REGION-BASED NEAR-LOSSLESS IMAGE COMPRESSION

Armando J. Pinho

Dep. Electrónica e Telecomunicações / IEETA
Universidade de Aveiro, 3810–193 Aveiro, Portugal
`ap@det.ua.pt` — `http://www.ieeta.pt/˜ap`

ABSTRACT

We present a near-lossless technique for the compression of images, which is based on the partitioning of the image into regions of constant intensity. The boundary information associated with the image partition is encoded with the method of the transition points. The compression of the intensities of the regions is based on the usual entropy encoding of the context-modeled prediction residuals. The experimental results show that this approach is able to provide significant compression improvements in images having sparse histograms, for small L_∞ errors.

1. INTRODUCTION

Traditionally, image coding techniques have been classified into one of two categories: lossless or lossy. Lossless methods are typically chosen for applications where small image details can be of paramount importance, such as in medical and space imaging or in remote sensing. On the other hand, lossy methods are required in situations where significant compression ratios are sought. This is the case, for example, in mobile applications or in digital photography, for which, generally, losing some image detail is not too problematic.

Recently, a third class of techniques has been generating a great interest among image coding researchers. It is positioned in between the other two, motivating the name by which it is generally known, "near-lossless image coding" (the designation "L_∞ constrained coding" is also frequently used). In L_∞ constrained image coding a tight numerical bound on the errors is assured. This contrasts with most of the lossy techniques, which are optimized for minimizing the L_2 norm of the error and, therefore, cannot guarantee some pre-defined maximum absolute error in the reconstructed image.

The most popular approach for near-lossless image compression is based on DPCM associated with more or less elaborated predictors and error modeling schemes (see, for example, [1–4]). However, some other approaches have been proposed, such as those based on vector quantization [5, 6], on wavelets [7, 8], or on some hybrid combinations [9].

Mostly motivated by two reasons, region-based image coding has been addressed only in the context of low and very low bit-rates. One of such reasons is that traditional encoding approaches, and specially those relying on block-based fixed-size segmentation, do not behave well at low bit-rates. Therefore, since region-based methods resort on more natural concepts such as contours and textures, they generally yield more satisfactory visual results at low bit-rates. The other reason has been an impediment for the use of region-based techniques at higher bit-rates: generally, the encoding of the boundaries of the regions is too costly in terms of bit-rate, immediately discouraging any attempts in that direction. However, we believe that recent results on the compression of high complexity contours [10] are able to provide the means to remove this limitation. Is precisely this path that we explore in this work.

In this paper, we present a L_∞ constrained technique for the compression of images, which is based on the partitioning of the image into regions of constant intensity. The boundary information associated with the image partitions is encoded with the method of the transition points [10], which provides a very efficient compression of high complexity contour maps. The compression of the (constant) intensities of the regions is based on the usual entropy encoding of the context-modeled prediction residuals. The experimental results show that this approach is able to provide significant compression improvements in images having sparse histograms, if the L_∞ error is kept small.

2. THE REGION-BASED ENCODER

The first step that is performed by the encoder, which is also the only step introducing distortion, requires an uniform quantization of the image, i.e.,

$$\tilde{g}(r,c) = \left\lfloor \frac{g(r,c)}{2\delta+1} \right\rfloor (2\delta+1) + \delta$$

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

where $g(r,c)$ and $\tilde{g}(r,c)$ denote, respectively, the original and quantized pixel values of the image, and δ denotes the maximum quantization error, i.e.,

$$|g(r,c) - \tilde{g}(r,c)| \leq \delta, \quad \forall(r,c)$$

Based on the quantized image, regions of constant intensity are determined. The set of these regions forms a partition of the image and, therefore, represents the original image within the pre-defined L_∞ error of δ. At this stage the information is separated into two distinct components: (1) a contour map, representing the boundaries of the regions; (2) a set of intensities, representing the interior of the regions.

The encoding of the contour map is performed using the method of the transition points [10–12]. This technique is very efficient, particularly for high complexity contours, allowing lossless compression of arbitrary contour maps. In this work, we used the version of the encoder described in [10], which relies on a four-symbol adaptive context-based arithmetic encoder that calculates contexts in the (binary) domain of the contour map.

The encoding of the intensities of the regions relies on the usual scheme of prediction followed by context modeling and entropy coding of the prediction residuals. However, there is a fundamental difference that characterizes the encoder that we describe here, in comparison to other methods also based on the same principles, such as LOCO-I [2] or CALIC [13]: each encoding step addresses an entire region instead of a single pixel.

