Embedding Companders in JPEG Compression

Neri Merhav
HP Laboratories Israel*
HPL-98-141
August, 1998

image compression,
JPEG,
companding

The use of a compander (*comp*ressor-exp*ander*) in conjunction with scalar uniform high-resolution quantization, so as to obtain an overall effect of an improved nonuniform quantization, has been widespread for many years and embedded in standards of digitization and compression of audio waveforms (μ-law in the U.S., and A-law in Europe). However, there is almost no reported work on its usefulness in still image and motion picture coding. In this work, we are making an attempt to embed companders in the JPEG algorithm and our results show considerable improvement (10-25%) in compression ratio for large PSNR's. Strictly speaking, the new scheme is not standard compliant. However, with suitable preprocessing and postprocessing it can be implemented in a standard compliant manner. Also, the results of this research might be useful for the quantization part of future standards such as JPEG 2000.

Internal Accession Date Only

*Hewlett-Packard Laboratories Israel, Technion City, Haifa 32000, Israel

1 Introduction

Companding is a very simple and well-known technique for non-uniform scalar quantization using a high resolution uniform quantizer (A/D converter) [1], [9], [11], [10], [6]. The idea is to apply a memoryless non-linearity (compressor) before the uniform quantizer and the inverse non-linearlity (expander) after, so as to obtain an overall effect of non-uniform quantization in a simple manner.

Theoretically, it is known [5], [6] that for a given input probability density function (PDF) $p(x)$, the best nonlinearity (in the sense of minimum square error distortion for a given number of quantization levels) before the quantization $c(x)$ is, in the high resolution limit, proportional to $\int_0^x p^{1/3}(u)du$ for $x > 0$, with an antisymmetric extension for $x < 0$. The constant of proportionality is chosen so as to maintain the dynamic range.

If the PDF is unknown, a common practice is to use a logarithmic compressor $c(x)$ because it leads robust quantization in the sense that the sensitivity of the SNR to the input PDF is drammatically reduced. Indeed, logarithmic companders have been widely used with great success. Moreover, they have been standardized in audio PCM conversion both in the U.S. (μ-law) and in Europe (A-law).

When entropy coding is allowed, namely, the optimum tradeoff between distortion and output entropy is saught, it turns out that the best compander is linear in the high resolution limit. In other words, companding becomes altogether superfluous as uniform quantization is asymptotically the best scalar quantization. Optimal entropy coding, however, requires knowledge of the statistics at the quantizer output, and in case of mismatch between the entropy coder and the those statistics, it is no longer true that linear companding is best.

A natural extension of the scalar compander-quantizer is the multi-dimensional one (see, e.g., [3], [4] and more recent studies [7], [8]). It turns out that for memoryless sources and additive distortion measures, there is no interaction between the different dimensions, and hence optimal multi-dimensional companding simply breaks to separate scalar compandings in each dimension.

With this background on companding and its great usefulness in audio compression, it appears somewhat surprising that there is almost no reported work on the use of companders in image and video compression. An exception is a recent work of Bakshi and Fuhrmann [2] who used the JPEG compression standard preceded by a μ-law compressor and followed by a μ-law expander and tested its improvement from a psychovisual point of view.

Instead of applying the compander non-linearity in the pixel domain of the raw input image, and the inverse non-linearity on the decoded image, as suggested in [2], in this work, we take an alternative approach, which seems to us more natural: We apply the the non-linearities directly in the domain in which the uniform scalar quantization of JPEG takes place, namely, the DCT domain (see fig. 1). From the theoretical point of view, this configuration is appealing, because if the DCT indeed removes correlation (and hence in the Gaussian case, also statistical dependency), then as explained above, the best multidimensional companding is given by separate scalar compandings of each DCT component. Strictly speaking, however, there is a practical difficulty with this approach because it is no longer standard compliant. Nonetheless, it is easy to see that such a scheme can be implemented by a standard JPEG algorithm if we allow suitable preprocessing on the input image and postprocessing on the output image. The preprocessing would include DCT operation on 8×8 blocks, followed

by scalar companding of each coefficient, followed in turn by IDCT. The postprocessing will then consist of the inverse operations in reverse order (see fig. 2).

From the theoretical point of view, the main emphasis in this work is to investigate optimal companding functions which take into account the particular method of entropy coding used in JPEG. This analysis will give us some guidelines about properties of good companding functions, which will turn out to be quite different from those tradionally used in audio digitization like μ-law and A-law. Experimental results show considerable improvment over the ordinary JPEG algorithm in the high resolution regime.

