

MSE optimal bit-rate allocation in JPEG2000 part 2 compression applied to a 3-D data set

Olga M. Kosheleva ^a, Sergio D. Cabrera ^a, Bryan E. Usevitch ^a, Alberto Aguirre ^a, and
Edward Vidal, Jr. ^b

^a Dept. of Electrical and Computer Engineering, University of Texas at El Paso, El Paso, TX,
79968-0523, USA

^b United States Army Research Laboratory, WSMR, NM, USA 88002-5501

ABSTRACT

A bit rate allocation (BRA) strategy is needed to optimally compress three-dimensional (3-D) data on a per-slice basis, treating it as a collection of two-dimensional (2-D) slices/components. This approach is compatible with the framework of JPEG2000 Part 2 which includes the option of pre-processing the slices with a decorrelation transform in the cross-component direction so that slices of transform coefficients are compressed. In this paper, we illustrate the impact of a recently developed inter-slice rate-distortion optimal bit-rate allocation approach that is applicable to this compression system. The approach exploits the MSE optimality of all JPEG2000 bit streams for all slices when each is produced in the quality progressive mode. Each bit stream can be used to produce a rate-distortion curve (RDC) for each slice that is MSE optimal at each bit rate of interest. The inter-slice allocation approach uses all RDCs for all slices to optimally select an overall optimal set of bit rates for all the slices using a constrained optimization procedure. The optimization is conceptually similar to Post-Compression Rate-Distortion optimization that is used within JPEG2000 to optimize bit rates allocated to codeblocks. Results are presented for two types of data sets: a 3-D computed tomography (CT) medical image, and a 3-D meteorological data set derived from a particular modeling program. For comparison purposes, compression results are also illustrated for the traditional log-variance approach and for a uniform allocation strategy. The approach is illustrated using two decorrelation transforms (the Karhunen Loeve transform, and the discrete wavelet transform) for which the inter-slice allocation scheme has the most impact.

Keywords: bit-rate allocation, rate-distortion optimization, JPEG2000, 3-D data compression, volumetric data compression, medical image compression, meteorological data compression

1. INTRODUCTION

JPEG2000 Part 2 has the capability to compress 3-D data by treating it as separate 2-D slices.¹ Alternatively, the separate slices can be transform coefficients obtained from the data that has undergone a decorrelation transformation as a pre-processing step in the third (cross-component or cross-slice) direction (see Figure 1 describing this approach to 3-D compression). Typically, the Karhunen-Loève Transform (KLT) or the Discrete Wavelet Transform (DWT) are used before the JPEG2000 coder is applied. The inverse KLT or inverse DWT are used after the JPEG2000 decoder. The question that arises is how to optimally allocate bits to the separate slices since JPEG2000 does not require that any specific method be used. In more precise terms, the problem is: given a desired average bit rate, assign bit rates to the individual slices so that a distortion metric, such as the Mean Squared Error (MSE), is minimized for the whole 3-D data cube. It can be shown that the solution of this (direct) problem also gives a solution to the inverse problem: given the upper bound on the MSE distortion metric, find the optimal bit rate allocation for the slices such that the average bit rate is minimized. The traditional approach uses a standard high rate approximation to the MSE rate-distortion, that in effect leads to bit rate allocation based on the logarithms of the variances of the corresponding slices.^{2,3}

Further author information: (Send correspondence to Sergio D. Cabrera)
E-mail: cabrera@ece.utep.edu, Telephone: (915)747-5470, FAX (915)747-7871

We recently⁴⁻⁶ solved the problem of inter-slice rate-distortion optimal bit rate allocation for 3-D JPEG2000 compression. The approach is called here the Rate-Distortion Optimal (RDO) method since it is inspired by the Post-Compression Rate-Distortion (PCRD) approach used within JPEG2000.⁷ In JPEG2000, PCRD optimization is used to select the optimal truncation points for the bit streams of the individual code-blocks⁷ associated with a single 2-D image or slice. Similar to the PCRD intra-slice allocation scheme, the RDO approach to inter-slice bit allocation gives the smallest MSE solution. Since this approach makes use of actual (experimentally obtained) rate-distortion curves, its main disadvantage is implementation complexity. We have devised other more efficient approaches⁴⁻⁶ not discussed here. Next we give a brief discussion of the JPEG2000 baseline coder.

