Extending RD-OPT with Global Thresholding for JPEG Optimization

Viresh Ratnakar
University of Wisconsin-Madison
Computer Sciences Department
Madison, WI 53706
Phone: (608) 262-6627
Email: ratnakar@cs.wisc.edu

Miron Livny
University of Wisconsin-Madison
Computer Sciences Department
Madison, WI 53706
Phone: (608) 262-0856
Email: miron@cs.wisc.edu

Abstract

In this paper, the RD-OPT algorithm for optimizing Discrete Cosine Transform quantization tables [RL95] is extended to incorporate global thresholding. Performance gains are possible by zeroing off some DCT coefficients in DCT-based image compression. We describe a global thresholding scheme, in which the zeroing thresholds for coefficients can be arbitrarily different from those determined by the quantization table. Unlike local thresholding [CR95], the zeroing decisions are not made separately for each image block. This simplifies the use of thresholding, is easier to optimize and is almost as effective as local thresholding.

1 Introduction

The Discrete Cosine Transform (DCT) [ANR74] lies at the heart of most commonly used lossy image and video compression schemes [PM93, MP91]. The extent of compression achieved depends upon the coarseness of quantization of the transform coefficients. Most DCT-based compression schemes use uniform scalar quantization of the coefficients, as determined by a table of quantizers. The RD-OPT algorithm [RL95] is a technique to optimize these quantization tables in an image-specific manner. RD-OPT has the advantages of being fast, and optimizing jointly over a wide range of compression rates, as compared to other DCT quantization table optimization techniques such as [WG93, MS93].

Within the framework of compression via DCT and quantization tables, it is possible to improve performance (i.e., the rate-distortion tradeoff), by selectively setting some coefficients to zero. This allows finer quantizer scales for the coefficients that are retained, and does not add any complexity to the decoder. A technique for making the zeroing decisions can be found in [CR95], which is an extension of the quantization table optimization scheme of [WG93]. However, making the zeroing decisions on a per-coefficient basis in large images is computationally expensive. Further, the algorithm in [CR95] needs to be rerun, every time a new rate of compression is desired. We extend RD-OPT here, to determine global *Threshold Tables* that decide the zeroing cutoff levels (or thresholds) for each DCT coefficient. We also present performance results and analysis of global thresholding.

2 Global Thresholding

The 8×8 DCT-based image compression process divides the image into 8×8 blocks, transforms each block using the DCT, quantizes the coefficients, and stores them using variable-length entropy coding such as Huffman coding. Quantization of coefficients is typically specified by a table of 64 quantizer scales, called the quantization table. For an image block f, if $\hat{f}$ represents the block of DCT coefficients, and Q is the quantization table, then the quantized coefficients, $\hat{f}_Q$ are calculated as:

$$\hat{f}_Q[n] = \hat{f}[n] /\!\!/ Q[n], \quad 0 \leq n \leq 63.$$

Here $/\!\!/$ represents division followed by rounding to the nearest integer[1].

If the quantization table entry $Q[n]$ is q, then the n^{th} DCT coefficient gets quantized to zero when its absolute value is less than $q/2$. Global thresholding allows us to increase this zeroing threshold above $q/2$. Let $T[64]$ be a table of non-negative real numbers, called the *threshold table.* We define global thresholding as follows: if for any image block f, $|\hat{f}[n]| < T[n]$, then set $\hat{f}[n]$ to zero in the quantization step. The combined result of quantization and thresholding is represented as the block $\hat{f}_{Q,T}$. Define

$$x /\!\!/\!\!/ (q,t) = \begin{cases} 0 & \text{if } |x| < t \\ x /\!\!/ q & \text{otherwise.} \end{cases}$$

Then, $\hat{f}_{Q,T}[n] = \hat{f}[n] /\!\!/\!\!/ (Q[n], T[n])$. We refer to this step as quantization by the pair (Q, T).

The table T need not be included with the compressed image, as the decompressor does not need to know the thresholds. The decompressor simply multiplies each quantized/thresholded coefficient by the quantization table entry, to calculate its approximation of the original DCT coefficients.