2 High Resolution Analysis

Consider the following model. A block of n independent random variables $X = (X_1, ..., X_n)$, distributed according to a joint probability density function (PDF) $p(X) = \prod_{i=1}^{n} p_i(X_i)$, is fed into an encoder that works as follows. First, each component X_i, $1 \leq i \leq n$, undergoes a memoryless nonlinearity (compressor) $c_i(\cdot)$, whose output $Y_i = c_i(X_i)$ is uniformly finely quantized with quantization step size Δ_i, to a codeword $\hat{Y}_i$. The block of codewords $\hat{Y} = (\hat{Y}_1, ..., \hat{Y}_N)$ is then losslessly encoded to its entropy $H(\hat{Y})$. At the decoder side, an entropy decoder reconstructs $\hat{Y}$, whose components $\hat{Y}_i$ undergo the corresponding inverse nonlinearities $\hat{X}_i = e_i(\hat{Y}_i) = c_i^{-1}(\hat{Y}_i)$ to yield the reconstructed source $\hat{X} = (\hat{X}_1, ..., \hat{X}_N)$.

A few basic assumptions will be the following.

1. The marginal PDF's p_i are all symmetric about the origin.
2. The dynamic range of the input is $[-A, A]$, $A < \infty$, for all i and the compressor functions all preserve this range, i.e., they map $[-A, A]$ onto itself.
3. The compression functions $c_i(\cdot)$ are antisymmetric, differentiable, and monotonically increasing, with $c_i(0) = 0$ and $c_i(A) = A$.
4. For every i, $1 \leq i \leq n$, the range $[-A, A]$ is evenly divided into N_i quantization levels, each of size $\Delta_i = 2A/N_i$.
5. The number N_i will be assumed odd, so that the quantization would be symmetric about the origin with a codeword at zero.
6. The quantization bins are so small that the PDF p_i is nearly constant therein (high resolution approximation).

This probabilistic model is an idealization that may serve as a theoretical ground for image and video compression standards, like JPEG, MPEG, H.261, H.263, etc. The input block X designated a block of DCT coefficients whose length N is given by the transform size ($8 \times 8 = 64$ in the case of JPEG). Assuming that the DCT decorrelates the data (being a good approximation to the KLT), then in the Gaussian case, it also creates statistical independence. The marginal densities of the different coefficients $p_i(\cdot)$, however, might be different (lower order coefficients normally have larger variances). In the standard compression algorithms, the DCT coefficients are uniformly quantized separately and then entropy-coded using run-length coding. Here we add the additional ingredient of companding.

Since the entropy encoder is based on run-length coding that takes advantage of the fact that many of the quantized coefficients $\hat{Y}_i$ are zero, we will treat separately the encodings of zeroes and non-zero values of these coefficients. Let Z_i denote the indicator function of the event that $\hat{Y}_i = 0$. Let $Z = (Z_1, ..., Z_n)$.

First observe that since Z is a deterministic function of $\hat{Y}$, the entropy of $\hat{Y}$ can be written as follows.

$$\begin{aligned} H(\hat{Y}) &= H(\hat{Y}, Z) \\ &= H(Z) + H(\hat{Y}|Z) \\ &= H(Z) + \sum_{i=1}^{n} H(\hat{Y}_i|Z_i) \end{aligned} \tag{1}$$

where the last line follows from independence. Now,

$$\begin{aligned} H(\hat{Y}_i|Z_i) &= \Pr\{Z_i = 0\} H(\hat{Y}_i|Z_i = 0) + \Pr\{Z_i = 1\} H(\hat{Y}_i|Z_i = 1) \\ &= \Pr\{Z_i = 0\} H(\hat{Y}_i|Z_i = 0) \end{aligned} \tag{2}$$

where the last line follows from the fact that $H(\hat{Y}_i|Z_i = 1) = 0$ because if $Z_i = 1$, then $\hat{Y}_i = 0$ by definition. Combining eqs. (1) and (2), we get

$$H(\hat{Y}) = H(Z) + \sum_{i=1}^{n} \Pr\{Z_i = 0\} H(\hat{Y}_i|Z_i = 0). \tag{3}$$

The first term on the r.h.s. of the last equation represents the part of run-length coding, whereas the second term is interpreted as entropy coding of the non-zero coefficients.

The idea now is that it might pay off to make the probabilities $\Pr\{Z_i = 1\}$ relatively large so as to decrease both the first and the second term, and hence to decrease the overall output entropy $H(\hat{Y})$. There is, however, a tradeoff between this entropy and the distortion associated with high probability of $Z_i = 1$.