The predictor used by our encoder is the same as the fixed predictor of JPEG-LS [14,15], also known as the "MED" predictor. It uses the values of three (causal) pixels, a, b and c, to produce the following estimate of x (see Fig. 1a):

$$\hat{x} = \begin{cases} \min(a,b), & \text{if } c \geq \max(a,b) \\ \max(a,b), & \text{if } c \leq \min(a,b) \\ a+b-c, & \text{otherwise.} \end{cases}$$

In our case, x represents the first pixel (in raster scan order) of the region being processed (see Fig. 1b). Moreover, it is also seen as the pixel that represents the intensity value of that region. It is important to notice that we always have $x \neq a$ and $x \neq b$. This is because a and b belong to neighbor regions of the region being processed and, by definition, two neighbor regions cannot share the same intensity level.[1]

The prediction residuals, $e = x - \hat{x}$, are encoded by an adaptive context-based arithmetic encoder. The determination of contexts (a total of 144) is based on the following four factors:

1. The activity level surrounding x. This is given by $(|a-c|+|b-c|+|b-d|)/3$. The values resulting from this measure are quantized into eight quantization regions: 0, 1, 2, 3, $\{4,5,6\}$, $\{7,\ldots,14\}$, $\{15,\ldots,30\}$ and $\{31,\ldots\}$.

2. The mean of the absolute residuals occurred in the (already processed) neighbor regions. These values are quantized into three levels: 0, $\{1,2,3\}$ and $\{4,\ldots\}$.

3. The sign (negative, zero or positive) of the mean of the residuals occurred in the (already processed) neighbor regions.

4. A binary parameter indicating if the predicted value corresponds to an impossible intensity for the region being processed (remember that a region cannot have the same intensity as one of its neighbor regions).

[1] Note also that, because we assume 4-connected neighborhoods, it is possible to have $x = c$ or $x = d$.

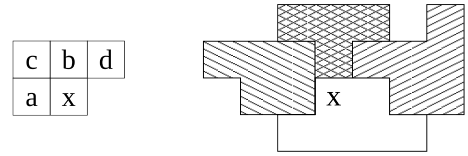

(a) (b)

Fig. 1. (a) Causal context used both for the generation of predictions and for context modeling. (b) Example of a region being processed (the white one) and some already processed neighbor regions (those painted with patterns). The "x" marks the first pixel (in raster scan order) of the region being processed, and represents the intensity of that region.

The total number of symbols handled by the arithmetic encoder is given by the number of possible values of the prediction residuals, i.e., $2N-1$, where N denotes the number of different intensity values contained in the (uniformly quantized) image. However, for a given prediction, $\hat{x}$, the residuals are limited to values in the set $\{-\hat{x},\ldots,N-1-\hat{x}\}$, assuming that $x \in \{0,\ldots,N-1\}$. Moreover, and also for a given $\hat{x}$, all residuals conducting to intensities known to belong to neighbor regions cannot occur. Therefore, for each symbol being encoded the arithmetic encoder is adapted in order to take advantage of these two aspects.

Note that, to be able to use the information related to region neighboring, the intensity encoder needs to know the boundary information associated with the image partition. Therefore, the boundary information should be encoded first.

3. EXPERIMENTAL RESULTS

In this section we present compression results of the encoder described in the preceding section, and compare these results with two state-of-the-art techniques: the JPEG-LS

Image	JPEG-LS			Improved CALIC [4]			Region-based		
	$\delta = 1$	$\delta = 3$	$\delta = 7$	$\delta = 1$	$\delta = 3$	$\delta = 7$	$\delta = 1$	$\delta = 3$	$\delta = 7$
air2.r	2.837	2.056	1.439	2.670	**1.861**	**1.268**	**2.367**	2.026	1.410
cafe.cyan	3.412	2.403	1.615	3.218	**2.185**	**1.430**	**2.219**	2.219	1.582
chart.l	1.138	0.840	0.615	0.972	0.668	**0.457**	**0.813**	**0.637**	0.468
cmpnd2.b	0.996	0.724	0.491	0.826	**0.564**	**0.362**	**0.754**	0.632	0.412
finger.raw	4.030	2.894	1.946	**3.822**	**2.665**	**1.767**	4.070	2.946	2.087
hotel.y	2.872	1.877	1.111	**2.697**	**1.686**	**0.953**	3.004	1.921	1.150
us.raw	1.643	1.142	0.792	**1.596**	**1.093**	0.718	1.718	1.136	**0.680**

Table 1. This table compares the compression performance of the region-based image coding method with the JPEG-LS standard and also with the improved version of CALIC for near-lossless compression described in [4]. Compression values, in bits per pixel, are given for L_∞ errors of 1, 3 and 7. The best result for each combination of image and δ are displayed in boldface.

standard [14] and the improved version of CALIC for near-lossless image compression [4]. To facilitate the comparison, we used the same set of images[2] of [4]. We also used the same set of values for δ, i.e., 1, 3 and 7.

Table 1 summarizes the experimental results, displaying compression performances in terms of bits per pixel. The results concerning JPEG-LS were obtained using version 2.1 of the codec developed by the Signal Processing & Multimedia Group at the University of British Columbia[3] (which is based on HP's implementation of JPEG-LS[4]). The results relative to the improved near-lossless CALIC are those published in [4]. The best compression ratio for each combination of image and δ are displayed in Table 1 in boldface.

