A CONTEXT-BASED RECURSIVE NONLINEAR INTERPOLATION FOR NEAR-LOSSLESS CODING OF X-RAY IMAGES

Bruno Aiazzi, Stefano Baronti, Franco Lotti
"Nello Carrara" IROE - CNR, Via Panciatichi, 64, I-50127 Florence, Italy
e-mail: {baronti lotti}@iroe.fi.cnr.it

Luciano Alparone
Dept. of Electron. Engin., Univ. of Florence, Via S. Marta, 3, I-50139 Florence, Italy
e-mail: alparone@cosimo.die.unifi.it

ABSTRACT

This paper addresses quality issues in Medical Image compression and proposes an approach to achieve *near-lossless* storage of digitized X-ray plates. An image is normalized to the standard deviation of its noise, which is estimated in an unsupervised fashion. The resulting bitmap is encoded without any further loss. The compression algorithm proposed is based on a two-stage recursive interpolation exploiting nonseparable median filtering a on a quincunx grid. The advantage is twofold: interpolation is performed from all error-free values and is unlikely to occur across edges, thereby reducing the coding cost of the outcome residuals. In addition, classification based on spatial context is employed to improve entropy coding. The scheme outperforms other established methods when applied to X-ray images.

1 INTRODUCTION

Image compression is gaining the attention of an ever-increasing audience in the field of Medical Imaging because of the huge amount of data to be archived and/or transmitted [1]. Compression methods can be either reversible (i.e. *lossless*) or irreversible (*lossy*), depending on whether images are exactly reconstructed after decoding or some errors are introduced. When lossless coding is mandatory, however, compression ratios greater than 2 or 3 can hardly be obtained, mainly because of the acquisition and digitization noise.

The evaluation of the maximum allowable distortion within the constraint of an efficient source coding is an open problem: quality must be preserved according to both objective and subjective criteria [2]. A viable approach consists of defining some *regions of interest* (ROI) [3]: compression will be lossless inside ROIs, and lossy outside. Another approach is to accept lossy compression and to adjust quality based on diagnostic trials performed, for instance, through *receiver operating characteristic* (ROC) curves [4]. Both the above approaches must be supervised by an expert, to identify ROIs in the former case, and to assess diagnostic quality in the latter. They also assume that the original digitized data possess a quality adequate for diagnoses.

Noteworthy for Medical Images are those lossy methods which allow to settle "a priori" the reconstruction error, not only on the whole, but also as a local measurement. *Near-lossless* compression is concerned when reconstruction errors are upper bounded, besides being small on an average [4, 5, 6]. In this way, notwithstanding the loss of semantic (e.g. diagnostic) information may be considered to be negligible at all, a decrement in bit-rate is obtained from the loss of statistical information given by the noisy image component.

The creation of digital archives starting from existing analog formats is a parallel problem in the medical field. This demand is specially pressing for X-ray images that are usually digitized from analog plates with a radiometric resolution of 12 bits, in order to preserve subtle variations of grey levels that can aid to diagnoses. High-quality scanning systems, however, can hardly provide SNRs comparable to those of 12 significant bits (> 70 dB). In addition, instrument set-up (e.g. radiometric calibration, spot size, sampling rate) are critical for performances. Therefore, an image digitized with 12 bits/pixel usally exhibits a number of significant levels (i.e. noise-free) which is lower than that expected.

Trivially, performances are penalized when noisy data are being lossessly compressed. A feasible solution is to estimate the noisiness of the data and remap them onto a more suitable number of levels, before coding. Performance penalty will be mitigated, and diagnostic quality thoroughly preserved. In fact, the outcome images will be virtually lossless, if the new quantization step size is comparable to the standard deviation of the noise.

In this paper, a method for estimating the number of significant grey levels in digital images is proposed. It makes use of information coming from the distribution of local variance of grey levels, and is totally unsupervised. Once the standard deviation of the noise has been assessed, only *significant* grey levels are considered for lossless coding. A hierarchical compression based on median filtering and context coding provides results among the best reported in the most advanced literature.

