

Multispectral Image Coding

Daniel Tretter
Hewlett Packard Laboratories
Palo Alto, CA 94304
tretter@bpl.hp.com
(650) 857-4727

Nasir Memon
Brooklyn Polytechnic University
Brooklyn, NY 11201
memon@poly.edu
(718) 260-3970

Charles A. Bouman
School of Electrical and Computer Engineering
Purdue University
West Lafayette, IN 47907
bouman@ecn.purdue.edu
(765) 494-0340

1 Introduction

Multispectral images are a particular class of images that require specialized coding algorithms. In multispectral images, the same spatial region is captured multiple times using different imaging modalities. These modalities often consist of measurements at different optical wavelengths (hence the name multispectral), but the same term is sometimes used when the separate image planes are captured from completely different imaging systems. Medical multispectral images, for example, may combine MRI, CT, and X-ray images into a single multilayer data set [10]. Multispectral images are three-dimensional data sets in which the third (spectral) dimension is qualitatively different from the other two. Because of this, a straightforward extension of two-dimensional image compression algorithms is generally not appropriate. Also, unlike most two-dimensional images, multispectral data sets are often not meant to be viewed by humans. Remotely sensed multispectral images, for example, often undergo electronic computer analysis. As a result, the quality of decompressed images may be judged by a different criterion than for two-dimensional images.

The most common example of multispectral images are conventional RGB color images, which contain three spectral image planes. The image planes represent the red, green, and blue color channels, which all lie in the visible range of the optical band. These three spectral images can be combined to produce a full color image for viewing on a display. However, most printing systems use four colors, typically cyan, magenta, yellow and black (CMYK) to produce a continuous range of colors. More recently, many high fidelity printing systems have begun to use more than four colors to increase the printer gamut, or range of printable colors. This is particularly common in photographic printing systems.

In fact, three colors are not sufficient to specify the appearance of an object under varying illuminants and viewing conditions. To accurately predict the perceived color of a physical surface, we must know the reflectance of the surface as a function of wavelength. Typically, spectral reflectance is measured at 31 wavelengths ranging from 400 to 700 nanometers; however, experiments indicate that the spectral reflectances of most physical materials can be accurately represented with 8 or fewer spectral basis functions [13]. Therefore, some high fidelity image capture systems collect and store more than three spectral measurements at each spatial location or pixel in the image [13]. For example, the VASARI imaging system developed at the National Gallery in London employs a seven-channel multispectral camera to capture paintings [20]. At this time, color image representations with more than four bands are only used in very high quality and high cost systems. However, such multispectral representations may become more common as the cost of hardware decreases and image quality requirements increase.

Another common class of multispectral data is remotely sensed imagery. Remote sensing consists of capturing image data from a remote location. The sensing platform is usually an aircraft or satellite, and the scene being imaged is usually the earth's surface. Since the sensor and the target are so far apart, each pixel in the image can correspond to tens or even hundreds of square meters on the ground. Data gathered from remote sensing platforms is normally not meant primarily for human viewing. Instead, the images are analyzed electronically to determine factors such as land use patterns, local geography, and ground cover classifications. Surface features in remotely sensed imagery can be difficult to distinguish with only a few bands of data. In particular, a larger number of spectral bands are necessary if a single data set is to be used for multiple purposes. For instance, a geographical survey may require different spectral bands than a crop study. Older systems, like the French SPOT or American thematic mapper, use only a handful of spectral bands. More modern systems, however, can incorporate hundreds of spectral bands into a single image [17]. Compression is important for this class of images both to minimize transmission bandwidth from the sensing platform to a ground station and to archive the captured images.

Some medical images include multiple image planes. Although the image planes may not actually correspond to separate frequency bands, they are often still referred to as multispectral images. For example, magnetic resonance imaging (MRI) can simultaneously measure multiple characteristics of the medium being imaged [11]. Alternatively, multispectral medical images can be formed from different medical imaging modalities such as MRI, CT, and X-ray [10]. These multimodal images are useful for identifying and diagnosing medical disorders.

Most multispectral compression algorithms assume that the multispectral data can be represented as a two-dimensional image with vector-valued pixels. Each pixel then consists of one sample from each image plane (spectral band). This representation requires all spectral bands to be sampled at the same resolution and over the same spatial extent. Most multispectral compression schemes also assume the spectral bands are perfectly registered, so each pixel component corresponds to the same exact location in the scene. For instance, in a perfectly registered multispectral image, a scene feature that covers only a single image pixel will cover exactly the same pixel in all spectral bands. In actual physical systems, registration can be a difficult task, and misregistration can severely degrade the resulting

compression ratio or decompressed image quality. Also, although the image planes may be resampled to have pixel values at the same spatial locations, the underlying images may not be of the same resolution.

As with monochrome image compression, multispectral image compression algorithms fall into two general categories: lossless and lossy. In lossless compression schemes, the decoded image is identical to the original. This gives perfect fidelity but limits the achievable compression ratio. For many applications, the required compression ratio is larger than can be achieved with lossless compression, so lossy algorithms are used. Lossy algorithms typically obtain much higher compression ratios but introduce distortions in the decompressed image. Popular approaches for lossy image coding are covered in chapters 5.2 - 5.5 of this volume, while lossless image coding is discussed in chapters 5.1 and 5.6. Lossy compression algorithms attempt to introduce errors in such a way as to minimize the degradation in output image quality for a given compression ratio. In fact, the rate distortion curve gives the minimum bitrate (and hence maximum compression) required to achieve a given distortion. If the allowed distortion is taken to be zero, the resulting maximum compression is the limit for lossless coding. The limit obtained from the theoretical rate distortion curve can be useful for evaluating the effectiveness of a given algorithm. While the bound is usually computed with respect to mean squared error (MSE) distortion, MSE is not a good measure of quality in all applications.

Most two-dimensional image coding algorithms attempt to transform the image data so that the transformed data samples are largely uncorrelated. The samples can then be quantized independently and entropy coded. At the decoder, the quantized samples are recovered and inverse transformed to produce the reconstructed image. The optimal linear transformation for decorrelating the data is the well-known Karhunen-Loeve (KL) transform. Since the KL transformation is data dependent, it requires considerable computation and must be encoded along with the data so it is available at the decoder. As a result, frequency transforms such as the discrete cosine transform (used in JPEG) or the wavelet transform are used to approximate the KL transform along the spatial dimensions. In fact, it can be shown that frequency transforms approximate the KL transform when the image is a stationary 2-D random process. This is generally a reasonable assumption since, over a large ensemble of images, statistical image properties should not vary significantly with spatial position. A large number of frequency transforms can be shown to approach the optimal KL transform as the image size approaches infinity, but in practice, the discrete cosine and wavelet transforms approach this optimal point much more quickly than many other transforms, so they are preferred in actual compression systems.

Multispectral images complicate this scenario. The third (spectral) dimension is qualitatively different from the spatial dimensions, and it generally cannot be modeled as stationary. The correlation between adjacent spectral bands, for example, can vary widely depending on which spectral bands are being considered. In a remotely sensed image, for instance, two adjacent infrared spectral bands might have consistently higher correlation than adjacent bands in the visible range. The correlation is thus dependent on absolute position in the spectral dimension, which violates stationarity. This means that simple frequency transforms along the spectral dimension are generally not effective. Moreover, we will see that most multispectral compression methods work by treating each spectral band differently.

