

Band Ordering in Lossless Compression of Multispectral Images

(Extended Abstract)*

Stephen R. Tate†

Department of Computer Sciences
University of North Texas
P.O. Box 13886
Denton, TX 76203-6886

E-mail: srt@cs.unt.edu

1 Introduction

Multispectral satellite images require enormous amounts of space, and with NASA's project EOS (the Earth Observing System) data will be generated at an unprecedented rate. The estimates are that over a terabyte (10^{12} bytes) of data will be generated every day by the EOS satellites, most of it being multispectral image data. Largely due to this fact, a lot of attention has recently been focussed on compression of multispectral images [3, 8, 11]. However, most of the compression methods that exploit spectral as well as spatial redundancy have been lossy compression algorithms, and for archival storage and for certain applications it is important to use lossless compression in order to preserve all of the data that is collected.

In this paper, we study lossless compression of multispectral images. Spectral redundancy is extracted by coding each band of the multispectral image by making use of a second "prediction band". In much the same way that standard single-image lossless compression is separated into the two separate components of prediction and coding, we divide the lossless compression of multispectral images into three components: band ordering, prediction, and coding. The new stage, band ordering, refers to selecting a permutation of the bands in which the bands that are coded first act as good predictors for the later bands in the ordering. The band ordering phase is independent of the other phases of the compressor, and the computational problems associated with this phase are identified and studied in this paper. In particular, given any particular predictor and coder, we give an efficient algorithm for computing the optimal ordering of the bands for compression performance. For an n -band multispectral image, our algorithm runs in $O(n^2)$ time, which is a vast improvement over the obvious $\Omega(n!)$ solution. In timing tests performed on a DEC

*The full version of this paper is available as University of North Texas Technical Report N-93-003.

†This research was supported by NASA subcontract 550-63 of prime contract NAS5-30428, and a substantial part of the experiments were conducted when the author was affiliated with Duke University and visiting the CESDIS office at the Goddard Space Flight Center.

Alpha workstation, our algorithm computed the optimal ordering from the input matrices (see Section 2) of a 12 band image in 1/100 of a second, whereas the naive algorithm took 87 minutes.

Unfortunately, in some cases the optimal ordering may not be a practical archival format for large numbers of bands (see Section 3 for more information on this), so we consider the natural restriction on band orderings that overcomes this problem. We refer to the restricted ordering as the *optimal compression order with bounded prediction*. Unfortunately, this new problem is shown to be NP-hard except in the most trivial case, so is computationally infeasible.

Next, we examine various possibilities for the prediction phase for multispectral data. Most predictors for image data use some form of linear prediction, and most of these use fixed coefficients for the prediction (or a small set of fixed coefficients as in lossless JPEG compression). With multispectral images it is vital that the linear prediction coefficients be chosen adaptively, and different methods for doing this are discussed in Section 4.

The coder for multispectral images is the simplest part, and is discussed in Section 5. We investigate the use of various configurations of an arithmetic coder, using data within both the current band and the previous, "prediction band", as context for the coder.

Finally, in Section 6 we present some experimental results using the techniques and algorithms presented in this paper on real multispectral data obtained from NASA. The experiments show that the compression methods developed in this paper can have a substantial practical impact on compression performance. In addition, our experiments show that a fixed band ordering can achieve almost optimal performance on a large class of data; therefore, using a pre-computed band ordering can give a very fast, high-performance compression system for multispectral images.

2 Terminology and Notation

Just as single-image data can be represented by a two-dimensional array $S[x, y]$, the data of a multispectral image may be represented as a three-dimensional array, with entry $M[b, x, y]$ representing the pixel from in band b , row x , column y . In traditional single-image compression, when coding pixel $S[x, y]$ the previously coded neighboring pixels may be used to predict the value of the current pixel. Likewise, when coding pixel $M[b, x, y]$ for a multispectral image, the previously coded neighboring pixels may be used for prediction, but we also allow neighboring pixels to come from a previously coded band. In other words, for any band b , we choose a *prediction band* $p(b)$ and use pixel values from band $p(b)$ in the encoding of $M[b, x, y]$.

