NEAR-LOSSLESS COMPRESSION OF MEDICAL IMAGES

M. Das[1], D. L. Neuhoff[2], and C. L. Lin[1]
[1]Dept. of Electrical Eng., Oakland University, Rochester, MI 48309
[2]Dept. of Electrical Eng. and Comp. Sc., University of Michigan, Ann Arbor, MI 48109

ABSTRACT

This paper studies the characteristic properties of a specific class of near-lossless image compression schemes which consists of a lossless coder followed by a uniform scalar quantizer. Three specific instances of such schemes are investigated; namely, differential pulse code modulation, hierarchical interpolation, and two-dimensional space-varying multiplicative autoregressive coders. The compression gains attainable with such schemes are studied and results of experiments conducted on several medical images are presented.

1. INTRODUCTION

The past two decades have witnessed a great deal of work in image coding, with a recently rising crescendo of activity. As a result of this, a multitude of different image coding techniques have emerged [1],[2]. All these techniques can be broadly categorized into two classes; namely, lossy and lossless. A lossless scheme typically achieves a compression ratio of the order of two only, but will allow exact recovery of the original image from the compressed version; a lossy scheme will not allow exact recovery, but can attain much higher compression ratios, e.g., twenty or more.

As against strictly lossless or lossy compression, this paper investigates the performance of a class of near-lossless image compression schemes that attempts to improve the compression efficiency of lossless coders by using a uniform scalar quantizer to quantize their residuals. The compression gain attainable with such schemes is investigated and the results are experimentally validated on several medical images by comparing the first order entropies of the residuals.

The organization of this paper is as follows. Section II introduces three near-lossless coding schemes based on differential pulse code modulation (DPCM) [4], hierarchical interpolation (HINT) [4], and two-dimensional space-varying multiplicative autoregressive (2-D SMAR) coders [5]. Section III presents some comparative experimental results and finally, some concluding remarks are given in Section IV.

2. NEAR-LOSSLESS CODING OF DIGITIZED IMAGES

There exists a variety of methods by which a lossless coder can be converted to a near-lossless one, which will confine the reconstruction errors for individual pixels within some predefined limits. For example, one may assure a reconstruction error of either 0, or 1 by simply dropping the least significant bit of the image. However, this clearly results in only a one bit gain in the average bit rate. Similarly, dropping the two least significant bits assures a gain of only two bits, while allowing errors to lie within [0,3]. Therefore, techniques that allow a better tradeoff between compression gain and reconstruction error size are preferred.

In this paper, we study a well known class of near-lossless coders that result from uniform scalar quantization of a lossless coder's residuals. For this purpose, we will focus on three types of predictive lossless coders, which are known to perform better than others [4],[5]. The methods chosen are: i) DPCM, which is a nonadaptive predictive coder, ii) HINT, which is a hierarchical predictive coder, and iii) 2-D SMAR, which is an adaptive predictive coder. The underlying ideas are first illustrated using DPCM method, and then near-lossless coders based on HINT and SMAR are discussed.

2.1. A Near-Lossless Coding Scheme Based on DPCM Technique

Two-dimensional differential pulse code modulation (DPCM) is a simple and well known technique for image data compression [3], which exploits the correlation of a pixel value with its four nearest causal neighbors. A lossless DPCM image compression technique [4] can be easily be converted to a near-lossless one as follows.

Assuming $\{f(i,j), 1 \le i \le L, 1 \le j \le L\}$ denote the pixel values of the original digitized image, the predicted value, $\tilde{f}(i,j)$, of $f(i,j)$ is given by

$$\tilde{f}(i,j) = \rho\hat{f}(i,j-1) + \rho\hat{f}(i-1,j) - \rho^2\hat{f}(i-1,j-1), \qquad (1a)$$

where $\hat{f}(.,.)$ denotes the reconstructed value of $f(.,.)$, and ρ, the correlation coefficient, is usually chosen to be 0.95.

