

Rate Bounds on SSIM Index of Quantized Images

Sumohana S. Channappayya, *Member, IEEE*, Alan Conrad Bovik, *Fellow, IEEE*, and Robert W. Heath, Jr., *Senior Member, IEEE*

Abstract—In this paper, we derive bounds on the structural similarity (SSIM) index as a function of quantization rate for fixed-rate uniform quantization of image discrete cosine transform (DCT) coefficients under the high-rate assumption. The space domain SSIM index is first expressed in terms of the DCT coefficients of the space domain vectors. The transform domain SSIM index is then used to derive bounds on the average SSIM index as a function of quantization rate for uniform, Gaussian, and Laplacian sources. As an illustrative example, uniform quantization of the DCT coefficients of natural images is considered. We show that the SSIM index between the reference and quantized images fall within the bounds for a large set of natural images. Further, we show using a simple example that the proposed bounds could be very useful for rate allocation problems in practical image and video coding applications.

I. INTRODUCTION

THE mean squared error (MSE) is a popular metric in the design of algorithms ranging from image quality assessment to quantization to restoration. The popularity of the MSE can be attributed to two main reasons: amenability to analysis and a lack of competitive perceptual distortion metrics. The importance of designing image processing algorithms optimized for perceptual quality measures, as opposed to the MSE, has been long recognized [7], [13]. Image coding algorithms that are optimized for perceptual distortion measures have been proposed by several authors and have become a part of image coding standards [10], [23], [24], [26], [43]. These algorithms use distortion measures such as those based on models of the human visual system (HVS), variants of the MSE such as the weighted MSE, or other empirical measures of quality.

Advances in understanding the HVS [29], and improved modeling of the statistics of natural scenes [36] have fostered the development of powerful image quality assessment (IQA) algorithms. These new IQA algorithms include, among others, the structural similarity (SSIM) index [39], or Wang-Bovik index, and its variants [41], [42], the Visual Information Fidelity Criterion (VIF) [30] and the Visual Signal to Noise Ratio (VSNR) [4]. These IQA algorithms outperform the measures currently used

in perceptual image coding algorithms such as the just-noticeable distortion (JND) metric [22], in addition to being better than the MSE in measuring perceptual image quality [32]. Image processing algorithms that optimize for these new IQA algorithms are only recently being developed [5], [6], [40]. The emergence of these IQA algorithms present the possibility of improved design of image coding algorithms.

In this paper, we analyze the relation between the quantization rate and the distortion measured by that SSIM index. The SSIM index is considered as opposed to other recent IQA algorithms since its performance is very competitive while being highly amenable to analysis. Nevertheless, a closed form relation between rate and SSIM index for fixed rate uniform quantization is extremely difficult, if not impossible, to derive. Instead, we derive upper and lower bounds on the SSIM index as a function of quantization rate (under a high rate assumption). It is shown that the SSIM index between the reference and quantized versions of the input falls within these bounds not only for uniform, Gaussian, and Laplacian sources, but also for natural images. The usefulness of the bounds in a practical image coding scenario is demonstrated using a simple rate allocation example.

A. Related Work

A brief overview of IQA algorithms is presented, followed by a discussion of image coding algorithms whose design is influenced by the properties of the HVS. The human eye is the ultimate receiver of all visual information. It is, therefore, natural to take into account the properties of the human visual system (HVS) in the design of an objective perceptual distortion measure. Several full-reference IQA algorithms have incorporated important properties of the HVS in their design. Yet, it is important to note that the HVS is still only weakly understood. This makes image quality assessment (even full-reference) a very challenging task that is still being actively researched.

The inadequacy of global MSE as a measure of image quality is well documented [8], [13], [21], [30], [31], [34], [37], [38], [39], [45]. In the context of image coding, it has been shown that incorporating models for several important aspects of the HVS such as the contrast sensitivity function (CSF), eye movement, edge masking etc. in the design, in addition to using distortion measures such as frequency weighted mean squared error, result in substantial quality gains when compared to using only the MSE [13], [43]. We briefly discuss perceptual distortion measures, followed by a discussion of perceptually motivated image coding systems.

Lubin's mechanistic model for the HVS [21], Daly's visual difference predictor (VDP) [8], Teo and Heeger's normalization model for the visual cortex [34], and Winkler's model of the HVS for video stimulus [45] are all examples of IQA and VQA algorithms that explicitly model the HVS in their design. All these algorithms use some form of the l_p norm and error-

Manuscript received August 23, 2007; revised April 16, 2008. Published August 13, 2008 (projected). This work was supported in part by the Texas Advanced Technology Program under Grant 003658-0380-2003 and in part by the National Science Foundation under Grant 0728748. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Giovanni Poggi.

