

PROGRESSIVE LOSSLESS COMPRESSION OF MEDICAL IMAGES

Armando J. Pinho and António J. R. Neves

Signal Processing Lab, DETI / IEETA
University of Aveiro, 3810-193 Aveiro, Portugal
ap@ua.pt — an@ua.pt

ABSTRACT

This paper describes a lossless compression method for medical images that produces an embedded bit-stream, allowing progressive lossy-to-lossless decoding with L_∞ oriented rate-distortion. The experimental results show that the proposed technique produces better average lossless compression results than several other compression methods, including JPEG2000, JPEG-LS and JBIG, in a publicly available medical image database containing images from several modalities.

Index Terms— Medical image compression, lossless image coding, progressive transmission, finite-context models.

1. INTRODUCTION

It is well-known that most medical imaging modalities produce huge amounts of data. Moreover, for several reasons, it is frequently needed to store or transmit those images at the highest possible fidelity. The discussion around the question of whether or not all data present in a medical image should be preserved is a long-lasting one (see, for example, [1, 2]). In our opinion, it seems reasonable that for long-term archiving (from medical sources or from some other sources) images should be compressed using reversible algorithms. In fact, whereas an image compressed with a lossless algorithm can be re-compressed more efficiently in the future by a better algorithm without losing any information, the same is usually not true with lossy-compressed images.

Despite the large number of works regarding embedded image coding, only a few of them address distortion measures that do not rely on error averages and, particularly, on the L_2 -norm. The works of Avcibaş *et al.* [3, 4], Alecu *et al.* [5, 6, 7] and Krivoulets [8, 9] are remarkable exceptions, since they address the problem of generating embedded bit-streams that minimize the L_∞ -norm of the reconstruction error. Avcibaş *et al.* proposed an approach that relies on a predictive-based method that successively refines the probability density function (pdf) used to estimate each pixel and by restricting the region of support of the pdf to fixed size intervals, which have to be predefined before encoding [3, 4]. Almost simultaneously, Alecu *et al.* proposed a wavelet-based scheme that allows full L_∞ scalability [5, 6, 7]. This algorithm was compared with JPEG2000 in terms of L_∞ rate-distortion, showing better results [7]. Krivoulets proposed a lossy plus near-lossless layered compression scheme with embedded quantization of the difference signal, where the initial lossy layer is encoded with JPEG2000 [8, 9].

In this paper, we describe a progressive lossless compression method and we present the results of its use in the medical imaging area. This method is based on binary tree decomposition and

This work was supported in part by the Fundação para a Ciência e a Tecnologia (FCT).

finite-context modeling, producing a L_∞ -constrained embedded bit-stream [10]. We studied the performance of the method in a medical image test set collected by Starosolski, containing images from Computed Radiography (CR), Computed Tomography (CT), Magnetic Resonance (MR) and Ultrasound (US) modalities of several anatomical regions, bit depths and acquired with devices from several vendors [11]. We compared the results with those attained by the current image coding standards, namely JPEG2000, JBIG and JPEG-LS.

2. THE CODING APPROACH

2.1. Hierarchical organization of the intensity levels

The compression technique is based on a hierarchical organization of the intensity levels of the image. This organization is obtained by means of a binary tree. Each node of the binary tree, n , represents a certain subset, S^n , of the intensities of the image. The root node is associated with the complete set of image intensities, $\mathcal{I} = \{I_1, I_2, \dots, I_N\}$. Therefore, $S^n \subset \mathcal{I}$ and $S^1 \equiv \mathcal{I}$. Each node possesses a representative intensity, I^n , given by

$$I^n = \left\lfloor \frac{I_m^n + I_M^n}{2} \right\rfloor, \quad (1)$$

where I_m^n and I_M^n are, respectively, the smallest and largest intensities in S^n , and where $\lfloor x \rfloor$ denotes the largest integer less than or equal to x . Computing the value of I^n according to (1) leads to the smallest possible L_∞ reconstruction error when the intensities associated to node n (those in S^n) are substituted by I^n . This error is $\epsilon_\infty^n = I_M^n - I^n$.

