

DCT QUANTIZATION NOISE IN COMPRESSED IMAGES

Mark A. Robertson and Robert L. Stevenson

Laboratory for Image and Signal Analysis
University of Notre Dame
Notre Dame, IN 46556

ABSTRACT

In lossy image compression schemes utilizing the discrete cosine transform (DCT), quantization of the DCT coefficients introduces error in the image representation and a loss of signal information. At high compression ratios, this introduced error produces visually undesirable compression artifacts that can dramatically lower the perceived quality of a particular image. This paper provides a spatial domain model of the quantization error based on a statistical noise model of the error introduced when quantizing the DCT coefficients. The resulting theoretically derived spatial domain quantization noise model shows that, in general, the compression noise in the spatial domain is both correlated and spatially varying. This provides some justification for many of the ad hoc artifact removal filters that have been proposed. An accurate description of quantization noise is essential if one hopes to remove, or at least alleviate, the visibility of compression artifacts.

1. INTRODUCTION

The presence of artifacts in compressed images is a well known phenomenon, both in the scientific and consumer communities. In order to compensate for these artifacts, one must have a thorough understanding of them. This paper derives statistical descriptions of the quantization noise that, in addition to providing important insight into the origins of compression noise, can also be readily applied in a post processing algorithm to remove the artifacts.

Blocking is one of the most common forms of compression artifact, and many attribute its existence to the inability of the block DCT to include correlations between blocks. While there is truth to this, the blocking phenomenon can also be explained by considering the compression of a single block only. This is the approach taken here: this paper considers a single block of the image to which the DCT is applied, and derives results for that block. As shall be seen, results for a single block prove sufficient to explain blocking artifacts, as well as compression noise in general.

Section 2 provides the basics of the DCT and establishes notation used in the remainder of the paper. Sections 3 and 4 derive a statistical description of the noise introduced by quantization of DCT coefficients. A discussion of the work is provided in Section 5, where an explanation for the existence of blocking artifacts is proposed, including where one can expect their presence or absence. Concluding remarks are given in Section 6.

2. THE DISCRETE COSINE TRANSFORM

In one dimension, the signal to be transformed is $z[n], 0 \leq n < N$, which is also represented by the length- N column vector \mathbf{z} .

The elements of the DCT matrix $\mathbf{H} = \{H[k, n]\}$ are

$$H[k, n] = \begin{cases} \frac{1}{\sqrt{N}}, & k = 0, 0 \leq n < N, \\ \sqrt{\frac{2}{N}} \cos \left[\frac{\pi(2n+1)k}{2N} \right], & 1 \leq k < N, \\ & 0 \leq n < N. \end{cases} \quad (1)$$

The DCT of \mathbf{z} is then $\mathbf{y} = \mathbf{H}\mathbf{z}$. Since the DCT is a real unitary transform, $\mathbf{H}^{-1} = \mathbf{H}^T$ and the inverse DCT (IDCT) is described by $\mathbf{z} = \mathbf{H}^T \mathbf{y}$. If \mathbf{h}_k^T denotes the k^{th} row of \mathbf{H} , then

$$\mathbf{z} = \sum_{k=0}^{N-1} \mathbf{h}_k \mathbf{y}[k]. \quad (2)$$

Equation (2) gives a series representation of the vector \mathbf{z} in terms of the DCT basis-functions \mathbf{h}_k and the DCT coefficients $\mathbf{y}[k]$. Any one-dimensional (1-D) signal of length N can be represented by a sum of these N basis functions of different frequencies, where each basis function in the sum is scaled by the DCT coefficient $\mathbf{y}[k]$. Figure 1 shows the DCT basis functions for $N = 8$.