2 NOISE ASSESSMENT OF X-RAY IMAGES

A noise model that is mathematatically tractable and physically congruent with X-ray images is stated as

$$X(m,n) \;=\; f(m,n) + f(m,n)^{\gamma} \cdot u(m,n) \qquad (1)$$

in which $X(m,n)$ denotes the recorded noisy image at pixel position (m,n), $f(m,n)$ the *noise-free* image and $u(m,n)$ a random process, independent of $f(m,n)$, stationary and uncorrelated, with zero mean and variance σ_u^2. For $\gamma = 0$, the model (1) reduces to additive *signal-independent* noise and the variance of (1) yields

$$\sigma_X^2(m,n) \;=\; \sigma_f^2(m,n) + \sigma_u^2 \qquad (2)$$

On homogeneous areas, in which $\sigma_f(m,n) \cong 0$, $\sigma_X^2(m,n)$ equals the noise variance σ_u^2. This fact is no longer true when ensemble statistics are replaced with local spatial statistics and $\sigma_f(m,n) \neq 0$. It can be noticed, however, that on an average σ_X always exhibits a unimodal unbalanced distribution. The presence of signal spreads the histogram towards right, thereby increasing its *mean*, but without practically affecting its *mode*. Thus, the mode of σ_X yields an estimate of σ_u. Figure 2 depicts the situation for a $12-bit$ X-ray image.

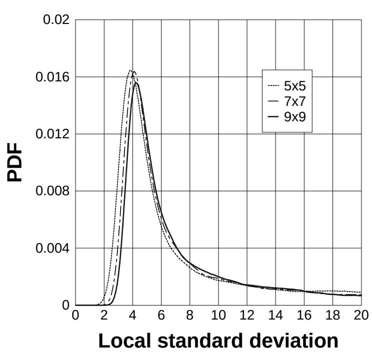

Figure 1: PDF of square root of local variance computed with several window size on a $12-bit$ X-ray chest image.

3 CONTEXT-BASED RECURSIVE NON-LINEAR INTERPOLATION

The idea of using nonlinear interpolation in a recursive fashion was first suggested as a variant of HINT [7]. The median of four pixels placed at the corners of a square or a diamond pattern, was used to predict the centre pixel. The median of a sample set with an even number of elements is defined as the average of the two middle values in the sorted set. In this way, interpolation is unlikely to be carried out *across* edges, but possibly *along* linear discontinuities. Unfortunately, the results of this *median-based* HINT were quite comparable to those of the standard scheme, on a test set of digitized X-ray images [7]. The classical linear HINT, therefore, remained unsurpassed for almost a decade [1].

3.1 Two-stage recursive nonlinear interpolation

Some experiments with separable interpolation, generalized [8] and two-stage bilinear [9], highlighted that the weak point of linear interpolation are edges: results may be partially adjusted by choosing a direction of larger correlation and carrying out interpolation along it in the first stage [10]. The idea of the present work is to replace linear interpolation with *median-based* nonlinear interpolation in the two-stage scheme [9, 10] which is forced to become nonseparable, as HINT was.

Let $\{X(m,n), m = 0,\dots,M-1; n = 0,\dots,N-1\}$, $M = p \times 2^K$, and $N = q \times 2^K$, be an integer-valued image, with p, q and K integers. The first interpolation stage at layer k yields an approximation of X, $\hat{X}$, as

$$\begin{aligned}\hat{X}(2^k(m+\tfrac{1}{2}), 2^k(n+\tfrac{1}{2})) \;&=\\ round[\mathcal{M}\{X(2^k m, 2^k n), X(2^k(m+1), 2^k n),&\quad (3)\\ X(2^k(m+1), 2^k(n+1)), X(2^k m, 2^k(n+1))\}]&\end{aligned}$$