2.1 Splitting distortion into DCT coefficients

The distortion between the original image blocks and the decompressed image blocks is the same as that between the original DCT coefficient blocks and the reconstructed (quantized and dequantized) coefficient blocks. We use this fact to split the distortion resulting from the use of any (Q, T) into 64 coefficient-wise component. We denote the contribution of the n^{th} coefficient (when quantized by (q, t)) to the total distortion by $D_n(q, t)$.

[1] For any $b > 0$, $a /\!\!/ b = \begin{cases} \lfloor \frac{a}{b} + 0.5 \rfloor & \text{if } a \geq 0 \\ -\lfloor \frac{-a}{b} + 0.5 \rfloor & \text{if } a < 0. \end{cases}$

2.2 Splitting rate into DCT coefficients

The degree of compression achieved is usually expressed in terms of the *rate* of the compressed image, which is the number of bits used per pixel:

$$\text{rate} = \frac{\text{size of compressed image in bits}}{\text{number of pixels in the image}}.$$

Low rates are achieved when the quantized blocks $\hat{f}_{Q,T}$ have similar entries (low entropy). The most common case is that of a coefficient being quantized to zero. The more zeros there are in $\hat{f}_{Q,T}$, the fewer bits it would take to store it. Thus, increasing the entries in Q and T tends to decrease the rate. We denote the rate resulting with the use of tables (Q, T) as $R(Q, T)$.

DCT has the nice property of being very close to the Karhunen-Loeve-Hotelling tranform, a transform that produces uncorrelated coefficients [ANR74]. The lack of correlation between coefficients allows the rate to be decomposed as a sum of contributions from individual coefficients. It has been shown in [RFVK94] that the coefficient-wise average of entropies of the quantized DCT coefficients is a very good estimate of the rate resulting from two-pass Huffman coding of runlengths. This allows us to approximate $R(Q, T)$ as a sum of rates of individual coefficients. Let $R_n(q, t)$ be defined as

$$R_n(q,t) = \frac{1}{64}\text{Entropy}\{(\hat{f}[n] /\!/\!/ (q,t))\},$$

Where the entropy is measured over all the blocks in the image[2]. Then

$$R(Q,T) \approx \sum_{n=0}^{63} R_n(Q[n], T[n]). \tag{1}$$

Thus, $R(Q, T)$ can be decomposed into a sum of contributions from individual coefficients, just like $D(Q, T)$.

2.3 Analysis of global thresholding

Consider a uniform scalar quantizer with step size q. Let the total number of samples be N, and of these, let kx be the number that get quantized to 0, and x be the number that get quantized to 1. To analyse global thresholding, we calculate the decrease in rate (empirical entropy) and the increase in distortion resulting from zeroing off a fraction ρx of the samples that were getting quantized to 1. If the mean distance of these samples from the value $q/2$ is δ, then it can be shown that:

$$\text{Decrease in rate, } \Delta R(\rho) = \frac{x}{N}\log_2\frac{(k+\rho)^{k+\rho}(1-\rho)^{1-\rho}}{k^k}, \text{ and}$$

[2] If $(\hat{f}[n] /\!/\!/ (q,t))$ takes the value v in a fraction $p_v > 0$ of all blocks $\hat{f}$, then this entropy is $-\sum_v p_v log_2 p_v$.

$$\text{Increase in distortion, } \Delta D(\rho) = \frac{x}{N} 2q\rho\delta.$$

The optimization problem of minimizing distortion while keeping the rate under some fixed budget is equivalent to minimizing the Lagrangian $R + \lambda D$ in the sense that solutions to the latter at each non-negative λ are solutions to the former for some rate budget. Then, for a given λ, the zeroing threshold (i.e., the fraction ρ) should be set so as to maximize the resulting drop in the Lagrangian, $\Delta R(\rho) - \lambda \Delta D(\rho)$. Typically, this drop increases for a while as ρ increases from 0 , and then decreases.

Note that the above results are general in the sense that they apply to uniform scalar quantization of arbitrary signals. Here, they have been applied to the DCT coefficient samples, for which k is large (as most coefficients get quantized to zero), and hence thresholding is especially effective.