During encoding (or decoding), a tree is constructed from the root node to the leaves, always choosing to expand the node that implies the highest reduction in the reconstruction error. In case of having several nodes leading to the same error, one is arbitrarily chosen. The only constraint is that the decoder picks the same node.

When node n is expanded, two subsets are formed by splitting S^n into S_l^n and S_r^n , such that $S_l^n = \{I \in S^n : I \leq I^n\}$ and $S_r^n = \{I \in S^n : I > I^n\}$. Therefore, all intensities $I \in S^n$ that are smaller or equal to the representative intensity, I^n , go to the set of the left node, whereas those that are larger go to the set of the right node. This procedure is repeated until expanding all nodes, i.e., until having a tree with N leaves (the number of image intensities).

At the decoder side, an identical binary tree is constructed. To do this, the decoder needs only to know the set of intensity values that occur in the image. This set can be efficiently communicated by sending the maximum intensity value, I_N , followed by a string of I_N bits, such that if the n^{th} bit of the string is one, then the intensity $n - 1$ exists in the image (if it is zero, then it does not exist).

It is interesting to note that two images sharing the same set of intensity values will have exactly the same binary tree and that it will be expanded exactly in the same node order, independently of how many times each intensity value is used and where it occurs in the image.

When node n is expanded, all pixels having intensity I^n have to change to one of the two new representative intensities, I_l^n or I_r^n . This can be seen as a region of arbitrary shape, containing zeros and ones, that needs to be communicated to the decoder. The shape of the region is known by the decoder (it corresponds to the position of the pixels in the image with current reconstruction value equal to I^n). However, the zeros and ones (corresponding to the positions of the new I_l^n and I_r^n intensities) need to be encoded.

2.2. Encoding of the binary masks

The encoding of these binary masks is of key importance to the final performance of the method, which is attained by a carefully chosen context modeling that drives a binary arithmetic encoder. [12]. The contexts are constructed based on a template where the context pixels (sixteen, at most) are numbered according to their distance to the encoding pixel. A particular context is represented using a sequence of bits, $b_1 b_2 \dots b_k$, where $b_i = 0$ if $|I(i) - I_l^n| \leq |I(i) - I_r^n|$ and is zero otherwise, and where $I(i)$ denotes the intensity of the pixel in the current reconstructed image corresponding to position i of the context template.

The value of k varies as encoding proceeds, because it is expected to have larger mask regions initially and smaller regions when $n \approx N$. Therefore, to avoid the problem of context dilution, smaller values of k are typically used when $n \approx N$. In this work, we present results based on three approaches regarding this context adaptation issue. One of them, the fastest, is based on a function that was fitted in order to predict the values of k [13]. We call this approach “Pred” and denote the value of the prediction by \hat{k} . The second approach uses \hat{k} as a starting point, and then picks the value of k given the lowest bitrate, for $k = \{\hat{k} - 2, \dots, \hat{k} + 2\}$. We call this approach, “Var2”. Finally, the third approach performs a full search for the best value of k . We call it “Full” and, as expected, it is the most demanding in terms of computational resources, but it also gives the best possible results for this context modeling setup.

3. EXPERIMENTAL RESULTS AND CONCLUSION

Table 1 presents lossless compression results, in bits per pixel, obtained using JPEG2000 [14, 15], JBIG [16, 17], JPEG-LS [18, 19] and the proposed L_∞ progressive compression algorithm. In this experiment, we used a publicly available medical image database composed of 48 images from CR, CT, MR and US modalities (twelve from each one). These images can be obtained from <http://sun.aei.polsl.pl/~rstaros/mednat/index.html>.

JPEG2000 lossless compression was obtained using version 5.1 of the JJ2000 codec with default parameters for lossless compression¹. JBIG compression was obtained using version 1.6 of the JBIG Kit package², with sequential coding (-q flag). Note that this flag disables progressive encoding inside a bit-plane. However, the L_∞ nature is maintained because the encoding is still done bit-plane by bit-plane, from the most significant to the least significant. JPEG-LS coding was obtained using version 2.2 of the SPMG JPEG-LS codec

with default parameters³.