Let $c_i'(\cdot)$ denote the derivative of $c_i(\cdot)$, and let $\delta_i(x) = \Delta_i / c_i'(x) = 2A/[N_i c_i'(x)]$ designate the approximate width of the inverse image of a companded quantization bin at the vicinity of a point $x \in [-A, A]$ in the input domain (before the compressor). This means that the compressor followed by a high resolution uniform quantizer induces a nonuniform quantizer whose quantization bin around x is about $\delta_i(x)$ (see fig. 3). For shorthand notation, let us denote $a_i \triangleq \delta_i(0)$. Now, observe that

$$q_i \triangleq \Pr\{Z_i = 1\} = \int_{-a_i}^{a_i} p_i(x) dx \tag{4}$$

and denote $\bar{q}_i = 1 - q_i = \Pr\{Z_i = 0\}$. Then,

$$\begin{aligned} H(\hat{Y}_i|Z_i = 0) &\approx -\int dx p_i(x|Z_i = 0) \log[p_i(x|Z_i = 0)\delta_i(x)] \\ &= -\int_{|x| \geq a_i} dx \frac{p_i(x)}{\bar{q}_i} \log\left[\frac{p_i(x)\delta_i(x)}{\bar{q}_i}\right] \\ &= -\int_{|x| \geq a_i} dx \frac{p_i(x)}{\bar{q}_i} \log p_i(x) - \int_{|x| \geq a_i} dx \frac{p_i(x)}{\bar{q}_i} \log \delta_i(x) + \log \bar{q}_i. \end{aligned} \tag{5}$$

Thus,

$$\begin{aligned}\Pr\{Z_i=0\}H(\hat{Y}_i|Z_i=0) &= -\int_{|x|\geq a_i} dx p_i(x)\log p_i(x) - \\ &\quad \int_{|x|\geq a_i} dx p_i(x)\log\delta_i(x) + \bar{q}_i\log\bar{q}_i, \end{aligned} \tag{6}$$

and so,

$$\begin{aligned}H(\hat{Y}_i) &= H(Z_i) + \Pr\{Z_i=0\}H(\hat{Y}_i|Z_i=0) \\ &= -q_i\log q_i - \int_{|x|\geq a_i} dx p_i(x)\log p_i(x) - \int_{|x|\geq a_i} dx p_i(x)\log\delta_i(x). \end{aligned} \tag{7}$$

On substituting $\delta_i(x) = 2A/[N_i c_i'(x)]$, we get

$$H(\hat{Y}_i) = -q_i\log q_i - \int_{|x|\geq a_i} dx p_i(x)\log p_i(x) - \bar{q}_i\log\frac{2A}{N_i} + \int_{|x|\geq a_i} dx p_i(x)\log c_i'(x). \tag{8}$$

As for the MSE distortion, since p_i is assumed nearly uniform within each quantization bin (cf. Assumption 6), then it contributes about $a_i^2/12$. Thus, the total MSE distortion D_i contributed by the ith component is

$$\begin{aligned}D_i &= \int_{-a_i}^{a_i} x^2 p_i(x)dx + \frac{1}{12}\int_{|x|\geq a_i}\delta_i^2(x)p_i(x)dx \\ &= \int_{-a_i}^{a_i} x^2 p_i(x)dx + \frac{A^2}{3N_i^2}\int_{|x|\geq a_i}\frac{p_i(x)}{[c_i'(x)]^2}dx \end{aligned} \tag{9}$$

Now observe that the first three terms in the expression for $H(\hat{Y}_i)$ (eq. (8)) as well as the first term in the expression for D_i (eq. (9)), depend on the choice of $c_i(\cdot)$ only via a_i, which asymptotically depends only on $c_i'(0)$. On the other hand, the remaining terms in eqs. (8) and (9) depend on the detailed behavior of $c_i'(\cdot)$ outside the interval $[-a_i, a_i]$. Thus for a given value of $c_i'(0)$, we can use the same method as in [6, pp. 153-154] to obtain optimal tradeoff between $\int_{|x|\geq a_i} dx p_i(x)/[c_i'(x)]^2$ and $\int_{|x|\geq a_i} dx p_i(x)\log c_i'(x)$, and reach the same conclusion that $c_i'(x)$ should be a constant outside the interval $[-a_i, a_i]$. Since inside this interval the quantizer output is zero by definition, the exact definition of $c_i(x)$ within this interval is immaterial as long as its absolute value is less than Δ_i (which guarantees the quantizer output to be zero). For example, $c_i(x)$ can be assumed linear in this interval, e.g.,

$$c_i(x) = \begin{cases} \frac{\gamma_i}{a_i}x & |x|\leq a_i \\ \left[\frac{A-\gamma_i}{A-a_i}(|x|-a_i)+\gamma_i\right]\mathrm{sgn}(x) & |x|>a_i \end{cases} \tag{10}$$

where γ_i is an arbitrary positive real smaller than Δ_i. Observe, however, that the choice of γ_i does not affect the quantization performance whenever $|x|\leq a_i$ because every x in this range is quantized to zero independenly of γ_i. For $|x|>a_i$, the choice of γ_i affects, most importantly, the slope of $c_i(\cdot)$. Higher slope and hence better quantization resolution is obtained for small γ_i. The limiting situation is, of course, when $\gamma_i = 0$, which leads to

$$c_i(x) = \begin{cases} 0 & |x|\leq a_i \\ \left[\frac{A}{A-a_i}(|x|-a_i)\right]\mathrm{sgn}(x) & |x|>a_i \end{cases} \tag{11}$$