This can be done by computing a KL transform across the spectral bands, using prediction filters which vary for each spectral band, or applying vector quantization methods which are trained for the statistical variation among bands.

Multispectral image compression algorithms can be roughly categorized by how they exploit the redundancies along the spatial and spectral dimensions. The simplest method for compressing multispectral data is to decompose the multispectral image into a set of monochrome images, and then to separately compress each image using conventional image compression methods. Other multispectral compression techniques concentrate solely on the spectral redundancy. However, the best compression methods exploit redundancies in both the spatial and spectral dimensions.

In [25], Tretter and Bouman categorized transform-based multispectral coders into three classes. These classes are important because they describe both the general structure of the coder and the assumptions behind its design. Figure 1 illustrates the structure of these three basic coding methods.

- **Spectral-spatial transform:** In this method, a KL-transform is first applied across the spectral components to decorrelate them. Then each decorrelated component is compressed separately using a transform based image coding method. The image coding method can be based on either block DCT's (see chapter 5.5 in this volume) or a wavelet transform (chapter 5.4). This method is asymptotically optimal if all image planes are properly registered and have the same spatial resolution.
- **Spatial-spectral transform:** In this method, a spatial transform (i.e. block DCT or wavelet transform) is first applied. Then the spectral components of each spatial-frequency band are decorrelated using a **different** KL transform. So for example, a different KL transform is used for each coefficient of the DCT transform or each band of the wavelet transform. This method is useful when the different spectral components have different spatial-frequency content. For instance, an infrared band may have lower spatial resolution than a visible band of the same multispectral image. In this case, the separate KL transforms result in better compression.
- **Complex spatial-spectral transform:** If, in addition, the image planes are not registered, the frequency transforms must be complex to retain phase information between planes. A spatial shift in one image plane relative to another (i.e. misregistration) corresponds to a phase shift in the frequency domain. In order to retain this information, frequency components must be stored as complex numbers. This method differs from the spatial-spectral method in that the transforms must be complex valued. This complex spatial-spectral transform has the advantage that it can remove the effect of misregistration between the spectral bands. However, because it requires the use of a DFT (instead of DCT), or a complex wavelet transform, it is more complicated to implement. If a real spatial-spectral transform is used to compress an image that has misregistered planes, much of the redundancy between image planes will be missed. The transform is unable to follow a scene feature as it shifts in location from one image plane to another, so the feature is essentially treated as a separate feature in each plane and is coded multiple times. As a result, the image will not compress well.

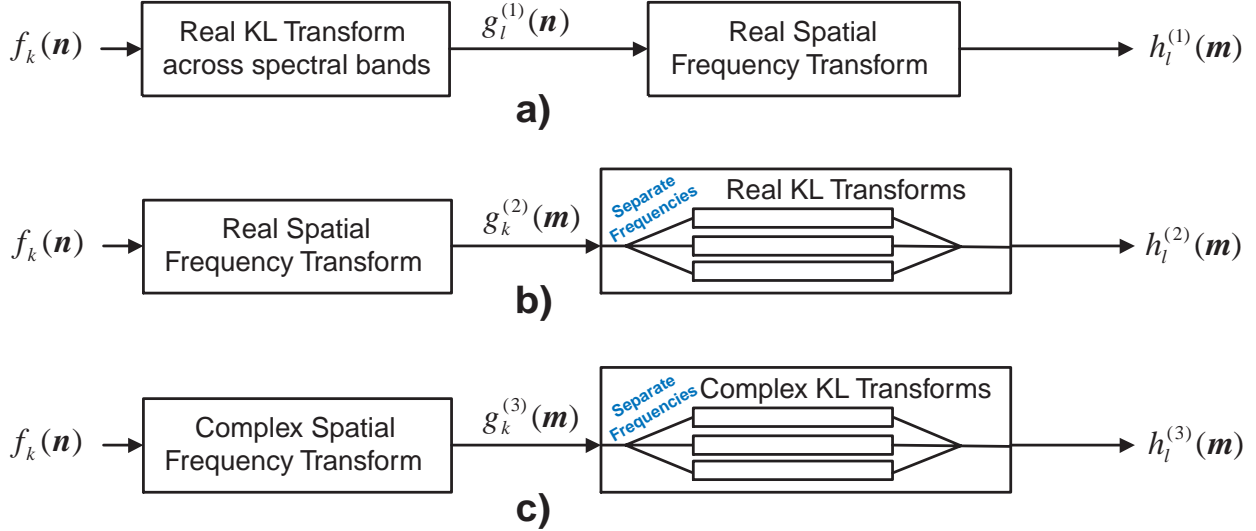


Figure 1: Three basic classes of transform coders for multispectral image compression. a) Spectral-spatial method: requires that all image planes be registered and of the same resolution; b) Spatial-spectral method: requires that all image planes be registered but not necessarily of the same resolution; c) Complex spatial-spectral method: does not require registration or planes of the same resolution.

In the following sections, we will discuss a variety of methods for both lossy and lossless multispectral image coding. The most appropriate coding method will depend on the application and system constraints.

2 Lossy Compression

Many researchers have worked on the problem of compressing multispectral images. In the area of lossy compression, most of the work has concentrated on remotely sensed data and RGB color images rather than medical imagery or photographic images with more than three spectral bands. For diagnostic and legal reasons, medical images are often compressed losslessly, and the high fidelity systems that use multispectral photographic data are still relatively rare.

Suppose we represent a multispectral image by $f_k(\mathbf{n})$, where

$$\mathbf{n} = (n_1, n_2), \quad n_1 \in \{0 \cdots M - 1\}, \quad n_2 \in \{0 \cdots N - 1\}, \quad \text{and } k \in \{0 \cdots K - 1\}.$$

In this notation, \mathbf{n} represents the two spatial dimensions and k is the spectral band number. In the development below, we will denote spectral band k by f_k , and $\mathbf{f}(\mathbf{n})$ will represent a single vector-valued pixel at spatial location $\mathbf{n} = (n_1, n_2)$.

Lossy compression algorithms attempt to introduce errors in such a way as to minimize the degradation in output image quality for a given compression ratio. To do this, algorithm designers must first decide on an appropriate measure of output image quality. Quality is often measured by defining an error metric relating the decompressed image to the original.

The most popular error metric is the simple mean squared error (MSE) between the original image and the decompressed image. Although this metric does not necessarily correlate well with image quality, it is easy to compute and mathematically tractable to minimize when designing a coding algorithm. If a decompressed two-dimensional $M \times N$ image $\hat{f}(\mathbf{n})$ is compared with the original image $f(\mathbf{n})$, the mean squared error is defined as


$$MSE = \sum_{n_1=0}^{M-1} \sum_{n_2=0}^{N-1} [f(n_1, n_2) - \hat{f}(n_1, n_2)]^2 .$$

For photographic images, quality is usually equated with visual quality as perceived by a human observer. The error metrics used thus often incorporate a human visual model. One popular choice is to use a visually weighted MSE between the original image and the decompressed image. This is normally computed in the frequency domain, since visual weighting of frequency coefficients is more natural than weightings in the spatial domain.