The results presented in Table 1 show that the proposed approach attains better lossless compression in the four groups of images. The largest difference occurs in the images from the CT and MR modalities. A considerable part of the gain that is attained is due to the sparse nature of the histogram of intensities of the images from these groups, as can be guessed from the number of levels indicated in the table. It is known that histogram sparseness originates loss of performance in predictive or transform based image compression techniques [20]. One of the simplest approaches for avoiding this loss of performance involves histogram packing (an one-to-one order preserving mapping) before compression [21]. However, histogram packing cannot be applied in lossy compression [22] and, therefore, the simultaneous use of histogram packing and progressive decoding is not possible. Regarding this issue, the proposed approach is clearly advantageous, because it is virtually immune to histogram sparseness and allows progressive decoding.

Figure 1 shows operational rate-distortion curves, in the L_∞ -norm sense, for images “cr.17218”, “ct.135960.001” and “us.19773”, showing the best behavior in the case of the proposed method. It can also be observed that, generally, JBIG provides better L_∞ rate-distortion than JPEG2000, but loses when the rates approach the lossless point. In fact, from the four images modalities, JBIG was only able to attain a better average lossless rate for US. Recall that the JPEG-LS standard does not allow progressive decoding.

4. CONCLUSION

In this work, we investigated the appropriateness of a progressive lossless image coding method based on binary tree decomposition and finite-context arithmetic coding applied to the compression of medical images. The technique produces an embedded bit-stream optimized for L_∞ -constrained decoding. In addition to its good performance, both in terms of lossless compression and L_∞ rate-distortion, it is immune to histogram sparseness, a characteristic not present in most predictive or transform based methods and that might considerably reduce the compression efficiency of those methods [20, 21, 22]. Three different approaches have been tested for context adaptation, to allow choosing between compression performance and compression according to convenience.

5. REFERENCES

- [1] S. Wong, L. Zaremba, D. Gooden, and H. K. Huang, “Radiologic image compression — a review,” *Proceedings of the IEEE*, vol. 83, no. 2, pp. 194–219, Feb. 1995.
- [2] D. A. Clunie, “Lossless compression of grayscale medical images - effectiveness of traditional and state of the art approaches,” in *Medical Imaging — Proc. SPIE*, 2000, pp. 74–84.
- [3] İ. Avcıbaşı, N. Memon, B. Sankur, and K. Sayood, “A progressive lossless/near-lossless image compression algorithm,” *IEEE Signal Processing Letters*, vol. 9, no. 10, pp. 312–314, Oct. 2002.
- [4] İ. Avcıbaşı, N. Memon, B. Sankur, and K. Sayood, “A successively refinable lossless image-coding algorithm,” *IEEE Transactions on Communications*, vol. 53, no. 3, pp. 445–452, Mar. 2005.

¹<http://jj2000.epfl.ch>.

²<http://www.cl.cam.ac.uk/~mgk25/jbigkit/>.

³The original web-site of this codec, <http://spmge.ece.ubc.ca>, is currently unavailable. However, it can be obtained from ftp://www.ieeta.pt/~ap/codecs/jpeg_ls_v2.2.tar.gz.