This function is noninvertible in the range $[-a_i, a_i]$, but considering the fact that expander operates on the quantized version of $c_i(x_i)$, and the quantized value in this range is zero anyhow, the natural expander would then be

$$e_i(y) = \begin{cases} 0 & y = 0 \\ \left[\frac{A-a_i}{A}|y| + a_i\right] \operatorname{sgn}(y) & y \neq 0 \end{cases} \tag{12}$$

In summary, under our assumptions on the source that generates the data, Assumptions 1–6 above, and the assumption on entropy coding, we have obtained a family of simple companders, henceforth referred to as *piecewise linear* companders, that are parametrized by the parameter a_i. This patrameter is subjected to optimal tuning, just like in μ-law and A-law. The characteristics of piecewise linear companders are, of course, very different from those of μ-law and A-law, however, they are intuitively appealing in view of the entropy coding that is being carried out. The choice of a_i controls the tradeoff between the efficiency of the run-length coding and the distortion within the range $[-a_i, a_i]$. One would like the quantizer to zero out small coefficients at the benefit of improved quantization for the significant coefficients.

3 Other Companders

It is virtually needless to say that the assumptions we made in the previous section are largely an idealization of the real situation. In practice, the DCT does not really break the dependency between the coefficients, but more importantly, it is also clear that the assumption regarding optimal entropy coding is not really met. Nonetheless, the above derivation suffices to serve as a guideline for our choice of several alternative companders to be examined in addition the piecewise linear compander. In view of the simple intuition behind the behavior of this compander and Assumptions 2 and 3 above, we will examine a few additional companding functions with the following properties:

1. $c_i'(0) = 0$.
2. $c_i'(x)$ is approximately constant for relatively large $|x|$.
3. $c_i(0) = 0;\ c_i(A) = A$.
4. $c_i'(x) \geq 0$ for all $x \in [-A, A]$.
5. $c_i(\cdot)$ is invertible (preferrably by a function with a closed-form expression).

Properties 1,2, and 4 suggest that for positive x,

$$c_i'(x) = K_i[1 - \epsilon_i(x)]$$

where K_i is a positive constant (chosen so as to satisfy Assumption 3) and $\epsilon_i(x)$ is a monotonically decreasing function with $\epsilon_i(0) = 1$ and $\epsilon_i(x) \approx 0$ for large x. Note that for the piecewise linear compander, we have

$$\epsilon_i(x) = \begin{cases} 0 & x \leq \delta_i \\ 1 & x > \delta_i \end{cases} \tag{13}$$

Two additional examples of such functions are

$$\epsilon_i(x) = \left[\frac{a_i}{x + a_i}\right]^{\rho}, \tag{14}$$

and

$$\epsilon_i(x) = e^{-a_i x} \tag{15}$$

where a_i and ρ are positive reals (companding parameters). The corresponding companders would then be

$$c_i(x) = \begin{cases} K_i \mathrm{sgn}(x) \left[|x| + a_i \ln \frac{a_i}{|x|+a_i}\right] & \rho = 1 \\ K_i \mathrm{sgn}(x) \left[|x| - \frac{a_i^{\rho}}{1-\rho}((|x| + a_i)^{1-\rho} - a_i^{1-\rho})\right] & \rho \neq 1 \end{cases} \tag{16}$$

and

$$c_i(x) = K_i \mathrm{sgn}(x) \left[|x| - \frac{1 - e^{-a_i|x|}}{a_i}\right], \tag{17}$$

respectively, where again, K_i is always chosen such that $c_i(A) = A$. The first compander above has an inverse (expander) with a closed form expression for $\rho = 1/2$ and $\rho = 2$. Specifically, for $\rho = 1/2$, the above general expression can be simplified to

$$c_i(x) = K_i \mathrm{sgn}(x) \left(\sqrt{|x| + a_i} - \sqrt{a_i}\right)^2 \tag{18}$$

whose inverse is

$$e_i(y) = \mathrm{sgn}(y) \left[\left(\sqrt{\frac{|y|}{K_i}} + \sqrt{a_i}\right)^2 - a_i\right], \tag{19}$$

and for $\rho = 2$, we have

$$c_i(x) = K_i \mathrm{sgn}(x) \left(|x| + \frac{a_i^2}{|x| + a_i} - a_i\right) \tag{20}$$

whose inverse is

$$e_i(y) = \mathrm{sgn}(y) \frac{|y| + \sqrt{|y|^2 + 4a_i K_i |y|}}{2K_i}. \tag{21}$$

In the next section, we will examine the performance of the piecewise linear companders as well as companders (18) and (20) on real images.