Some images are used for purposes other than viewing. Medical images may be used for diagnosis, and satellite photos are sometimes analyzed to classify surface regions or identify objects. For these images, other error metrics may be more appropriate since the measure of image quality is quite different. We will discuss this topic further with respect to multispectral images later in this chapter.

2.1 RGB color images

Lossy compression of RGB color images deserves special mention. These images are by far the most common type of multispectral image, and a considerable body of research has been devoted to development of appropriate coding techniques. Color images in uncompressed form typically consist of red, green, and blue color planes, where the data in each plane has undergone a nonlinear gamma correction to make it appropriate for viewing on a CRT monitor [22]. Typical CRT monitors have a nonlinear response, so doubling the value of a pixel (the frame buffer value), for instance, will increase the luminance of the displayed pixel, but the luminance will not double. The nonlinear response approximates a power function, so digital color images are usually prewarped using the inverse power function to make the image display properly. Different color imaging systems can have different definitions of red, green, and blue, different gamma curves, and different assumed viewing conditions. In recent years, however, many commercial systems are moving to the sRGB standard to provide better color consistency across devices and applications [2]. Before compression, color images are usually transformed from RGB to a luminance-chrominance representation. Each pixel vector $\mathbf{f}(\mathbf{n})$ is transformed via a reversible transformation to an equivalent luminance-chrominance vector $\mathbf{g}(\mathbf{n})$.



Two common luminance-chrominance color spaces used are YCrCb, a digital form of the YUV format used in NTSC color television, and CIELab [22]. YCrCb is obtained from





Figure 2: Detail illustrating JPEG compression artifacts (75 dpi). a) Original image data, b) JPEG compressed 30:1 using chrominance subsampling, c) JPEG compressed 30:1 using no chrominance subsampling.

sRGB via a simple linear transformation, while CIELab requires nonlinear computations and is normally computed using lookup tables.

The purpose of the transformation is to decorrelate the spectral bands visually so they can be treated separately. After transformation, the three new image planes are normally compressed independently using a two dimensional coding algorithm such as the ones described earlier in this chapter. The luminance channel Y (or L) is visually more important than the two chrominance channels, so the chrominance images are often subsampled by a factor of 2 in each dimension before compression [22]. Perhaps the most common color image compression algorithm uses the YCrCb (sometimes still referred to as YUV) color space in conjunction with chrominance subsampling and standard JPEG compression on each image plane. Many color devices refer to this entire scheme as JPEG even though the standard does not specify color space or subsampling. Most JPEG images viewed across the World Wide Web by browsers have been compressed in this way. Figure 2 illustrates the artifacts introduced by JPEG compression. Figure 2a shows a detail from an original uncompressed image, 2b illustrates the decompressed image region after 30:1 JPEG compression with chrominance subsampling, and 2c illustrates the decompressed image after 30:1 JPEG compression with no chrominance subsampling. Figures 2b and 2c both show typical JPEG compression artifacts; the reconstructed images have blocking artifacts in the smooth regions such as the back of the hand, and both images show ringing artifacts along the edges. However, the artifacts are much more visible in Figure 2c, which was compressed without chrominance subsampling. Since the chrominance components are retained at full resolution, a larger percentage of the compressed data stream is required to represent the chrominance information, so fewer bits are available for luminance. The additional artifacts introduced by using fewer bits for luminance are more visible than the artifacts caused by chrominance subsampling, so Figure 2c has more visible artifacts than Figure 2b.

One interesting approach to color image storage and compression is to use color palettization. In this approach, a limited palette of representative colors (usually no more than 256 colors) is stored as a lookup table, and each pixel in the image is replaced by an index into the table that indicates the best palette color to use to approximate that pixel. This is essentially a simple vector quantization scheme (vector quantization is covered in detail in

chapter 5.3). Palettization was first designed not for compression, but to match the capabilities of display monitors. Some display devices can only display a limited number of colors at a time, due to either a limited internal memory size or to characteristics of the display itself. As a result, images had to be palettized before display.



Palettization collapses the multispectral image into a single image plane, which can be further compressed if desired. Both lossy and lossless compression schemes for palettized images have been developed. The well-known GIF format, which is often used for images transmitted over the world wide web, is one example of this sort of image. As a compression technique, palettization is most useful for non-photographic images, such as synthetically generated images, which often only use a limited number of colors to begin with.

2.2 Remotely sensed multispectral images

Remotely sensed multispectral images have been in use for a long time. The Landsat 1 system, for example, was first launched in 1972. Aircraft based systems have been in use even longer. Satellite and aircraft platforms can gather an extremely large amount of data in a short period of time, and remotely sensed data is often archived so changes in the earth's surface can be tracked over long periods of time. As a result, compression has been of considerable interest since the earliest days of remote sensing, when the main purpose of compression was to reduce storage requirements and processing time [9, 16]. Although processing and data storage facilities are becoming increasingly more powerful and affordable, recent remote sensing systems continue to stress state of the art technology. Compression is particularly important for spaceborne systems where transmission bandwidth reduction is a necessity [9, 26]. Reviews of compression approaches for remotely sensed images can be found in [21, 28].

The simplest type of lossy compression for multispectral images, known as spectral editing, consists of not transmitting all spectral bands. Some sort of algorithm is used to determine which bands are of lesser importance, and those bands are not sent. Since this amounts to simply throwing away some of the data, such a technique is obviously undesirable. For one thing, the choice of bands to eliminate is strongly dependent on the information desired from the image. Since a variety of researchers may want to extract entirely different information from the same image, all of the bands may be needed at one time or another. As a result, a number of researchers have proposed more sophisticated ways to combine the spectral bands and reduce the spectral dimensionality while retaining as much of the information as possible.

As in two-dimensional image compression, many algorithms attempt to first transform the image data such that the transformed data samples are largely uncorrelated. For multispectral images, the spectral bands are often modeled as a series of correlated random fields. If each spectral band f_k is a two dimensional stationary random field, frequency transforms are appropriate across the spatial dimensions. However, redundancies between

spectral bands are usually removed differently. A variety of schemes use a KL or similar transformation across spectral bands followed by a two-dimensional frequency transform like DCT or wavelets across the two spatial dimensions [1, 5, 6, 7, 24, 25]. These **spectral-spatial transform methods** are of the general form shown in Figure 1a) and have been shown to be asymptotically optimal for a MSE distortion metric as the size of the data set goes to infinity when the following three assumptions hold [25].


- i) The spectral components can be modeled as a stationary Gaussian random field.
- ii) The spectral components are perfectly registered with one another.
- iii) The spectral components have similar frequency distributions (for instance, are of the same resolution as one another).

If *iii)* does not hold, a separate KL spectral transform must be used at every spatial frequency. Algorithms of this sort have been proposed by several researchers [1, 25, 27]. If assumption *ii)* does not apply either, a complex frequency transform must be used to preserve phase information if the algorithm is to remain asymptotically optimal [25]. However, the computational complexity involved makes this approach difficult, so it is generally preferable to add more preprocessing to better register the spectral bands. Some recent algorithms also get improved performance by adapting the KL transform spatially based on local data characteristics [6, 7, 24].