3 The RD-OPT algorithm extended for global thresholding

In this section, we describe the RD-OPT algorithm for optimizing quantization and threshold tables with respect to rate-distortion tradeoffs for a given image. This is very similar to optimizing the quantization tables alone as in [RL95], hence only the aspects pertaining to thresholding are described.

It is desirable to have low rate (high compression) and low distortion (high quality). However, varying (Q, T) has opposite effects on distortion and rate. The distortion $D(Q, T)$ tends to increase and the rate $R(Q, T)$ tends to decrease as the entries in Q and T are made larger. The tradeoff between $D(Q, T)$ and $R(Q, T)$ is different for different images.

RD-OPT takes an image I as input and optimizes quantization and threshold tables for a wide range of rates and distortions. Only integral entries are considered for threshold tables, but arbitrary precision is possible at the cost of higher complexity.

Recall that RD-OPT calculates the contributions of individual coefficients to the total rate and total distortion, and then runs a Dynamic Programming algorithm to minimize two sums [RL95]. For each coefficient n and for each possible quantizer scale q, the contribution to total rate of the n^{th} coefficient is calculated as $\mathtt{R}[n][q]$, and the contribution to total distortion as $\mathtt{D}[n][q]$. Thresholding can be included in a straightforward manner: the individual contributions of each coefficient are now calculated for each possible q, and for each possible threshold t. We use thresholds t that are multiples of 0.5, so that the DCT statistics (which are histograms of coefficient counts over bins of size 0.5) can be used to predict rate (entropy) and distortion accurately.

Then, for any (Q, T), the rate of compression is

$$R(Q, T) = \sum_{n=0}^{63} \mathtt{R}[n][Q[n]][T[n]],$$

and the distortion is

$$D(Q,T) = \sum_{n=0}^{63} \mathtt{R}[n][Q[n]][T[n]].$$

Algorithm RD-OPT

Input: An image I of width W and height H, with pixel values in the range $[0 \ldots M]$.

Output: Optimized DCT quantization tables Q and threshold tables T.

Step 1. Gather DCT statistics for the image.

Step 2. Use the statistics to calculate $\mathtt{R}[n][q][t]$ and $\mathtt{D}[n][q][t]$ for each possible (q,t).

Step 3. Use dynamic programming to optimize $R(Q,T)$ against $D(Q,T)$.

The details of these steps are exactly as in [RL95]. The only difference is the extra freedom in choosing thresholds along with quantizer scales.

3.1 Complexity

Let MAXRATE be the largest value of rate (after discretization by multiplying with a large integer constant). If all thresholds are allowed to be in the range $0 \ldots 255$, and all quantization table entries in the range $1 \ldots 255$, then each row of the dynamic program may require upto MAXRATE$\cdot 255 \cdot 256$ steps, which is prohibitively expensive. We reduce this with the following observations:

1. If the n^{th} coefficient has maximum value $C < 255$ in the entire image, $Q[n]$ need not be more than $2C+1$, and $T[n]$ need not be bigger than C. This is useful for higher-frequency coefficients.

2. For $Q[n] = q$, $T[n]$ is usually in the interval $[q/2, 2q]$. This is an empirical observation, but is very useful for reducing complexity.

3. If for the n^{th} coefficient the total number of $(Q[n], T[n])$ pairs possible (after the previous two reductions) is greater than MAXRATE, then it is better to precompute the best (q,t) pair for each of the MAXRATE possible (discretized) values of $\mathtt{R}[n][Q[n]][T[n]]$. In fact, $\mathtt{R}[n][Q[n]][T[n]]$ will typically span a much smaller range than MAXRATE. This usually results in two orders of magnitude reduction in complexity.