Finally, a general remark is in order. We mentioned earlier that the assumption on perfect entropy coding may not be realistic. For example, a more realistic assumption would be that the number of bits required for coding the quantized version of $c_i(x)$, is roughly a constant plus $\log c_i(x)$. (This results from applying universal codes for the integers on $\lceil c_i(x)/\Delta_i \rceil$, and it also reflects the behavior of the Huffman coding tables of JPEG). Thus, optimal tradefoff between rate and distortion, in the high resolution limit outside $[-\delta_i, \delta_i]$, can be formalized

as a Largrange minimization problem: Find a monotone function $c_i(\cdot)$ from $[0, A]$ onto itself (with an antisymmetric extension to $[-A, 0]$) that minimizes

$$\int_{\delta_i}^{A} p_i(x) \left[\log c_i(x) + \frac{\lambda}{[c_i'(x)]^2} \right] dx$$

with boundary conditions $c_i(0) = 0$ and $c_i(A) = A$. We are not aware of the existence of a closed-form expression to the solution of this problem. However, the discretized version of this problem

$$\sum_k p_i(k\Delta) \left[\log c_i(k\Delta) + \frac{\lambda'}{[c_i((k+1)\Delta) - c_i(k\Delta)]^2} \right], \quad \Delta > 0,$$

can be solved by dynamic programing. Nonetheless, by inspection of this Lagrange function, it is intuitively appealing that for small values of x (and hence also of $\log c_i(x)$), the derivative $c_i'(x)$ should be small at the benefit that for large values of x, it would be large in order to minimize the entire cost while keeping the boundary conditions. Again, this supports the above mentioned guidelines for the choice of good companders.

4 Experimental Results

We have examined the piecewise linear compander, the compander given in eq. (18), and the one given in eq. (20), all with $a_i = a$ for all i, i.e., the companders of all coefficients have the same paramter value, commonly denoted by a. We examined several natural black-and-white images: "Boats", "Gold", "Lena", 'Hotel", "Barbara", and "Zelda." In all our experiments, the compander of eq. (18) turned out to give the best results, and so, we will present results for this compander only.

Our experiments were conducted as follows: For each image, we compared the high resolution rate-distortion tradeoff that is obtained in two different ways. The first is by using the ordinary JPEG algorithm with the default quantization table for luminance scaled by the Q-factor, for several values of Q. The second method was to hold $Q = 0$, and to scan a range of values for the paramter a of the compander.

The results are summarized in Tables 1–6. Each table is organiized as follows: the left-hand part displays PSNR results (in dB) vs. compression ratio (CR) for ordinary JPEG as a function of the Q-factor, whereas the right-hand part does the same for companded JPEG as a function of the paramter a. Also, for the sake of convenience of comparing the results, each table caption indicates pairs of entries in the table corresponding to approximately the same PSNR. For these pairs of entries it is therefore fair to compare the compression ratios. As can be seen, the compander offers a considerably better rate-distortion tradeoff in the high resolution limit: For the same PSNR, it gives a better CR by 10–25%. The advantage of companding over ordinary JPEG decreases as we move to working points of more aggressive quantization. Nonetheless, the range where companding is advantageous is fairly wide.

6 References

[1] V. R. Algazi, "Useful approximations to optimum quantization," *IEEE Trans. Commun. Tech.*, COM–14, pp. 297–301, 1966.

[2] M. Bakshi and D. R. Fuhrmann, "Improving the visual quality of JPEG-encoded images via companding," *J. Electronics Imaging*, vol. 6, no. 2, pp. 189–197, April 1997.

[3] J. A. Buclew, "Companding and random quantization in several dimensions," *IEEE Trans. Inform. Theory*, vol. IT–27, no. 2, pp. 207–211, March 1981.

[4] J. A. Buclew, "A note on optimal multidimensional companders," *IEEE Trans. Inform. Theory*, vol. IT–29, no. 2, p. 279, March 1983.

[5] A. Gersho, "Principles of quantization," *IEEE Transactions on Circuits and Systems*, vol. CAS–25, no. 7, pp. 427–436, July 1978.

[6] N. S. Jayant and P. Noll, *Digital Coding of Waveforms.* Englewood Cliffs, NJ, Prentice-Hall, 1984.

[7] T. Linder, R. Zamir, and K. Zeger, "High-resolution source coding for non-difference distortion measures: multidimensional companding," preprint 1997.

[8] P. W. Moo and D. L. Neuhoff, "Optimal compressor functions for multidimensional companding of memoryless sources," preprint 1998.

[9] P. F. Panter and W. Dite, "Quantizing distortion in pulse-code modulation with non-uniform spacing," *Proc. IRE*, vol. 39, pp. 44–48, 1951.

[10] G. M. Roe, "Quantizing for minimum distortion," *IEEE Trans. Inform. Theory*, vol. IT–10, pp. 384–385.

[11] B. Smith, "Instantaneous conpanding fo quantizing signals," *Bell Syst. Tech. J.* , vol. 27, pp. 446–472, 1948.