Figure 3 shows the result of applying two different coding algorithms to a Thematic Mapper multispectral data set. The data set consists of bands 1, 4, and 7 from a Thematic Mapper image. Figure 3a shows a region from the original uncompressed data. The image is shown in pseudo-color, with band 1 being mapped to red, band 4 to green, and band 7 to blue. Figure 3b shows the reconstructed data after 30:1 compression using an algorithm from [25] that uses a single KL transform followed by a two-dimensional frequency subband transform across the two spatial dimensions (RSS algorithm), and 3c shows the reconstructed data after 30:1 compression using a similar algorithm that first applies the frequency transform and then computes a separate KL transform for each frequency subband (RSM algorithm). For this imaging device, band 7 is of lower resolution than the other bands, so assumption *iii)* does not hold for this data set. As a result, we expect the RSM algorithm to outperform the RSS algorithm on this data set. Comparing Figures 3b and 3c, we can see that 3c has slightly fewer visual artifacts than 3b. The mean squared error produced by the RSM algorithm was 27.44 for this image, compared with a mean squared error of 28.65 for the RSS algorithm. As expected, the RSM algorithm outperforms the RSS algorithm on this data set both in visual quality and mean squared error.

Rather than decorrelating the data samples by using a reversible transformation, some approaches use linear prediction to remove redundancy. The predictive algorithms are often used in conjunction with data transformations in one or more dimensions [9, 14]. For instance, spectral redundancy may be removed using prediction, while spatial redundancies are removed via a decorrelating transformation.


Correlation in the data can also be accounted for by using clustering or vector quantization (VQ) approaches, often coupled with prediction. A number of predictive VQ and



a



b



c

Figure 3: Detail illustrating transform-based compression on Thematic Mapper data (100 dpi). a) Original image data in pseudo-color, b) compressed 30:1 using RSS algorithm (single KL transform), c) compressed 30:1 using RSM algorithm (multiple KL transforms). The RSM algorithm gives better compression for this result, with a mean squared error of 27.44 vs. the RSS algorithm at 28.65.

clustering techniques have been proposed [4, 8, 9, 26]. As with predictive algorithms, VQ methods can be combined with decorrelating data transformations [1, 27].

Finally, some compression algorithms have been devised specifically for multispectral images, where the authors assumed the images would be subjected to machine classification. These approaches, which are not strongly tied to two-dimensional image compression algorithms, use parametric modeling to approximate the relationships between spectral bands [15]. Classification accuracy is used to measure the effectiveness of these algorithms.

Two popular approaches for lossy compression of remotely sensed multispectral images have emerged in recent years. One approach is based on predictive VQ, while the other consists of a decorrelating KL transform across spectral bands in conjunction with frequency transforms in the spatial dimensions. We discuss a representative algorithm of each type in more detail below to help expand upon and illustrate the main ideas involved in each approach.

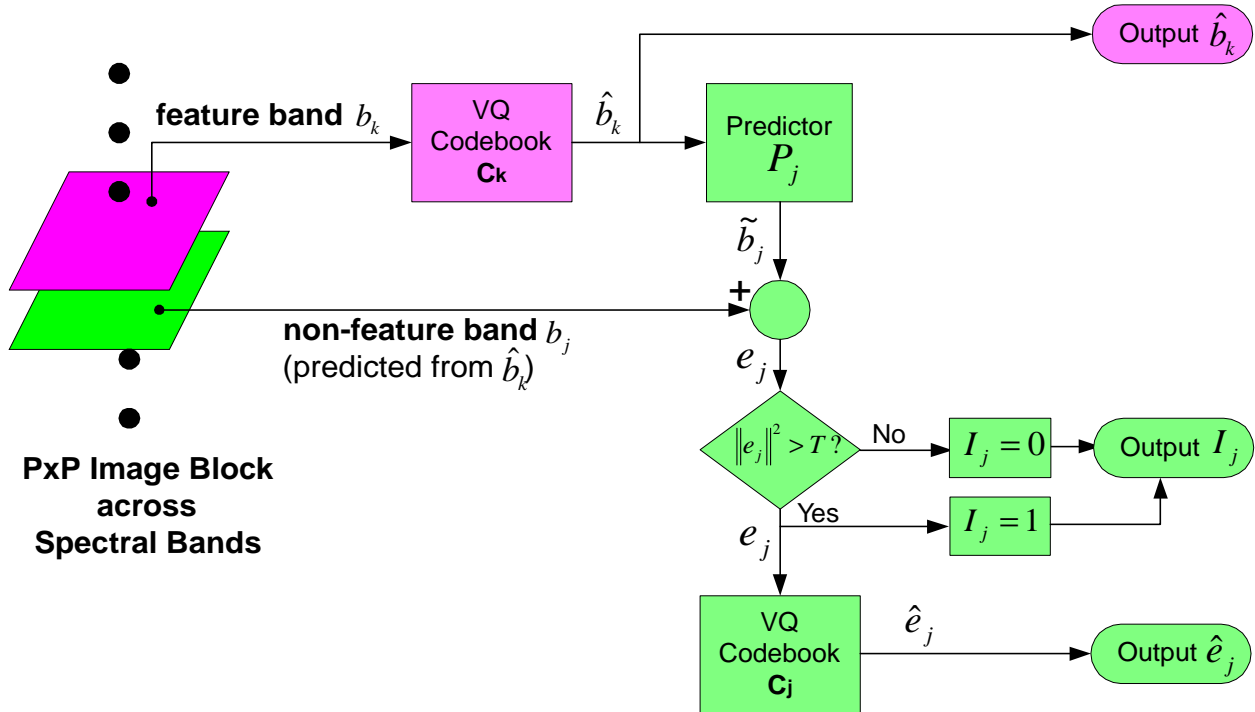


Figure 4: Gupta and Gersho’s feature predictive vector quantization scheme encodes each spectral block separately, where feature blocks are used to predict non-feature blocks, thus removing both spatial and spectral redundancy.

Gupta and Gersho propose a feature predictive vector quantization approach to the compression of multispectral images [8]. Vector quantization is a powerful compression technique, known to be capable of achieving theoretically optimal coding performance. However, straightforward VQ suffers from high encoding complexity, particularly as vector dimension increases. Thus, Gupta and Gersho couple VQ with prediction to keep the vector dimension manageable while still accounting for all of the redundancies in the data. In particular, VQ is used to take advantage of spatial correlations, while spectral correlations are removed using prediction. The authors propose several algorithm variants, but we only discuss one of them here.

