) / (NB^2)] \quad (3)$$

in bits per pixel. This equation suggests that the overhead bit rate increases approximately linearly with the number of spectral bands. The overhead bit rate is, however, inversely proportional to the square of the covariance update window size. Figure 14 shows the overhead bit rate for various number



17. Illustrative results from unsupervised and supervised classification.

of bands and sub-block sizes. The optimum block size for the terrain-adaptive approach is 64 x 64. For this choice of window size, the overhead bit rate is a negligible amount less than 0.04 bpp regardless of the number of bands in the set. Selection of a smaller window size results in a substantial increase in the overhead bit rate.

Look-up Table Approach to Reduce Overhead Bit Rate

For a smaller sub-block size, the overhead bit rate increases significantly. To alleviate this problem, we have to resort to a look-up-table approach. Here, a stored table of pre-calculated covariance matrices is searched to find the one that approximates the estimated covariance matrix of the current sub-block. This approach yields a substantial savings in the overhead bit rate since only the pointer to the table, instead of the entire covariance matrix itself, will be transmitted to the receiver.

The covariance LUT approach was not implemented in this study because the overhead bit rate was insignificant, less than 0.04 bpp, for block sizes of 64 x 64 or greater. The improvement in the spectral decorrelation efficiency as the result of smaller blocks was not significant enough to pursue implementation of the covariance LUT approach. The covariance LUT approach may, however, be useful for hardware implementation due to a potentially lighter computation requirement.

Experimental Results

The terrain-adaptive compression algorithm has been programmed for both, SUN and PC platforms. The algorithm is conveniently parameterized to accommodate various parameters including size of data, dynamic range, number of bands, desired bit rate and window size for terrain adaptation (covariance update window). The discussion in this section is mostly limited to the rate versus distortion results only. The complete experimental results including machine-based exploitation are discussed elsewhere [14, 15].

Rate vs. Distortion Results

Figure 15 shows the rate versus mean square and maximum error curves for the Military 1, Scene 1 test image. These results are representative of the performance of the compression system for a set of eleven multispectral images. This test image contains sharp edges which are difficult to compress with high fidelity at low rates. Thus, the rate versus distortion curve shown here is a conservative estimate of the algorithm performance. These results demonstrate, that on the average, each pixel value is off by 1 count in 256 at 0.8 bpp (10:1 compression ratio). Visual evaluation of the reconstructed imagery reveals that visually lossless performance is achieved in the 0.2 - 1.6 bpp range. (5:1 to 40:1 compression ratios). Above 1.6 bpp, the subtle difference between the original and the reconstructed images can only be observed using advanced machine based exploitation measures.

Figure 16 shows the difference images at 0.8 bpp for the Coast 2, Scene 2, and the Military 1, Scene 1 test images. The images show the difference between the original and the reconstructed images using the terrain-adaptive and the non-terrain-adaptive approaches. The difference images shown here have been linearly scaled up to the 0-255 range. The difference images for both approaches reflect a relatively uniform spatial distribution of the error. A homogeneous difference image is one of the important properties of any robust compression technique. The presence of any structure in the difference image reflects the loss of information in the compression process.

Classification

Figure 17 shows illustrative results from the unsupervised and supervised classification metrics [15]. The unsupervised classification yielded better than 95 percent correct classification, at 1.0 bpp and over 98 percent at 1.5 bpp. Supervised classification was more impacted by the compression, with over 87 percent of correct classification at 1.0 bpp and 90 percent at 1.5 bpp. The classification results for the DPCM-based approach [11] were between 5 to 20 percent worse, depending on the bit rate. Further details on statistical and machine-based metrics appear in [14] and [15]. Note that the terms supervised and unsupervised used therein refers to whether the number categories to be classified are known in advance.

Spectral Fidelity

In the bandwidth compression of multispectral imagery, preservation of the spectral resolution across bands is as essential as preserving the spatial resolution within each band. Based on the correlation coefficient matrix a useful tool for measuring the spectral fidelity of multispectral data was devised. The spectral fidelity is depicted by a convenient and very revealing 3-dimensional color surfaces plot of the correlation coefficient matrix. Any deviation from the original correlation coefficient matrix indicates the loss in spectral fidelity due to the compression process. The 3-D surface plot of the deviation from the original correlation matrix highlights the loss in spectral fidelity as the result of compression.

Experimental results indicate that the loss of spectral fidelity, as measured by the deviation from the original correlation coefficient matrix, is very insignificant, regardless of the image and the coding bit rate. Figure 18 shows the original scene and the loss of spectral fidelity for the Military 1, Scene 1 test image at compression ratios of 4:1, 8:1, and 16:1. Note that the maximum of the color bar scale is at 0.003, or 0.3 of one percent. At 4:1 CR, no loss of the spectral fidelity is observed. At 8:1 CR, a small loss of the spectral fidelity between bands 7 and bands 10 and 11. At 16:1 CR additional loss of the spectral fidelity is noticed between band 4 and bands 10 and 11.