Q	PSNR	CR	a	PSNR	CR
0	55.35	1.69	0	55.35	1.69
5	44.81	2.77	5	45.05	3.41
10	42.11	4.31	10	43.44	4.21
15	40.63	5.60	20	42.17	5.04
20	39.56	6.66	30	41.38	5.71
30	38.00	8.65	40	40.68	6.23
40	36.89	10.30	50	40.04	6.68
50	35.93	11.91	80	39.37	7.45
60	35.30	13.12	100	38.86	7.89
70	34.74	14.40	120	36.24	8.32
80	34.24	15.65	150	37.93	8.83

Table 1: PSNR[dB] vs. CR for the image "Boats." Entries corresponding to approximately the same PSNR: $Q = 5$ vs. $a = 5$, $Q = 10$ vs. $a = 20$, $Q = 15$ vs. $a = 40$, $Q = 20$ vs. $a = 80$, $Q = 30$ vs. $a = 150$.

Q	PSNR	CR	a	PSNR	CR
0	55.29	1.49	0	55.29	1.49
5	43.47	2.32	5	44.80	2.62
10	39.83	3.56	10	42.59	3.14
15	38.19	4.64	20	40.81	3.75
20	37.18	5.55	30	39.78	4.32
30	35.86	7.24	40	38.88	4.82
40	34.98	8.70	50	38.11	5.31
50	34.25	10.17	80	37.34	6.07
60	33.77	11.29	100	36.76	6.59
70	33.32	12.61	120	36.24	7.11
80	32.92	13.88	150	35.78	7.71

Table 2: PSNR[dB] vs. CR for the image "Gold." Entries corresponding to approximately the same PSNR: $Q = 10$ vs. $a = 30$, $Q = 15$ vs. $a = 50$, $Q = 20$ vs. $a = 80$, $Q = 30$ vs. $a = 150$.

Q	PSNR	CR	a	PSNR	CR
0	55.22	1.48	0	55.22	1.48
5	43.11	2.30	5	44.50	2.61
10	39.46	3.56	10	42.15	3.14
15	37.99	4.71	20	40.37	3.74
20	37.09	5.70	30	39.35	4.35
30	35.92	7.56	40	38.51	4.96
40	35.10	9.27	50	37.82	5.57
50	34.42	10.97	80	37.16	6.45
60	33.97	12.25	100	36.69	7.07
70	33.56	13.67	120	36.28	7.64
80	33.18	15.00	150	35.92	8.25

Table 3: PSNR[dB] vs. CR for the image "Lena." Entries corresponding to approximately the same PSNR: $Q = 10$ vs. $a = 30$, $Q = 15$ vs. $a = 50$, $Q = 20$ vs. $a = 80$, $Q = 30$ vs. $a = 150$.

Q	PSNR	CR	a	PSNR	CR
0	55.25	1.50	0	55.25	1.50
5	43.41	2.33	5	44.73	2.63
10	40.03	3.62	10	42.63	3.11
15	38.51	4.74	20	41.02	3.67
20	37.52	5.66	30	40.10	4.22
30	36.20	7.39	40	39.32	4.72
40	35.30	8.83	50	38.67	5.17
50	34.54	10.16	80	37.99	5.84
60	33.97	12.25	100	37.49	6.26
70	33.50	12.26	120	37.04	6.62
80	33.05	13.30	150	36.61	7.05

Table 4: PSNR[dB] vs. CR for the image "Hotel." Entries corresponding to approximately the same PSNR: $Q = 10$ vs. $a = 30$, $Q = 15$ vs. $a = 50$, $Q = 20$ vs. $a = 100$.

Q	PSNR	CR	a	PSNR	CR
0	55.24	1.48	0	55.24	1.48
5	44.24	2.31	5	44.98	2.58
10	40.83	3.43	10	43.11	3.03
15	38.99	4.32	20	41.64	3.54
20	37.68	5.05	30	40.75	3.97
30	35.79	6.43	40	39.97	4.28
40	34.46	7.60	50	39.28	4.55
50	33.36	8.73	80	38.53	5.04
60	32.75	9.59	100	37.96	5.32
70	31.86	10.50	120	37.42	5.57
80	31.22	11.43	150	36.96	5.88

Table 5: PSNR[dB] vs. CR for the image "Barbara." Entries corresponding to approximately the same PSNR: $Q = 10$ vs. $a = 30$, $Q = 15$ vs. $a = 80$, $Q = 20$ vs. $a = 120$.

Q	PSNR	CR	a	PSNR	CR
0	55.24	1.79	0	55.24	1.79
5	44.15	2.94	5	44.34	3.84
10	41.80	4.87	10	42.48	4.82
15	40.86	6.66	20	41.19	6.11
20	40.23	8.14	30	40.49	7.62
30	39.32	11.02	40	39.94	9.14
40	38.63	13.53	50	39.51	10.50
50	37.99	15.97	80	39.09	12.06
60	37.57	17.82	100	38.78	13.08
70	37.14	19.75	120	38.50	13.93
80	36.73	21.62	150	37.34	15.60

Table 6: PSNR[dB] vs. CR for the image "Zelda." Entries corresponding to approximately the same PSNR: $Q = 5$ vs. $a = 5$, $Q = 15$ vs. $a = 30$.