Gupta and Gersho begin by partitioning each spectral band k into $P \times P$ nonoverlapping blocks, which will each be coded separately. Suppose $b_k(\mathbf{m})$ is one such set of blocks, with

$$\mathbf{m} = (m_1, m_2), \quad m_1 \in \{0 \cdots P - 1\}, \quad m_2 \in \{0 \cdots P - 1\}, \quad \text{and } k \in \{0 \cdots K - 1\} .$$

Figure 4 illustrates the operation of the algorithm for the two types of spectral blocks b_k and b_j in this set of blocks. A small subset L of the K spectral bands is chosen as feature bands. The number of feature bands will be chosen based on the total number of spectral bands and the correlations among them. Vector quantization is used to code each feature band separately, as illustrated for block b_k in the figure. Each feature band is coded using a separate VQ codebook. Each of the $K - L$ non-feature bands is then predicted from one of the coded feature bands. The prediction is subtracted from the actual data values to get an error block e_j for each non-feature band. If the energy (squared norm) of the error block

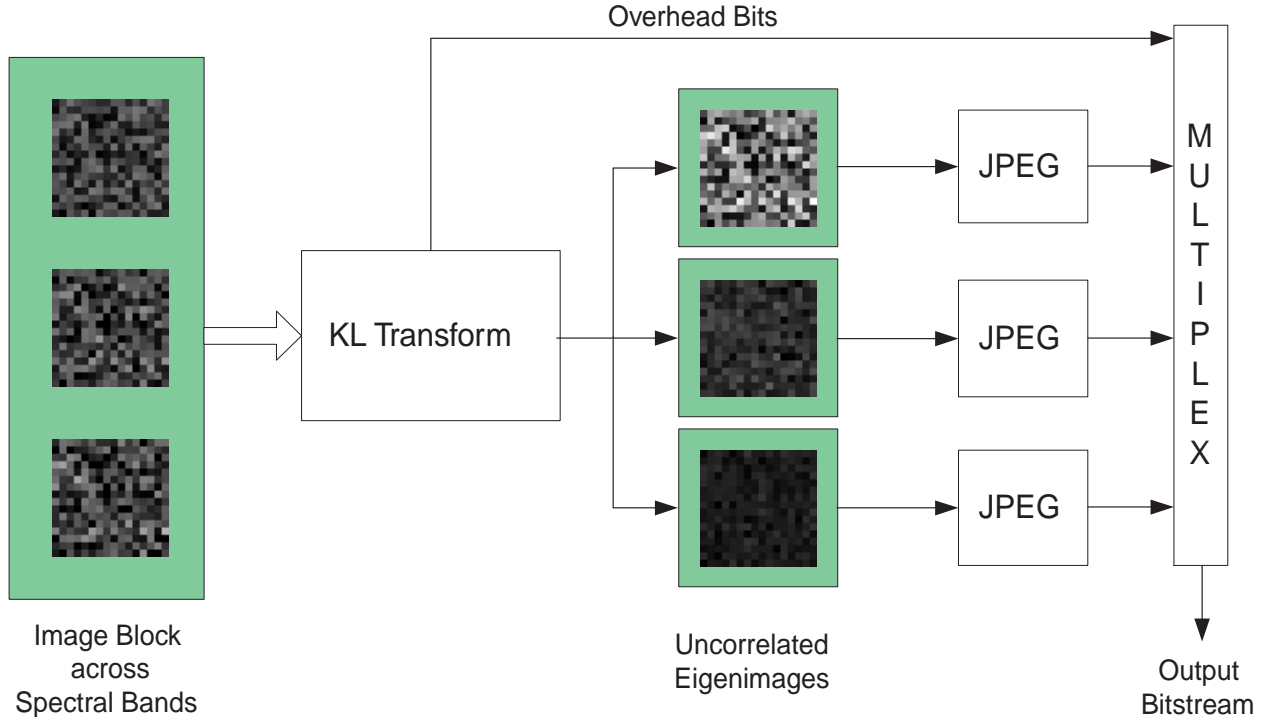


Figure 5: Saghri, Tescher, and Reagan have devised an algorithm that uses the KL transform to decorrelate image blocks across spectral bands, followed by JPEG to remove spatial redundancies. The KL transform concentrates much of the energy into a single band, improving coding efficiency. Note that the KL transform is data dependent, so the transformation matrix must be sent to the decoder as overhead bits.

exceeds a predefined threshold T , the error block is coded using yet another VQ codebook. This procedure is illustrated for block b_j in Figure 4. A binary indicator flag I_j is set for each non-feature band to indicate whether or not the error block was coded.

$$I_j = \begin{cases} 1 & \|e_j\|^2 > T \\ 0 & \text{else} \end{cases}$$

To decode, simply add any decoded error blocks to the predicted block for non-feature bands. Feature bands are decoded directly using the appropriate VQ codebook. Gupta and Gersho also derive optimal predictors P_j for their algorithm and discuss how to design the various codebooks from training images. See [8] for details.

One example of a transform-based compression system is proposed by Saghri, Tescher, and Reagan in [19]. Their algorithm uses the KL transform to decorrelate the data spectrally, followed by JPEG compression on each of the transformed bands. Like Gupta and Gersho, they begin by partitioning each spectral band into nonoverlapping blocks, which will be coded separately. A separate KL transform is computed for each spatial block, so different regions of the image will undergo different transformations. This approach allows the scheme to adapt to varying terrain in the scene, producing better compression results.

Figure 5 illustrates the algorithm for a single image block. In this example, the image consists of three highly correlated spectral bands. The KL transform concentrates much

of the energy into a single band, improving overall coding efficiency. Saghri, Tescher, and Reagan have designed their algorithm for use on board an imaging platform, so they must consider a variety of practical details in their paper. The KL transform is generally a real valued transform, so they use quantization to reduce the required overhead bits. Also, JPEG expects data of only 8 bits per pixel, so in order to use standard JPEG blocks, they scale and quantize the transformed bands (eigenimages) to 8 bits per pixel. This mapping is also sent to the decoder as overhead. The spatial block size is chosen to give good compression performance while keeping the number of overhead bits small.

The authors also discuss practical ways to select the JPEG parameters to get the best results. They use custom quantization tables and devise several possible schemes for selecting the appropriate quality factor for each transformed band.

Saghri, Tescher, and Reagan have found that their system can give about 40 : 1 compression ratios with visually lossless quality for test images containing eleven spectral bands. They give both measured distortion and classification accuracy results for the decompressed data to support this conclusion. They also examine the sensitivity of the algorithm to various physical system characteristics. They found that the coding results were sensitive to band misregistration, but were robust to changes in the dynamic range of the data, dead/saturated pixels, and calibration and preprocessing of the data. More details can be found in the paper [19].

Most coding schemes for remotely sensed multispectral images are designed for a MSE distortion metric, but many authors also consider the effect on scene classification accuracy [9, 19, 25]. For fairly modest compression ratios, MSE is often a good indicator of classification accuracy. System issues in the design of compression for remotely sensed images are discussed in some detail in [28].

2.3 Multispectral medical and photographic images

Relatively little work has been done on lossy compression for multispectral medical images and photographic images with more than three spectral bands. Medical images are usually compressed losslessly for legal and diagnostic reasons, although preliminary findings indicate that moderate degrees of lossy compression may not affect diagnostic accuracy. Since most diagnosis is performed visually, visual quality of decompressed images correlates well with diagnostic accuracy. Hu, Wang, and Cahill propose linear prediction algorithms for the lossy compression of multispectral MR images [11]. They compare MSE and visual quality of their algorithm versus several other common compression schemes and report preliminary results indicating that diagnostic accuracy is not affected by compression ratios up to 25:1. They do note that their results rely on the spectral bands being well registered, so a preprocessing step may be necessary in some cases to register the planes before coding to get good results.