When compressing a signal, the DCT coefficients are typically quantized rather than the actual signal \mathbf{z} . The quantized DCT coefficients are denoted as \mathbf{y}_q , with $\mathbf{y}_q = Q[\mathbf{y}]$, where $Q[\cdot]$ is the quantization operator. Quantization is a non-linear operation that results in a loss of information; only scalar quantization is considered here, where each element of \mathbf{y} is quantized individually. Scalar quantization is a many-to-one mapping that transforms intervals of real numbers $[q_i^k, q_{i+1}^k)$ to single real numbers. The superscript “ k ” accounts for the possibility of different quantization intervals for different frequency coefficients, and the subscript “ i ” indicates the i^{th} quantization level. Transform coefficients that are in these intervals are typically mapped to the midpoint of the interval, so that $y_q[k] = \frac{1}{2}(q_i^k + q_{i+1}^k)$ for $q_i^k \leq y[k] < q_{i+1}^k$.

The recovered signal \mathbf{z}_q is found by performing the inverse DCT on the quantized frequency values, $\mathbf{z}_q = \mathbf{H}^T \mathbf{y}_q$. Two quantities of interest in this paper are the quantization errors in both the spatial and the frequency domains. Spatial-domain error is represented by $\mathbf{e}_z = \mathbf{z}_q - \mathbf{z}$, and frequency-domain error by $\mathbf{e}_y = \mathbf{y}_q - \mathbf{y}$. Note that the quantization error in the spatial domain can be expressed as

$$\mathbf{e}_z = \mathbf{H}^T [\mathbf{y}_q - \mathbf{y}] = \sum_{k=0}^{N-1} \mathbf{h}_k (y_q[k] - y[k]). \quad (3)$$

Equation (3) gives the basis-function representation of the quantization error: the quantization error is the sum of the errors of each frequency coefficient multiplied by the corresponding DCT basis function.

The two-dimensional (2-D) DCT is simply a separable extension of the DCT in one dimension. Here, the matrix \mathbf{Z} is used




Fig. 1. One-dimensional DCT basis functions of length 8.

to represent the 2-D image block, with \mathbf{Y} , \mathbf{Z}_q , \mathbf{Y}_q , \mathbf{E}_z and \mathbf{E}_y defined analogous to the 1-D case.

As for the 1-D case, a basis-function representation for $N \times N$ image blocks can be written. Defining the $N \times N$ basis images $\mathbf{H}_{k,l} = \mathbf{h}_k \mathbf{h}_l^T$,

$$\mathbf{Z} = \sum_{k=0}^{N-1} \sum_{l=0}^{N-1} \mathbf{H}_{k,l} Y[k, l]. \quad (4)$$

The basis images in (4) are a staple of image processing text books, and will not be shown here.

Quantization of the 2-D DCT coefficients is performed as in the 1-D case, with a resulting spatial-domain quantization error of

$$\mathbf{E}_z = \sum_{k=0}^{N-1} \sum_{l=0}^{N-1} \mathbf{H}_{k,l} (Y_q[k, l] - Y[k, l]), \quad (5)$$

which gives the basis-image representation of the quantization error: the quantization error is the sum of the errors of each transform coefficient multiplied by each of the 2-D DCT basis images. A similar way of looking at the quantization error in compressed images has also been considered in [1].

On some occasions it will be convenient to use vector notation to represent images. When using this vector notation, the vector \mathbf{z} of length N^2 represents the image \mathbf{Z} , where $z[lN + k] = Z[k, l]$. This is equivalent to stacking the individual columns of \mathbf{Z} into a single vector. The matrix \mathcal{H} is then defined as the Kronecker product of \mathbf{H} with itself, $\mathcal{H} = \mathbf{H} \otimes \mathbf{H}$. The $N^2 \times N^2$ matrix \mathcal{H} is the 2-D DCT matrix for image blocks in vector format. If the various image matrices from before are represented in this stacked-column vector notation, the results can be summarized as $\mathbf{y} = \mathcal{H}\mathbf{z}$, $\mathbf{z} = \mathcal{H}^T \mathbf{y}$, and $\mathbf{z}_q - \mathbf{z} = \mathcal{H}^T (\mathbf{y}_q - \mathbf{y})$.