Handling the Sharpening Band

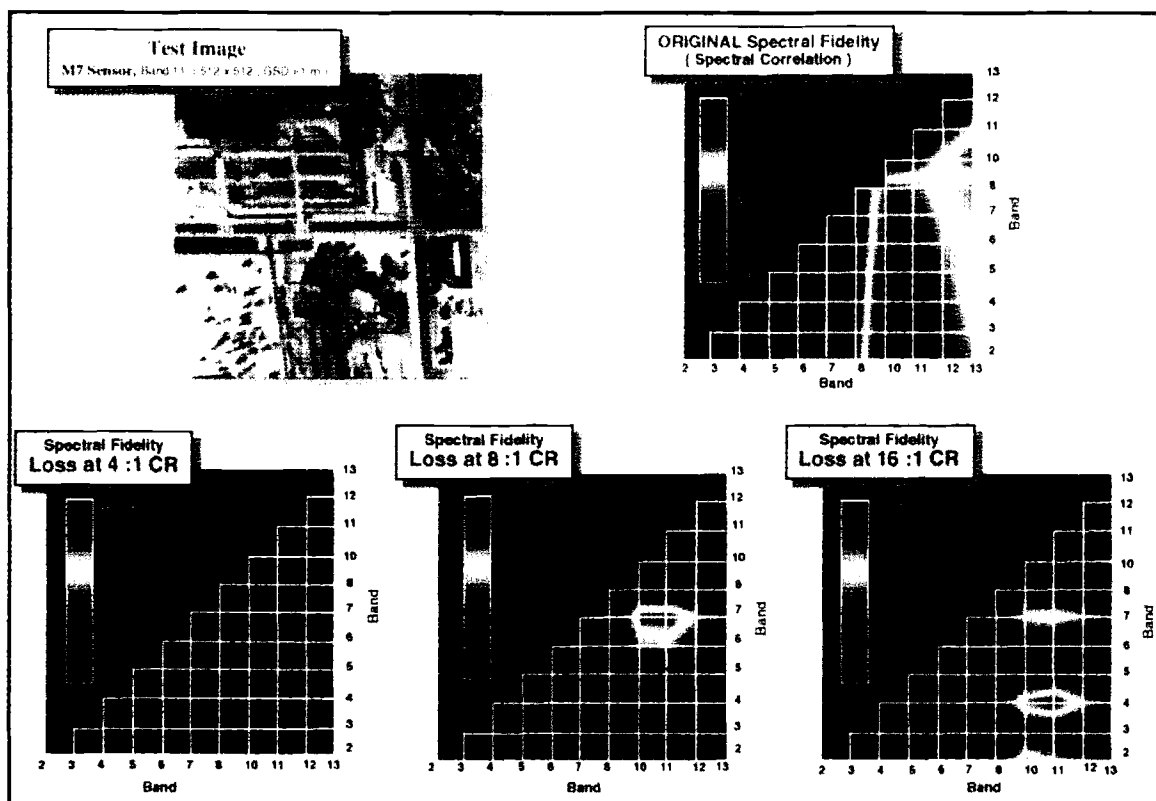
Typically, a panchromatic sharpening band of twice the resolution in both horizontal and vertical directions accompanies the multispectral imagery. We have developed a simple scheme to jointly compress the sharpening band along with the multispectral bands. The sharpening band is "fanned out" into four non-overlapping sub-sampled images equal in size to the multispectral images. That is, each of the four pixels in a 2 by 2 block of the original sharpening band forms one pixel in a subsampled image. The fanned out images are then regarded as additional bands in the compression process. This process ensures exploitation of the inherent correlation between the sharpening band and the multispectral bands. At the receiver, the reconstructed sub-sampled sharpening bands are "fanned in" into the reconstructed sharpening band.

Sensitivity Issues

Since the characteristics of future multispectral systems are yet to be determined, a simulation study was performed to assess the sensitivity of the compression system to various system parameters. The sensitivity issues considered were (1) impact of band misalignment; (2) dynamic range of data; (3) impact of calibration and pre-processing of the data, and (4) impact of dead/saturated pixels. Results of the experiments indicated that except for band misalignment, the performance of the compression system is not significantly impacted by the above parameters at moderate compression ratios (10:1). Band misalignment can be minimized prior to the spectral decorrelation process via some simple pre-processing. The brute-force technique is to reposition individual bands in order to maximize the inter-band correlation. This correlation maximization process need not to be carried out based on the entire image set. The correction magnitude for band positions may be derived from a small subset of the data (64 x 64 section).

