

Histogram packing, total variation, and lossless image compression

Paulo J. S. G. Ferreira and Armando J. Pinho*

Dept. de Electrónica e Telecomunicações / IEETA
Universidade de Aveiro
3810-193 Aveiro, PORTUGAL

{pjf,ap}@det.ua.pt, <http://www.ieeta.pt/~{pjf,ap}>

ABSTRACT

State-of-the-art lossless image compression methods, such as JPEG-LS, lossless JPEG-2000, and CALIC, perform considerably better on images with sparse histograms when a recently introduced preprocessing technique is used. Bitrate savings of up to 50% have been reported, but so far there is no firm theoretical foundation for this surprising performance. In this paper we address this issue, and attempt to explain how the preprocessing stage, which basically packs the histogram of the images, affects the image total variation, and as a result the ability of the compression algorithms to work more effectively.

1 INTRODUCTION

Most of the image compression techniques currently available were designed mainly for compressing continuous-tone natural images. However, the amount of images that nowadays falls outside this class is large and is also continuously increasing. In fact, in addition to natural content, numerous images of interest may also include other types of content, such as graphical and textual. Frequently, this kind of images do not use the complete set of available intensities (of colors or tones of gray), i.e., the histogram of intensities is sparse.

It has been demonstrated recently [1] that the compression efficiency of state-of-the-art lossless image compression methods can be substantially improved when the images to compress have sparse histograms. This is true for JPEG-LS [2, 3], lossless JPEG-2000 [4, 5], and CALIC [6].

So far, little is known concerning the theoretical justification for the improvement. This paper addresses this issue, and comments on the relation between histogram packing and the overall smoothness of the image, as measured by the variation norm, and the degree to which the variation norm controls the ability of image encoders to efficiently represent an image. A more detailed analysis of some of the points addressed in this paper can be found in the work [7].

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




Figure 1: Improving lossless image coding: the image is subject to a transformation that packs its histogram, and then coded. Applying the inverse mapping T^{-1} to the decoded image yields back the original.

2 THE PREPROCESSING TECHNIQUE

The technique introduced in [1] has been shown to considerably improve the performance of the state-of-the-art lossless image coding methods. It implies no modifications to the basic lossless encoders (JPEG-LS [2, 3], lossless JPEG-2000 [4, 5], or CALIC [6]). In fact, rather than modifying the codecs, the method modifies the original image, using one preprocessing and one post-processing stage, as shown in Fig. 1. Before coding, the image is subject to a transformation T which packs its histogram. To recover the original image, the inverse transformation T^{-1} is applied to the decoded image. Recent research shows that the method is also useful (often even more so) when the histogram is only “locally sparse”.

The packing and unpacking operations are relatively simple, and can be implemented efficiently. Clearly, packing is image-dependent, and therefore the unpacking stage needs some side information (basically, the packing table). This side information adds to the bitrate and should of course be accounted for.

However, the overall compression gain, including the overheads due to the side information, shows savings of up to 50% and above, as reported in [1]. Some results are shown in Table 1, for CALIC, JPEG-LS and lossless JPEG-2000. The improvements are surprising, considering that the methods mentioned represent the state-of-the-art in lossless image coding.

3 HISTOGRAM PACKING AND VARIATION

What is the impact of histogram packing on the norms through which the smoothness of an image is normally measured?

In the following discussion, in which Besov spaces and the total variation norm naturally appear, the term “ N -term non-linear approximation” will mean approximation using the N most significant coefficients of the expansion, whereas “ N -term linear approximation” will imply an approximation using a fixed, image-independent set of N coefficients. For details see [8, Chapter 9], for example.

For $\alpha > 0$, the α -class of an image compression algorithm is the set of all f satisfying

$$a_N(f) := \inf_{\tilde{f} \text{ has } N \text{ coefficients}} \|f - \tilde{f}\| = O(N^{-\alpha}),$$

as $N \rightarrow \infty$. In words, it is the set of images that lead to an approximation error norm that decreases asymptotically with $N^{-\alpha}$. It turns out that the α -class with respect to L_p norms for wavelet-based algorithms is a Besov space.

