Symposium

Electronic Image Capture and Publishing (EICP) '98

Conference on

Electronic Imaging: Processing, Printing and Publishing in Colour

chaired by: J. Bares (Xerox Corp., Webster, N.Y., USA)

ADAPTIVE COMPRESSION OF STILL IMAGES : AUTOMATING THE CHOICE OF ALGORITHM AND PARAMETERS

Frédéric GILBERT
ALCATEL ALSTHOM RECHERCHE
Route de Nozay
91460 MARCOUSSIS
FRANCE
Frederic.Gilbert@aar.alcatel-alsthom.fr

Fabien MOUTARDE
ECOLE DES MINES DE PARIS
60, Bd Saint-Michel
75272 PARIS Cedex 06
FRANCE
moutarde@cc.ensmp.fr

Abstract

This paper presents a new adaptive compression scheme for continuous-tone still images with two main benefits: 1°/ it allows to take advantage of a set of algorithms (such as JPEG, wavelets, fractals, ...) by switching between them according to the image content and target compression ratio; 2°/ it tunes automatically the algorithm-specific parameters (e.g. JPEG quantization factor) in order to meet the user compression objective (e.g. a target compression ratio). A performance evaluation test indicate the following: 1°/ our adaptive coder can succeed with probability ≈ 72% in selecting automatically the best technique (among JPEG, a wavelet algorithm, and a fractal one); 2°/ the difference between the obtained compression ratio and the target is lower than 10%; 3°/ the computing overhead due to algorithm selection and parameter choice is negligible. This adaptive compression technique has been proposed to the ISO standardisation committee currently working on the definition of the future "JPEG-2000" standard.

1. Motivation

For the last ten years, JPEG compression [1] has been the dominant algorithm used for lossy compression of continuous-tone still images. It generally provides very good results for compression ratios up to ≈ 20-30. However, new compression techniques based on wavelets [2] or fractals [3] have emerged, which can outperform JPEG, particularly at high compression ratios. The improvement that can sometimes be attained by these new algorithms is such that the standardisation committee ISO-IEC/JTC1/SC29/WG1 has started working in 1996 on the definition of a future standard (nick-named JPEG-2000) that could be based on, or include, one of these new techniques (cf. [8]).

Despite the claims by various authors that their particular version of wavelet or fractal compression always outperforms JPEG and others, it becomes clear that there is no such thing as the ultimate algorithm capable of the best performance for any image and any compression ratio; it is indeed a well-known fact that the behaviour of any compression technique may strongly vary from one image to another. Some authors have already explored the possibility to adapt one given algorithm to the content of the image to compress (e.g. modification of DCT quantization table for JPEG in [4], or change of wavelet filters in [5]). In this paper, we investigate a new approach consisting in switching between algorithms, rather than adapting one single technique: we propose and test a method to design a smart adaptive coder able to choose the best algorithm (among a predefined set) depending on the image content and on the target compression ratio.

Another requirement and motivation for our adaptive coder, which comes from ergonomic considerations, is to make it possible for the end-user to tune the compression of images by specifying only one easily understandable parameter (such as the compression ratio, or the percentage of maintained quality), no matter what algorithm is used.

2. Adaptive coder design

For a feasibility experiment, we took the three following elementary algorithms: the usual lossy (DCT-based) JPEG [1], a wavelet-based algorithm with zero-tree coding (EZW) [2], and a fractal-based algorithm (Iterated Function System, IFS) [3]. The software used are: the Independent Jpeg Group (IJG) implementation of JPEG, an Alcatel implementation of EZW, and a freeware implementation of the IFS by Y. Fisher.

The error criterion that we have adopted for choosing the best technique at a given compression ratio on a given image is the Quadratic Mean Error (QME). We also characterise the "complexity" of the image to compress by a complexity measure *Cx*, which is the mean absolute pixel-to-pixel horizontal variation (see formula (1) below); although not very sophisticated, this criteria proved to give a good report on the image local regularity, on its texturation and consequently on its "compressibility".

$$(1)\quad C_x = \frac{1}{(w-1)\times h}\sum_{i=1}^{h}\sum_{j=1}^{w-1}\left|x_{i,j} - x_{i,j+1}\right|$$

where x_{ij} = pixel value at line i and column j,
and (w,h) = image dimensions.