$m = 0,\dots,M/2^k - 1, n = 0,\dots,N/2^k - 1, 1 \le k \le K$. $\mathcal{M}\{\cdot,\cdot,\cdot,\cdot\}$ stands for *median* and *round*$[\cdot]$ for roundoff to integer. Interpolation *residuals* are defined as

$$\begin{aligned}D(2^k(m+\tfrac{1}{2}), 2^k(n+\tfrac{1}{2})) \;&=\\ X(2^k(m+\tfrac{1}{2}), 2^k(n+\tfrac{1}{2})) - \hat{X}(2^k(m+\tfrac{1}{2}), 2^k(n+\tfrac{1}{2}))&\quad (4)\end{aligned}$$

The second step of interpolation is twofold:

$$\begin{aligned}\hat{X}(2^k(m+\tfrac{1}{2}), 2^k n) \;&=\\ round[\mathcal{M}\{X(2^k m, 2^k n), X(2^k(m+\tfrac{1}{2}), 2^k(n-\tfrac{1}{2})),&\\ X(2^k(m+1), 2^k n), X(2^k(m+\tfrac{1}{2}), 2^k(n+\tfrac{1}{2}))\}]&\\ \hat{X}(2^k m, 2^k(n+\tfrac{1}{2})) \;&=\\ round[\mathcal{M}\{X(2^k m, 2^k n), X(2^k(m+\tfrac{1}{2}), 2^k(n+\tfrac{1}{2})),&\\ X(2^k m, 2^k(n+1), X(2^k(m-\tfrac{1}{2}), 2^k(n+\tfrac{1}{2}))\}]&\quad (5)\end{aligned}$$

The outcome residuals will be

$$\begin{aligned}D(2^k(m+\tfrac{1}{2}), 2^k n) \;&=\\ X(2^k(m+\tfrac{1}{2}), 2^k n) - \hat{X}(2^k(m+\tfrac{1}{2}), 2^k n)&\\ D(2^k m, 2^k(n+\tfrac{1}{2})) \;&=\\ X(2^k m, 2^k(n+\tfrac{1}{2})) - \hat{X}(2^k m, 2^k(n+\tfrac{1}{2}))&\quad (6)\end{aligned}$$

Samples outside subscript ranges are skipped. Thus, the median of three pixels is used on frame borders. Starting from the $p \times q$ icon at level K, the procedure is repeated for $k = K-1,\dots,1$ to yield multiple sets of integer coefficients. Figure 2 shows the two-stage interpolation.

(a) (b)

Figure 2: Example of two-stage nonseparable median-based interpolation. First stage (a): the median of 4 neighbouring samples (dots) yields an approximation (triangle); the interpolation error (difference between a triangle and its underlying dot) is encoded. Second stage (b): interpolation is carried out from 4 dots placed at the corners of a diamond to yield an approximation (square); interpolation errors are encoded as well.

Figure 3 shows the interpolation errors arranged into subband-like regions: for each resolution, those yielded by (4) lie along the main diagonal; those coming from (6) along the co-diagonal.

At the decoder, the $p \times q$ icon is interpolated onto the quincunx grid by median filtering, the first set of residuals (4) is added to such reconstructed samples to yield all error-free pixels lying on the quincunx lattice (dots in Fig. 2(b)). These samples are interpolated according to (5) to produce approximations that are refined to reversibility by using the decoded residuals of the second stage (6). This procedure is repeated decrementing k by one, until the full-size image is retrieved for $k = 1$.

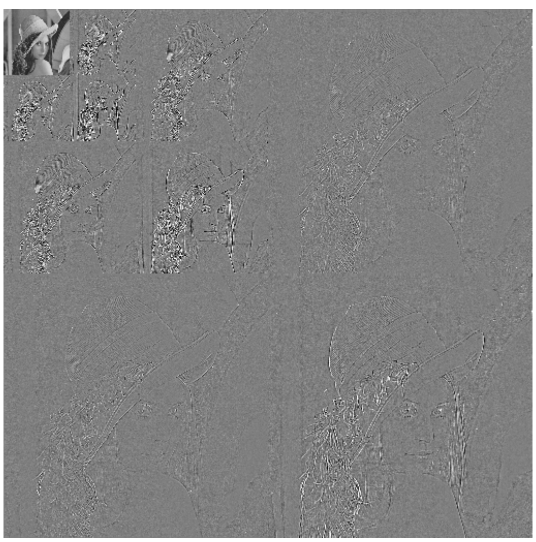

Figure 3: Icon and interpolation errors with $K = 3$.