S. S. Channappayya was with the Department of Electrical and Computer Engineering, The University of Texas at Austin, Austin TX 78712-0240 USA. He is now with the PacketVideo Corporation, San Diego, CA 92121 USA (e-mail: sumohana@gmail.com).

A. C. Bovik and R. W. Heath, Jr. are with the Department of Electrical and Computer Engineering, The University of Texas at Austin, Austin, TX 78712-0240 USA (e-mail: bovik@ece.utexas.edu; rheath@ece.utexas.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TIP.2008.2001400

pooling to arrive at a measure of quality. These measures perform consistently better than the MSE in measuring the perceptual quality of images and videos [32].

In a departure from earlier quality measures, Wang and Bovik [37] proposed the universal image quality index (UQI). The UQI differs from the earlier philosophy in that there is no explicit HVS modeling and the error is no longer measured using an l_p norm. The idea behind UQI is to measure the distortion of three image features (locally)—luminance, contrast, and correlation between the reference and the distorted image. The correlation between UQI and the mean opinion score (MOS) of subjective studies is significantly better than the correlation between MSE and MOS [37]. In addition to better perceptual correlation, this metric is intuitive, computationally efficient, and analytically amenable. Wang *et al.*, [39] proposed an improvement to the UQI in a measure called the structural similarity (SSIM) index. The SSIM index introduces stabilizing constants to handle the instability issues associate with the UQI. A detailed discussion of the SSIM index is presented in Section II-A. Some of the other recent IQA algorithms include the information fidelity criterion (IFC) [31], and its improved version, the visual information fidelity criterion [32]. Both these measures use empirical statistical models for natural scenes in their design [33], [36]. A performance analysis of these recent algorithms can be found in [30]. Another recent IQA algorithm is the visual signal to noise ratio (VSNR) [4]. A detailed discussion of the state-of-the-art image quality metrics can be found in [38].

The importance of optimizing image codecs for perceptual distortion measures has also been long recognized. A natural way to design such codecs is to use HVS models (such as the aforementioned ones) in the bit-allocation process. Mannos and Sakrison [23] provided the first-ever analysis of rate versus a perceptual distortion measure. The distortion measure used here is a weighted mean squared error and the optimal weighting function is determined empirically so that it maximizes perceptual quality. The general approach to the design of a perceptually optimal encoder is that the image is decomposed using transforms such as the DCT or wavelets. Visual models [21], [8] are used to mask the coefficients in such a way that the visibility of quantization errors is minimized. Examples of such codecs include Nill's perceptually weighted cosine transform approach [24], Eggerton and Srinath's perceptually weighted quantization of DCT coefficients subject to an entropy constraint [10], the Safranek–Johnston perceptual image coder [28], Watson's work on perceptually optimized DCT quantization matrices [44], Buccigrossi and Simoncelli's image codec based on the statistics of wavelet coefficients [2], Hontsch and Karam's adaptive coder with perceptual distortion control [16], Chandler and Hemami's dynamic contrast-based image coder [3], and Liu *et al.*'s JPEG2000 compliant encoding with a perceptual distortion control mechanism [20]. This list is nonexhaustive but covers several significant contributions made towards the design of perceptually optimal image codecs.

Eckert and Bradley [9] provide a thorough review of perceptually optimized image coding techniques. Pappas and Safranek [25] summarize several popular image quality measures with particular reference to those used for image compression. The flavor of our work presented here is most closely related to [23] and [44] in that we analyze the relation between the quantization rate and a popular and successful perceptual distortion measure (SSIM index).

B. Proposed Work

Next, we discuss the nature and relevance of the proposed work, followed by a brief outline of the rest of the paper. The MSE between a random variable (RV) X and its quantized version $Q(X)$ at a given quantization rate for fixed rate uniform scalar quantization is well known and is approximated as $\text{MSE}(X, Q(X)) \approx \Delta^2/12$ [12], [15], where Δ is the quantization step size. This is valid under the high resolution assumption, where the rate is large, and the contribution of the overload region is ignored. This result holds well for most practical scenarios in image and video coding and has been used to estimate the MSE (PSNR) of quantized images as a function of quantization rate (e.g., Sabir *et al.* [27]). Currently, a similar relation between the SSIM index and quantization rate does not exist. As with the MSE, such a relation would be very useful in bit-allocation problems in image and video coding. In a broader context, such a relation (between rate and SSIM index) could be used in the design of image and video codecs that can guarantee a desired level of perceptual quality.

