

A COMPARATIVE STUDY OF L_∞ -DISTORTION LIMITED IMAGE COMPRESSION ALGORITHMS

Hannes Hartenstein, Ralf Herz and Dietmar Saupe

Institut für Informatik, Universität Freiburg, Am Flughafen 17, 79110 Freiburg, Germany
E-mail: hartenst,herz,saupe@informatik.uni-freiburg.de

ABSTRACT

In L_∞ -distortion limited compression each single pixel value is only changed by maximal $\pm\tau$ grey values. In this paper we present a theoretical framework for L_∞ -distortion limited compression that covers several recently proposed methods. The basics of each of these methods are described. We give a comparison of coding results for the Lenna test image, a coronary angiogram, and a Landsat image. Results are reported for various tolerances. Standard DPCM is used as a reference. While this paper gives an overview over various algorithms, the main purpose is to indicate what level of compression can be expected when limiting the error in L_∞ -distortion sense.

1. A L_∞ -DISTORTION LIMITED COMPRESSION FRAMEWORK

In many applications, for example medical imagery, SAR imagery, or numerical weather simulations, the large amount of data to be stored or transmitted asks for data compression. Since lossless coding usually gives a compression ratio of at most 4:1, lossy coding methods have to be employed when higher compression ratios are needed. Most lossy compression schemes operate by minimizing some *average* error measure such as the root mean square error. However, in error critical applications such as medical imagery or target recognition, such average error measures are inappropriate. Instead, there is usually a need for a guarantee that a single pixel has not been changed by more than a certain tolerance (which may depend on the pixel location). Thus, the error in *each* pixel has to be controlled.

In this paper we consider a L_∞ -distortion limited compression scheme with global tolerance τ . For such an encoding method the code for a one-dimensional signal $s \in Z^n$ represents a reconstruction signal $\hat{s} \in \mathcal{N}_\tau(s)$, where

$$\mathcal{N}_\tau(s) := \{t \in Z^n \mid \|t - s\|_\infty \leq \tau\},$$

$$\|t - s\|_\infty = \max_{i \in \{1..n\}} |t_i - s_i|.$$

For $\tau = 0$ this leads to lossless compression. If τ is small, the term 'near-lossless coding' appears to be




Figure 1: Each left-to-right path is an element of $\mathcal{N}_\tau(s)$ with $s = (3, 3, 2, 1, 3, 5, 7, 6, 5, 4)$.

justified. $\mathcal{N}_\tau(s)$ can be seen as the set of all left-to-right paths in a trellis as depicted in Figure 1.

Which of the $(2\tau + 1)^n$ elements of $\mathcal{N}_\tau(s)$ can be coded most efficiently? All coding methods described below use a lossless coding strategy C and try to determine, at least approximately or heuristically, the element of $\mathcal{N}_\tau(s)$ that can be coded most efficiently using C . Mathematically, let \mathcal{C} be the set of lossless coding methods. For example, $C \in \mathcal{C}$ could be a 0-order entropy coder. Then the coding problem for this particular C is:

$$\begin{aligned} &\text{find } t^* \in \mathcal{N}_\tau(s) \text{ such that} \\ &\ell(C(t^*)) = \min_{t \in \mathcal{N}_\tau(s)} \ell(C(t)) \end{aligned} \quad (*)$$

where $\ell(\cdot)$ gives the length of the code or some estimate thereof.

In the following sections we give short descriptions of several L_∞ -based compression methods. Most of these methods were implemented and tested with the images given in Figure 2.

2. QUANTIZATION VS PRECONDITIONING

In the problem formulation of the previous section the signal to be coded can be modified in each component independently of the other components. Thus, for a signal s and $i \neq j$ it is possible that $s(i) = s(j)$ but $\hat{s}(i) \neq \hat{s}(j)$. In other words, \hat{s} is in general not simply

the result of using a quantizer q for the values of s . We refer to \hat{s} as a *preconditioned* version of s in contrast to a *quantized* version. The emphasis in this paper is on *preconditioning*.