Our smart adaptive coder is composed of an adaptive module and of the 3 simple compression routines. The adaptive module (cf figure 1) can be further decomposed in the complexity computation module, and a selector function. The selector chooses the algorithm, and determines the value(s) to use for its parameter(s). The inputs of the selector are: 1°/ the compression objective (target compression ratio or target reconstruction quality); 2°/ the complexity measure *Cx* of the image.

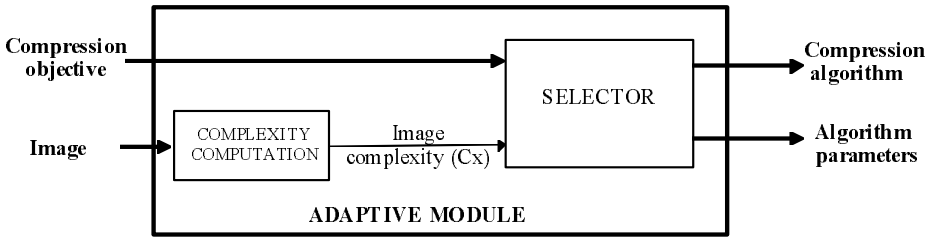

Figure 1: structure of the adaptive module

When the compression objective is a reconstruction quality, it is given as a percentage of maintained quality Q between 0 and 100, with Q=100 corresponding to QME=0 (lossless compression), and Q=0 corresponding to maximal degradation. The target quality Q defines in fact a target QME given by:

$QME = QME_{\max} \times (1 - \frac{Q}{100})$ where QME_{max} is the QME corresponding to maximal degradation.

The selector itself is obtained through an experimental calibration process described below, which consists in testing all algorithms in competition on a database of images and for various compression ratios, in order to build a performance model of each algorithm.

3. Adaptive coder calibration

During a learning phase, the 3 elementary algorithms have been systematically applied to a calibration database of 50 images. This calibration database was chosen to be representative of the type of images we wanted to apply the adaptive coder to: in the present case a mixture of aerial, satellite and "natural" 8-bit greyscale images. Each algorithm was tested on each image with varying values of its parameters so as to cover evenly the parameter space.

For each compression experiment, we stored: 1°/ the complexity measure *Cx* of the image to compress; 2°/ the algorithm parameter value, 3°/ the obtained compression rate: 4°/ the reconstruction error (QME). These data were then used to automatically build, for each algorithm, a 3D abacus relating the QME, the compression ratio, and the image complexity. The abacuses are obtained by applying smoothing splines to the experimental data points. An example of one of these $QME=F_{algo}$(compression_ratio, image_complexity) abacuses is presented in figure 2.

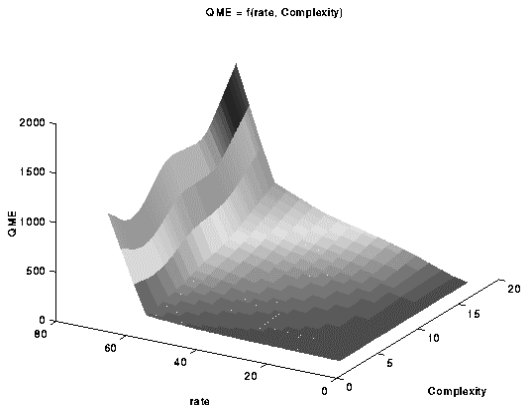

Figure 2: QME = f(compression_ratio, image_complexity) for JPEG

For EZW, compression ratio or QME is directly tunable with the native parameters, but this is *not* the case for JPEG and for the IFS. However, we wanted the selector of the adaptive coder to be able to find automatically the good parameter value for a given compression objective *for any of the 3 elementary algorithms*. We thus also built automatically (from the same experimental data,

and also with smoothing splines), four other 3D abacuses: 2 for JPEG and 2 for the IFS. They provide respectively, for each algorithm, the compression ratio and the QME as a function of image complexity and parameter value. An example of one of these compression_ratio=f(image_complexity, parameter) is shown in figure 3.

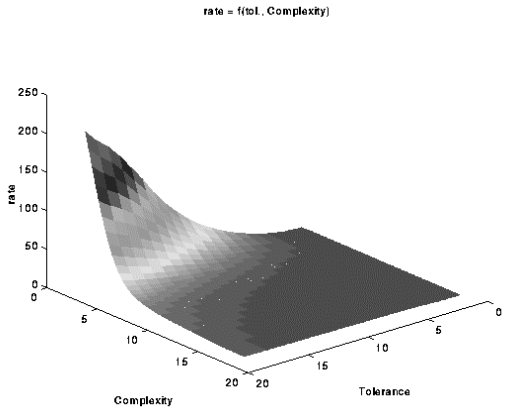

Figure 3: compression_ratio= f(tolerance_parameter, image_complexity) for Fisher's IFS

Finally, an empirical relation between QME_{max} and the complexity was also determined, in order to simplify the processing of quality-based compression objective: QME_{max} is approximated by $QME_{\max} \approx \alpha \times C_x + \beta$, where Cx is the image complexity and α and β are constants determined experimentally on the calibration image database.