4 Using the Lagrangian instead of Dynamic Programming

Instead of using Dynamic Programming to optimize $R(Q,T)$ against $D(Q,T)$, one can use the Lagrangian minimization approach. Given a $\lambda > 0$, we find Q and T

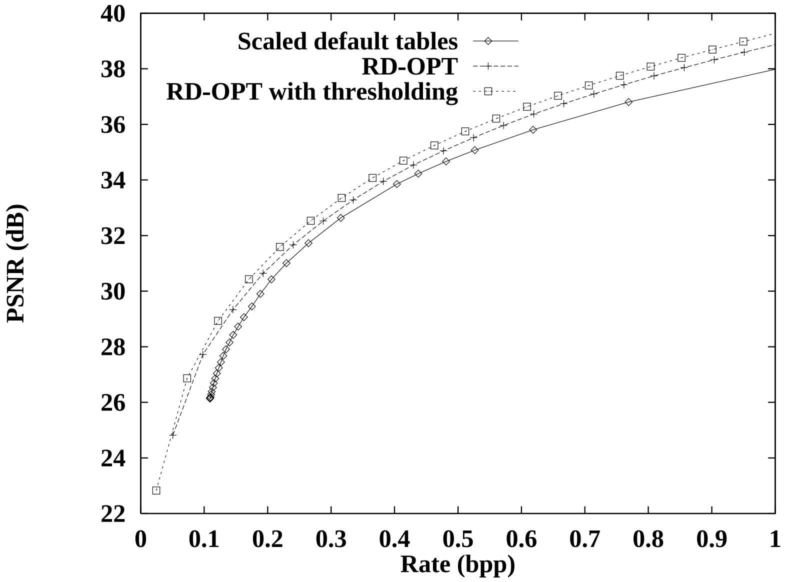

Figure 1: Performance results for *Lena*

that minimize $R(Q,T) + \lambda D(Q,T)$. This is readily done by choosing each $q = Q[n]$ and $t = T[n]$ so as to minimize $\mathtt{R}[n][q][t] + \lambda\mathtt{D}[n][q][t]$. This has the advantage that for a given q, not all t values need be tried: starting at $t = q/2$, t needs only to be increased until the Lagrangian stops decreasing. With the Lagrangian approach, however, a search for λ needs to be done when a target rate or distortion is specified.

5 Performance

Figure 1 shows the PSNR-Rate plots for JPEG compression of Lena using various quantization strategies. RD-OPT with global thresholding gives a PSNR gain of about 2dB at most bit rates. This is comparable to the results with local thresholding in [CR95]. For example, at 1.0 bits per pixel, RD-OPT with global thresholding gives a PSNR of 39.3 dB, as compared to 39.6 dB for [CR95]. The advantage here is that with one execution of RD-OPT, quantization and threshold tables can be obtained for a wide range of rates, while other optimization strategies such as [WG93, MS93, CR95] need to be used repeatedly, every time a new rate/quality setting is needed.

Figure 2 and Figure 3 show the PSNR-Rate plots for two scientific images, *Brain* and *CloudySky*, respectively. *Brain* is a 665×810 grayscale image of a cat's brain, while *CloudySky* is a 1024×512 grayscale image of a canopy with cloudy sky in the background. For both these images, global thresholding along with quantization table optimization results in an improvement of 1-2 dB.

Note that many perceptually weighted distortion measures can be used in RD-OPT [RL95], and this remains true for the thresholding extension.