From (3) and (5), the spatial-domain error introduced by image compression is dependent on two key factors: the DCT basis functions, and the frequency-domain error introduced by quantization of the DCT coefficients. Statistics of the spatial-domain quantization noise given statistics of the frequency-domain quantization noise are examined in the next section. Statistics of the noise in the frequency domain are considered in Section 4.

3. QUANTIZATION NOISE STATISTICS

Pixel values in images are now viewed as *random* rather than as deterministic, which will enable the determination of expectations for the quantization error. The problem is first considered in one dimension, after which the results are extended to two dimensions.

Note that for a large variety of images the DCT has a decorrelating tendency, i.e., the resulting frequency coefficients are approximately uncorrelated [2]. With mild symmetry conditions on




Fig. 2. One-dimensional basis functions of the spatial-domain quantization error variance for sequences of length 8.

the pdf of these DCT coefficients [3], the DCT-domain quantization errors are uncorrelated random variables. The covariance matrix $\mathbf{K}_{e_y} = E[(\mathbf{y}_q - \mathbf{y})(\mathbf{y}_q - \mathbf{y})^T]$ is then diagonal, with its N non-zero elements equal to the quantization noise of the individual frequency-domain coefficients, $\sigma_{e_y}^2[k]$.

Suppose for the moment that \mathbf{K}_{e_y} is known; it is then possible to determine statistics of the error in the spatial domain. Of primary interest here is the covariance

$$\mathbf{K}_{e_z} = E[(\mathbf{z}_q - \mathbf{z})(\mathbf{z}_q - \mathbf{z})^T] = \mathbf{H}^T \mathbf{K}_{e_y} \mathbf{H}. \quad (6)$$

Equation (6) includes information about the correlation of the spatial-domain error sequence, but another quantity of interest is the variance of the individual spatial-domain errors. The variance of $e_z[n]$ is found as $\sigma_{e_z}^2[n] = \mathbf{K}_{e_z}[n, n]$, or in summation notation

$$\sigma_{e_z}^2[n] = \sum_{k=0}^{N-1} H^2[k, n] \mathbf{K}_{e_y}[k, k] = \sum_{k=0}^{N-1} H^2[k, n] \sigma_{e_y}^2[k], \quad (7)$$

where $H^2[k, n]$ is the square of the $(k, n)^{th}$ element of \mathbf{H} . A new transformation matrix \mathbf{M} is then defined, with elements $M[k, n] = H^2[k, n]$. Letting \mathbf{m}_k^T denote the k^{th} row of \mathbf{M} ,

$$\sigma_{e_z}^2 = \sum_{k=0}^{N-1} \mathbf{m}_k \sigma_{e_y}^2[k]. \quad (8)$$

Here, $\sigma_{e_z}^2$ is the vector representing the error variance of \mathbf{z}_q . The vectors \mathbf{m}_k form a basis for the error variance of \mathbf{z}_q : the error variance can be found by summing each error variance basis function scaled by the corresponding frequency-domain error variance. Error variance basis functions for sequences of length 8 are shown in Fig. 2. Transform-domain coefficients with large quantization error variances cause their corresponding basis functions to contribute more to the overall spatial-domain error variances. In general, this leads to spatial-domain error variances that vary with n , the spatial location. Also, note from (6) that in general the spatial-domain quantization noise is correlated.

To represent the covariance of the noise in two dimensions, it is convenient to use the stacked-column notation introduced earlier, where the random vectors \mathbf{z} , \mathbf{y} , \mathbf{z}_q , \mathbf{y}_q , \mathbf{e}_z and \mathbf{e}_y are formed by stacking the columns of their respective random images. As before, assume for the moment that the covariance matrix of the 2-D frequency-domain quantization noise is a known diagonal matrix \mathbf{K}_{e_y} , whose N^2 diagonal elements are taken from the individual frequency-domain error variances $\sigma_{e_y}^2[k, l]$. The covariance matrix of the spatial-domain error can then be written as

$$\mathbf{K}_{e_z} = E[(\mathbf{z}_q - \mathbf{z})(\mathbf{z}_q - \mathbf{z})^T] = \mathcal{H}^T \mathbf{K}_{e_y} \mathcal{H}. \quad (9)$$




Fig. 3. Two-dimensional DCT error variance images for 8×8 image blocks. DC corresponds to the top-left corner. Basis image values are ordered with 0 at black, and increase as the image gets brighter.