Nevertheless, the problem of finding the quantization function such that the quantized version of s has minimal 0-order entropy can be solved in $\mathcal{O}(n)$ time using dynamic programming [1]. It is also shown that for a tolerance $\tau > 0$ the entropy savings are at most $\log_2(2\tau + 1)$ per pixel.

3. ENTROPY-CODED DPCM BASED METHODS

The entropy coding of the prediction residuals of a DPCM scheme is a standard method for lossless compression. It can easily be modified to serve as an L_∞ distortion based compression method.

DPCM1. The signal is uniformly quantized with quantization bin size $2\tau + 1$. Thus, a quantized version of the original signal is computed. Then the residuals of a linear predictor are entropy coded. The disadvantage of this method is that for larger τ there are only a few different grey levels leading to 'plateau effects'.

DPCM2. No a priori grey value reduction is performed, but the prediction error of the DPCM scheme is uniformly quantized to match the desired tolerance τ . When the predictor coefficients are not integer values, this method does not coincide with the method *DPCM1* and does not show the plateau effects. Results for several medical images are reported in [2].

In the above mentioned methods, there is actually no mechanism to minimize the entropy of the error sequence. When we use a lossless predictive coder followed by an entropy coder, the optimization problem (*) asks for the path in the trellis whose corresponding residual sequence has minimum entropy. We conjecture that this optimization problem is NP-hard. Note that the complexity depends on the signal length n and the tolerance τ .

We applied genetic algorithms (GA) [3, 4] to solve this optimization problem for a signal $s \in Z^n$ and a tolerance τ .

GA. In our setting a chromosome c is a word of length n over the alphabet $\{-\tau, \dots, 0, \dots, \tau\}$ and represents the signal $s + c \in \mathcal{N}_\tau(s)$. The genetic operations are 2-point crossover and mutation. We use roulette wheel parent selection. The evaluation of a chromosome is given by the entropy of the distribution of prediction residuals of $\hat{s} = s + c$. For the fitness function we use exponential ranking. Large tests for the determination of suitable parameters were performed.

The results obtained with the *GA* approach are rather disappointing. For example, as a signal s a line of the image Lenna was taken. The entropy of that signal is 7.0, after quantization with tolerance $\tau = 2$, and after prediction the sequence can be coded with an

entropy of 3.1. The solution found with the *GA* only gave an entropy of 3.9. Thus, the *GA* is not even able to beat the method *DPCM1*.

The minimum-entropy constrained-error DPCM (*MECE*) of [5] is another method that tries to minimize the entropy of the prediction residual sequence. It uses an iterative optimization method that arrives at a local optimum.

MECE. Assume that an ideal entropy coder is given for a fixed residual distribution. To find the optimal element of $\mathcal{N}_\tau(s)$ for this coder one has to solve a shortest path problem. This can easily be done via Dynamic Programming. Now, using an entropy coder that is optimal for the actual residual distribution will give a decrease in entropy. These two steps are performed iteratively until a stopping criterion is matched.

For images a two dimensional 3-tap predictor is used and the images are coded row by row. The results can be further improved by using a 1-order entropy coder with a certain number of contexts.

In the above mentioned methods the predictor and the contexts are fixed. Of course, it would be advantageous to include the choice of predictor coefficients and context into the optimization problem; clearly, this makes the problem even more complicated. A sophisticated method that uses adaptive context modeling to correct prediction biases is the L_∞ -constrained CALIC [6]. The converse problem of determining a predictor such that the prediction residuals have minimum entropy was investigated for lossless coding in [7].

4. PIECEWISE LINEAR CODING

Piecewise linear coding (*PCL*) is a generalization of Run Length Encoding. It is also called fan-based coding; for an extensive overview see [8]. In piecewise linear coding a signal is split into segments each of which can be described by a linear function. Each segment then is coded by the length of the segment and the slope parameter. The constant additive part of the function is implicitly given by the previous segment; only for the first segment the initial signal value has to be coded. For example, the signal in Figure 1 is represented as 3(1,0)(2,-1)(3,2)(3,-1).