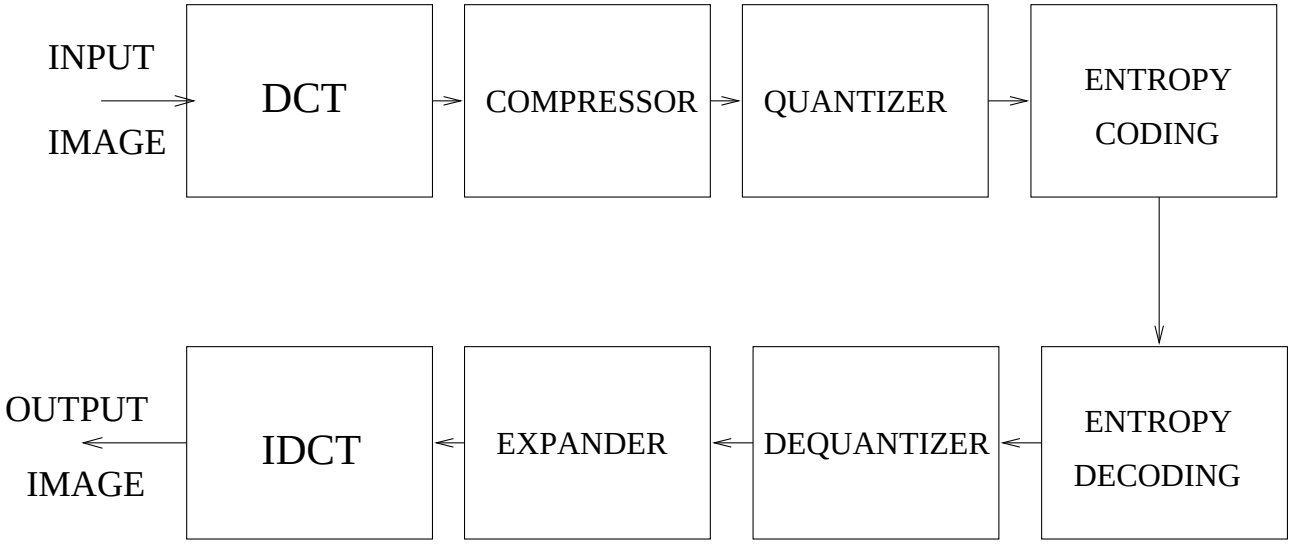

Figure 1: DCT domain companded JPEG codec.

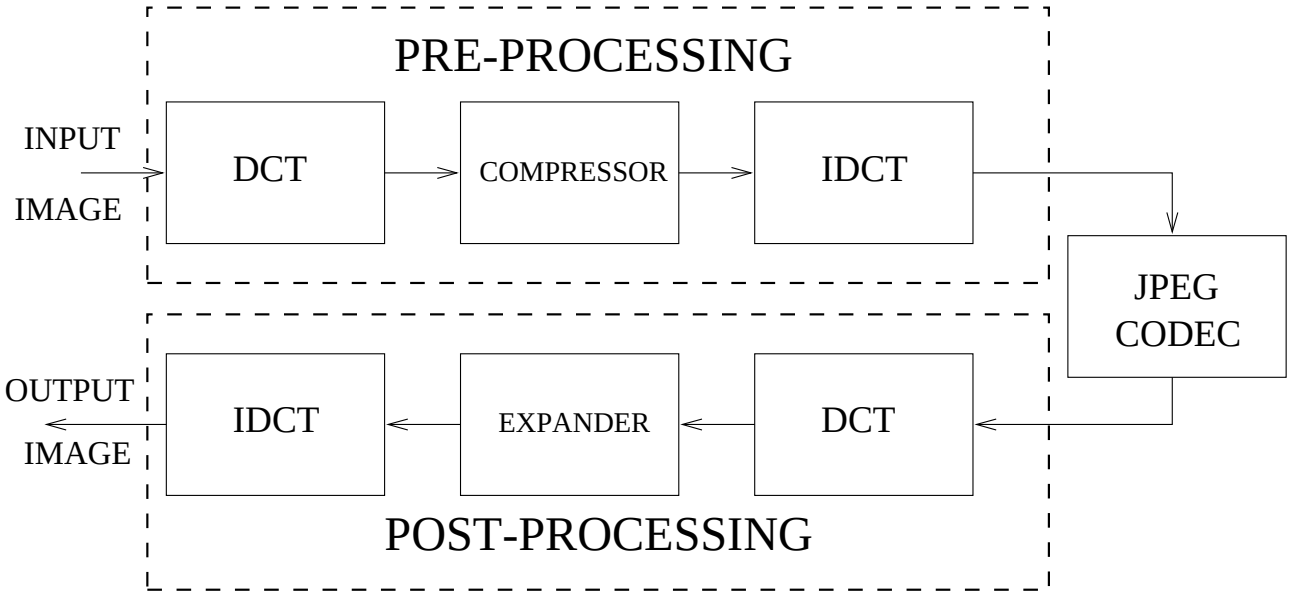

Figure 2: Standard compliant implementation.