The rate of decay of the error between an image f and a compressed representation \tilde{f} as the bitrate increases is related to the smoothness of f in terms of these Besov spaces [9]. Images that belong to these spaces can be near-optimally compressed (within the class of stable transform-based possibly nonlinear compression methods) using wavelet-based techniques.

There is a connection between total variation and Besov norms. Roughly speaking, the total (univariate) variation norm is bounded by two Besov norms, and the space of (univariate) functions of bounded variation is embedded in the corresponding Besov spaces. The (nonlinear) approximation error is bounded by

$$E_n(N) \leq CV(f)^2 \frac{1}{N^2},$$

where C is a constant. The decay N^{-2} obtained with wavelet-based approximations cannot be improved by any nonlinear approximation calculated in an orthonormal basis [10]. It is in this sense that wavelets are optimal for approximating bounded variation functions.

For images of bounded variation, the linear approximation error satisfies

$$E_l(N) \leq CV(f) \|f\|_{\infty} \frac{1}{N^{1/2}},$$

and the nonlinear approximation error is given by

$$E_n(N) \leq CV(f)^2 \frac{1}{N}.$$

Note how, in both cases, the error decreases with the variation $V(f)$ of the image, or its square. Clearly, if the image is preprocessed as shown in Fig. 1, and if the variation of the preprocessed image becomes smaller than the variation of the original image, the encoder will in principle be able to

CALIC

Image	Normal		Off-line packing		
	Size	bps	Size	bps	%
benjerry	6,094	1.743	4,758	1.361	21.9
books	23,229	3.264	11,673	1.640	49.7
cmpnnd	68,704	1.397	60,668	1.234	11.7
cmpndn	55,564	1.130	49,471	1.006	11.0
gate	24,700	3.244	18,959	2.490	23.2
music	2,340	1.519	1,445	0.938	38.2
netscape	19,915	2.603	12,917	1.688	35.1
sea_dusk	1,736	0.088	1,417	0.072	18.4
sunset	77,049	2.006	70,348	1.831	8.7
winaw	34,045	0.925	20,159	0.547	40.8
yahoo	7,131	2.101	6,934	2.043	2.8
Total	320,507	—	258,749	—	19.3

JPEG-LS

Image	Normal		Off-line packing		
	Size	bps	Size	bps	%
benjerry	6,707	1.919	4,881	1.396	27.2
books	39,859	5.601	13,396	1.882	66.4
cmpnnd	71,469	1.454	62,431	1.270	12.6
cmpndn	58,639	1.193	51,619	1.050	12.0
gate	27,656	3.632	20,718	2.721	25.1
music	4,534	2.943	1,747	1.134	61.5
netscape	21,249	2.777	13,191	1.724	37.9
sea_dusk	4,061	0.206	3,479	0.176	14.3
sunset	83,552	2.175	75,412	1.963	9.7
winaw	48,189	1.309	20,102	0.546	58.3
yahoo	8,822	2.600	8,401	2.476	4.8
Total	374,737	—	275,377	—	26.5

Lossless JPEG-2000

Image	Normal		Off-line packing		
	Size	bps	Size	bps	%
benjerry	14,076	4.027	9,664	2.765	31.3
books	43,859	6.164	15,318	2.152	65.1
cmpnnd	114,362	2.326	98,767	2.009	13.6
cmpndn	107,596	2.189	92,594	1.883	13.9
gate	32,916	4.323	24,316	3.193	26.1
music	8,457	5.491	3,180	2.064	62.4
netscape	30,769	4.022	17,887	2.338	41.9
sea_dusk	8,214	0.417	5,894	0.299	28.2
sunset	119,031	3.099	106,013	2.760	10.9
winaw	84,913	2.307	33,757	0.917	60.2
yahoo	13,782	4.062	13,001	3.832	5.7
Total	577,975	—	420,391	—	27.3

Table 1: Performance of CALIC, JPEG-LS and lossless JPEG-2000 with and without histogram packing. The compression gains include the overhead due to the side-information necessary to invert the packing transformation.