- [5] A. Alecu, A. Munteanu, P. Schelkens, J. Cornelis, and S. Dewitte, "Wavelet-based L-infinite scalable coding," *Electronics Letters*, vol. 38, no. 22, pp. 1338–1340, Oct. 2002.
- [6] A. Alecu, A. Munteanu, P. Schelkens, J. Cornelis, and S. Dewitte, "Wavelet-based fixed and embedded L-infinite-constrained image coding," *Journal of Electronic Imaging*, vol. 12, no. 3, pp. 522–538, July 2003.
- [7] A. Alecu, A. Munteanu, J. Cornelis, and P. Schelkens, "Wavelet-based scalable L-infinity-oriented compression," *IEEE Trans. on Image Processing*, vol. 15, no. 9, pp. 2499–2512, Sept. 2006.
- [8] A. Krivoulets, "A method for progressive near-lossless image compression," in *Proc. of the IEEE Int. Conf. on Image Processing, ICIP-2003*, Barcelona, Spain, Sept. 2003, vol. 2, pp. 185–188.
- [9] A. Krivoulets, "Progressive near-lossless coding of medical images," in *Proc. of the 3rd Int. Symp. on Image and Signal Processing, ISPA-03*, Sept. 2003, pp. 202–207.
- [10] A. J. Pinho and A. J. R. Neves, "L-infinity progressive image compression," in *Proc. of the Picture Coding Symposium, PCS-07*, Lisbon, Portugal, Nov. 2007.
- [11] R. Starosolski, "Performance evaluation of lossless medical and natural continuous tone image compression algorithms," in *Medical Imaging — Proc. SPIE*, Sept. 2005, pp. 116–127.
- [12] T. C. Bell, J. G. Cleary, and I. H. Witten, *Text compression*, Prentice Hall, 1990.
- [13] A. J. Pinho and A. J. R. Neves, "A context adaptation model for the compression of images with a reduced number of colors," in *Proc. of the IEEE Int. Conf. on Image Processing, ICIP-2005*, Genova, Italy, Sept. 2005, vol. 2, pp. 738–741.
- [14] ISO/IEC, *Information technology - JPEG 2000 image coding system*, ISO/IEC International Standard 15444-1, ITU-T Recommendation T.800, 2000.
- [15] D. S. Taubman and M. W. Marcellin, *JPEG 2000: image compression fundamentals, standards and practice*, Kluwer Academic Publishers, 2002.
- [16] ISO/IEC, *Information technology - Coded representation of picture and audio information - progressive bi-level image compression*, International Standard ISO/IEC 11544 and ITU-T Recommendation T.82, Mar. 1993.
- [17] D. Salomon, *Data compression - The complete reference*, Springer, 4th edition, 2007.
- [18] ISO/IEC, *Information technology - Lossless and near-lossless compression of continuous-tone still images*, ISO/IEC 14495-1 and ITU Recommendation T.87, 1999.
- [19] M. J. Weinberger, G. Seroussi, and G. Sapiro, "The LOCO-I lossless image compression algorithm: principles and standardization into JPEG-LS," *IEEE Trans. on Image Processing*, vol. 9, no. 8, pp. 1309–1324, Aug. 2000.
- [20] A. J. Pinho, "On the impact of histogram sparseness on some lossless image compression techniques," in *Proc. of the IEEE Int. Conf. on Image Processing, ICIP-2001*, Thessaloniki, Greece, Oct. 2001, vol. III, pp. 442–445.
- [21] A. J. Pinho, "An online preprocessing technique for improving the lossless compression of images with sparse histograms," *IEEE Signal Processing Letters*, vol. 9, no. 1, pp. 5–7, Jan. 2002.
- [22] P. J. S. G. Ferreira and A. J. Pinho, "Why does histogram packing improve lossless compression rates?," *IEEE Signal Processing Letters*, vol. 9, no. 8, pp. 259–261, Aug. 2002.