Multispectral photographic images with more than three spectral bands are still relatively rare. Most work with these images concentrates on determining the appropriate number and placement of the spectral bands and deriving the mathematical techniques to use the additional bands for improved color reproduction. Many researchers must design and build their own image capture systems as well, so considerable attention and effort is being spent on image capture issues. As a result, little work has been published on lossy coding of such

data, although we expect this to be an interesting area for future research. One interesting approach that could be the first stage in a lossy compression algorithm is the multispectral system proposed by Imai and Berns, which combines a low spatial resolution multispectral image with a high resolution monochrome image [12]. The monochrome image contains lightness information, while color information is obtained from the lower resolution multispectral image. Since human viewers are most sensitive to high frequencies in lightness, these images can be combined to give a high resolution color image with little or no visible degradation resulting from the lower resolution of the multispectral data. In this way, the approach of Imai and Berns is analogous to the chrominance subsampling often done during RGB color image coding.

3 Lossless Compression

Because of the difference in goals, the best way of exploiting spatial and spectral redundancies for lossy and lossless compression is usually quite different. The decorrelating transforms used for lossy compression usually cannot be used for lossless compression as they often require floating point computations which result in loss in data when implemented with finite precision arithmetic. This is especially true for “optimal” transforms like the KL transform and the DCT transform. Also, techniques based on vector quantization are clearly of little utility for lossless compression. Furthermore, irrespective of the transform used, there is often a significant amount of redundancy that remains in the data after decorrelation, the modeling and capturing of which constitutes a crucial step in lossless compression.

There are essentially two main approaches used for lossless image compression. The first is the traditional DPCM approach based on prediction followed by context modeling of prediction errors. The second and more recent approach is based on reversible integer wavelet transforms followed by context modeling and coding of transform coefficients. For a detailed description of these techniques and specific algorithms that employ these approaches, the reader is referred to accompanying chapters in this volume on lossless image compression (Chapter 5.1) and Wavelet-Based Coding (Chapter 5.4). In the rest of this section, we focus on how techniques based on each of these two approaches can be extended to provide lossless compression of multispectral data.

3.1 Predictive Techniques

While extending a predictive technique to exploit inter-band correlations, the following new issues arise -

- *Band ordering.* In what order does one encode the different spectral bands. This is related to the problem of determining which band(s) are the best to use as reference band(s) for predicting and modeling intensity values in a given band.
- *Inter-band prediction.* How to best incorporate additional information available from pixels located in previously encoded spectral bands to improve prediction?

- *Inter-band error modeling.* How to exploit information available from prediction errors incurred at pixel locations in previously encoded spectral bands to better model and encode the current prediction error?

We examine typical approaches that have been taken to address these questions in the rest of this sub-section.

3.1.1 Band Ordering

In [29], Wang *et. al.* analyzed correlations between the seven bands of LANDSAT TM images, and proposed an order, based on heuristics, to code the bands that result in the best compression. According to their studies, bands 2, 4, and 6 should first be encoded by traditional intra-band linear predictors optimized within individual bands. Then pixels in band 5 are predicted using neighboring pixels in band 5 as well as those in bands 2, 4, and 6. Finally, bands 1, 3, and 7 are coded using pixels in the local neighborhood as well as selected pixels from bands 2, 4, 5, and 6.

If we restrict the number of reference bands that can be used to predict pixels in any given band, then Tate [23] showed that the problem of computing an optimal ordering can be formulated in graph theoretic terms, admitting an $O(N^2)$ solution for an N -band image. He also observed that using a single reference band is sufficient in practice as compression performance does not improve significantly with additional bands. Although significant improvements in compression performance were demonstrated, one major limitation of this approach is the fact that it is two-pass. An optimal ordering and corresponding prediction coefficients are first computed by making an entire pass through the data set. This problem can be alleviated to some degree by computing an optimal ordering for different types of images. Another limitation of the approach is that it re-orders entire bands. That is, it makes the assumption that spectral relationships do not vary spatially. The optimal spectral ordering and prediction coefficients will change spatially depending on the characteristics of the objects being imaged.

In the remainder of this sub-section, for clarity of exposition, we assume that the image in question has been appropriately re-ordered, if necessary, and simply use the previous band as the reference band for encoding the current band. However, before we proceed, there is one further potential complication that needs to be addressed. The different bands in a multi-spectral image may be represented one pixel at a time (pixel interleaved), one row at a time (line interleaved), or an entire band at a time (band sequential). Since the coder needs to utilize at least one band (the reference band) in order to make compression gains on other bands, buffering strategy and requirements would vary with the different representations and should be taken into account before adopting a specific compression technique. We assume this to be the case in the remainder of this sub-section and discuss prediction and error modeling techniques for lossless compression of multi-spectral images, irrespective of the band ordering and pixel interleaving employed.

3.1.2 Inter-band Prediction

Let Y denote the current band and X the reference band. In order to exploit inter-band correlations, it is easy to generalize a DPCM-like predictor from two-dimensional to three-dimensional sources. Namely, we predict the current pixel $Y[i, j]$ to be

$$\hat{Y}[i, j] = \sum_{a,b \in N_1} \theta_{a,b} Y[i - a, j - b] + \sum_{a',b' \in N_2} \theta'_{a',b'} X[i - a', j - b'], \quad (1)$$

where N_1 and N_2 are appropriately chosen neighborhoods that are causal with respect to the scan and the band interleaving being employed. The coefficients $\theta_{a,b}$ and $\theta'_{a',b'}$ can be optimized by standard techniques to minimize $\|Y - \hat{Y}\|$ over a given multispectral image. In [18] Roger and Cavenor performed a detailed study on AVIRIS¹ images with different neighborhood sets and found that a 3rd order spatial-spectral predictor based on the immediate two neighbors $Y[i, j - 1], Y[i - 1, j]$ and the corresponding pixel $X[i, j]$ in the reference band is sufficient and larger neighborhoods provide very marginal improvements in prediction efficiency.

Since the characteristics of multispectral images often vary spatially, optimizing prediction coefficients over the entire image can be ineffective. Hence, Roger and Cavenor [18] compute optimal predictors for each row of the image and transmits them as side information. The motivation for adapting predictor coefficients a row at a time has to do with the fact that an AVIRIS image is acquired in a line interleaved manner and a real time compression technique would have to operate under such constraints. However, for offline compression, say for archival purposes, this may not be the best strategy as one would expect spectral relationships to change significantly across the width of an image. A better approach to adapt prediction coefficients would be to partition the image into blocks, and compute optimal predictors on a block-by-block basis.