The compression methods described in this paper all code the pixels of a band in scan order. At any point in time, the pixel $M[b, x, y]$ that the compression algorithm is predicting or coding is called the "current pixel", which comes from the "current band" (band b). As mentioned above, we also designate a separate band $p(b)$ as the "prediction band" for the current band. Just as in single-image compression, all pixel values used in prediction must be coded before the current pixel, including pixel values from the prediction band. Thus we may use the prediction band's pixel location corresponding to the current pixel (i.e., $M[p(b), x, y]$) if and only if band $p(b)$ is coded before band b . Clearly, the ordering of the bands determines what bands may be selected as the prediction band for the current band, and determining an ordering

of the bands that maximizes compression over all possible orderings is an important problem.

For any given predictor/coder pair, and two bands b_1 and b_2 , we can define the following two values:

B_{b_1, b_2} is the compressed size of band b_2 if band b_1 is used for prediction, but $M[b_1, x, y]$ is *not* used in the prediction of $M[b_2, x, y]$. In other words, this is the compressed size for band b_2 using b_1 for prediction, even if band b_2 is encoded *before* band b_1 .

A_{b_1, b_2} is the compressed size of band b_2 if band b_1 is used for prediction, and the current pixel location from band b_1 is used in the prediction. This is the compressed size possible for band b_2 if band b_2 is encoded *after* band b_1 .

We will represent the A and B values for a multispectral image as matrices. In the next section, we will discuss various ways of choosing a band ordering when given matrices A and B .

3 Determining Band Compression Order

In this section, we examine the following problem: In what order should we compress the bands of a multispectral image to achieve the best compression? For example, if band 2 acts as a good predictor for band 1, substantial savings may be achieved by compressing band 2 first. We will show that, given the matrices as defined in the previous section, the optimal ordering of the bands can be found very efficiently. For multispectral data with a small number of bands, this result shows that optimal ordering is both simple and practical.

However, for data with a large number of spectral bands (such as the 210 band AVIRIS data), the optimal ordering can be computed, but may not be what is desired. In particular, given a compressed file containing all 210 bands of AVIRIS data, it may be desirable to be able to extract one particular band from the compressed data. Unfortunately, the band dependencies of the optimal ordering may be a single, long chain, which would mean that *all* bands would have to be uncompressed in order to uncompress the last single band in the chain. A good solution to this problem would be to partition the bands into small sets, each containing at most b bands, and then when compressing one band, the only other bands that may be used for prediction are those within the same set. With this method at most b bands would have to be uncompressed in order to extract a single band from the compressed file. For a particular partition size bound b , we will call the optimal solution to this problem the *optimal compression order with bounded prediction*.

We will show that finding the optimal compression order with bounded prediction is easy if $b = 2$, but is NP-hard for any $b \geq 3$.

3.1 Unconstrained Optimal Ordering. In this section we consider the problem of finding an optimal compression order, with no constraints on how many bands are necessary for uncompression of a single band. This ordering is the best possible ordering for encoding bands, and is practical when either there are a small number of bands (such as the 5 bands of CZCS data or the 7 bands of TM data), or if extraction of a single band is never needed (an archive of AVIRIS data may require uncompression of all 210 bands when it is accessed).

We will transform the problem of finding an optimal compression ordering into a problem on weighted directed graphs which has a known efficient solution. The directed graph $G = (V, E)$ is constructed so that each edge $E_{i,j}$ has a corresponding weight $W_{i,j}$ that represents the compression savings attainable by compressing band i before band j . To construct this graph, define $B_{\min,j}$ to be

$$B_{\min,j} = \min_{1 \leq i \leq n} B_{i,j}.$$

$B_{\min,j}$ represents the minimum compressed size of band j if the current pixel location is not used from any other band. In particular, a compressed size of $B_{\min,j}$ is attainable for band j regardless of the order in which the bands are compressed.