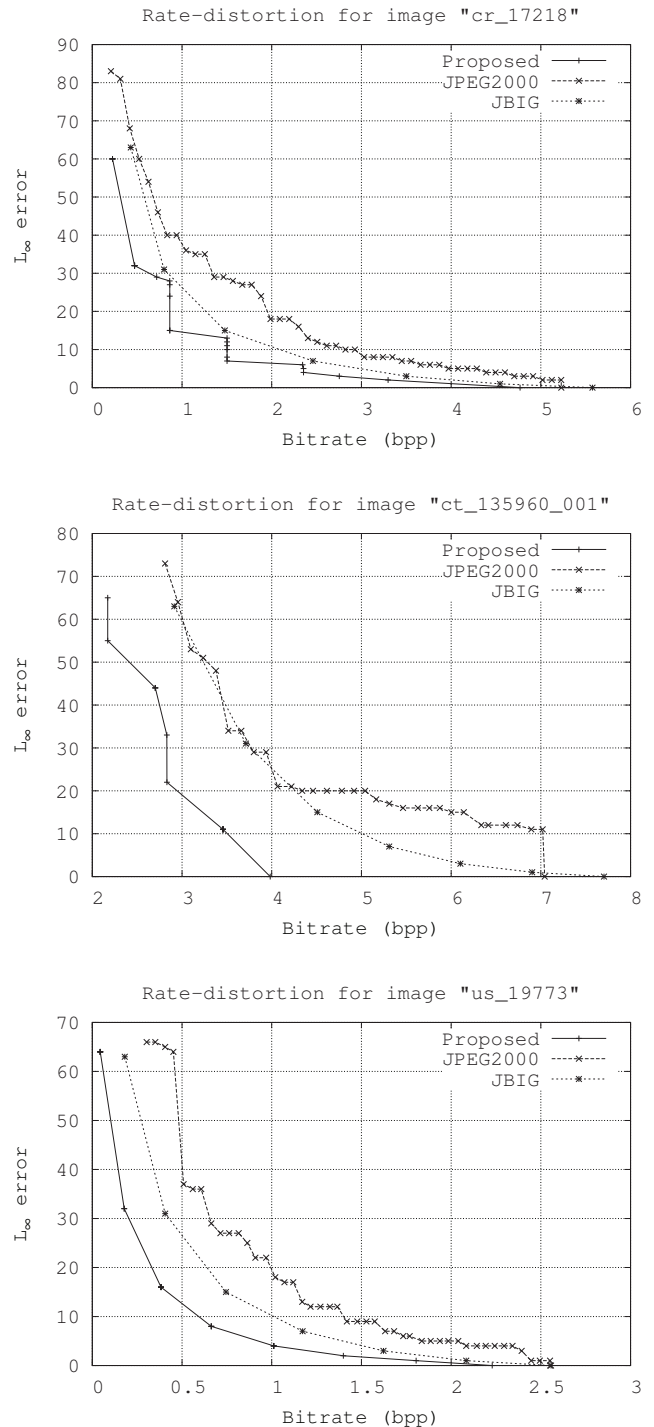


Fig. 1. Operational L_∞ -norm rate-distortion curves of the "cr.17218", "ct.135960.001" and "us.19773" images for the proposed method ("Pred" version of context adaptation), for JPEG2000 and for JBIG.

Table 1. Lossless compression results, in bits per pixel, obtained with the proposed method, with JPEG2000, JBIG and JPEG-LS. The images have been taken from a publicly available database and cover four medical imaging modalities: Computed Radiography (CR), Computed Tomography (CT), Magnetic Resonance (MR) and Ultrasound (US). The “Prev”, “Var2” and “Full” versions of the proposed method refer to three different approaches for context adaptation (additional details in the main text).