The JPEG2000 Part 1 baseline or simply JPEG2000⁷ has brought a new paradigm to image compression standards. It provides among other advantages superior low bit-rate performance, bit rate scalability and progressive transmission by quality or resolution. Quality scalability is achieved by dividing the wavelet transformed image into codeblocks B_i . Each codeblock is encoded into an embedded representation yielding minimum distortion $D_i(n_i)$ as a function of bit rate $R_i(n_i)$ for each given truncation point n_i . After having encoded all codeblocks, a post-processing operation determines where each code-block's embedded stream should be truncated in order to meet a pre-defined bit-rate (or distortion) bound for the whole image. This bitstream rescheduling module is referred to as the Tier 2. It establishes a multi-layered representation of the final bitstream, guaranteeing an optimal performance at several bit rates or resolutions. The Tier 2 component optimizes the truncation process, and tries to reach the desired bit-rate while minimizing the introduced distortion, utilizing Lagrangian rate allocation principles. The procedure is known as Post-Compression Rate-Distortion (PCRD) optimization⁷ and the basic theory behind it is extensively discussed in literature.⁸ The following is a brief description of the corresponding optimization problem solved using the Lagrange multipliers approach.

Assuming that the overall distortion metric is additive, i.e.,

$$D = \sum_{i=1}^N D_i(n_i), \quad (1)$$

it is desired to find the optimal selection of bit stream truncation points n_i^λ such that the overall distortion metric is minimized subject to a constraint $R = \sum_i R_i(n_i) \leq R^{max}$. To solve the problem, the Lagrange multiplier method is used. It leads to the unconstrained optimization problem

$$Q = (D(\lambda) + \lambda \cdot R(\lambda)) \rightarrow \min. \quad (2)$$

The resulting objective function Q depends on N variables n_i , but can be represented as a sum of N terms

$$Q_i = D_i(n_i) + \lambda \cdot R_i(n_i).$$

Therefore, to minimize the sum, we must find, for each code-block i , truncation point n_i^λ that minimizes the corresponding term Q_i . This approach follows from the general result proven in.⁸ The determination of the N optimal truncation points for any given λ is performed efficiently based on the experimental information about rate-distortion dependence collected during generation of each code-block's embedded bit-stream. Only convex hull rate-distortion points, (i.e. the largest set of truncation points for which the corresponding distortion-rate slopes are strictly decreasing in magnitude as bit-rate increases) are used in the optimization. Basically, the algorithm finds the truncation points n_i^λ , where each rate-distortion slope $S_i(n_i^\lambda) = \Delta D_i(n_i^\lambda) / \Delta R_i(n_i^\lambda)$ is closest to the fixed λ .

Since this algorithm has to be repeated for several values of λ , the distortion-rate slopes are pre-calculated, and only the set of acceptable truncation points N_i (defining the convex hull) is stored together with the corresponding rates and distortion-rate slopes.

As mentioned earlier, the PCRD optimization approach is the motivation for the inter-slice RDO method described in this paper. The RDO bit allocation method is evaluated and compared here by applying it to 3-D Meteorological (Met) data and 3-D Medical data. The specific Met data set used was generated by the

Battlescale Forecast Model (BFM), which is an analytical model developed by the U. S. Army Research Lab.⁹ The Medical data set used in this study is a helical CT scan of the chest area.

In all our experiments, the default pre-processing is the application of the KLT in the cross-slice direction. Results from this paper will show that for the Met data, the RDO approach allows us to achieve significant improvement in MSE compared to the uniform and the traditional (based on logarithms of variances) inter-slice rate allocation approaches. Similarly, experiments on Medical data show the same trend for KLT pre-processing. In addition, using DWT pre-processing can sometimes produce slightly better results than with KLT pre-processing when both are used in conjunction with the RDO allocation approach.