Once all the abacuses and parameters are constructed, the selector algorithm is the following:

1°/ Choice of compression algorithm

If the compression objective is a compression ratio, look-up in the 3 abacuses for the 3 QMEs predicted for the (compression_ratio, image_complexity) value corresponding to the image and target compression ratio, and choose the algorithm for which the predicted QME is the lowest.
If the compression objective is a target maintained quality Q, first convert it to a target QME by the formula (2) below:

$$(2) \quad QME = (\alpha \times C_x + \beta) \times (1 - \frac{Q}{100})$$

Then, read the three (QME, compression_ratio, image_complexity) abacuses (corresponding to the 3 different elementary algorithms) as ratio=f(complexity, QME) in order to compare the predicted compression ratios, and select the compression technique for which it is biggest.

2°/ Parameter determination

If the algorithm can be natively parameterized with the compression objective (which is the case for the EZW wavelet technique), the parameter value is exactly the compression objective.
If there is no way to tune directly the compression objective (as is the case, for instance, when there is a target compression rate to reach with JPEG compression), the appropriate "parameter" abacus (depending on the algorithm and on the type of compression objective) is used to evaluate the parameter value which, given the image complexity, is expected to respect the compression objective.

4. Results

We have tested the adaptive coder described above on a set of 50 images (independent of the images used during the calibration process, but of the same nature, i.e. 8-bit greyscale aerial, satellite and natural images), and with various compression objectives for each test image.

The global success rate of the automatic optimum algorithm selection is 72.5%. A more detailed analysis of the performance consists in evaluating separately the confidence level that can be placed on the selector decision depending on what is this decision (see table 1 below).

Adaptive coder choice:	JPEG	EZW	IFS
Best algorithm in reality:	JPEG best: 62 % EZW best: 38 % IFS best: 0 %	JPEG best: 23 % EZW best: 77 % IFS best: 0 %	never chosen

Table 1: each column corresponds to (image, compression objective) for which the adaptive coder selected a given algorithm; the percentages shown present the algorithm that was actually the best (and thus the one that should have been chosen). A perfect adaptive coder would have 100% on the first diagonal, and 0% everywhere else.

The last column of table 1 cannot be computed because in our tests, the fractal algorithm is never chosen by our adaptive coder. It is also never the best algorithm (as can be seen on the last line of table 1), which may be due to the nature of our images, or to the systematically sparse domain search that we had to use for the IFS in order to compress in a reasonable time. The main other information from table 1 is that the choice of JPEG by the selector is significantly less reliable than the choice of EZW.

The conditional success and failure rates can also be computed in the reverse way, as the probability that a given algorithm is selected knowing that it is the best in reality: this corresponds to what is called the "confusion matrix", which is the usual performance measure of classifiers (cf. table 2).

Best algorithm in reality:	JPEG	EZW
Adaptive coder choice:	JPEG selected: 55 % EZW selected: 45 %	JPEG selected: 20 % EZW selected: 80 %

Table 2: in each column, one can read the conditional success and failure rates when the actual best algorithm is fixed (success rates are on the first diagonal of the confusion matrix). IFS algorithm is omitted, since it is never the best and never chosen.

From table 2, one can see that the adaptive coder tends to under-estimate JPEG when it is in fact the best algorithm, and thus makes more mistakes in this case than in the case when EZW is the best algorithm.

In a first analysis of the 27.5% of cases when the selector fails to choose the best algorithm, we found that they correspond to cases for which average QME difference between JPEG and EZW is less important (by a factor of 2) than when the selection is successful. This means that the automatic algorithm selection tends to fail more when the quality difference between algorithms is less clear, which is rather normal and difficult to avoid.

We also noticed in this experiment that JPEG is generally the best algorithm for low compression ratios, but is outperformed by wavelets for higher compression ratios; the switching limit depends on the image complexity, but is around 20 (see figure 4). This particular result is in contradiction with some other papers comparing JPEG and wavelets, in particular with [6] in which wavelets *always* outperform JPEG (even at low compression ratios), but the test images used in [6] are not the same, and their wavelet algorithm is quite different from our EZW.