3.2 Context-based coding

A recently proposed image compression scheme [11] has proven itself to be more efficient than DPCM schemes, usually employed for (near-)lossless compression [5]. Multiresolution decorrelation was provided by the S-transform (ST) [1], which is a dyadic Haar wavelet yielding integer valued coefficients and working on 2×2 image blocks, up to a *root* level K. The superior coding performance depends on the use of context, as well as of an empirically adjusted prediction across scale (*P-prediction*) to enhance the decorrelation capability of ST, which is indeed rather poor. Prediction errors of ST coefficients are encoded. To futher improve performances, *context-based* coding was introduced following a trend recently established for lossless compression of Medical Images [12]. Prediction error are mapped into new values (magnitudes) in such a way that larger absolute prediction errors are associated to larger magnitudes. Each magnitude is then context-coded through a function which takes into account of the average magnitudes of adjacent previously encoded coefficients as well as of the *parent* pixel. The sign and least significant bits of the prediction error, whose number is variable and increases with the magnitude, are finally encoded.

In this paper, the median-based interpolation, instead of ST, is embedded in the context-based coding scheme.

4 EXPERIMENTAL RESULTS

Trials were carried out on a variety of test images to assess the lossless compression capability of the proposed *context-based recursive nonlinear interpolation* (CRNI). Such standard 8-bit pictures as *Lenna, Bridge, Harbour, Baboon*, and *Peppers* were used for experiments. Coding results were compared with those of lossless JPEG, HINT, GRI [10], as well as of S&P, with coding method II for best results. Gross bit-rate in bits per pixel (*bpp*), with all overhead and side information, were considered for all the schemes compared throughout. Table 1 shows

Table 1: Gross bit-rates (*bpp*) of test images for several lossless compression methods: JPEG denotes best predictor. HINT and GRI use built-in Huffman codebooks, S&P and CRNI use classified arithmetic coding.

IMAGE	HINT	GRI	CRNI	S&P	JPEG
Lenna	4.71	4.44	4.21	4.18	4.71
Harbour	5.60	5.16	4.80	4.73	5.23
Baboon	6.74	6.49	6.31	6.19	6.63
Bridge	6.09	5.41	5.40	5.58	5.76
Peppers	4.69	4.62	4.34	4.37	4.77

that, for all the test images, CRNI and S&P are always far superior to the other methods, which are not based on context, and on an average the two schemes are comparable to each other. The efficiency of CRNI is greater than that of S&P, as appears from a new set of trials, carried out on medical images, whose results are reported in Figure 4 as a scatterplot. A test set of 25 1024×1024 X-ray chest images, originally scanned at 12 *bit/pixel*, have been remapped onto fewer greylevels through rescaling by their noise standard deviation. As an example, the test image used for the plot of Fig. 2 was normalized to its estimated standard deviation of noise $\sigma_u = 4.1$. This operation has the effect of reducing the code rate from 5.67 *bpp* to 3.67 *bpp*. Although the images are differently textured, CRNI always outperforms S&P (all the 25 stars lie below the diagonal line). Incidentally, the application of a *P-prediction* (see [11]) in CRNI is unessential: it has been noticed that the bit-rates in the two cases always differ by less than 0.1%.

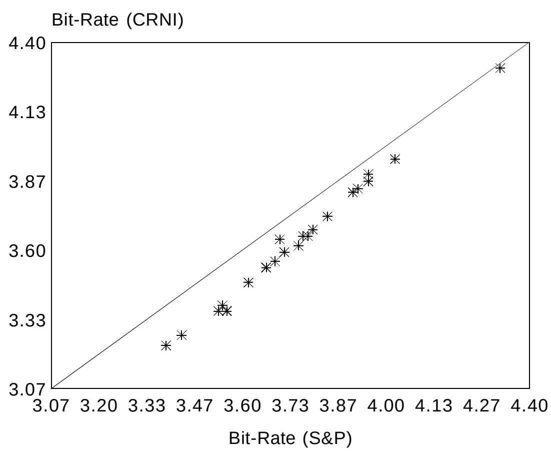

Figure 4: Bit-Rates of CRNI plotted versus those of S&P for the test set of 25 remapped X-ray chest images.