The remainder of the paper proceeds as follows. In Section 2 we present a brief description of the 3-D data sets used and of the software tools that we have developed for multi-slice compression. Section 3 describes the pre-processing schemes and Section 4 contains general descriptions of the relevant optimization problems and Lagrange Multiplier techniques. Section 5 describes how rate-distortion curves are obtained for each slice and the algorithms used in solving for the optimal bit rate allocation. Experimental results on Met data and conclusions are provided in Section 6.

2. DESCRIPTION OF DATA AND SOFTWARE TOOLS

A model being used today at the Army Research Lab. to generate Met data is the Battlescale Forecast Model (BFM).⁹ Tables 1 shows the range and units of the six Met variables that make up our test data for this and other studies.^{4, 10}

Table 1. Meteorological components

Component	Max	Min	Range	Units
U (wind)	32.9	-18.5	51.4	m/sec
V (wind)	25.9	-18.5	44.4	m/sec
W (wind)	0.37	-0.45	0.82	m/sec
T (temperature)	333.26	270.91	62.35	deg K
Wv (water vapor)	13.81	0.027	13.783	$\frac{\text{gr of water}}{\text{Kg of air}}$
P (pressure)	1019.1	247	772.1	millibars

The BFM uses physics models and actual measured data as boundary conditions to produce floating point data values on a user specified 3-D spatial grid. The BFM data set available for use in this study consists of a 3-D array ("cube") of data for each of six physical variables. For a specific variable *Met* we will call the refer to the corresponding cube as $Met(Z, X, Y)$ which is of dimensions 64x129x129. The first dimension is the vertical height Z , and the other two are X and Y for the two horizontal spatial variables.

The medical data set used in this and other¹¹ studies is a helical CT scan of the chest area (from an anonymous patient), obtained from a HiSpeed CT/i machine from GE Medical Systems at Toronto Hospital, West Div. It is composed of 109 cross-sectional views of 512 by 512 pixels each. The pixel depth is 16 bits with a resolution of 0.676 mm/pixel.

Various software packages were used to perform this study. We have developed a special tool denoted CompressMD to perform, among other tasks, data compression experiments for 3-D data sets in the form illustrated in figure 1 . This software tool under our development has as its main purpose to enable the user to perform JPEG2000 data compression experiments on collections of images and 2-D slices from multi-dimensional data sets (up to 4-D) with appropriate pre- and post-processing . The tool is written in JAVA and can run on many platforms. Other CompressMD features relevant to this and future similar studies include:

- Computation and tabulation of several standard compression distortion metrics.

- An implementation of the Karhunen Love and the discrete wavelet cross-component transforms for pre- and post-processing compatible with JPEG2000 Part 2.
- Various inter-slice rate allocation algorithm.
- Support for TIFF, BMP, PGX and PGM image formats and NetCDF data file formats.
- Visualization of data in 2-D with a user friendly graphical user interface.
- The ability to use different JPEG2000 baseline coders and decoders.