Conclusions

We presented an implementable three dimensional terrain-adaptive transform based bandwidth compression technique for multispectral imagery. The algorithm exploits the inherent spectral and spatial correlations in the data. The compression technique is based on Karhunen-Loeve transformation for spectral decorrelation followed by the standard JPEG algorithm for coding the resulting spectrally decorrelated eigen images. The algorithm is conveniently parameterized to accommodate reconstructed image fidelities ranging from near-lossless at about 5:1 CR to visually lossy beginning at about 40:1 CR. The novelty of this technique lies in its unique capability to adaptively vary the characteristics of the spectral decorrelation transformation as a function of the variation of the local terrain. The spectral and spatial modularity of the algorithm architecture allows the JPEG to be replaced by alternate spatial coding procedure. The significant practical advantage of this proposed approach is that it is leveraged on



18. Spectral fidelity.

the standard and highly developed JPEG compression technology.

Acknowledgment

This research was supported in part by the LANDSAT Program Office, SAFSP, Office of the Secretary of the Air Force.

John A. Saghri, Andrew G. Tescher, and John T. Reagan are members of the technical staff at Lockheed Palo Alto Research Laboratories, Palo Alto, CA. Dr. Tescher may be reached at Lockheed Missiles & Space Co. Inc., Dept. 9780/251, 3251 Hannover St., Palo Alto, CA 94304.

References

1. T.J. Lynch, *Data Compression Techniques and Applications*, Wadsworth, Inc., Lifetime Learning Publications, Belmont, CA, 1985.
2. R.J. Clarke, *Transform Coding of Images*, Academic Press, Inc., 1990.
3. N.S. Jayant and P. Noll, *Digital Coding of Waveforms*, Prentice-Hall, Inc, 1984.
4. M. Rabbani and P.W. Jones, *Digital Image Compression Techniques*, SPIE Optical Engineering Press, Vol. TT7, 1991.
5. W.K. Pratt, *Digital Image Processing*, John Wiley & Sons, Inc., 2nd Edition, 1991.
6. J.A. Saghri and A.G. Tescher, "Near-Lossless Bandwidth Compression for Radiometric Data," *Optical Engineering*, Vol. 30, No. 7, Jul. 1991.
7. J.A. Saghri, A.G. Tescher, and J.T. Reagan, "Three-Dimensional Transform Coding of Multispectral Data," *Proc. of Twenty Seventh Asilomar Conf on Signals, Systems & Computer*, Nov. 1993.
8. J. A. Saghri, A. G. Tescher, and J. T. Reagan, "Terrain-Adaptive Multispectral Bandwidth Compression," *Proc. of the Industry Workshop of Data Compression Conference*, Snowbird Utah, April 1994.
9. B.R. Epstein, R. Hingorani, J.M. Shapiro, and M. Czigler, "Multispectral KLT-Wavelet Data Compression for Landsat Thematic Mapper Images," *Proceeding of Data Compression Conference*, Mar. 1992.
10. S. Gupta and A. Gersho, "Feature Predictive Vector Quantization of Multispectral Images," *IEEE Trans. on Geoscience and Remote Sensing*, Vol. 30, No. 3, May 1992.
11. B. Brower, B. Gandhi, D. Couwenhoven and C. Smith, "ADPCM for Advanced LANDSAT Downlink Applications," *Proc. of Twenty Seventh Asilomar Conf on Signals, Systems & Computers*, Nov. 1993.
12. A. K. Rao and S. Bihargava, "Multispectral Data Compression using Inter-band Prediction," *Proc. of the Industry Workshop of Data Compression Conference*, Snowbird Utah, April 1994.
13. S. Gupta, A. Gersho and A. G. Tescher, "Recent Trends in Compression of Multispectral Imagery," *Proc. of Twenty Seventh Asilomar Conf on Signals, Systems & Computer*, Nov. 1993.
14. S. S. Shen, J.E. Lindgren, and P.M. Payton, "Effects of Multispectral Image Compression on Machine Exploitation," *Proc. of Twenty Seventh Asilomar Conf on Signals, Systems & Computers*, Nov. 1993.
15. K. Riehl Jr. and C.D. Erdman, "An Assessment of the Effects of Data Compression on Multispectral Imagery," *Proc. of the Industry Workshop of Data Compression Conference*, Snowbird Utah, April 1994.
16. K. Gutttag, R. J. Gove, and J. R. Van Aken, "A Single-Chip Multiprocessor for Multimedia: The MVP," *IEEE Computer Graphics & Applications*, Vol. 12, No. 6, Nov. 1992.
17. P. Wayner, "Digital Video Goes Real Time," *BYTE*, pp 108-112, Jan. 1994.