In the case that $\ell(\cdot)$ counts the number of segments the optimization (*) can be solved in $\mathcal{O}(n^2)$ -time via Dynamic Programming [9]. In [10, 11] a suboptimal greedy method that works in linear time is proposed for the same optimization problem. Essentially, it works as follows. The image is transformed into a 1-dimensional signal, e.g., by a Hilbert-Peano scan. Then the linear segments are successively determined: starting at the endpoint of the last determined segment, the new segment is chosen to be the one of greatest possible length. Finally, an 0-order entropy coder is applied to the list of segment lengths and segment slopes. Better results can be obtained when the length of the 0-order entropy code is minimized in place of




Figure 2: The 512×512 8-bit test images Lenna, Angio, Landsat.

	$\tau = 0$	$\tau = 1$		$\tau = 2$		$\tau = 3$		$\tau = 4$		$\tau = 5$		$\tau = 10$		$\tau = 15$	
	bpp	bpp	rms	bpp	rms	bpp	rms	bpp	rms	bpp	rms	bpp	rms	bpp	rms
DPCM1	4.6	3.0	0.8	2.3	1.4	1.9	2.0	1.6	2.6	1.4	3.2	0.9	6.1	0.7	8.9
DPCM2	4.5	3.0	0.8	2.3	1.4	1.9	2.0	1.7	2.6	1.5	3.2	1.0	5.9	0.8	8.7
MECE	4.6	3.1	0.8	2.5	1.4	2.0	1.9	1.7	2.5	1.4	3.0	0.7	5.3	0.5	7.6
PLC	5.0	4.2	0.8	3.5	1.5	3.0	2.2	2.5	2.8	2.2	3.3	1.3	5.9	0.9	8.5
DLVQ	4.7	3.2	0.8	2.5	1.4	2.0	2.0	1.7	2.5	1.3	2.9	0.7	4.9	0.6	6.8
SPIHT	4.5	3.0	0.8	2.3	1.4	1.9	1.8	1.5	2.1	1.2	2.3	0.6	3.2	0.4	3.9

Table 1: Compression results for the Lenna image (bpp \doteq bits per pixel, rms \doteq root mean square error).

	$\tau = 0$	$\tau = 1$		$\tau = 2$		$\tau = 3$		$\tau = 4$		$\tau = 5$		$\tau = 10$		$\tau = 15$	
	bpp	bpp	rms	bpp	rms	bpp	rms	bpp	rms	bpp	rms	bpp	rms	bpp	rms
DPCM1	3.9	2.5	0.8	1.9	1.4	1.5	1.9	1.3	2.5	1.2	3.1	0.8	6.0	0.6	9.0
DPCM2	3.9	2.6	0.9	2.0	1.4	1.7	2.0	1.5	2.6	1.3	3.1	1.0	6.0	0.9	8.9
MECE	3.9	2.6	0.8	2.0	1.3	1.7	1.9	1.5	2.5	1.2	3.0	0.5	5.2	0.3	7.0
PLC	4.5	3.9	0.8	3.3	1.5	2.8	2.3	2.3	2.9	2.0	3.5	0.9	5.9	0.5	7.9
DLVQ	4.2	2.6	0.8	1.9	1.4	1.5	2.0	1.3	2.5	1.0	3.0	0.5	4.8	0.4	6.5
SPIHT	4.0	2.5	0.8	1.9	1.4	1.6	1.7	1.2	1.9	1.0	2.2	0.3	4.2	0.1	4.9

Table 2: Compression results for the Angio image.

	$\tau = 0$	$\tau = 1$		$\tau = 2$		$\tau = 3$		$\tau = 4$		$\tau = 5$		$\tau = 10$		$\tau = 15$	
	bpp	bpp	rms	bpp	rms	bpp	rms	bpp	rms	bpp	rms	bpp	rms	bpp	rms
DPCM1	4.8	4.6	0.8	4.1	1.4	3.3	1.9	3.2	2.5	3.0	3.1	2.1	5.8	1.7	8.8
DPCM2	5.7	4.5	0.8	3.8	1.4	3.3	2.0	3.0	2.6	2.7	3.1	1.9	6.0	1.4	8.7
MECE	4.8	4.1	0.7	4.0	1.4	3.5	2.1	3.1	2.7	3.0	3.2	2.0	5.8	1.4	8.7
PLC	4.6	5.1	0.8	5.1	1.4	4.7	2.3	4.6	2.9	4.4	3.5	3.2	7.1	2.3	10.1
DLVQ	6.0	4.4	0.8	3.7	1.4	3.2	2.0	2.9	2.5	2.6	3.1	1.7	6.0	1.2	8.6
SPIHT	5.9	4.6	0.8	3.9	1.4	3.4	2.0	3.0	2.5	2.7	3.1	1.9	5.9	1.4	7.2