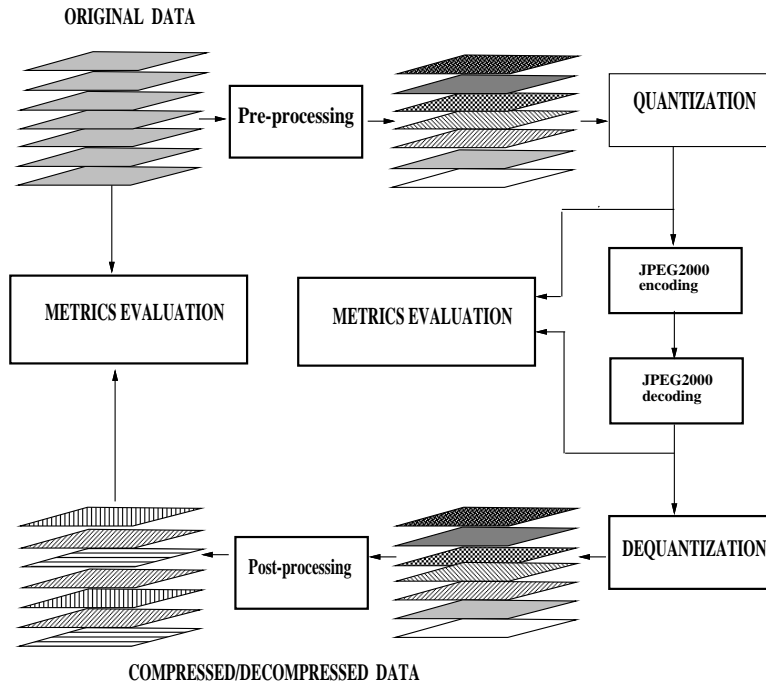


Figure 1. Diagram illustrating 3-D compression on a per-slice basis using JPEG2000.

3. PRE-PROCESSING SCHEMES

In the simplest approach to 3-D compression using a 2-D coder, the volume of data is processed as a set of independent slices without the pre- and post-processing. The disadvantage of this approach is that intra-slice redundancies are left unexploited. In order to take advantage of the third dimension correlation, the JPEG2000 Part 2 standard¹ allows for a cross-component (slice) transform such as the KLT or the DWT which are used in this paper.

Let us first consider KLT pre-processing of N slices. To establish notation we consider the vertical direction vectors at each (x, y) , $\vec{I}(x, y) = (I(x, y, 1), \dots, I(x, y, N))$.

In terms of vector space representation, these values are the coefficients in an N -dimensional space using the standard basis. In the KLT transform representation we are selecting a different orthonormal basis $\vec{e}_1, \dots, \vec{e}_M$ and finding the following representation (after subtracting the mean vector):

$$\vec{I}(x, y) - \vec{I}_0 = a_1(x, y) \cdot \vec{e}_1 + \dots + a_N(x, y) \cdot \vec{e}_N, \quad (3)$$

where vectors $\vec{e}_1, \dots, \vec{e}_N$ are the eigenvectors of the covariance matrix¹² of the ensemble of N -vectors with mean \vec{I}_0 .




Figure 2. Diagram illustrating DWT pre-processing.

Similarly, for the discrete wavelet transform the mean is first subtracted from each data slice before the DWT is computed. Figure 2 illustrates the transform slices for $N = 16$ after 1, 2, and 3 levels of decomposition. It is clear that in this case, the slices of transform coefficients have a natural grouping that is not present in the KLT case. In our examples, the 9/7 Daubechies kernel is used and 5 levels of decomposition are performed.

In both cases the transform coefficient slices are then quantized in order to process them with a JPEG2000 coder, since a baseline implementation is designed to work with fixed point precision data. In our examples, the transformed slices are quantized to $b = 16$ bits using a simple uniform quantizer that ensures a mapping from 0 to 0. Defining $absmax = \max\{|\min(T_z)|, |\max(T_z)|\}$, the quantizer is defined as:

$$Y_z(x, y) = \frac{2^b - 1}{2 \cdot absmax} \cdot (T_z(x, y) + absmax) - 2^{b-1}, Q_z(x, y) = \lfloor Y_z(x, y) + 0.5 \rfloor,$$

4. APPLICABLE OPTIMIZATION TECHNIQUES

The problem of inter-slice optimal bit rate allocation becomes most important when we apply the pre-processing with a decorrelation transform in JPEG2000 Part 2. The diagram in Fig. 3 shows that in most cases, the KLT transformed data slices will have a very large dynamic range of variances or energies organized in decreasing order from first (top) to last (bottom). It is clear that an optimal rate allocation strategy will also need to produce a widely varying, generally decreasing set of bit rates for the KLT slices. If no decorrelation is done, the data domain slices are expected to have similar variances, requiring similar bit rates as illustrated in the figure. Because the KLT concentrates most of the energy into relatively few slices, these first few slices will be allocated significantly more bits than the desired average. On the other hand, the last few slices may actually be completely discarded since zero bits will be allocated to them.