In this paper, upper and lower bounds on the SSIM index as a function of quantization rate are proposed. Fixed rate uniform quantization under the high-rate assumption is considered. Since the discrete cosine transform (DCT) is commonly used in the transform coding of images, our analysis is carried out in the DCT domain. It is shown for a large set of natural images, in addition to uniform, Gaussian, and Laplacian sources, that the SSIM index between the reference and quantized versions of the input lie within the proposed bounds. We demonstrate the usefulness of the proposed bounds in a practical image coding scenario with a simple rate allocation example.

The paper is organized as follows. Section II presents a brief overview of the SSIM index, and uniform quantization, and formulates the SSIM index versus rate problem. Bounds on the SSIM index as a function of quantization rate are presented in Section III, followed by a discussion of its properties. In Section IV, we present results and highlight the usefulness of the proposed bounds, followed by concluding remarks in Section V.

II. PROBLEM FORMULATION

In this section, we provide an overview of the SSIM index, and express the space domain SSIM index in terms of the DCT coefficients of the space domain vectors. We briefly discuss uniform quantizers, and then the notation used in the sequel. Finally, the expression for the average SSIM index as a function of quantization step size is presented, and the motivation to bound this expression is discussed.

A. Structural Similarity Index

The most general form of the metric that is used to measure the structural similarity between two signal vectors \mathbf{x} and \mathbf{y} (both in \mathbb{R}^N) is

$$\text{SSIM}(\mathbf{x}, \mathbf{y}) = [l(\mathbf{x}, \mathbf{y})]^\alpha [c(\mathbf{x}, \mathbf{y})]^\beta [s(\mathbf{x}, \mathbf{y})]^\gamma. \quad (1)$$

The term $l(\mathbf{x}, \mathbf{y}) = (2\mu_x\mu_y + C_1)/(\mu_x^2 + \mu_y^2 + C_1)$ compares the luminance of the signals, $c(\mathbf{x}, \mathbf{y}) = (2\sigma_x\sigma_y + C_2)/(\sigma_x^2 + \sigma_y^2 + C_2)$ compares the contrast of the signals, and $s(\mathbf{x}, \mathbf{y}) = (\sigma_{xy} + C_3)/(\sigma_x\sigma_y + C_3)$ measures the structural correlation of the signals. The quantities μ_x, μ_y are the sample means of \mathbf{x}

and \mathbf{y} respectively, σ_x^2, σ_y^2 are the sample variances of \mathbf{x} and \mathbf{y} respectively, and σ_{xy} is the sample cross-covariance between \mathbf{x} and \mathbf{y} . The constants C_1, C_2, C_3 are used to stabilize the metric for the case where the means and variances become small. The parameters $\alpha > 0$, $\beta > 0$, and $\gamma > 0$, are used to adjust the relative importance of the three components. We use the following simplified form of the SSIM index in our work (where $\alpha = \beta = \gamma = 1$, and $C_3 = C_2/2$)

$$\text{SSIM}(\mathbf{x}, \mathbf{y}) = \left(\frac{2\mu_x\mu_y + C_1}{\mu_x^2 + \mu_y^2 + C_1} \right) \left(\frac{2\sigma_{xy} + C_2}{\sigma_x^2 + \sigma_y^2 + C_2} \right). \quad (2)$$

In image quality assessment, image blocks from the reference and distorted image constitute the vectors \mathbf{x} and \mathbf{y} respectively. The average of the SSIM values across the image (also called mean SSIM or MSSIM) gives the final quality measure. The design philosophy of the SSIM index is to acknowledge the fact that natural images are highly structured, and that the measure of structural correlation (between the reference and the distorted image) is important for deciding overall visual quality. Further, the SSIM index measures quality locally and is able to capture local dissimilarities better, unlike global quality measures such as MSE (and, hence, PSNR). Though (2) has a form that is more complicated than MSE, it remains analytically tractable. These features make the SSIM index attractive to work with.

B. Measuring SSIM index From DCT Coefficients

The DCT is widely used in the transform coding of images and videos and is central to several popular image (JPEG) and video coding standards (MPEG-x) [26], [11]. Highly efficient software and hardware implementations of the DCT form the core of several of these standards. The DCT is popular due to its energy compaction property, combined with efficient implementations. These reasons motivate us to perform our analysis in the DCT domain. The SSIM index in (2) is defined in the space domain, however. In the following, we derive a simple yet useful formula for measuring the SSIM index between two vectors from their DCT coefficients. Similar expressions can be obtained for the Fourier transform, as well.