Table 3: Compression results for the Landsat image.

the number of segments [9]. However, in this case the global optimum can no longer be achieved by linear programming alone. Instead, an iterative procedure similar to that in [5] can be applied leading to a locally optimal solution.

5. DISTORTION-LIMITED VQ

The *DLVQ* method proposed by [12] is a multistage vector quantization approach. The image is partitioned,

e.g., in 4×4 blocks. A block is coded by a predictive vector quantizer. If there is a component with a distortion above the threshold τ , the error residual block is coded with another codebook and a second index has to be sent. If the error tolerance is still not satisfied, a scalar quantizer is applied to the second-order residuals which finally guarantees that the error tolerance criterion is satisfied.

The special feature of [12] is a codebook design algorithm that reduces the number of large distortion




Figure 3: Error images for Lenna and $\tau = 10$. The absolute values of the errors are scaled such that white indicates an error of 10, black stands for exact pixel coding.

occurrences. It is based on a modified L_∞ -distortion:

$$\|t - s\|_{mod} = \begin{cases} 0 & \max_{1 \leq i \leq n} |s_i - t_i| \leq \tau \\ \max_{1 \leq i \leq n} |s_i - t_i| - \tau & \text{otherwise} \end{cases}$$

The centroid computation in the Generalized Lloyd Algorithm for the L_∞ -distortion and for the above modification is done using a specially designed iterative optimization.

6. OTHER METHODS

Any lossy compression method can be turned into a lossless scheme by adding a stage of lossless error residual coding [13]. Thus, likewise a lossy method can also be extended to a near-lossless coding method. The problem is to efficiently encode the positions where a refinement to the lossy scheme has to be done. The *DLVQ* can be viewed as a method of this type.

We have tested the 'lossy plus near-lossless residual' option using the *SPIHT* scheme of [14]. The images are coded with the *SPIHT* method for various bit rates; then, the residual images are coded via DPCM

(for a fixed tolerance τ). That combination is chosen which has minimized the total bit rate. Of course, better adapted predictors and context selectors will improve the performance.

When using transform coding without near-lossless residual coding for L_∞ -distortion based compression, the problem is to determine how the transformed coefficients can be modified such that after the application of the inverse transform the tolerance criterion is not violated. In [15] this is studied in the context of the wavelet transform. The maximal quantization step size is determined for a uniform quantizer in the wavelet domain, given a tolerance τ .

7. RESULTS & CONCLUSIONS

For our empirical tests we have implemented the two basic DPCM methods, the minimum-entropy constrained-error (*MECE*) approach, the optimal piecewise linear coding (*PLC*) that minimizes the number of runs, the distortion-limited VQ (*DLVQ*), and the *SPIHT plus DPCM* scheme. The programs were applied to the three test images shown in Figure 2 for various L_∞ -tolerances $\tau \geq 0$. The results are given in Tables 1-3. All bit per pixel (bpp) numbers are entropy estimates. Besides the bpp value the error dis-

tribution (Figure 3) is another criterion for comparison. The *DPCM1* and *DPCM2* methods give almost the same results, but the *DPCM2* coded images give better visual impression. The entropy minimization of *MECE* shows little effect for small error tolerances. For high tolerances, we see some significant gains, but these gains come with high computing times. For the *DPCM1* and the *MECE* schemes the simple planar predictor was employed. For the *DPCM2* method the coefficients of the 3-tap predictor were chosen to minimize the squared error of the residuals. The contexts were selected as in [5].

The *PLC* strategy gives good results only for the relatively smooth angiogram at high tolerances. For the Landsat image increasing the tolerance actually can increase the bpp. This is because the optimization criterion is not the entropy of the resulting parameters.