0-7803-2431-5/95 $4.00 © 1995 IEEE

Next the prediction error, d(i,j), and its quantized value, $\hat{d}(i,j)$, are computed as

$$d(i,j) = f(i,j) - R[\tilde{f}(i,j)], \quad (1b)$$

$$\hat{d}(i,j) = Q[d(i,j)], \quad (1c)$$

where R[.] denotes the rounding operator to the nearest integer, and Q[.] is a center-clipping quantizer whose input-output relationship is dictated by the maximum allowable reconstruction error (MARE). For example, if MARE is ±1, the input-output relationship of Q is chosen as,

$$Q[k] = m \text{ for integer } k \in [m-1,m+1], \quad (2a)$$

where m = 0, ±3, ±6, etc. Similarly in general if MARE is ±n, choose Q as,

$$Q[k] = m \text{ for integer } k \in [m-n,m+n], \quad (2b)$$

where m = 0, ±(2n+1), ±(4n+1), etc.

Finally, the reconstructed pixel value, $\hat{f}(i,j)$, of f(i,j) is obtained as

$$\hat{f}(i,j) = R[\tilde{f}(i,j)] + \hat{d}(i,j), \quad (3)$$

and at the end, after repeating the above steps for each pixel in a recursive fashion, the sequence $\{\hat{d}(i,j)\}$ is entropy coded and transmitted. At the receiver, the reconstructed image is obtained recursively by first computing $\tilde{f}(i,j)$ from (1a), rounding it and then adding the same to the decoded value of $\hat{d}(i,j)$, as shown in (3).

The above compression scheme is a near-lossless one because by choosing "n" in equation (2b) to be small, a nearly perfect version of the original image can be reconstructed at the receiver end. For instance, if the original image is 8-bit deep, the peak-signal-to-noise-ratio (PSNR) of the reconstructed image is given by

$$PSNR = 20*\log_{10}(255/n), \quad (4)$$

which equals 48.13 db for n=1 and 42.11 db for n=2. In practice, reconstructed images having PSNR greater than or equal to 35 db are found to be hardly distinguishable from the original ones.

Next, the above ideas are extended to HINT and SMAR coders below.

2.2. A Near-Lossless Coding Scheme Based on HINT Method

The hierarchical interpolation (HINT) [4] is a lossless coding scheme that begins with a low-resolution version of the original image, P_0, and successively generates the higher resolutions P_k, $1 \le k \le 4$, using interpolations. The lowermost resolution, P_0, is entropy coded and transmitted first. Thereafter, in a hierarchical fashion, the interpolation scheme is used to generate estimates of the unknown pixel values of P_k by calculating the average of its four nearest neighbors that are provided by P_{k-1}. The estimates are rounded to their nearest integers and then subtracted from the true pixel values. The difference signals pertaining to each of the higher resolutions, D_k, $1 \le k \le 4$, are also entropy coded and transmitted.

As in the case of DPCM, the above scheme can be converted to a near-lossless one by first quantizing the residuals, D_k, and calculating the reconstructed version, $\hat{P}_k$, of P_k, before carrying out interpolation to generate the next higher resolution. In this case then the quantized residuals, $\hat{D}_k$, $1 \le k \le 4$, are encoded and transmitted. The quantizer used is same as the center-clipping one described in the previous Section.

2.3. A Near-Lossless Coding Scheme Based on 2-D SMAR Coders

2-D SMAR is an adaptive predictive coder that has been found to be very efficient for lossless compression [5]. Such predictors are most easily described using delay polynomials (D-transforms). In this notation, an image f is represented as a polynomial

$$F(D_1,D_2) = \sum_{i,j} f(i,j)\, D_1^{-i} D_2^{-j}, \quad (5)$$

where $f(i,j)\, D_1^{-i} D_2^{-j}$ represents the fact that the pixel at location (i,j) has value f(i,j). Also, the operation of a linear predictor,

$$y(i,j) = \sum_{k,l} p(k,l)\, f(i-k,j-l), \quad (6)$$

may be described by

$$Y(D_1,D_2) = F(D_1,D_2)\, P(D_1,D_2), \quad (7)$$

where

$$Y(D_1,D_2) = \sum_{i,j} y(i,j)\, D_1^{-i} D_2^{-j}, \quad (8)$$

and

$$P(D_1,D_2) = \sum_{i,j} p(i,j)\, D_1^{-i} D_2^{-j}. \quad (9)$$

In MAR [6],

$$P(D_1,D_2)=1-A(D_1,D_2), \quad (10)$$

where $A(D_1,D_2)$ is the product of a number of low degree polynomials. This yields a high degree polynomial (and a large prediction context) described by just a few parameters. For example, we have had excellent success with the 4x4 nonsymmetric half-plane (NSHP) predictor [6] characterized by