In the *direct* optimization problem, we are given a target average bit rate for the collection of N slices as $R_t = 1/N \cdot \sum_{z=1}^N R_z$. and we want to find, among all possible non-negative bit rate allocations $\{R_1, \dots, R_N\}$ having the desired average, a specific one for which the mean squared error (MSE) distortion attains the smallest possible value. In the *inverse* optimization problem, we are given the required bound on the MSE, and we want to find, among all bit rate allocations that guarantee this compression quality, the allocation $\{R_1, \dots, R_N\}$ for which the average bit rate is the smallest. From the mathematical point of view, the solution of the inverse problem is easy to deduce from the direct problem, and therefore we will focus on the direct optimization problem.

Defining $MSE_z(R_z)$ as the Rate-Distortion Curve (RDC) of the slice with index z , solving the direct constrained optimization problem using the Lagrange multipliers method requires the solution of an unconstrained optimization problem with an objective function:

$$Q \triangleq \sum_{z=1}^N MSE_z(R_z) + \lambda \cdot \sum_{z=1}^N R_z = \sum_{z=1}^N (MSE_z(R_z) + \lambda \cdot R_z) \quad (4)$$

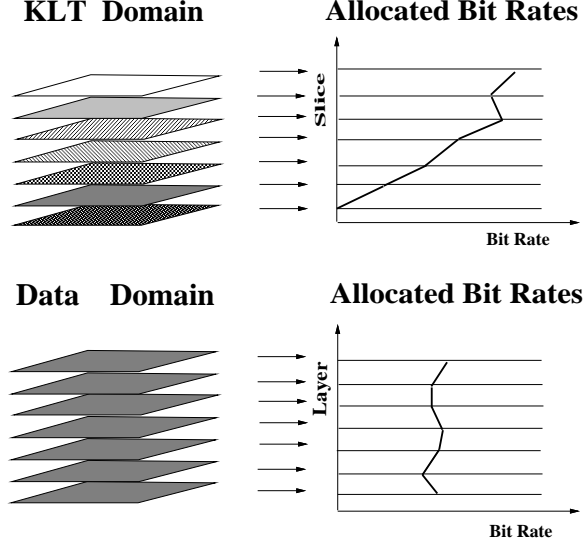


Figure 3. Diagram illustrating the expected nature of optimal bit rate allocation schemes for the KLT and data domains.

As shown in⁸ minimization of Q can be accomplished through individual minimization of each entry

$$(MSE_z(R_z) + \lambda \cdot R_z)$$

of the sum (4). Differentiating and setting to zero, the desired value R_z^o can be obtained as a solution of $MSE'_z(R_z) + \lambda = 0$. Thus, the optimal values R_z^o are such points where all the slopes of the rate-distortion curves are equal $-MSE'_1(R_1^o) = \dots = -MSE'_N(R_N^o) = \lambda$. To deal with nonnegative slopes and λ , the slope is redefined to be the negation of the standard slope. Since MSE is preserved under the (orthonormal) KLT transform, we can use the same approach in the KLT or in the data domain. Thus, we need to have RDCs for each slice; once they are determined it is straightforward to find the optimal values R_z^o . For DWT we need to normalize the subbands prior to the optimization so that the MSE is closely preserved as in the KLT case. Equivalently, we use a weighted MSE in the wavelet domain to approximate the MSE in the data domain.

In the proposed approach the generation of (discrete) RDCs is done experimentally by acquiring pairs $(R_z, MSE_z(R_z))$ for each slice. Then using the RDO approach analogous to the PCRD of JPEG2000, we generate a MSE optimal inter-slice bit rate allocation. To generate experimental data, we select for each slice z several different increasing bit rates $R_z(1), \dots, R_z(T_z)$ in a pre-defined acceptable range (from $RMIN_z$ to $RMAX_z$) and obtain the corresponding $MSE_z(R_z)$ at each point using a JPEG2000 coder and decoder. Once the RDCs are determined, the optimal bit rates R_z^o are determined by finding on all the curves, points of the same slope such that the corresponding bit rate average is the desired target rate.