The $N^2 \times N^2$ matrix $\mathbf{K}_{\mathbf{e}_z}$ contains correlation information of the spatial-domain error image, but the individual error variances of the N^2 locations are also of interest. These individual variances are the diagonal elements of $\mathbf{K}_{\mathbf{e}_z}$. Switching back to 2-D notation, they can be written in a form similar to (7),

$$\sigma_{\mathbf{E}_z}^2[m, n] = \sum_{k=0}^{N-1} \sum_{l=0}^{N-1} \mathbf{H}^2[l, m] \mathbf{H}^2[k, n] \sigma_{\mathbf{E}_y}^2[k, l]. \quad (10)$$

Defining a set of images $\mathbf{M}_{k,l} = \mathbf{m}_k \mathbf{m}_l^T$, with elements $M_{k,l}[m, n] = \mathbf{H}^2[l, m] \mathbf{H}^2[k, n]$, allows the formation of a basis-image representation for the quantization error variance of the DCT in two dimensions,

$$\mathbf{\Lambda} = \sum_{k=0}^{N-1} \sum_{l=0}^{N-1} \mathbf{M}_{k,l} \sigma_{\mathbf{E}_y}^2[k, l]. \quad (11)$$

Here, $\mathbf{\Lambda}$ is the $N \times N$ array of error variances in the spatial domain. Thus the variance of the spatial-domain errors due to quantization with the 2-D DCT consists of a sum of the error basis images defined above, where each basis image is scaled by the error variance of its corresponding 2-D DCT coefficient. Figure 3 shows the error variance basis images for $N = 8$. Note that the error variance image will in general be spatially varying in m and n , as discussed previously for the 1-D case, and from (9) will in general be correlated.

4. DCT-DOMAIN QUANTIZATION ERROR

In the previous section, the autocovariance matrix of the spatial-domain quantization noise was derived, and depended on the frequency-domain quantization noise, which was assumed known. Expressions for these missing quantities are provide here. For the sake of brevity, focus is limited to the 2-D case.

A simple model for frequency-domain quantization noise would be a uniform distribution. The original random variable

$y[k, l]$ is quantized as $y_q[k, l] = Q[y[k, l]]$. If a realization $y_q[k, l]$ is observed such that $y_q[k, l]$ lies in the quantization interval $[q_i^{k,l}, q_{i+1}^{k,l})$, then the uniform model simply says that the quantization error $(y_q[k, l] - y[k, l])$ is distributed uniformly in the interval $[q_i^{k,l} - y_q[k, l], q_{i+1}^{k,l} - y_q[k, l])$. From an intuitive point of view, the uniform model says that there is no prior knowledge about the value of the original frequency coefficient $y[k, l]$ —it is equally likely anywhere within the quantization interval. For the uniform model, the DCT-domain quantization error can then be expressed as

$$\sigma_{\mathbf{E}_y}^2[k, l] = \frac{(q_{i+1}^{k,l} - q_i^{k,l})^2}{12}. \quad (12)$$

Using the uniform model makes analysis simple, for the variance of a uniform random variable is easy to calculate. However, although the uniform model makes analysis much simpler, it is not necessarily always appropriate. For example, many images (or portions of images) are quite often smooth, i.e., they have most of their energy concentrated at low frequencies, with little energy at high frequencies. This situation provides an example where the uniform model fails: if a quantized high-frequency coefficient of a “smooth” image is observed to be zero, then it makes little sense to assume that the original coefficient varied uniformly throughout the entire quantization interval. Rather, it makes more sense to assume that the original coefficient was quite close to zero, and in fact did not vary far from zero.

Simulations have also shown that using a uniform noise model sometimes results in poor estimates of quantization noise. Rather than assuming uniform DCT-domain quantization noise, one can incorporate prior knowledge of the DCT coefficients, which can yield more accurate results. Start by considering an image model in the spatial domain, which will lead to the desired result in the frequency domain.