$$A(D_1,D_2) = (1+a_1D_2^{-1})(1+a_2D_1^{-1})(1+a_3D_1^{-1}D_2)(1+a_4D_1^{-1}D_2^{-1}), \quad (11)$$

which uses only four coeffcients, but results in the context of 11 pixels shown in Figure 1. It is further shown in [5] that this type of predictor works better when it follows the simple, nonadaptive predictor

$$b(i,j) = 0.25*[f(i-1,j)+f(i,j-1)+f(i-1,j-1)+f(i-1,j+1)]. \quad (12)$$

That is, the MAR predictor is applied to the residuals of the above predictor. The first-stage of the predictor can be thought of as removing the space-varying mean.

The SMAR near-lossless coder works in a block-by-block fashion. At the transmitter end, the coding for each block proceeds as follows. First, the coefficients of $A(D_1,D_2)$ are estimated using the so called recursive pseudolinear regression (RPLR) method [6],[7], and the predictor, $P(D_1,D_2)$ to be utilized for the block is calculated from the estimated 2-D SMAR model, $\hat{A}(D_1,D_2)$, as

$$P(D_1,D_2) = 1 - \hat{A}(D_1,D_2). \quad (13)$$

Next, the coding proceeds recursively in the following manner:

calculate approximate space-varying mean:

$$\tilde{b}(i,j) = R[(1/4)(\hat{f}(i,j-1) + \hat{f}(i-1,j-1) + \hat{f}(i-1,j) + \hat{f}(i-1,j+1))], \quad (14a)$$

calculate zero-mean pixel:

$$y(i,j) = f(i,j) - \tilde{b}(i,j), \quad (14b)$$

predict y(i,j) and calculate residual:

$$\tilde{y}(i,j) = \sum_{i,j} p(k,l)\, \hat{y}(i-k,j-l), \quad (14c)$$

$$y_r(i,j) = R[\tilde{y}(i,j)], \quad (14d)$$

$$d(i,j) = y(i,j) - y_r(i,j), \quad (14e)$$

quantize residual: $\hat{d}(i,j) = Q[d(i,j)], \quad (14f)$

reconstruct y(i,j) and f(i,j) :

$$\hat{y}(i,j) = y_r(i,j) + \hat{d}(i,j), \quad (14g)$$

$$\hat{f}(i,j) = \hat{y}(i,j) + \tilde{b}(i,j). \quad (14h)$$

Finally, the quantized residual sequence, $\{\hat{d}(i,j)\}$, is entropy coded using an optimal encoder and transmitted.

At the receiver, the reconstructed image is obtained recursively by first computing $\tilde{b}(i,j)$ from (14a), next calculating $\hat{y}(i,j)$ from (14c),(14d) and (14g), and finally computing $\hat{f}(i,j)$ from (14h).

Next, the compression gain attainable using the above near-lossless coders is discussed above

2.4. Expected Compression Gain

Assume that the residuals of the underlying lossless coder consist of integer values over some interval, [-I,+I], and have a probability distribution, p(i), $-I \le i \le I$. Next, suppose the residuals are quantized to allow a maximum reconstruction error of $\pm n$. For the purpose of an approximate analysis, if we assume that p(i+k), $-n \le k \le n$ are approximately equal, then the near-lossless coding achieves a compression rate of about $\log_2(2n+1)$ bits less than lossless coding. For example, if n = 1, this is a savings of $\log_2 3$ = 1.58 bits/pixel, and if n=2, this is a savings of 2.3 bits.

Finally, some experimental results are given below.

3. EXPERIMENTAL RESULTS

The performance of the three near-lossless coding schemes presented above was evaluated using four X-rays, R1-R4, and six MRI images, MRI1-MRI6. All images are 512x512 and digitized to 256 gray levels. For SMAR experiments, the model used was 4x4 NSHP and the block sizes were chosen as 32x32. The comparative performance of different schemes was evaluated solely on the basis of the attainable bit rates, as measured by the first order entropies of the quantized residuals.