5. FINDING THE RATE-DISTORTION OPTIMAL ALLOCATION FROM THE RATE-DISTORTION CURVES

To find the optimal bit rate allocation using the RDO method, we first have to find a set of experimentally determined set of points to form a rate-distortion curve for each slice. We then need to select a set of feasible points that define the vertices of the convex hull for each of these RDCs. For the convex hull, the slope of the RDC varies monotonically with respect to the bit rate. Figure 4 shows an example of a convex hull construction (points from RDC that violate the convexity are not included and those that remain are implicitly connected by a straight line).

As discussed in Section 4, the optimal bit rate allocation R_1^o, \dots, R_N^o (for a given target average bit rate Rt) will be such that the corresponding slopes on all the rate-distortion curves will have the same value λ . Therefore, we can solve the problem by searching through different λ and finding the one that leads to the bit rates whose average equals Rt .




Figure 4. U, KLT, example of convex hull.

Since we have only finitely many bit rates on each convex hull, there are only finitely many possible slope values. Thus, for a given λ , generally we will not be able to find the value t_z for which the slope is exactly equal to λ ; the best we can do is to find the neighboring index values t_z and $t_z + 1$ (on the convex hull) such that λ is between the slopes at these two points. The search is done using the bisection algorithm¹³ to find the optimal point t_z for each slice.

To complete the description of this algorithm, we must describe how to select the optimal value of the Lagrange multiplier λ^o . We must select it in such a way that the corresponding average bit rate $R(\lambda^o) = 1/N \cdot \sum_{z=1}^N R_z(\lambda^o)$ is equal to the given target average bit rate Rt .

As we have mentioned earlier, the solution should be such that λ^o is the value of the slope on the RDC of slice z corresponding to the optimal bit-rate value R^o_z for that slice. The slope (magnitude) decreases as the bit rate R_z increases; thus, when λ decreases, the corresponding values $R_z(\lambda)$ increase and therefore, the average bit rate $R(\lambda)$ increases as well. So, the dependence $R(\lambda)$ is monotonically decreasing and again, we use bisection to find the value λ^o for which $R(\lambda^o) = Rt$. We start with $[\underline{\lambda}, \bar{\lambda}]$, where $\underline{\lambda} = 0$ and $\bar{\lambda}$ is the largest slope magnitude that appears on any RDC (clearly, this interval is guaranteed to contain the optimal value of λ).

At each stage of the bisection algorithm, the interval $[\underline{\lambda}, \bar{\lambda}]$ is halved and the algorithm is as follows:^{4, 5}

- first, we find a midpoint

$$\lambda_m \stackrel{\text{def}}{=} \frac{\underline{\lambda} + \bar{\lambda}}{2}; \quad (5)$$

- for this λ_m , we find the bit rates for each slice $R_1(\lambda_m), \dots, R_N(\lambda_m)$ (also using bisection to match slopes on all slices), and then compute the average bit rate $R(\lambda_m)$;
- if $R(\lambda_m) > R$, this means that the desired value of λ is larger than λ_m ; thus, we can take a half-interval $[\lambda_m, \bar{\lambda}]$ as the new interval that is guaranteed to contain λ ;
- if $R(\lambda_m) < R$, this means that the desired value of λ is smaller than λ_m ; thus, we can take a half-interval $[\underline{\lambda}, \lambda_m]$ as the new interval that is guaranteed to contain λ .

We reach convergence and stop the iterations when the bisection interval stops changing.

