Smoothing DCT Compression Artifacts

A. J. Ahumada, Jr. and R. Horng
NASA Ames Research Center, Moffett Field, CA

ABSTRACT

Image compression based on quantizing the image in the discrete cosine transform (DCT) domain can generate blocky artifacts in the output image. It is possible to reduce these artifacts and RMS error by adjusting measures of block edginess and image roughness, while restricting the DCT coefficient values to values that would have been quantized to those of the compressed image. We also introduce a DCT coefficient amplitude adjustment that reduces RMS error.

INTRODUCTION

Lossy image compression in the DCT domain is achieved by the quantization of the DCT coefficients. The quantization of a single coefficient in a single block causes the reconstructed image to differ from the original image by an error image proportional to the associated basis function in that block. Our goal here is to try to reduce the blockiness without reducing the accuracy of the reconstruction.

The sum of squared differences between adjacent block edge pixels is a measure of blockiness that tends to increase with the amount of compression. A similar measure taken away from the block edge provides an estimate of the blockiness in the original image. Lowering the blockiness value to this estimate, while limiting coefficient change to the quantization level, can reduce both apparent blockiness and RMS (root mean square) error. Weighting spatial frequency components by their spatial frequency and then summing their squared values give a measure of image roughness. Some improvement results from rereducing the within-block image roughness if reducing the blockiness has increased it. This method has been presented earlier.[1] Here we also describe a DCT amplitude adjustment procedure that itself reduces RMS error and improves the performance of the smoothing algorithm.

The goal of this project is the development of algorithms for improving the quality of images that have already been compressed by a JPEG-like scheme in which the image is divided up into blocks, each block is converted to DCT coefficients, and these coefficients are then quantized. In JPEG there follows a stage of lossless encoding that we ignore for the present purpose. We assume that we have the quantized coefficients and the matrix of quantization values. Our problem is to find an image that is more like the original image than the image obtained simply by performing the inverse DCT on the quantized coefficients.

Fig. 1 shows an original image. Fig. 2 shows that image after quantization and restoration without any de-blocking. Fig. 3 shows the corresponding images after our de-blocking algorithm is applied. Our method can be described as follows:

1) Adjust the amplitudes of the DCT coefficients to reduce the RMS error.
2) Measure the blockiness of the image and estimate how blocky it should be.
3) Lower the blockiness to the estimate.
4) Ensure that all DCT coefficients quantize to those of the compressed image.
5) If the within-block roughness of the image has increased, restore it to its original value.
6) Ensure that all DCT coefficients quantize to those of the compressed image.

In the rest of the paper, we make this description more precise, show some quantitative results for this image. and relate our work to that of others. We conclude that if the quantization is strong enough to generate significant block artifacts, our method gives moderate de-blocking and a small decrease in the RMS image error.

THEORY

The DCT image transform

The discrete cosine transform (DCT) has become a standard method of image compression[2,3]. Typically the image is divided into 8×8-pixel blocks, which are each transformed into 64 DCT coefficients. The DCT coefficients $I_{u,v}$, of an $N \times N$ block of image pixels $i_{x,y}$, are given by

$$I_{u,v} = \sum_{x=0}^{N-1} \sum_{y=0}^{N-1} i_{x,y}\ c_{x,u}\, c_{y,v}\ ,\quad u,v = 0, N-1, \qquad (1a)$$

where

$$c_{x,u} = \alpha_u \cos\left(\frac{\pi u}{2N}[2x+1]\right), \qquad (1b)$$

and

$$\alpha_u = \begin{cases} \sqrt{1/N}\ , & u = 0 \\ \sqrt{2/N}\ , & u > 0 \end{cases}. \qquad (1c)$$

DCT coefficient quantization

In JPEG quantization[2,3] a coefficient is quantized by the operation

$$S_{u,v} = \text{Round}\left(\frac{I_{u,v}}{Q_{u,v}}\right). \quad (2)$$

The compressed image contains both the $S_{u,v}$ for all the blocks and the $Q_{u,v}$. To retrieve the image, first the DCT coefficients are restored (with their quantization error) by

$$\hat{I}_{u,v} = S_{u,v}\, Q_{u,v}\,, \quad (3)$$

where $Q_{u,v}$ denotes the quantizer step size used for coefficient $I_{u,v}$. The blocks of image pixels are reconstructed by the inverse transform:

$$\hat{i}_{x,y} = \sum_{u=0}^{N-1} \sum_{v=0}^{N-1} \hat{I}_{u,v}\, c_{x,u}\, c_{y,v}\,, \quad (4)$$

which for this normalization is the same as the forward transform. Our goal is to find better estimates of these coefficients.