Of course, if the bands are ordered such that band i is encoded before band j , then the current pixel location from band i can be used to predict the current pixel of band j . In some cases, the resulting compressed size may be substantially smaller than the size that was attainable without considering ordering. We define the edge weights of our graph in terms of the savings possible by encoding band j using the current pixel location from band i for prediction. In particular,

$$W_{i,j} = \begin{cases} 0 & \text{if } A_{i,j} \geq B_{\min,j}, \\ B_{\min,j} - A_{i,j} & \text{otherwise.} \end{cases}$$

The problem of finding an optimal compression order is equivalent to finding a maximum weight directed spanning forest of G . The resulting spanning forest defines a partial order on the vertices, and since the spanning forest is maximum weight, the partial order represents the maximum compression savings of any band ordering, or the optimal compression order of the bands of the multispectral image. It is known that a maximum weight directed spanning forest (also known as an "optimal branching") can be found in $O(|V| \log |E| + |E|)$ time on sparse graphs and $O(|V|^2)$ time on dense graphs [13, 2], so it follows that an optimal compression order can be found in $O(n^2)$ time. We summarize the result of this section in the following theorem.

Theorem 3.1 Given values $A_{i,j}$ and $B_{i,j}$ for all pairs of bands (i, j) from an n -band multispectral image, the optimal compression ordering can be computed in $O(n^2)$ time.

3.2 Optimal Ordering with Bounded Prediction. As mentioned at the beginning of this section, the optimal ordering is not always practical. In particular, if it is necessary to be able to extract and uncompress a single band of the image efficiently, the optimal compression ordering does not always make this possible. Consider a case where the directed spanning tree obtained as in the previous subsection is actually a long directed path. In a 210 band data set, it would be possible for band 210 to be coded using band 209 for prediction, band 209 to be coded using band 208, etc. Thus, in order to extract band 210, all of the 209 other bands would have to be uncompressed in addition to band 210.

For data sets with a large number of bands, a reasonable approach would be to partition the bands into small sets (containing, say, b bands each), and then compress each set using its optimal compression ordering. In this case, if a single band needs to be extracted, only the other bands in the same set as the requested band will possibly need to be uncompressed. We call such an ordering the *optimal compression order with bounded prediction*, and prove the following theorem regarding the complexity of computing such an ordering.




Figure 1: Neighboring pixels used in prediction phase.

Theorem 3.2 Given values $A_{i,j}$ and $B_{i,j}$ for all pairs of bands (i,j) from an n -band multispectral image, the problem of finding an optimal compression order with bounded prediction is NP-hard if $b \geq 3$, but can be computed in $O(n^3)$ time if $b = 2$.

Proof: The NP-hardness result is proved by reducing the PARTITION INTO PATHS OF LENGTH TWO to a decision problem dealing with the existence compression order with bounded prediction. The algorithm for $b = 2$ is by a reduction to maximum weight graph matching. Details of these proofs can be found in the full paper. ■

4 Prediction for Multispectral Images

The vast majority of lossless image compression algorithms are made up of a prediction stage in which pixel values are predicted, and a coding stage in which the difference between the actual pixel value and the predicted value is coded (for recent examples, see [9, 15]). The predictor typically uses a linear function to predict the current pixel value, based on the values of previously coded pixels (the “neighborhood” of the current pixel). In compression algorithms that code the pixels in scan-order, the neighborhood is usually a subset of the 4-neighborhood shown in Figure 1(b), where the current pixel is marked with the solid circle. For example, we may predict the current pixel by computing the average of its two nearest neighbors: $(A + C)/2$. The lossless JPEG scheme has 8 different possible linear prediction functions of A , B , and C , and is free to use any of those eight [15].

Since the relationship between different bands of a multispectral image is not known ahead of time, we must find a good prediction function by examining the image data. In particular, in our experiments, we consider both 2-neighborhoods and 4-neighborhoods for pairs of bands (see Figure 1), and use the actual image data to determine the coefficients of the linear function of neighborhood values that minimizes the variance of the prediction error. Determining the coefficients in this way is similar (although not exactly the same) to standard linear prediction methods (see for example [10]), and can be done using least squares approximation algorithms.

The experiments of this paper were done with a predictor that finds linear coefficients for the 2-neighborhood and 4-neighborhood of the current pixel location, where values from both the current band and the prediction band are used in the prediction function. Both the 2-neighborhood and 4-neighborhood are then used for a second pass on the data using the computed coefficients, and whichever resulted in a smaller compressed size is reported. The program computes the coefficients both with and without the current pixel location in the prediction band, and the resulting compressed sizes are used for the A and B matrix values, respectively.