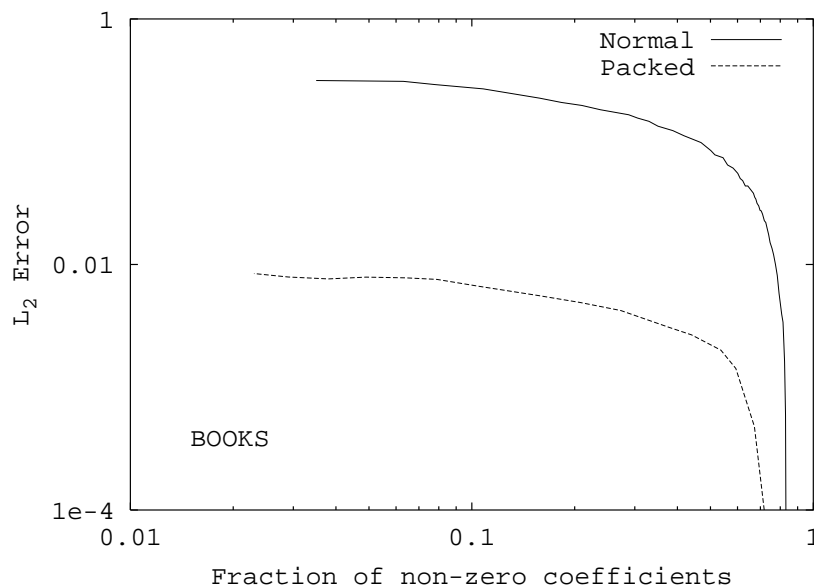


Figure 2: L_2 error as a function of the fraction of coefficients N , for the original and packed image (wavelet-based nonlinear approximation). As explained in the text, the packed image has a smaller variation and as a result it can be compressed using less coefficients, for the same desired L_2 error.

perform more efficiently. In other words, it becomes possible to achieve a certain fixed approximation error with less coefficients N , provided that $V(f)$ is reduced accordingly.

The compression curves on logarithmic plots are expected to differ by constants in zones where the mentioned asymptotic results are meaningful (note that when dealing with finite images one cannot have $N \rightarrow \infty$). This is confirmed in Fig. 2.

This establishes a link between the transformation T that is the key to the preprocessing technique described in Fig. 1 and the concept of total variation. It suggests that investigating the effect of packing on the variation norm might help in understanding more completely the reasons that lead to the performance demonstrated in Table 1 and [1]. Additional evidence confirming this behaviour is presented in Fig. 3, where both the reduction of image variation and also of compression size can be noticed when step-by-step packing is performed on the histograms.

It can be shown that histogram packing does indeed reduce the total variation of an image (see [7] for a more detailed discussion). A simple heuristic argument is presented here: let α be some constant greater than one. Consider an image approximately equal to α times the indicator function of some domain D , with a smooth boundary. Let the length of the boundary of D be L . Then, the variation of the image will be proportional to α (and L). The histogram of the image will be sparse, since it will contain intensity values close to zero and close to α . Packing the histogram will reduce α and as a result the variation of the image.

4 CONCLUSIONS

We have shown that the preprocessing technique reported in [1] and illustrated in the block diagram of Fig. 1 can be

understood in terms of its effect on the image total variation. The packing transformation reduces the total variation of the image, yielding an image of smaller total variation, easier to compress. Results that confirm the performance of the method (Table 1) and the conclusions taken have been presented.