DCT coefficient amplitude adjustment

The standard method of restoring the coefficient, using Eq. 3, is equivalent to replacing each coefficient by the center of the quantization interval in which the original coefficient falls. The distribution of the non-DC coefficients for a given u,v peaks at zero and decreases monotonically. For quantization intervals not including zero, the distribution of the original coefficients is denser at the end of the interval closer to zero. The mean of the distribution is the minimum mean squared error reconstructor. For simplicity, we model the distribution of absolute amplitudes as exponential with mean μ. We estimate μ by the mean of the $|S_{u,v}|$. We then replace $S_{u,v}$ by

$$S_{u,v} - 0.5 + \mu - \frac{e^{-1/\mu}}{1-e^{-1/\mu}}\,, \ \text{if } S_{u,v} > 0\,,$$

$$S_{u,v}\,, \ \text{if } S_{u,v} = 0\,, \quad (5)$$

$$S_{u,v} + 0.5 - \mu + \frac{e^{-1/\mu}}{1-e^{-1/\mu}}\,, \ \text{if } S_{u,v} < 0\,.$$

The dotted line in Fig. 4 shows the RMS improvement for the image in Fig. 1 as a function of the level of quantization, which ranged from 5 to 100 in steps of 5. A constant quantization matrix was used. For moderately high levels of quantization, the amplitude adjustment was not as effective. Comparing the predicted means of the interval distribution with the actual means, we find that Eq. 5 overestimates the desired correction when the mean of the $|S_{u,v}|$ is a small fraction. This is probably caused by the poor fit of the exponential near zero, where the actual distribution is flat.

A global measure of blockiness

Suppose i_1 and i_2 are the image values of two pixels that are next to each other in the same row or column, but are in different blocks. We assume that the blockiness of the compressed image is related to the fact that before compression, the values of i_1 and i_2 were usually similar, but they have been made more different by the quantization. We define the edge variance E to be sum of the squared differences for all such pixel pairs.

$$E = \sum (i_1 - i_2)^2, \quad (6)$$

The block edge variance E is our measure of image blockiness.

We estimate the desired value of the edge variance by computing the same measure for the pixels just inside the edge on either side and taking the average. If this estimate is less than the edge variance, we attempt to reduce the edge variance to this value. This reduction is done in the direction of the gradient of edge variance and may not be completely achieved if the minimum reduction in this direction is above the next-to-edge variance.

Adjusting the edge variance in this way only alters the edge pixels. The problem has been reduced at the boundary, but a new problem has been created inside the blocks. We attempt to reduce this problem by monitoring a measure of image roughness in the blocks.

A global measure of intra-block roughness

Our measure of the intra-block roughness is the sum of squares of the DCT coefficients weighted by their spatial frequency,

$$R = \sum_{u,v} (u^2 + v^2)\, {I_{u,v}}^2, \quad (7)$$

summed over all blocks. By weighting each component $I_{u,v}$ by its spatial frequency $\sqrt{u^2+v^2}$, we obtain a measure closely related to the total edge variance inside the block. If this measure increases after reducing the edge variance, we attempt to return it to its original value by changing it along its gradient.

Smoothing results

The solid line in Fig. 4 shows the ratio of the RMS error in the smoothed picture to that of the compressed image. It shows that the smoothing usually improved the accuracy of image restoration: it improves the RMS error except in the case that the quantization is very low and the the next-to-edge estimate is also low. Fortunately, in this case, the image will only be slightly changed and there would be no apparent need for de-blocking.