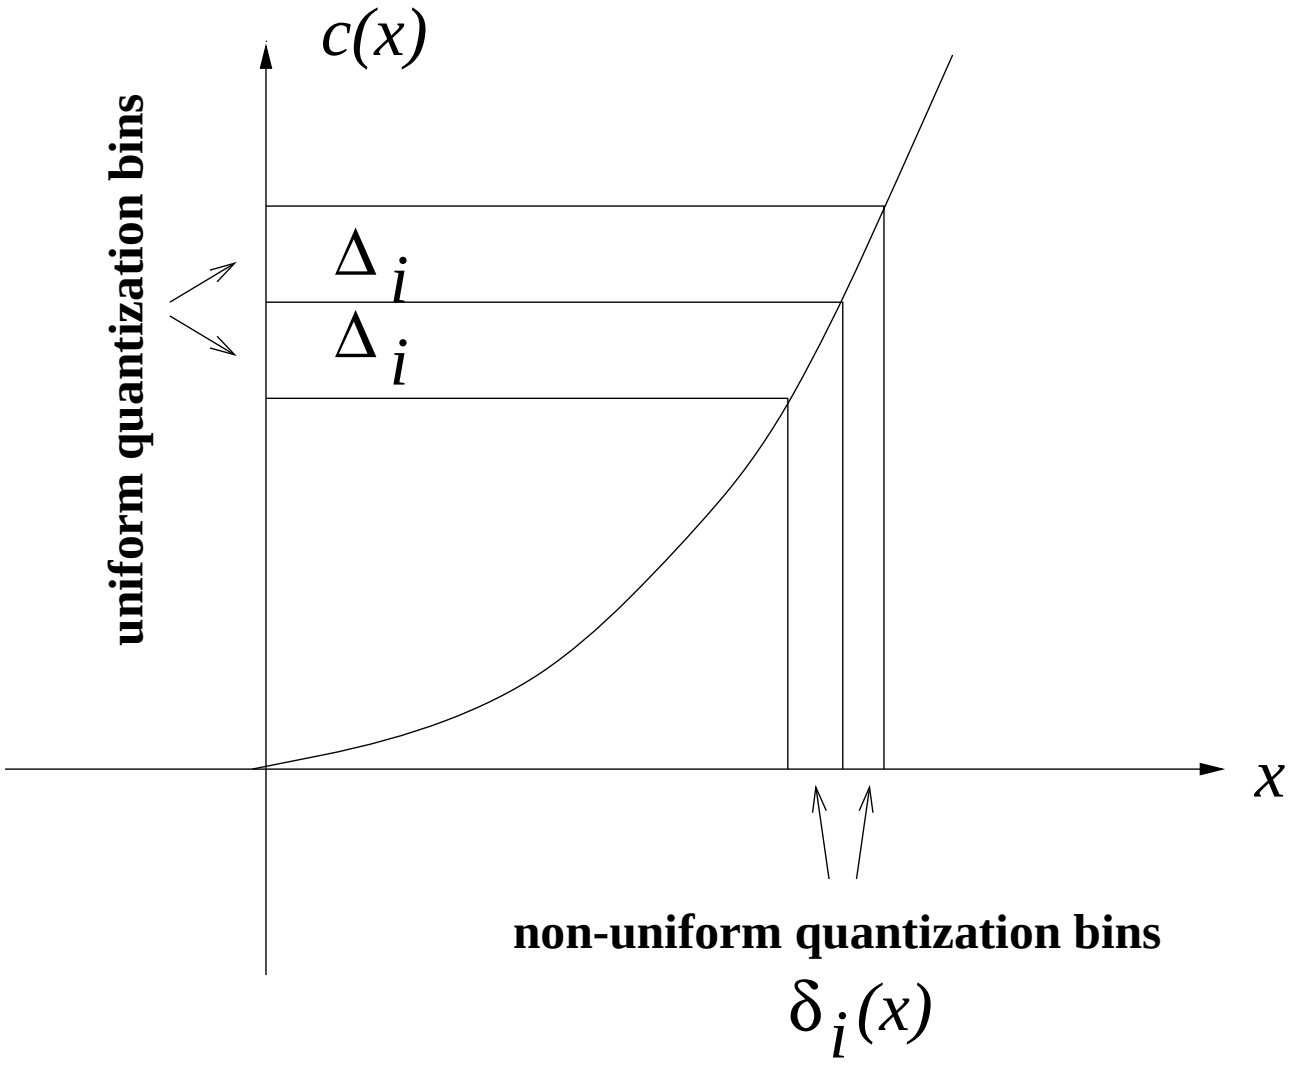

Figure 3: Compander and quantization levels.