A commonly used model for an image block \mathbf{Z} makes use of the separable autocorrelation function

$$E[Z[m, n] Z[i, j]] = \sigma_{\mathbf{Z}}^2 \rho_1^{|m-i|} \rho_2^{|n-j|}, \quad (13)$$

where the parameters ρ_1 and ρ_2 are the one-step correlations in the vertical and horizontal dimensions, and $\sigma_{\mathbf{Z}}^2$ is the variance of the image block \mathbf{Z} . It is assumed that the mean has been subtracted so that $E[\mathbf{Z}] = \mathbf{0}$. Using such a model, the variance of the $(k, l)^{th}$ DCT coefficient can be expressed as [2]

$$\sigma_{\mathbf{Y}}^2[k, l] = \sigma_{\mathbf{Z}}^2 [\mathbf{H} \mathbf{K}_1 \mathbf{H}^T]_{k,k} [\mathbf{H} \mathbf{K}_2 \mathbf{H}^T]_{l,l}, \quad (14)$$

where \mathbf{K}_x are defined as

$$\mathbf{K}_x = \begin{bmatrix} 1 & \rho_x & \rho_x^2 & \cdots & \rho_x^{N-1} \\ \rho_x & 1 & \rho_x & \cdots & \rho_x^{N-2} \\ \rho_x^2 & \rho_x & 1 & \cdots & \rho_x^{N-3} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \rho_x^{N-1} & \rho_x^{N-2} & \rho_x^{N-3} & \cdots & 1 \end{bmatrix}. \quad (15)$$

From (14) the variance, or energy, of each DCT coefficient is known provided that ρ_1 , ρ_2 and $\sigma_{\mathbf{Z}}^2$ are known. In practice, these parameters can be estimated from the observed data, or assumed known *a priori*.

The nature of the distributions of the DCT coefficients is of high importance for determination of frequency-domain quantization noise. There is precedence for using a Laplacian model for DCT coefficients in images and video [4], and that is the approach

taken here. The Laplacian parameter for the $(k, l)^{th}$ coefficient is found as

$$\lambda_{k,l} = \sqrt{\frac{2}{\sigma_Y^2[k,l]}}, \quad (16)$$

and the distribution of $Y[k, l]$ becomes

$$p_{Y[k,l]}(y) = \frac{\lambda_{k,l}}{2} \exp\{-\lambda_{k,l}|y|\}. \quad (17)$$

If the quantized DCT coefficient $Y_q[k, l] = y_q$ is observed to lie in the interval $[q_i^{k,l}, q_{i+1}^{k,l})$, then the distribution of $Y[k, l]$ given $Y_q[k, l] = y_q$ is zero outside of the quantization interval and $\gamma p_{Y[k,l]}(y)$ inside the quantization interval, where γ is a constant that ensures the distribution integrates to unity. Thus,

$$\sigma_{E_y}^2[k, l] = \gamma \int_{q_i^{k,l}}^{q_{i+1}^{k,l}} (y_q - y)^2 p_{Y[k,l]}(y) dy. \quad (18)$$

Equations (12) and (18) give two possible methods of determining the frequency-domain quantization noise. The method of (12) uses a uniform distribution for the DCT coefficients, which is equivalent to assuming no prior knowledge of the coefficients. Alternatively, the equation in (18) was derived based on a prior model for the image.

Simulations have suggested that by themselves, neither (12) nor (18) accurately predict the spatial-domain quantization noise. For quantized DCT coefficients that are observed to be zero, the Laplacian model indeed works well. However, when non-zero quantized DCT coefficients are observed, the Laplacian model seems to be inappropriate and the uniform model works better. Using the uniform model for non-zero coefficients can be viewed as approximating the tails of the Laplacian with a uniform distribution, rather than the exponentially decaying tail. Using different distributions depending on whether or not the observed quantized coefficient is zero essentially says that the energy of the DCT coefficients is expected to have concentrations near zero, and when away from zero the energy is expected to be more evenly distributed.