DISCUSSION

The present problem is a special case of the general problem discussed by Wu and Gersho[4], optimal decoding of an image under the assumption of constrained encoding. They formalize the problem as that of finding the image that minimizes the average value of a distortion function. In an earlier paper[5], they applied this concept to the derivation of optimal additive contribution to the block for each possible level of each DCT coefficient. Their (NLI) decoder gave a 0.7 dB improvement in mean square error on a diverse 23 image training set and about 0.5 dB improvement on new images. They report apparent reduction in blockiness, but since the method was restricted to within blocks, it does not directly attack the problem. Our amplitude adjustment is an application of this strategy to the single DCT coefficient amplitudes. The use of the exponential distribution model removes the dependence on amplitude and represents the effects of coefficient indices and quantization level to be represented by a single parameter easily estimated from the quantized data.

Stevenson[6,7] has analyzed this problem from the maximum *a posterior* (MAP) point of view. The goal is to find the image that maximizes the probability of the image given the quantized image. Using a non-Gaussian Markov random field model for the image distribution, the resulting solution is the minimum of a roughness function similar to ours with the squaring operation replaced by a Huber operation in the space domain. The quantization constraint is also enforced. The method appears to strongly reduce blocking for the Lena picture, JPEG compressed at 30 to 1, but no quantitative measures are reported. The Huber function is reported to be better than the squaring operation at allowing edges in the original image to persist through the smoothing.

Our methods and results are similar to the iterative projection method of Yang, Galatsanos, and Katsaggelos[8], They also use edge variance and the quantization constraint. They compute separate horizontal and vertical edge variances and force them to their correct values in the original image by a weighted averaging of edge pixels. They iterate these two constraints in conjunction with the quantization constraint and range constraints in both the space and DCT domains. Since the constraints are projections onto convex sets, iterating them is guaranteed to terminate, since the original image is a solution. They report a 1 dB improvement in RMS error of reconstruction and strong apparent reduction in the blockiness for the 256×256 Lena image when the PSNR for the original reconstruction was 27.9 dB. Our method differs from their method mainly in the addition of the within block smoothness constraint and the estimation of the edge variance.

SUMMARY

We have presented a method of estimating DCT coefficients from their quantized values and the quantization matrix, which are both included in the JPEG[2] standard compressed image file. This method of image reconstruction can reduce blockiness and RMS error in DCT quantized images. It includes a simple method of DCT coefficient amplitude adjustment that reduces the RMS error itself.

ACKNOWLEDGMENTS

We appreciate the assistance of A. Watson and J. Solomon. This work was supported by NASA RTOP Nos. 199-06-31 and 505-64-53 and NASA Cooperative Agreement NAC 2-307 with Stanford University.

REFERENCES

1. A. J. Ahumada, Jr., R. Horng, "De-blocking DCT compressed images", B. Rogowitz and J. Allebach, eds., *Human Vision, Visual Processing, and Digital Display V*, Conference Proceedings vol. 2179, SPIE, Bellingham, WA, 1994.
2. G. Wallace, "The JPEG still picture compression standard", *Communications of the ACM*, vol. 34, no. 4, pp. 31-44, 1991.
3. W. B. Pennebaker, J. L. Mitchell, *JPEG Still Image Data Compression Standard*, van Nostrand Reinhold, New York, 1993.
4. S. Wu, A. Gersho, "Enhanced video compression with standardized bit stream syntax", IEEE ISCASSP Proceedings, vol. I, pp. 103-106, 1993.
5. S. Wu, A. Gersho, "Enhancement of transform coding by nonlinear interpolation", Visual Communications and Image Processing '91: Visual Communication vol. 1605, SPIE, Bellingham, WA, pp. 487-498, 1991.
6. Robert L. Stevenson, "Reduction of coding artifacts in transform image coding", IEEE ISCASSP Proceedings, vol. V, pp. 401-404, 1993.
7. T. P. O'Rourke, R. L. Stevenson, "Improved image decompression for reduced transform coding artifacts", In S. Rajala and R. Stevenson, eds., *Image and Video Processing II*, Conference

Proceedings vol. 2182, SPIE, Bellingham, WA, 1994.
8. Y. Yang, N. Galatsanos, A. Katsaggelos, "Iterative Projection algorithms for removing the blocking artifacts of block-DCT compressed images", IEEE ISCASSP Proceedings, vol. V, pp. 405-408, 1993.

Figure 1. The original image.

Figure 2. The image after standard quantization.

x

Figure 3. The image after smoothing.

Figure 4. The ratio of RMS error after amplitude adjustment to that after standard quantization (dotted line) and the same ratio after smoothing (solid line).