Computing an optimal least-square multispectral predictor for different image segments does not always improve coding efficiency despite the high computational costs involved. This is because frequently changing prediction coefficients incur too much side information (high model cost) especially for color images which have only 3 or 4 bands. In view of this, Wu and Memon [30] propose an adaptive inter-band predictor that exploits relationships between local gradients among adjacent spectral bands. Local gradients are an important piece of information that can help resolve uncertainty in high activity regions of an image, and hence improve prediction efficiency. The gradient at the pixel being currently coded is known in the reference band but missing in the current band. Hence, the local waveform shape in the reference band can be projected to the current band to obtain a reasonably accurate prediction, particularly in the presence of strong edges. Although there are several ways in which one can interpolate the current pixel on the basis of local gradients in the reference band, in practice the following *difference based inter-band interpolation*, works well.

$$\hat{Y}[i, j] = \frac{Y[i - 1, j] + (X[i, j] - X[i - 1, j]) + Y[i, j - 1] + (X[i, j] - X[i, j - 1])}{2} \quad (2)$$

¹Airborne Visible InfraRed Imaging Spectrometer. AVIRIS is a world class instrument in the realm of Earth Remote Sensing. It delivers calibrated images in 224 contiguous spectral bands with wavelengths from 400 to 2500 nanometers (nm)

Wu and Memon also observed that performing inter-band prediction in an unconditional manner does not always give significant improvements over intra-band prediction and sometimes leads to a degradation in compression performance. This is due to the fact that correlation between bands varies significantly in different regions of the image depending on the objects present in that specific region. Thus it is difficult to find an inter-band predictor that works well across the entire image. Hence, they propose a switched inter/intra band predictor that performs inter-band prediction only if the correlation in the current window is strong enough; otherwise intra-band prediction is used. More specifically, they examine the correlation $\text{Cor}(X_w, Y_w)$ between the current and reference band in a local window w . If $\text{Cor}(X_w, Y_w)$ is high then inter-band prediction is performed else intra-band prediction is used. Since computing $\text{Cor}(X_w, Y_w)$ for each pixel can be computationally expensive they give simple heuristics to approximate this correlation. They report that switched inter/intra band prediction gives significant improvement over optimal predictors using inter or intra band prediction alone.

3.1.3 Error modeling and coding

If the *residual image* consisting of prediction errors is treated as an source with independent identically distributed (i.i.d.) output, then it can be efficiently coded using any of the standard variable length entropy coding techniques, like Huffman coding or arithmetic coding. Unfortunately, even after applying the most sophisticated prediction techniques, the residual image generally has ample structure which violates the i.i.d. assumption. Hence, in order to encode prediction errors efficiently we need a model that captures the structure that remains after prediction. This step is often referred to as *error modeling*. The error modeling techniques employed by most lossless compression schemes proposed in the literature, can be captured within a *context modeling* framework. In this approach, the prediction error at each pixel is encoded with respect to a conditioning state or *context*, which is arrived at from the values of previously encoded neighboring pixels. Viewed in this framework, the role of the error model is essentially to provide estimates of the conditional probability of the prediction error, given the context in which it occurs. This can be done by estimating the probability density function by maintaining counts of symbol occurrences within each context or by estimating the parameters of an assumed probability density function. The accompanying chapter on lossless image compression (chapter 5.1) gives more details on error modeling techniques. Here we look at examples of how each of these two approaches have been used for compression of multispectral images.

An example of the first approach used for multispectral image compression is provided in [18] where Roger and Cavenor investigate two different variations. First they assume prediction errors in a row belong to a single geometric probability mass function (pmf) and determine the optimal Rice-Golomb code by an exhaustive search over the parameter set. In the second technique they compute the variance of prediction errors for each row and based on this utilize one of eight pre-designed Huffman codes. An example of the second approach is provided by Tate [23] who quantizes the prediction error in the corresponding location in the reference band and uses this as a conditioning state for arithmetic coding. Since this involves estimating the pmf in each conditioning state, only a small number of states (4 to 8) are used. An example of a hybrid approach is given by Wu and Memon

[30] who propose an elaborate context formation scheme that includes gradients, prediction errors and quantized pixel intensities from the current and reference band. They estimate the variance of prediction error within each context and based on this estimate they select between one of eight different conditioning states for arithmetic coding. In each state they estimate the pmf of prediction errors by keeping occurrence counts of prediction errors.

Another simple technique for exploiting relationships between prediction errors in adjacent bands that can be used in conjunction with any of the above error modeling techniques, follows from the observation that prediction errors in neighboring bands are correlated and just taking a simple difference between the prediction error in the current and reference band can lead to a significant reduction in the variance of the prediction error signal. This in turn leads to a reduction in bit rate produced by a variable length code like a Huffman or an arithmetic code. The approach can be further improved by conditioning the differencing operation based on statistics gathered from contexts. However, it should be noted that the prediction errors would still contain enough structure to benefit from one of the error modeling and coding techniques described above.

3.2 Reversible transform based techniques

An alternative approach to lossless image compression, that has emerged recently is based on subband decomposition. There are several advantages offered by a subband approach for lossless image compression. The most important of which is perhaps the natural integration of lossy and lossless compression that becomes possible. By transmitting entropy coded subband coefficients in an appropriate manner, one can produce an embedded bit stream that permits the decoder to extract a lossy reconstruction at a desired bit-rate. This enables progressive decoding of the image that can ultimately lead to lossless reconstruction. The image can also be recovered at different spatial resolutions. These features are of great value for specific applications in remote sensing and “network-centric” computing in general. Although quite a few subband based lossless image compression have been proposed in recent literature, there has been very little work on extending them to multispectral images. Bilgin, Zweig and Marcellin [3] extend the well known Zero-tree algorithm for compression of multispectral data. They perform a 3D dyadic subband decomposition of the image and encode transform coefficients by using a zero tree structure extended to three dimensions. They report an improvement of 15 to 20% over the best 2D lossless image compression technique.

3.3 Near-lossless compression

Recent studies on AVIRIS data have indicated that the presence of sensor noise limits the amount of compression that can be obtained by any lossless compression scheme. This is supported by the fact that the best results reported in the literature on compression of AVIRIS data seem to be in the range of 5 - 6 bits per pixel. Increased compression can be obtained with lossy compression techniques which have been shown to provide very high compression ratios with little or no loss in *visual fidelity*. Lossy compression, however, may not be desirable in many circumstances due to the uncertainty of the effects caused by

lossy compression on subsequent scientific analysis that is performed with the image data. One compromise then is to use a *bounded distortion* (or *nearly-lossless*) technique which guarantees that each pixel in the reconstructed image is within $\pm k$ of the original.

Extension of a lossless predictive coding technique to a near lossless one can be done in a straight forward manner by prediction error quantization according to the specified pixel value tolerance. In order for the predictor at the receiver to track the predictor at the encoder, the reconstructed values of the image need then to be used to generate the prediction at both the encoder and the receiver. More specifically, the following uniform quantization procedure leads us to a near-lossless compression technique.

$$Q[x] = \left\lfloor \frac{x + k}{2k + 1} \right\rfloor (2k + 1), \quad (3)$$

where x is the prediction error, k is the maximum reconstruction error allowed in any given pixel and $\lfloor \cdot \rfloor$ denotes the integer part of the argument. At the encoder, a label l is generated according to

$$l = \left\lfloor \frac{x + k}{2k + 1} \right\rfloor. \quad (4)$$

This label is encoded, and at the decoder the prediction error is reconstructed according to

$$\hat{x} = l \times (2k + 1). \quad (5)$$

Near-lossless compression techniques can yield significantly higher compression ratios as compared to lossless compression. For example, ± 1 near-lossless compression can usually lead to reduction in bit rates by about 1 to 1.3 bits per pixel.

4 Conclusion

In applications such as remote sensing, multispectral images were first used to store the multiple images corresponding to each band in a optical spectrum. More recently, multispectral images have come to refer to any image formed by multiple spatially registered scalar images, independent of the specific manner in which the individual images were obtained. This broader definition encompasses many emerging technologies such as multimodal medical images and high fidelity color images. As these new sources of multispectral data become more common, the need for high performance multispectral compression methods will increase.