The following modification of the algorithm was done to avoid problems with transform coefficient slices where zero bits are expected to be allocated to the smallest energy slices. For low target average bit rates, the corresponding λ values encountered near convergence will be large. Due to energy compaction, some of the RDCs of the lowest energy slices do not have sufficiently large slopes (this is true mostly when the KLT is used, see Fig. 5). Therefore, we allocate zero bits to these slices and they are removed from consideration in the allocation algorithm. Clearly, this corresponds to completely discarding these slices except possibly for their mean values which were previously saved. Some details of a typical KLT domain bit rate allocation result for the Met variable U are illustrated in Fig. 6.




Figure 5. Diagram illustrating KLT-specific problem for low bit rates.

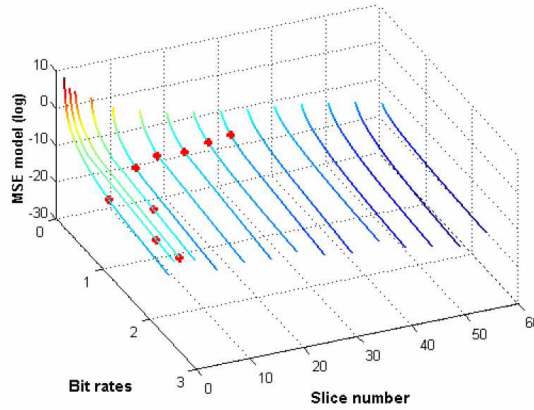


Figure 6. Optimal bit rate allocation for U variable.

6. RESULTS AND CONCLUSIONS

The experiments were performed for several approaches:(1) the traditional (TRAD) approach, based on the logarithms of variances, (2) the RDO approach, providing optimal bit rate allocation, and (3) one bit rate (OBR) approach which assigns equal bit rates to all slices. Clearly, the traditional approach is used without the standard constraint which forces bit rates to be integers. JPEG2000 can achieve most small-enough non-negative fractional bit rates using the 9/7 irreversible DWT with a small enough step size in the quantizer parameter.⁷

Table 2. Met variables T, Wv : RMSE for each bit rate (BR) in % of the full amplitude range

BR	T TRAD	T OBR	T RDO	WV TRAD	WV OBR	WV RDO
0.3	0.26	0.29	0.12	0.49	0.23	0.1
0.5	0.14	0.22	0.06	0.19	0.16	0.05
0.7	0.1	0.19	0.04	0.13	0.13	0.04
0.9	0.08	0.16	0.03	0.08	0.10	0.03
1.1	0.06	0.14	0.03	0.06	0.08	0.02
1.3	0.05	0.12	0.02	0.05	0.07	0.02
1.5	0.04	0.11	0.02	0.04	0.06	0.01

One of the major purposes of these experiments was to show that by using the RDO approach we can improve the MSE at low bit rates. Tables 2, 3 show the results for the T, Wv, U and V components of BFM Meteorological data. All the results for Meteorological data are shown in percentages of the total amplitude range of the specific data cube (see Table 1 describing this). As we can see from the results in these tables, the optimal RDO bit rate allocation provides significant improvement in MSE for small bit rates. For higher bit rates the RDO is slightly

Table 3. Met variables U, V : RMSE for each bit rate (BR) in % of the full amplitude range

BR	U TRAD	U OBR	U RDO	V TRAD	V OBR	V RDO
0.3	0.31	0.09	0.04	0.51	0.17	0.08
0.5	0.16	0.06	0.03	0.24	0.11	0.05
0.7	0.09	0.04	0.02	0.13	0.08	0.04
0.9	0.07	0.04	0.02	0.09	0.07	0.03
1.1	0.05	0.03	0.01	0.06	0.05	0.02
1.3	0.03	0.02	0.01	0.05	0.05	0.02
1.5	0.03	0.02	0.01	0.03	0.04	0.01

better than the traditional approach. Another interesting observation is that the OBR approach is sometimes better than the TRAD approach at low bit rates.

The MSE results for the 3-D CT medical data are shown in Figs. 7, 8, and 9. Note again that the RDO bit rate allocation provides significant improvement in MSE for all the small bit rates shown. For higher bit rates, the RDO is slightly better than the traditional approach results. Figure 7 also shows that for this data, the OBR approach always gives the worst results. Evaluation of the absolute distortion can be done using these figures and the fact that this 16-bit 3-D data set has an amplitude range of 4760.