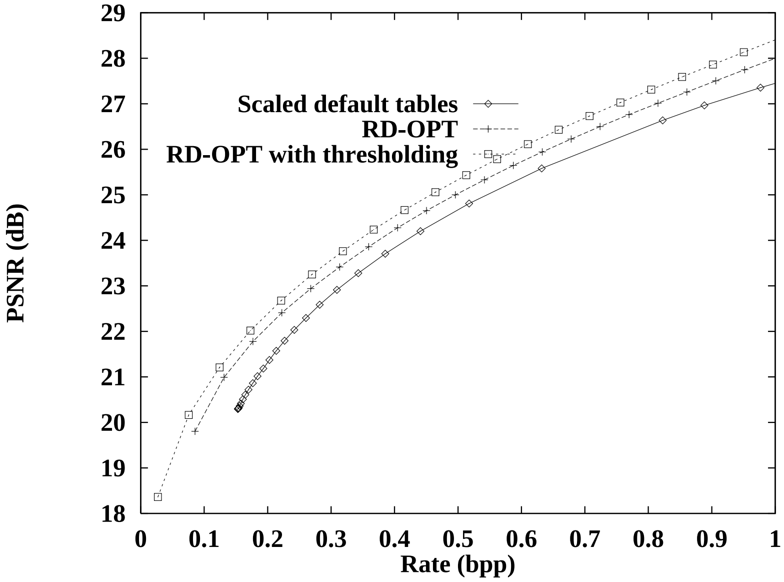

Figure 2: Performance results for *Brain*

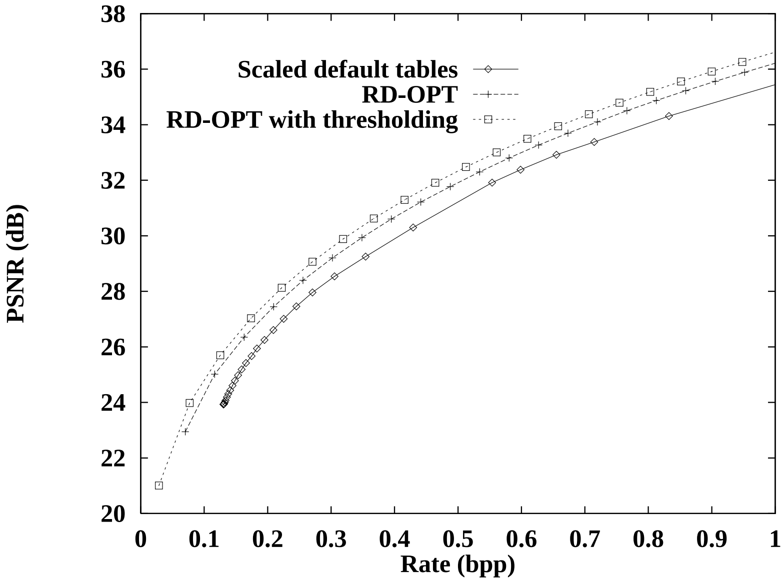

Figure 3: Performance results for *CloudySky*

6 Conclusion

We have introduced the notion of global thresholding, which can be used to improve the performance of DCT-based image compression. Global thresholding provides a conceptually simple way to exploit the advantages of thresholding possible (within the JPEG syntax, for example), without the complexity of local thresholding which makes zeroing decisions on a per-coefficient basis. It gives almost equally high quality improvements. In addition, global thresholding can be easily incorporated into most existing JPEG compressors, by giving them the threshold table as a parameter. This separates the optimization itself from the actual compression.

References

[ANR74] Ahmed, N., Natarajan, T., and Rao, K. R. Discrete Cosine Transform. *IEEE Trans. Computers*, C-2390-3, Jan. 1974.

[CR95] Crouse, M. and Ramchandran, K. JPEG optimization using an entropy-constrained quantization framework. *Proceedings of Data Compression Conference*, pages 342–351, 1995.

[MP91] MPEG I draft: Coding of Moving Pictures and associated audio for digital storage, 1991. Document ISO/IEC-CD-11172.

[MS93] Monro, D. M. and Sherlock, B. G. Optimum DCT Quantization. *Proceedings of Data Compression Conference*, pages 188–194, 1993.

[PM93] Pennebaker, W. B. and Mitchell, J. L. *JPEG Still Image Data Compression Standard*. Van Nostrand Reinhold, New York, 1993.

[RFVK94] Ratnakar, V., Feig, E., Viscito, E., and Kalluri, S. Runlength encoding of quantized DCT coefficients. *IBM RC 19693 (87318) 8/5/94 (Also in Proceedings of SPIE '95)*, 1994.

[RL95] Ratnakar, V. and Livny, M. RD-OPT: An Efficient Algorithm For Optimizing DCT Qunatization Tables. *Proceedings of Data Compression Conference (Also, Technical Report 1257, Dept of Computer Sciences, UW-Madison)*, pages 332–341, 1995.

[WG93] Wu, S. and Gersho, A. Rate-constrained picture-adaptive quantization for JPEG baseline coders. *Proc. Inter. Conf. Acoustics, Speech and Signal Processing*, 5:389—392, April 1993.