The results for the *DLVQ* approach look very promising. The best results are obtained with a first codebook of size 64 or 128 and a second one of size 512.

The *SPIHT plus DPCM* scheme shows a very good performance for the images Lenna and Angio; especially for larger τ there is the additional advantage that the root mean square error is much smaller compared to the other methods. For lossless coding the *DPCM* alone is as good as the combination with the *SPIHT* scheme. But using the lossless coding method of [16] one gets an improvement of about 0.3 bpp.

What compression is gained by allowing a small per pixel tolerance compared to lossless coding? For the Landsat image a small tolerance does not significantly improve coding results. Even large tolerances lead to only small compression ratios. For the other images allowing a grey value tolerance of ± 1 results in an additional 33% size reduction over lossless encoding, a tolerance of ± 5 leads to an additional 66% size reduction.

8. ACKNOWLEDGMENT

We thank Bettina Heck and Maximillian Thiel for helpful discussions and programming work, and Bernd Triebfürst for providing us various Landsat images.

9. REFERENCES

- [1] Gormish, M. J., Gill, J. T., *Discrete minimum entropy quantization*, IEEE Trans. Image Processing, Vol. 4, No. 9, Sept. 1995, pp. 1314–1317
- [2] Das, M., Neuhoﬀ, D. L., Lin, C. L., *Near-lossless compression of medical images*, Proceedings of the International Conference on Acoustics, Speech, and Signal Processing (ICASSP), Detroit, 1995, pp. 2347–2350
- [3] Tang, K. S., Man, K. F., Kwong, S., He, Q., *Genetic algorithms and their applications*, IEEE Signal Processing Magazine, 1996
- [4] Davis, L. (ed), *Handbook of Genetic Algorithms*, van Nostrand Reinhold, New York, 1991
- [5] Ke, L., Marcellin, M. W., *Near-lossless image compression: minimum entropy, constrained-error DPCM*, Proceedings of the International Conference on Image Processing (ICIP), Washington, D. C., 1995
- [6] Wu, X., *L_∞ -constrained high-fidelity image compression via adaptive context modeling*, Proceedings of the Data Compression Conference, Snowbird, Utah, 1997, pp. 91–100
- [7] Bowen-Wright, R., Sayood, K., *Lossless compression by simulated annealing*, Proceedings of the Data Compression Conference, Snowbird, Utah, 1995, p. 452
- [8] Rosenberg, C. J., *A lossy image compression algorithm based on nonuniform sampling and interpolation of the image intensity surface*, Master Thesis, Massachusetts Institute of Technology, 1990
- [9] Saupe, D., *Optimal piecewise linear image coding*, submitted
- [10] Pinciroli, F., et al., *Some experiments in compressing angiocardigraphics images according to the Peano-Hilbert scan path*, Computer methods and programs in biomedicine, vol. 43, No. 3-4, 1994, pp. 247–253
- [11] Bhaskaran, V., Natarjan, B. K., Konstantinides, K., *Optimal piecewise-linear compression of images*, Proceedings of the Data Compression Conference, Snowbird, Utah, 1993, pp. 168–177
- [12] Hahn, P. J., Mathews, V. J., *Distortion-limited vector quantization*, Proceedings of the Data Compression Conference, Snowbird, Utah, 1996, pp. 340–348
- [13] Rabbani, M., Jones, P. W., *Digital Image Compression Techniques*, SPIE Press, 1991
- [14] Said, A., Pearlman, W. A., *A new fast and efficient implementation of an image codec based on set partitioning in hierarchical trees*, IEEE Transactions on Circuits and Systems for Video Technology, vol. 6, June 1996, pp. 243–250
- [15] Karray, L., Rioul, O., Duhamel, P., *L_∞ -coding of images: a confidence interval criterion*, Proceedings of the International Conference on Image Processing (ICIP), Austin, Texas, 1994, pp. 888–892
- [16] Said, A., Pearlman, W. A., *An image multiresolution representation for lossless and lossy compression*, Visual Communications and Image Processing, Proceedings SPIE 2094, Nov. 1993, pp. 664–674