In this chapter, we have described some of the current methods for both lossless and lossy coding of multispectral images. Effective methods for multispectral compression exploit the redundancy across spectral bands while also incorporating more conventional image coding methods based on spatial dependencies of the image data. Importantly, spatial and spectral redundancy differ fundamentally in that spectral redundancies generally depend on the specific choices and ordering of bands and are not subject to the normal assumptions of stationarity used in the spatial dimension.

We described some typical examples of lossy image coding methods. These methods use either a Karhunen-Loeve (KL) transform or prediction to decorrelate data along the spectral

dimension. The resulting decorrelated images can then be coded using more conventional image compression methods. Lossless multispectral image coding necessitates the use of prediction methods because general transformations result in undesired quantization error. For both the lossy and lossless compression, adaptation to the spectral dependencies is essential to achieve the best coding performance.

References

- [1] F. Amato, C. Galdi, and G. Poggi. Embedded zerotree wavelet coding of multispectral images. In *Proc. IEEE Int. Conf. Image Processing I*, pages 612–615, Santa Barbara, California, Oct. 1997.
- [2] M. Anderson, R. Motta, S. Chandrasekar, and M. Stokes. Proposal for a standard default color space for the internet – sRGB. In *Proc. Fourth Color Imaging Conference: Color Science, Systems, and Applications*, pages 238–246, Scottsdale, Arizona, Nov. 1996.
- [3] A. Bilgin, G. Zweig, and M. W. Marcellin. Efficient lossless coding of medical image volumes using reversible integer wavelet transforms. In *Proceedings of the Data Compression Conference*, pages 428–437. IEEE Computer Society, 1998.
- [4] G. R. Canta and G. Poggi. Kronecker-product gain-shape vector quantization for multispectral and hyperspectral image coding. *IEEE Trans. on Image Processing*, 7(5):668–678, May 1998.
- [5] B. R. Epstein, R. Hingorani, J. M. Shapiro, and M. Czigler. Multispectral KLT-wavelet data compression for Landsat Thematic Mapper images. In *Proceedings of Data Compression Conference*, pages 200–208, Snowbird, Utah, Mar. 1992.
- [6] G. Fernandez and C. M. Wittenbrink. Coding of spectrally homogeneous regions in multispectral image compression. In *Proc. IEEE Int. Conf. Image Processing II*, pages 923–926, Lausanne, Switzerland, Sept. 1996.
- [7] M. Finelli, G. Gelli, and G. Poggi. Multispectral-image coding by spectral classification. In *Proc. IEEE Int. Conf. Image Processing II*, pages 605–608, Lausanne, Switzerland, Sept. 1996.
- [8] S. Gupta and A. Gersho. Feature predictive vector quantization of multispectral images. *IEEE Trans. on Geoscience and Remote Sensing*, 30(3):491–501, May 1992.
- [9] A. Habibi and A. S. Samulon. Bandwidth compression of multispectral data. In *Proc. SPIE Conf. on Efficient Transmission of Pictorial Information*, pages 23–35, 1975. vol. 66.
- [10] S. K. Holland, A. Zavaljevsky, A. Dhawan, R. S. Dunn, and W. S. Ball. Multispectral magnetic resonance image synthesis, 1998. CHMC Imaging Research Center website – <http://www.chmcc.org/departments/misc/ircmsmris.htm>.

- [11] J.-H. Hu, Y. Wang, and P. T. Cahill. Multispectral code excited linear prediction coding and its application in magnetic resonance images. *IEEE Trans. on Image Processing*, 6(11):1555–1566, Nov. 1997.
- [12] F. H. Imai and R. S. Berns. High-resolution multi-spectral image archives: A hybrid approach. In *Proc. Sixth Color Imaging Conference: Color Science, Systems, and Applications*, pages 224–227, Scottsdale, Arizona, Nov. 1998.
- [13] T. Keusen. Multispectral color system with an encoding format compatible with the conventional tristimulus model. *Journal of Imaging Science and Technology*, 40(6):510–515, November/December 1996.
- [14] J. Lee. Quadtree-based least-squares prediction for multispectral image coding. *Optical Engineering*, 37(5):1547–1552, May 1998.
- [15] C. Mailhes, P. Vermande, and F. Castanie. Spectral image compression. *Journal of Optics (Paris)*, 21(3):121–132, 1990.
- [16] N. Pendock. Reducing the spectral dimension of remotely sensed data and the effect on information content. In *Proc. of International Symposium on Remote Sensing of Environment*, pages 1213–1222, 1983. no. 17.
- [17] J. A. Richards. *Remote Sensing Digital Image Analysis*. Springer-Verlag, Berlin, Germany, 1986.
- [18] R. E. Roger and M. C. Cavenor. Lossless compression of AVIRIS images. *IEEE Transactions on Image Processing*, 5(5):713–719, 1996.
- [19] J. A. Saghri, A. G. Tescher, and J. T. Reagan. Practical transform coding of multispectral imagery. *IEEE Signal Processing Magazine*, 12(1):32–43, Jan. 1995.
- [20] D. Saunders and J. Cupitt. Image processing at the national gallery: The vasari project. Technical Report 14:72, National Gallery, 1993.
- [21] K. Sayood. Data compression in remote sensing applications. *IEEE Geoscience and Remote Sensing Society Newsletter*, pages 7–15, Sept. 1992.
- [22] G. Sharma and H. J. Trussell. Digital color imaging. *IEEE Trans. on Image Processing*, 6(7):901–932, July 1997.
- [23] S. R. Tate. Band ordering in lossless compression of multispectral images. In *Proceedings of the Data Compression Conference*, pages 311–320. IEEE Computer Society, 1994.
- [24] A. G. Tescher, J. T. Reagan, and J. A. Saghri. Near lossless transform coding of multispectral images. In *Proc. of IGARSS'96 Symposium*, pages 1020–1022, May 1996. vol. 2.
- [25] D. Tretter and C. A. Bouman. Optimal transforms for multispectral and multilayer image coding. *IEEE Trans. on Image Processing*, 4(3):296–308, Mar. 1995.

- [26] Y. T. Tse, S. Z. Kiang, C. Y. Chiu, and R. L. Baker. Lossy compression techniques of hyperspectral imagery. In *Proc. of IGARSS'90 Symposium*, pages 361–364, 1990.
- [27] J. Vaisey, M. Barlaud, and M. Antonini. Multispectral image coding using lattice vq and the wavelet transform. In *Proc. IEEE Int. Conf. Image Processing II*, Chicago, Illinois, Oct. 1998.
- [28] V. D. Vaughn and T. S. Wilkinson. System considerations for multispectral image compression designs. *IEEE Signal Processing Magazine*, 12(1):19–31, Jan. 1995.
- [29] J. Wang, K. Zhang, and S. Tang. Spectral and spatial decorrelation of Landsat-TM data for lossless compression. *IEEE Transactions on Geoscience and Remote Sensing*, 33(5):1277–1285, 1995.
- [30] X. Wu, W. K. Choi, and N. D. Memon. Context-based lossless inter-band compression. In *Proceedings of the IEEE Data Compression Conference*, pages 378–387. IEEE Press, 1998.