5 Coding for Multispectral Images

The coder has the responsibility of turning the predicted values into an output bitstream. Various approaches have been taken in different compression systems, including Huffman coding, Rice coding, and arithmetic coding (or variants such as Q-coding). A good explanation of coding in general, including Huffman coding and arithmetic coding, can be found in [1]; Rice coding is described in [12], which is similar to Golomb coding as described in [6]. Huffman coding performs poorly on small alphabets or when the prediction phase produces very small (or low entropy) errors, and Rice coding performs poorly when the errors are not generated according to a Laplacian distribution. While the distribution of errors is usually close to Laplacian, it was discovered in our experiments that CZCS data produced very markedly non-Laplacian error distributions (the other sensors did produce prediction errors that were more or less Laplacian).

Our final choice for a coder was an adaptive arithmetic coder, with three neighboring prediction errors used as context for the coder (the two nearest neighbors in the current band, and the current pixel location from the prediction band when possible). While the actual prediction errors could easily be used for neighboring pixels in the same band, the prediction error for the current location in the prediction band was estimated by using the simple single-image linear prediction formula $X - \lfloor (A + C)/2 \rfloor$, where X is the prediction band pixel in the current location, and A and C are neighboring pixels in the prediction band as defined by Figure 1(a). The prediction errors were bucketed (similar to [14]) into “zero”, “small” (1-2), “medium” (3-8), and “large” (over 9), with the sign of the error also used in the context. Since the context is computed from 3 prediction errors with 7 possible values each, the coder used 343 contexts.

The actual prediction errors are encoded by first encoding a tag denoting negative, positive, or zero error, and then encoding the error magnitude two bits at a time. This method substantially reduces the amount of storage required for maintaining the error frequency statistics, since frequency tables only need to be maintained for a 4 symbol alphabet (the 2-bit blocks). For example, in an image with 16 bit pixels, the error can range from -65,535 to +65,535, so directly coding the prediction error would require 131,071 frequency counts per context (for a total of almost 4.5 million counts over all contexts). On the other hand, keeping statistics only for the error sign and the two-bit blocks of the error magnitude requires only 35 frequency counts per context (for a total of 12,005 frequency counts over all contexts). It is very significant to note that when the errors are Laplacian distributed (the majority of cases), encoding the errors in this significantly more efficient way is fully as efficient as keeping separate frequency counts for all possible error values. Furthermore, when the errors are not Laplacian distributed (such as for the CZCS data), this method greatly outperforms true Laplacian coders such as the Rice coder, and performs almost as well as the arithmetic coder with full frequency counts.

6 Experimental Results

We performed many experiments using the algorithms described in this paper and real multispectral satellite data. Our experiments show that band reordering can substantially improve compression performance for certain data sets, and that using a fixed, pre-computed band ordering works well for entire classes of data, giving a




Figure 2: Compression ratio for some sample data files

very fast and high-quality compression system for multispectral images. The data for our experiments came from several types of sensors: images "neworl", "ridgely", and "washdc" come from the 7-band Thematic Mapper (TM); images "czcs" and the series "czcs1-7" come from the 6-band Coastal Zone Color Scanner (CZCS); images "avhrr1" and "avhrr2" are from the 6-band Advanced Very High Resolution Radiometer (AVHRR); and image "aviris" comes from the 210-band Airborne Visible Infrared Imaging Spectrometer (AVIRIS).

To compute the A and B matrices, the matrix entries were computed in parallel on a CM-5 parallel computer. Once these values were computed, the A and B matrices were fed into implementations of the algorithms described in section 3 for finding the optimal band ordering, and for finding the optimal paired ordering. The resulting compression ratio for the various sample images is shown in Figure 2. The compression ratio using our prediction and coding phases, but with no inter-band prediction, is shown as "Independent Coding" (meaning that the bands are coded independently of other bands). In addition, the compression ratios using lossless JPEG (with each band coded independently of the others) and the UNIX compress command are included for comparison purposes. The performance of lossless JPEG compares well with the best lossless single-image compressors. The UNIX compress results are included since many installations currently choose this method for compressing their multispectral data sets, and these results emphasize how poor this choice is.