The bit rates attainable using different coding techniques are summarized in Table 1. The bit rates presented for HINT and DPCM are same as the first order entropies of their residual signals, whereas those for SMAR include both first order entropies of the quantized residuals and the average number of bits per pixel used for transmitting the blockwise model coefficients. Finally, Table 2 depicts the actual compression gains attainable using SMAR as a near-lossless coder rather than a lossless one.

It is clearly seen that for all the images, both HINT and SMAR outperform DPCM in terms of attainable compression gains. This is, of course, according to our expectation because DPCM uses a much simpler image model compared to the other two. Whereas such a simple model works reasonably good for lossless compression [4], it does not perform as well in the near-lossless case. Finally, as shown in Table 2, the actual compression gains attainable by SMAR compare well with the expected ones discussed in Section 2.

4. CONCLUSION

A near-lossless compression scheme based on incorporation of a uniform quantizer into a lossless coder is studied in this paper. The performance of the scheme is tested on three lossless coders, namely, DPCM, HINT, and SMAR. Whereas all the techniques offer significant compression gains with very little loss of information, HINT and SMAR perform superior compared to DPCM. It is believed that the proposed method can be useful for near-lossless compression of medical and other kinds of images.

REFERENCES

[1] A. K. Jain, "Image Data Compression: A Review," Proceedings of the IEEE, Vol. 69, pp. 349-389, March 1981.

[2] M. Rabbani and P. W. Jones, *Digital Image Compression Techniques*, SPIE Optical Engineering Press, 1991.

[3] A. K. Jain, *Fundamentals of Digital Image Processing*, Prentice Hall, Englewood Cliffs, 1989.

[4] P. Roos et al, "Reversible Intraframe Compression of Medical Images," IEEE Transactions on Medical Imaging, Vol. 7, No. 4, pp. 328-336, 1988.

[5] M. Das and S. Burgett, "A Lossless Medical Image Compression Scheme Based on Two- Dimensional Space-Varying Multiplicative Autoregressive Models," Oakland University Technical Report No. TR-92-ESE-06-02, to appear in IEEE Trans. on Image Processing.

[6] M. Das and S. Burgett, "Adaptive Predictive Coding of Images Using Two-Dimensional Multiplicative Autoregressive Models," Signal Processing, Vol. 31, No. 2, pp. 121-132, 1993.

[7] L. Ljung and T. Soderstrom, *Theory and Practice of Recursive Identification*, MIT Press, Cambridge, Massachusetts, 1983.

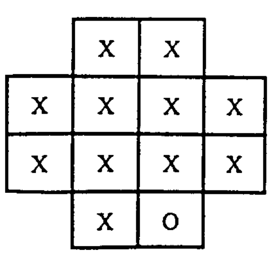

Figure 1. Support Region of NSHP(4x4) model
"X" is a pixel in the support
"O" is the pixel to be predicted

Images	Ave. bit rates (in bits/pixel) for max. error of ± 1			Ave. bits rates (in bits/pixel) for max. error of ± 2		
	HINT	DPCM	SMAR	HINT	DPCM	SMAR
R1	1.15	1.40	1.07	0.68	1.12	0.59
R2	1.19	1.74	1.09	0.75	1.44	0.58
R3	1.17	1.48	1.06	0.66	1.23	0.50
R4	1.28	1.59	1.20	0.76	1.28	0.66
MRI1	2.35	2.37	2.27	1.88	1.90	1.85
MRI2	1.79	1.76	1.62	1.37	1.38	1.25
MRI3	2.92	3.05	2.76	2.27	2.41	2.10
MRI4	2.37	2.53	2.25	1.78	1.95	1.68
MRI5	2.03	2.15	1.91	1.50	1.66	1.40
MRI6	3.57	3.75	3.43	2.88	3.06	2.72

Table 1. Entropies of different coding schemes

Images	Ave. bit rates (in bits/pixel) using SMAR		
	Lossless	Near-lossless (max. error of ± 1)	Near-lossless (max. error of ± 2)
R1	2.28	1.07	0.59
R2	2.38	1.09	0.58
R3	2.40	1.06	0.50
R4	2.55	1.20	0.66
MRI1	3.33	2.27	1.85
MRI2	2.56	1.62	1.25
MRI3	4.26	2.76	2.10
MRI4	3.64	2.25	1.68
MRI5	3.16	1.91	1.40
MRI6	4.99	3.43	2.72

Table 2. Compression gains using SMAR