The preprocessing method cannot be expected to work, in general, for lossy compression, that is, when the lossless encoder and decoder pair depicted in Fig. 1 are replaced by a lossy codec. In fact, coding and then decoding the preprocessed image Tf leads only to an approximation of Tf ; let that approximation be denoted by g . In general, g may contain intensity values that were not present in Tf , and which consequently do not belong to the range of the packing function. As a result, T^{-1} cannot be applied, and the effects of packing cannot be entirely compensated for.

This does not prevent the existence of other reversible transformations T , which when used as shown in Fig. 1 may lead to overall compression gains. As far as we know, this is an open question.

References

- [1] A. J. Pinho, "An on-line pre-processing technique for improving the lossless compression of images with sparse histograms", *IEEE Signal Processing Letters*, 2002 (in press).
- [2] ISO/IEC 14495-1 and ITU Recommendation T.87, *Information technology - Lossless and near-lossless compression of continuous-tone still images*, 1999.
- [3] M. J. Weinberger, G. Seroussi, and G. Sapiro, "The LOCO-I lossless image compression algorithm: principles and standardization into JPEG-LS", *IEEE Trans.*

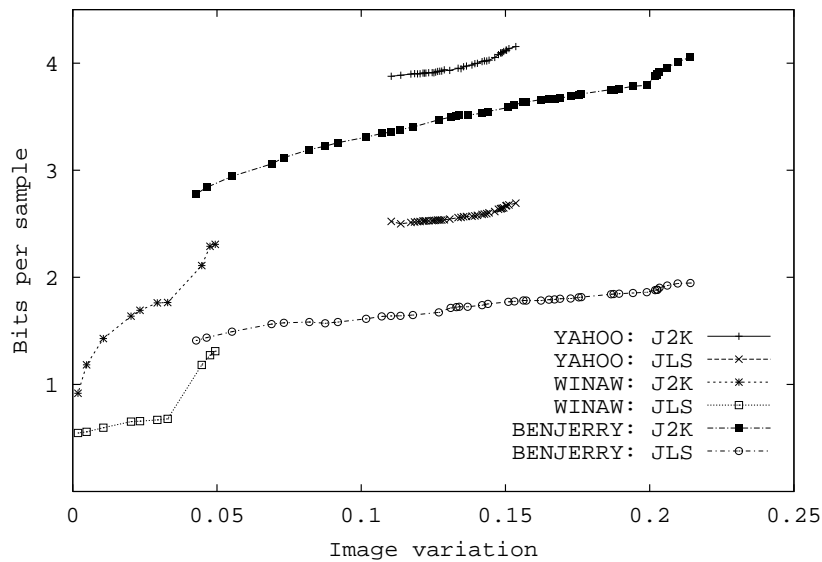


Figure 3: Effect of histogram packing on the variation of three images, and the corresponding bitrates, obtained using lossless JPEG-2000 (J2K) and JPEG-LS (JLS). The dots mark the packing steps. Each step reduces the variation and the bitrate. The bitrate accounts for the side-information necessary to invert the packing transformation.

on Image Processing, vol. 9, no. 8, pp. 1309–1324, Aug. 2000.

[4] ISO/IEC International Standard 15444–1, ITU-T Recommendation T.800, *Information technology - JPEG 2000 image coding system*, 2000.

[5] M. W. Marcellin, M. J. Gormish, A. Bilgin, and M. P. Boliek, “An overview of JPEG-2000”, in *Proc. of the Data Compression Conf., DCC-2000*, Snowbird, Utah, Mar. 2000, pp. 523–541.

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

[7] P. J. S. G. Ferreira and A. J. Pinho, “Why does histogram packing improve lossless compression rates?”, preprint, 2002.

[8] S. G. Mallat, *A wavelet tour of signal processing*, Academic Press, San Diego, 1998.

[9] R. A. DeVore, B. Jawerth, and B. J. Lucier, “Image compression through wavelet transform coding”, *IEEE Trans. on Information Theory*, vol. 38, no. 2, pp. 719–746, Mar. 1992.

[10] R. A. DeVore, “Nonlinear approximation”, *Acta Numer.*, pp. 51–150, 1998.