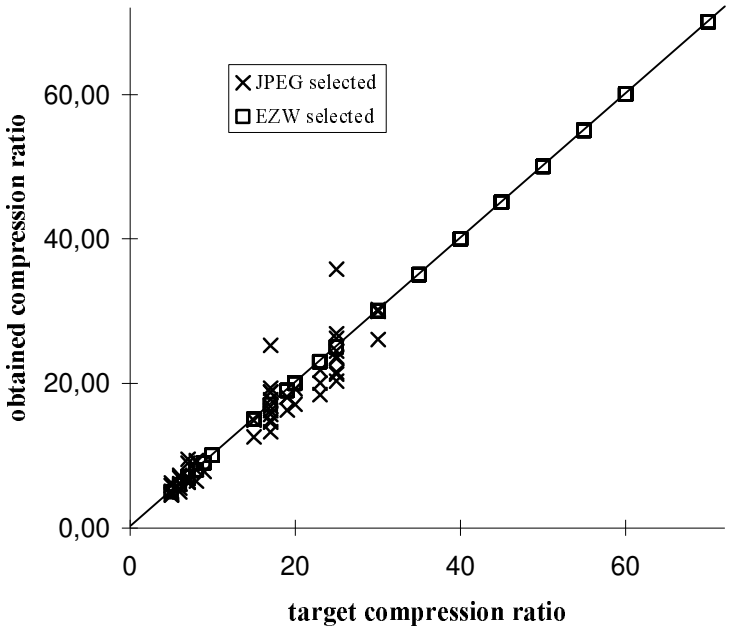

Figure 4: comparison of target and obtained compression ratios for some of the test images and various ratios; cases when the adaptive coder selected JPEG and EZW are plotted differently, so that one can also see how the target compression ratio influence the algorithm choice.

Another important result about the adaptive coder is the good precision with which it fulfils the compression objectives: the maximum error between the target value and the obtained one is about 10% in compression ratio and 20% in quality for JPEG (for EZW, the errors are negligible because it can be directly parameterized with the compression objective). As an illustration, we provide plots comparing target and obtained compression ratios and QMEs for some of the test results (cf. figures 4 and 5).

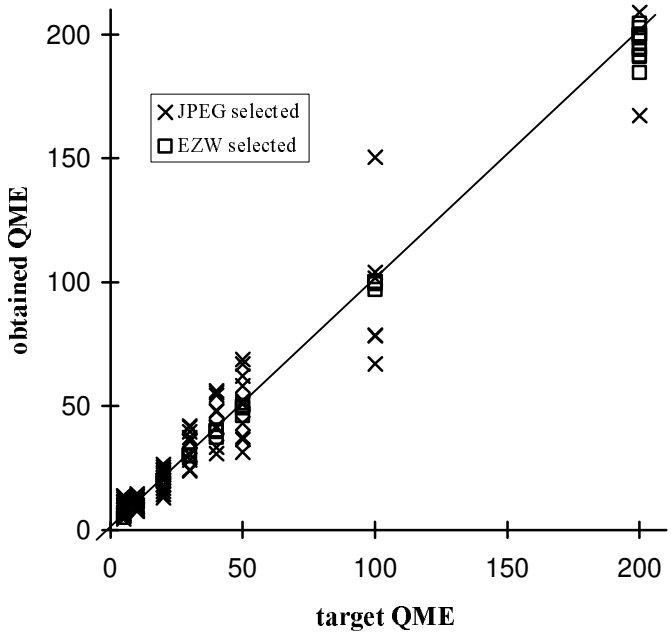

Figure 5: comparison of target and obtained QME for various test images and QMEs.

Finally, it is important to note that the computation overhead corresponding to the selector, as well as the computation of the image complexity *Cx*, is negligible compared to the compression itself, so that the total computation time of the adaptive coder is nearly that of the selected algorithm.

5. Discussion

First, the mean performance of the algorithm selection (≈ 72%) is rather good, but the error rate is still too high when JPEG is selected or should be selected. We are currently working to improve the automatic selection performance. Among the possible modifications for improvement, we are mainly considering: 1°/ changing the smoothing technique used to construct the abacuses from experimental data in the calibration phase; 2°/ using as error measure some kind of normalized QME, so that the error=f(compression_ratio) data used for calibrating the abacuses would have lower image-to-image dispersion; 3°/ changing the complexity measure (e.g. similar to *Cx*, but bi-dimensional version); 4°/ using a more specialised calibration database (distinguish three separate categories for natural, aerial and satellite images, and eventually others – fingerprints, people faces, ...).