As can be seen from Figure 2, the effect of band ordering seems to rely heavily on the source of the data. In particular, large gains were seen by selecting the optimal band ordering for the AVIRIS and CZCS data, moderate gains were seen in the TM data sets, and almost nothing was gained by band-ordering in the AVHRR data. What

this suggests is that the individual bands of the AVHRR data are more independent than the bands of the other sensors, which is supported by the technical specifications of the sensor.

It is interesting to note that for data sets in which band ordering made a difference, the performance of paired ordering is only a moderate improvement over independent coding. This implies that the band dependencies are of a more global nature than pairs of bands. In light of our proof that computing the optimal partitioning into blocks of size three or more is NP-hard, it seems that close-to-optimal performance is not possible without using the absolute optimal ordering and risking long dependence chains. To examine the importance of selecting a good partitioning, we divided the 210 bands of AVIRIS data into blocks of 5 adjacent bands per block, without regard for the content of those bands. In other words, the blocks consisted of bands 1-5, bands 6-10, bands 11-15, etc. Within each block, the bands were optimally ordered, but there is no reason to believe that the chosen blocks bear any resemblance to the optimal partitioning. The results are shown in the following table.

Blocking Method	Compression Ratio
No blocking (optimal order)	3.53
Blocks of 5 adjacent bands	3.39
Optimal partitioning of 2-band blocks	3.21
Block size 1 (independent coding)	2.98

It can be seen from this table that simply picking blocks of adjacent bands for the partitioning seems to work well for AVIRIS data. While the compression ratio does not quite match the performance of the optimal ordering, the added benefit of being able to efficiently extract a single band from the compressed archive may be worth sacrificing some compression.

The next experiment we tried was to determine how much the band ordering depends on the sensor, and how much it depends on the individual data sets. For this test, we obtained seven independent data sets from a single sensor (CZCS) and ran all of the previous tests, in addition to the following test: for any particular data set, the optimal compression ordering for the *other* six data sets is computed, and then that ordering is used for the data set being tested. The results, with the last type of test labeled as "Pre-computed Order", are shown in Figure 3.

As can be seen from the figure, in most cases the order computed from the other images gave a compression ratio that was competitive with the optimal order computed specifically for that image. This suggests that, at least for CZCS data, the band ordering is more a function of the sensor than of the individual image; therefore, a reasonable approach to a production compression package would be to compute the optimal ordering from a large, representative set of images, and hard code this band ordering into the compressor. To compress a multispectral image with this approach, no band-ordering has to be computed, and the *A* and *B* matrices do not need to be computed, so the compression is very efficient. The results of this last experiment show that this approach should yield almost optimal compression.

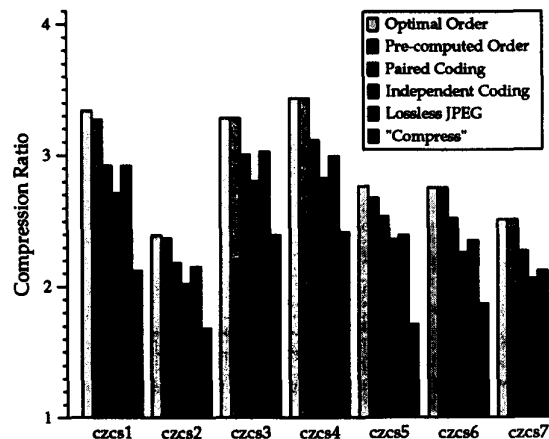


Figure 3: Compression ratio for sample CZCS data files

7 Conclusions

In this paper, we have extensively studied the benefits and problems associated with reordering the bands of a multispectral image for performing lossless compression. Under a reasonable model of inter-band prediction, we have shown how to efficiently compute the optimal compression band ordering. In addition, we have formalized the restrictions that arise when bands need to be extracted individually from a compressed archive, and have shown that computing the optimal ordering under these restrictions is NP-hard, except in the most simple case.