5. INTERPRETATION AND IMPLICATIONS

Consider the 1-D error variance basis functions of Fig. 2. If the DCT quantization errors were *equal* for each frequency, then the basis sum of (8) would result in *equal* variance at each spatial location. However, in general the DCT quantization errors will not be equal at each frequency. For example, for relatively smooth signals more signal energy will be contained in the low-frequency coefficients, with decreasing energy as the frequency increases. The high-frequency coefficients are quantized to zero, but since they have little energy they contribute very little quantization noise. However, the low-frequency coefficients contribute considerable quantization noise. Thus, one can see that the sum of (8) will have higher weights for the low-frequency components, and a quick consultation with Fig. 2 suggests that this will result in *higher error variance for locations near the boundaries*. This simple argument accounts for the presence of blocking artifacts. Note also that errors in the high-frequency components contribute little to the error right at the block boundaries.

On the other hand, consider signals that *do* contain significant high-frequency content, such as textured regions of an image. For these cases the situation described above is reversed, and it is easy to construct examples for which pixels near block boundaries have

less error variance than pixels near the middle. In these cases one should trust boundary pixels more than interior pixels!

The 2-D case is a ready extension of the case in one dimension, with similar arguments for the existence or absence of blocking artifacts in one dimension. While not as obvious as the 1-D case, the error variance basis images of Fig. 3 exhibit the same properties as the basis functions in Fig. 2: high error at block boundaries for low-frequency components, and low error at block boundaries for high-frequency components.

This analysis gives theoretic justification for many of the ad hoc post processing algorithms that attempt to remove blocking artifacts. Many of these algorithms perform selective smoothing at block boundaries, where more smoothing is applied in regions that are deemed “smooth,” and less or no smoothing in regions that are “not so smooth.” Filtering block edges like this makes perfect sense in light of the preceding discussion on the presence or absence of blocking artifacts.

Extensive tests have been performed for DCT compression noise in one dimension, and the results support the discussions of this paper—quantization noise in the spatial dimension is undoubtedly spatially varying and correlated, with a strong dependence on the quantization intervals and on the original signal statistics, namely the model parameters ρ_1 , ρ_2 and σ^2 . Experimental results in two dimensions have, while less complete than in one dimension, also supported the discussions of this paper.

6. CONCLUSION

This paper has provided a statistical description of the spatial-domain quantization noise in images compressed using the DCT, and has considered both the variance of the quantization error at individual pixels, as well as the correlations between them. The presence of blocking artifacts became easily explicable in this context, as was the absence of blocking artifacts. It was shown that in many cases one can actually expect pixel values near block boundaries to be *more* accurate than those in the block interior. Although not derived explicitly for the case of video compression, the results of this paper are easily applied to the quantization of a motion compensated residual as well. An understanding of compression noise is crucial if one expects to be able to compensate appropriately for such noise, making the concepts discussed in this paper valuable for those in the compressed image and video post-processing fields. In particular, the statistical descriptors of compression noise derived in this paper are readily applicable in a probabilistic image restoration formulation [3].

7. REFERENCES

- [1] H. A. Peterson, A. J. Ahumada, and A. B. Watson, “The visibility of DCT quantization noise,” in *J. Marreale, ed., Digest of Technical Papers, Society for Information Display*, Playa del Rey, CA, 1993, vol. 24, pp. 942–945.
- [2] A. K. Jain, *Fundamentals of Digital Image Processing*, Prentice-Hall, Inc., Englewood Cliffs, NJ, 1989.
- [3] M. A. Robertson, *High-Quality Reconstruction of Digital Images and Video from Imperfect Observations*, Ph.D. dissertation, University of Notre Dame, Apr. 2001.
- [4] E. Y. Lam and J. W. Goodman, “A mathematical analysis of the DCT coefficient distributions for images,” *IEEE Trans. on Image Proc.*, vol. 9, no. 10, pp. 1661–1666, Oct. 2000.