Also, it could also be argued that using EQM as quality measure for departing algorithms does not reflect the visual quality, which is indeed true (at least for a small EQM difference). For this reason, we have tried to apply our technique with an error measure derived from psycho-visual properties (the Image Quality Measure, IQM, proposed in [7]), but this raised several problems: first, the computation cost of the IQM is much higher than that of the EQM, and secondly, some numerical stability problems of this measure made it impossible to obtain good and usable IQM=f(ratio,complexity) abacuses so that it was not possible to simply transpose what we had done with the QME.

One can wonder if our adaptive compression scheme could be used also for colour images. Everything presented in this paper was done for greyscale images, but it could easily be modified for treating colour images by treating separately each colour component that is compressed separately by the elementary algorithm (usually Y, Cb and Cr in the case of JPEG). However, it could also be interesting to try to define a global complexity measure for colour images, so as to treat them globally (instead of component by component).

Finally, we considered in this paper only the possibility to switch between algorithm globally for each image. But it might be interesting to design some segmentation procedure that could be combined with the switching technique to compresss separately different portions of a single image with different algorithms and parameters. This could allow to optimize even more the image compression, and also to compress different parts of the image with different qualities depending on their relative interests.

6. Conclusion

For obtaining the best image quality at a given compression ratio, it can be necessary to choose among various algorithms depending on the image content, and on the target compression ratio. This motivated us to design a multi-algorithm adaptive coder able to select automatically the best algorithm (among a predefined set) for compressing a given image with a given compression objective. Our tests indicate that such a coder can indeed have a reasonable success rate in automatically selecting the best algorithm (at least on a particular kind of images) for a computation cost overhead nearly negligible (except for the initial calibration, which is normally to be done only once for all).

Moreover, our adaptive coder also provides a unified user-oriented parametrization of image compression: whatever the algorithm actually applied and its native parameters, the selector allows to tune the compression with either a target compression ratio or a target percentage of quality preservation.

This multi-algorithm adaptive approach seems very promising, and we proposed it to the ISO standardisation working group that is currently working on the next generation of still image compression standard named "JPEG-2000" (cf. [9]). This does not mean that the future standard should normalise the algorithm-switching technique, but that JPEG-2000 should provide support for easy use of multiple elementary algorithms, and possibly for separate compression of different portions of an image with different algorithms and parameter values.

Note: the adaptive compression technique described in this paper is covered by a patent filed by Alcatel (french patent application number 97 03212).

Bibliography

[1] ISO-IEC/JTC1/SC29/WG1, "*Digital compression and coding of continuous-tone still images: requirements and guidelines*", ISO/IEC 10918-1, February 1994.

[2] J. Shapiro "*Embedded Image Coding Using Zerotrees of Wavelet Coefficients*", IEEE Trans. on Signal Processing, Vol. 41, n°12, Dec. 1993.

[3] Y. Fisher (ed.), "*Fractal image compression: theory and applications*", Springer Verlag, New York, 1995.

[4] A.B. Watson, "*Visually optimal DCT quantization matrices for individual images*", proceedings of Data Compression Conference '93, held in Snowbird (USA), March 1993.

[5] O. Egger, "*Region representation using non-linear techniques with application to image and video coding*", PhD thesis, Ecole Polytechnique Fédérale de Lausanne, April 1997.

[6] V.R. Algazi & R.R. Estes Jr, "*Comparative performance of wavelet and JPEG coders at high quality*", proceedings of SPIE conference on "Very High Resolution and Quality Imaging, held in San Jose (USA), February 1997.

[7] N.B. Nill & B.H. Houzas, "*Objective Image quality Measures derived from digital image power spectra*", Optical Engineering, Vol. 31, No 4, pp 813-825, April 1992.

[8] ISO-IEC/JTC1/SC29/WG1, "*Call for contribution for JPEG 2000 (JTC 1.29.14, 15444): Image Coding System*", ISO/IEC JTC1/SC29/WG1/N.505, March 21st, 1997.

[9] F. Gilbert, "*Adaptive compression mode for a JPEG 2000 toolbox architecture*", ISO/IEC JTC1/SC29/WG1/N.533, June 6th, 1997.