Figure 9 shows results for the RDO approach using the two different pre-processing schemes discussed earlier. The computationally lighter DWT preprocessing gives results that are slightly better than the KLT pre-processing which also requires additional overhead since the basis vectors need to be saved and provided to the decoder. The compression distortion analysis given here emphasizes the relative behavior of the inter-slice bit rate allocation methods discussed.




Figure 7. Medical Data results for KLT pre-processing, low bit rates.

Acknowledgment

We would like to thank Jorge Melchor, Jr. for his assistance in the preparation of this manuscript. We gratefully acknowledge the input from Dr. T. Henmi of the U. S. Army Research Lab. for generating and providing the




Figure 8. Medical Data results for KLT pre-processing, high bit rates.




Figure 9. Medical data results comparing DWT pre-processing with KLT pre-processing.

Meteorological data used in this study. We also thank Katitza Gerdau of Ximis Inc. for providing us the 3-D medical image data.

REFERENCES

1. *ITU-T Recommendation T.801 | ISO/IEC 15444-1:2001 Information technology – JPEG2000 image coding system – Part 2: Extensions*, 2000.
2. A. Gersho and R. M. Gray, *Vector Quantization and Signal Compression*, Kluwer, Boston, 1992.

3. M. Pal, C. Brislawn, and S. Brumby, "Feature extraction from hyperspectral images compressed using the JPEG2000 standard," *Proc. of the IEEE Southwest Symposium on Image Analysis and Interpretation*, pp. 168–172, April 2002.
4. Olga M. Kosheleva, "Task-Specific Metrics and Optimized Rate Allocation Applied to Part 2 of JPEG2000 and 3-D Meteorological Data", Ph.D. Dissertation, University of Texas at El Paso, August 2003.
5. Olga M. Kosheleva, Bryan E. Usevitch, Sergio D. Cabrera, and Edward Vidal, Jr., "Rate Distortion Optimal Bit Allocation Methods for Volumetric Data Using JPEG2000," submitted to *IEEE Transactions on Image Processing*, January 2004.
6. Olga Kosheleva, Bryan Usevitch, Sergio Cabrera, and Edward Vidal, Jr., "MSE Optimal Bit Allocation in the Application of JPEG2000 Part 2 to Meteorological Data," *Proceedings of the 2004 IEEE Data Compression Conference DCC'2004*, Snowbird, Utah, p. 546, March 23–25, 2004.
7. D. Taubman and M. Marcellin, *JPEG2000 Image Compression Fundamentals, Standards and Practice*. Kluwer, Boston, Dordrecht, London, 2002.
8. H. E. Everett III, "Generalized lagrange Multiplier method for solving problems of optimum allocation of resources," *Operations Research*, vol. 11, pp. 399–417, 1963.
9. T. Henmi and J. R. Dumais, *Description of the Battlescale Forecast Model*, Army Research Lab, May 1998), Technical Report.
10. Olga Kosheleva, Alberto Aguirre, Sergio D. Cabrera, and Edward Vidal, Jr., "Assessment of KLT and Bit-Allocation Strategies in the Application of JPEG2000 to the Battlescale Forecast Meteorological Data", *Proc. IEEE International Geoscience and Remote Sensing Symposium IGARS'03*, Toulouse, France, July 21–25, 2003.
11. Alberto Aguirre, Sergio D. Cabrera, Aldo Lucero, Edward Vidal, Jr. and Katitza Gerdau, "Compression of Three-Dimensional Medical Image Data based on JPEG 2000," *Proceedings of 17th IEEE Symposium on Computer-Based Medical Systems*, Bethesda, MD, June 2004.
12. S. G. Mallat, *A wavelet tour of signal processing*. Academic Press, 1997.
13. T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms*, MIT Press, Cambridge, MA and Mc-Graw Hill Co., N.Y., 2001.