The results presented in Table 1 show that the region-based approach is able to provide higher compression for small δ and for some types of images. Particularly, it outperformed the other two methods at $\delta = 1$ for the images "air2", "cafe", "chart" and "cmpnd2", offering improvements from 8.7% ("cmpnd2") to 31% ("cafe") over the second best results. Taking into consideration the size of the images, an overall (calculated over the seven test images) compression gain of 22.6% is attained by the region-based method over the improved near-lossless CALIC, for $\delta = 1$. For $\delta = 3$ and $\delta = 7$ the region-based method provides worse results in most of the cases.

4. CONCLUSIONS

This paper addresses, to the best of our knowledge for the first time, the question of L_∞ constrained image compression based on the segmentation of the uniformly quantized image into a set of constant intensity regions. Two information components are generated and encoded: (1) the boundaries of the regions; (2) the intensities of the regions.

The results that we obtained with this method show that it can be very effective in images having sparse histograms, such as "air2", "cafe", "chart" or "cmpnd2", and when the L_∞ error is kept small. Moreover, due to its "flat region" foundation, this method seems to be appropriate for encoding images containing significant areas of text, graphics or, more generally, man-made items.

The encoder described in this paper is still at an experimental stage. Most of the issues related to the part responsible for intensity coding need to be further and more systematically studied. This applies, particularly, to the prediction stage and to the residual context modeler. Nevertheless, and even at its current developing stage, the encoder that we described in this paper was able to provide sufficiently encouraging results for motivating further study.

5. ACKNOWLEDGEMENT

The author would like to thank the following people, for providing publicly available software: Alistair Moffat[5] *et al.* (arithmetic codec, which was used and adapted by the author in the implementation described in this paper); The Signal Processing & Multimedia Group at the University of British Columbia (SPMG / JPEG-LS V.2.1 codec).

6. REFERENCES

[1] K. Chen and T. V. Ramabadran, "Near-lossless compression of medical images through entropy-coded DPCM," *IEEE Trans. on Medical Imaging*, vol. 13, no. 3, pp. 538–548, Sept. 1994.

[2] These images were obtained from `ftp://www.cipr.rpi.edu/pub/image2/jpeg_cont_tone`.

[3] `http://spmg.ece.ubc.ca`.

[4] `http://www.hpl.hp.com/loco`.

[5] `http://www.cs.mu.oz.au/~alistair`.

[2] M. J. Weinberger, G. Seroussi, and G. Sapiro, “LOCO-I: A low complexity, context-based, lossless image compression algorithm,” in *Proc. of the Data Compression Conf., DCC’96*, Snowbird, Utah, Mar. 1996, pp. 140–149.

[3] L. Ke and M. W. Marcellin, “Near-lossless image compression: minimum-entropy, constrained-error DPCM,” *IEEE Trans. on Image Processing*, vol. 7, no. 2, pp. 225–228, Feb. 1998.

[4] X. Wu and P. Bao, “L_∞ constrained high-fidelity image compression via adaptive context modeling,” *IEEE Trans. on Image Processing*, vol. 9, no. 4, pp. 536–542, Apr. 2000.

[5] V. J. Mathews and P. J. Hahn, “Vector quantization using the L_∞ distortion measure,” *IEEE Signal Processing Letters*, vol. 4, no. 2, pp. 33–35, Feb. 1997.

[6] C. Zhu and Y. Hua, “Image vector quantization with minimax L_∞ distortion,” *IEEE Signal Processing Letters*, vol. 6, no. 2, pp. 25–27, Feb. 1999.

[7] M. J. Gormish, E. L. Schwartz, A. Keith, M. Boliek, and A. Zandi, “Lossless and nearly lossless compression for high quality images,” in *Very High Resolution and Quality Imaging II — Proc. of the SPIE*, 1997, vol. 3025, pp. 62–70.

[8] L. Karray, P. Duhamel, and O. Rioul, “Image coding with an L^∞ norm and confidence interval criteria,” *IEEE Trans. on Image Processing*, vol. 7, no. 5, pp. 621–631, May 1998.

[9] X. Wu and P. Bao, “Near-lossless image compression by combining wavelets and CALIC,” in *Proc. of the 31st Asilomar Conf. on Signals, Systems, and Computers, 1997*, 1998, vol. 2, pp. 1427–1431.

[10] A. J. Pinho, “Adaptive context-based arithmetic coding of arbitrary contour maps,” *IEEE Signal Processing Letters*, vol. 8, no. 1, Jan. 2001.

[11] A. J. Pinho, “A method for encoding region boundaries based on transition points,” *Image and Vision Computing*, vol. 16, no. 3, pp. 213–218, Mar. 1998.

[12] A. J. Pinho, “A JBIG-based approach to the encoding of contour maps,” *IEEE Trans. on Image Processing*, vol. 9, no. 5, pp. 936–941, May 2000.

[13] X. Wu and N. Memon, “Context-based, adaptive, lossless image coding,” *IEEE Trans. on Communications*, vol. 45, no. 4, pp. 437–444, Apr. 1997.

[14] ISO/IEC 14495–1 and ITU Recommendation T.87, *Information technology - Lossless and near-lossless compression of continuous-tone still images*, 1999.

[15] 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.