Image	Rows × Cols	Depth	Levels	JPEG2000	JBIG	JPEG-LS	Proposed		
							Pred	Var2	Full
cr_17218	1792 × 2392	12	2068	5.228	5.575	5.221	4.765	4.735	4.715
cr_17220	2048 × 2500	12	3186	3.712	4.118	3.785	3.745	3.723	3.704
cr_17222	2392 × 1792	12	2939	4.493	4.878	4.546	4.539	4.509	4.489
cr_4503	2010 × 1670	10	256	4.800	4.164	4.733	2.781	2.771	2.758
cr_4507	1760 × 1760	10	1024	2.158	2.432	2.188	2.191	2.174	2.162
cr_4509	2140 × 1760	10	882	4.194	4.571	4.236	3.924	3.909	3.894
cr_pacem_1	1910 × 1716	16	24180	11.186	11.665	10.903	9.635	9.602	9.586
cr_pacem_2	1965 × 1531	16	28627	10.736	11.305	10.537	9.515	9.470	9.456
cr_rtg_jb	746 × 612	16	3280	11.223	11.698	11.029	6.792	6.767	6.759
cr_siem_01_02	2128 × 1744	10	913	5.317	5.677	5.242	5.083	5.065	5.052
cr_siem_14_02	2368 × 1760	10	638	2.981	3.219	2.916	2.462	2.452	2.445
cr_slim_1	2031 × 1866	16	26539	11.046	11.518	10.759	9.664	9.626	9.610
Average	–	–	–	5.845	6.151	5.782	5.175	5.151	5.136
ct_135960_001	512 × 512	16	2442	7.043	7.704	6.766	3.946	3.849	3.826
ct_135960_005	512 × 512	16	2806	7.009	7.670	6.706	4.040	3.922	3.896
ct_17	512 × 512	12	1883	4.879	5.507	4.599	4.194	4.159	4.153
ct_27154	512 × 512	12	1300	2.739	3.106	2.600	2.100	2.033	2.025
ct_29513	340 × 340	12	2570	5.259	5.678	4.829	4.631	4.543	4.530
ct_29920	512 × 512	12	1723	4.879	5.438	4.617	4.008	3.970	3.962
ct_3030	691 × 512	16	778	11.690	12.227	11.493	5.224	5.200	5.189
ct_3071	512 × 512	16	1696	9.406	9.659	9.033	4.983	4.951	4.927
ct_4006	512 × 512	16	2100	11.444	11.898	11.290	6.437	6.413	6.396
ct_4087	512 × 512	16	1731	11.704	12.273	11.535	6.429	6.406	6.392
ct_4165	512 × 512	16	1735	12.166	12.644	12.010	6.842	6.817	6.798
ct_tk_kl_piers0021	512 × 512	16	2644	8.893	9.101	8.573	5.386	5.343	5.317
Average	–	–	–	8.334	8.822	8.089	4.874	4.825	4.809
mr_2321	512 × 512	16	894	11.287	12.209	11.337	5.532	5.519	5.512
mr_2331	512 × 512	16	893	11.348	12.200	11.438	5.651	5.636	5.628
mr_2337	512 × 512	16	1047	8.398	9.013	8.295	4.165	4.137	4.130
mr_2371	512 × 512	16	1415	8.181	8.589	8.135	4.153	4.122	4.115
mr_2412	512 × 512	16	1300	10.900	11.797	10.888	5.573	5.559	5.557
mr_2807	256 × 256	16	1858	12.555	13.597	12.366	8.555	8.466	8.463
mr_2882	512 × 512	16	501	1.852	1.957	1.725	1.005	0.995	0.990
mr_2896	512 × 512	16	604	9.648	9.964	9.347	4.431	4.410	4.398
mr_6624	256 × 256	16	795	12.027	10.683	12.265	6.706	6.689	6.684
mr_6706	256 × 256	16	1088	12.680	13.470	12.548	7.309	7.284	7.279
mr_6774	512 × 512	16	1799	10.743	11.249	10.645	5.533	5.520	5.518
mr_6837	256 × 256	16	1055	11.300	12.160	11.117	6.370	6.339	6.337
Average	–	–	–	9.389	9.940	9.321	4.809	4.789	4.783
us_19773	480 × 640	8	256	2.552	2.557	2.277	2.220	2.202	2.193
us_27704	480 × 640	8	249	3.493	3.238	3.110	2.838	2.818	2.792
us_27743	480 × 640	8	246	3.663	3.388	3.232	2.928	2.909	2.883
us_28279	480 × 640	8	250	3.090	2.619	2.552	2.295	2.278	2.266
us_28282	480 × 640	8	247	3.070	3.266	2.783	2.784	2.763	2.747
us_28289	480 × 640	8	254	2.866	2.226	2.339	1.978	1.968	1.944
us_28322	480 × 640	8	213	3.515	3.428	3.283	2.916	2.897	2.881
us_28329	480 × 640	8	213	3.940	3.716	3.557	3.139	3.120	3.101
us_28348	480 × 640	8	217	3.629	3.164	3.117	2.627	2.616	2.601
us_3393	476 × 640	8	218	2.926	3.048	2.584	2.505	2.485	2.471
us_3403	484 × 584	8	256	2.762	2.572	2.524	2.313	2.282	2.245
us_3405	476 × 640	8	197	2.102	1.635	1.608	1.440	1.425	1.394
Average	–	–	–	3.138	2.908	2.750	2.501	2.483	2.462