5 CONCLUSIONS

CRNI is a *non-expansive* pyramid employing median filtering for interpolation and noise feedback (roundoff of interpolated values) to achieve perfect reconstruction of integer valued images. Two sets of residuals are produced at the encoder for each layer. Interpolation both at the encoder and at the decoder is carried out in two steps, interleaved with the encoding/decoding of the two sets of residuals, in order for all error-free pixels to be used for interpolation. A context-based strategy is adopted for encoding the residuals, to take advantage of the self-similarity of coefficients at different scales. A coding scheme is borrowed from a work by Said & Pearlman, whose superior use of context and joint use of arithmetic coding make it a winner. Nevertheless, the overall compression is comparable, as proven on standard grayscale images, and significantly enhanced on a test set of $12 - bits$ X-ray images, which have been remapped onto approximately one thousand of greylevels in order to exhibit noise with variance one.

ACKNOWLEDGEMENT

Work partially supported by Ministry of University and Scientific Research, nationwide project on Telemedicine.

References

[1] S. Wong, L. Zaremba, D. Gooden, and H. K. Huang, "Radiologic Image Compression -A Review," *Proc. of the IEEE*, **83**(2), 194–219, 1995.

[2] P. C. Cosman, R. M. Gray, and R. A. Olshen, "Evaluating Quality of Compressed Medical Images: SNR Subjective Ratings and Diagnostic Accuracy," *Proc. of the IEEE*, **82**(6), 919–932, 1994.

[3] J. Ström and P. C. Cosman, "Medical image compression with lossless regions of interest," *Signal Processing*, **59**(2), 155–171, 1997.

[4] B. Aiazzi, L. Alparone, S. Baronti, G. Chirò, F. Lotti, M. Moroni, "A pyramid-based error-bounded encoder: An evaluation on X-ray chest images," *Signal Processing*, **59**(2), 73–187, 1997.

[5] K. Chen and T. V. Ramabadran, "Near-Lossless Compression of Medical Images Through Entropy-Coded DPCM," *IEEE Trans. Medical Imaging*, **13**(3), 538–548, 1994.

[6] B. Aiazzi, L. Alparone, S. Baronti, and F. Lotti, "Lossless Image Compression by Quantization Feedback in a Content-Driven Enhanced Laplacian Pyramid," *IEEE Trans. Image Processing*, **6**(6), 831–843, 1997.

[7] P. Roos, M. A. Viergever, M. C. A. Van Duke, and J. H. Peters, "Reversible Intraframe Compression of Medical Images," *IEEE Trans. Medical Imaging*, **7**(4), 328–336, 1988.

[8] B. Aiazzi, L. Alparone, and S. Baronti, "A Reduced Laplacian Pyramid for Lossless and Progressive Image Communication," *IEEE Trans. Communications*, **44**(1), 18–22, 1996.

[9] A. Abrardo, L. Alparone, F. Bartolini, "Encoding-interleaved hierarchical interpolation for lossless image compression," *Signal Processing*, **56**(2), 321–328, 1997.

[10] B. Aiazzi, P. S. Alba, L. Alparone, S. Baronti, "Three-Dimensional Lossless Compression Based on a Separable Generalized Recursive Interpolation," *Proc. IEEE-ICIP'96*, **I/III**, 85–88, 1996.

[11] A. Said and W. A. Pearlman, "An Image Multiresolution Representation for Lossless and Lossy Compression," *IEEE Trans. Image Processing*, **5**(9), 1303–1310, 1996.

[12] T. V. Ramabadran and K. Chen, "The Use of Contextual Information in the Reversible Compression of Medical Images," *IEEE Trans. Medical Imaging*, **11**(2), 185–195, 1992.