Our experiments show that for certain data sources (such as CZCS or AVIRIS satellite imagery), the band ordering algorithms of this paper can substantially improve compression performance. Interestingly, some data sources (such as the AVHRR data) don't benefit from the techniques of this paper at all, implying that the bands are independent, and can be treated separately (as if the data was not a multispectral image at all). By computing the band ordering as described in this paper, and evaluating the compression results, it is easy to determine if band re-ordering is going to have a substantial benefit for a particular data set.

While computing the optimal band ordering may be efficient (i.e., polynomial time), it is certainly not fast, requiring that all pairs of bands be tested as prediction-coding pairs. An alternative is to compute a fixed band ordering for a class of data, and use that fixed ordering when coding all data sets in that class. For instance, a fixed "CZCS ordering" could be pre-computed and used for all CZCS data sets. This has the benefit of being extremely fast (similar in speed to simply ignoring inter-band correlation); furthermore, our experiments have shown that for CZCS data, the results are almost indistinguishable from the optimal ordering results. Thus, this last approach of using a fixed band reordering seems to be a very promising lossless multispectral compression method.

8 Acknowledgements

The author would like to thank Doreen Revis for her help with various aspects of the CM-5, Jim Tilton and Manohar Mareboyana for supplying a wide range of test data, Gene Feldman for supplying the CZCS data, and Mary James for supplying the AVHRR data.

References

- [1] T. C. Bell, J. G. Cleary, and I. H. Witten. *Text Compression*, Prentice Hall, Englewood Cliffs, NJ, 1990.
- [2] P. M. Camerini, L. Fratta, and F. Maffioli. "A Note on Finding Optimum Branchings," *Networks*, Vol. 9, pp. 309-312, 1979.
- [3] B. R. Epstein, R. Hingorani, J. M. Shapiro, and M. Czigler. "Multispectral KLT-Wavelet Data Compression for Landsat Thematic Mapper Images," *Proc. of the Data Compression Conference (DCC '92)*, pp. 200-208, 1992.
- [4] H. N. Gabow. "An Efficient Implementation of Edmonds' Algorithm for Maximum Matching on Graphs," *JACM*, Vol. 23, No. 2, pp. 221-234, 1976.
- [5] H. N. Gabow and R. E. Tarjan. "Faster Scaling Algorithms for General Graph-Matching Problems," *JACM*, Vol. 38, No. 4, pp. 815-853, 1991.
- [6] R. G. Gallager and D. C. Van Voorhis. "Optimal Source Codes for Geometrically Distributed Integer Alphabets," *IEEE Transactions on Information Theory*, Vol. IT-21, pp. 228-230, March 1975.
- [7] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*, W. H. Freeman and Company, New York, 1979.
- [8] S. Gupta and A. Gersho. "Feature Predictive Vector Quantization of Multispectral Images," *IEEE Transactions on Geoscience and Remote Sensing*, Vol. 30, No. 3, May 1992.
- [9] P. G. Howard and J. S. Vitter. "New Methods for Lossless Image Compression Using Arithmetic Coding," *Information Processing and Management*, Vol. 26, No. 6, pp. 765-779, Nov. 1992.
- [10] J. Makhoul. "Linear Prediction: A Tutorial Review," *Proceedings of the IEEE*, Vol. 63, No. 4, pp. 561-580, April 1975.
- [11] T. Markas and J. Reif. "Multispectral Image Compression Algorithms," *Proc. of the Data Compression Conference (DCC '93)*, pp. 391-400, 1993.
- [12] R. F. Rice. "Some Practical Universal Noiseless Coding Techniques", *JPL Publication 79-22*, Jet Propulsion Laboratory, Pasadena, CA, Mar. 1979.
- [13] R. E. Tarjan. "Finding Optimum Branchings," *Networks*, Vol. 7, pp. 25-35, 1977.
- [14] S. Todd, G. Langdon, J. Rissanen. "Parameter Reduction and Context Selection for Compression of Gray-scale Images", *IBM J. Res. Develop.*, Vol. 29, No. 2, pp. 188-193, March 1985.
- [15] G. K. Wallace. "The JPEG Still Compression Standard", *Communications of the ACM*, Vol. 34, No. 4, pp. 30-44, April 1991.