In order to measure the SSIM index from DCT coefficients, the space domain mean, variance, and cross correlation are expressed in terms of DCT coefficients. The DCT of a vector $\mathbf{x} \in \mathbb{R}^N$ is [18]

$$X(k) = \sum_{i=0}^{N-1} \lambda(k) \cos\left(\frac{(2i+1)\pi k}{2N}\right) x(i) \quad (3)$$

$$\text{where } \lambda(k) = \begin{cases} \sqrt{\frac{1}{N}}, & \text{if } k = 0 \\ \sqrt{\frac{2}{N}}, & k \neq 0. \end{cases}$$

The DCT is a unitary transform and obeys the Parseval's theorem [35]. Using this property and (3), the following relations between space domain mean, variance, cross correlation, and the DCT coefficients are established

$$\mu_x = \frac{\sum_{i=0}^{N-1} x(i)}{N} = \frac{X(0)}{\sqrt{N}}, \text{ from (3)} \quad (4)$$




Fig. 1. Three-bit uniform quantizer. Shown are the quantization intervals and quantization levels. The figure also shows the notation used in the paper.

$$\sigma_x^2 = \frac{\sum_{i=0}^{N-1} x(i)^2 - N\mu_x^2}{N-1} = \frac{\sum_{k=1}^{N-1} X(k)^2}{N-1} \quad (5)$$

$$\sigma_{xy} = \frac{\sum_{i=0}^{N-1} x(i)y(i) - N\mu_x\mu_y}{N-1} = \frac{\sum_{k=1}^{N-1} X(k)Y(k)}{N-1}. \quad (6)$$

Substituting the space domain mean, variance, and cross correlation terms in the definition of SSIM (2) with the expressions in (4)–(6)

$$\text{SSIM}(\mathbf{x}, \mathbf{y}) = \left(\frac{2\frac{X(0)Y(0)}{N} + C_1}{\frac{X(0)^2 + Y(0)^2}{N} + C_1} \right) \times \left(\frac{2\frac{\sum_{k=1}^{N-1} X(k)Y(k)}{N-1} + C_2}{\frac{\sum_{k=1}^{N-1} X(k)^2 + Y(k)^2}{N-1} + C_2} \right). \quad (7)$$

This expression can be particularly useful when perform quality assessment of JPEG compressed images without having to decompress the images to the space domain (for computing the SSIM index from nonoverlapping blocks). We use the DCT domain expression for the SSIM index in the following analysis.

C. Uniform Quantization

Uniform quantization [1], [15] is the earliest, simplest, and most common form of quantization. It is used in a range of audio, image, and video coding applications [11], [26] mainly due to its simplicity. While other forms of quantization are well studied [15], asymptotic analysis of the relation between rate and distortion (mean squared error) for fixed-rate uniform quantization of symmetric sources with infinite support was reported only as recently as 2001 by Hui and Neuhoff [17]. We use results from this work in our implementation.

A uniform quantizer is illustrated in Fig. 1, with its granular region highlighted. The following notation is used in our analysis. The range of the granular region is denoted by $2L$, the number of quantization levels $N = 2^R$, where R is the quantization rate. The quantizer step size is denoted by $\Delta = (2L)/(N)$. The quantization levels are denoted by $y_i (0 \leq i \leq N-1)$, with $y_i = -L + (i+1/2)\Delta$. An interval in the granular region is denoted by $S_i = (y_i - \Delta/2, y_i + \Delta/2]$.

The relation between SSIM index and rate is derived under the high rate assumption and includes contributions only from the granular region. We assume that the DCT coefficients are independent [19], and that they are quantized at different step sizes Δ_i [26]. In the sequel, we use the term rate and quantization step size interchangeably (for notational convenience) since they are related by $\Delta = (2L)/(2^R)$, where $2L$ and R are as defined above.

D. Relation Between SSIM Index and Quantization Rate

Let $\mathbf{X} = [X_0 \ X_1 \ \dots \ X_{N-1}]^T$ denote a random vector composed of N DCT coefficients. In the sequel, we assume that the elements X_0, X_1, \dots, X_{N-1} are independent and have a joint density $f_{\mathbf{X}}(\mathbf{x}) = f(x_0)f(x_1)\dots f(x_{N-1})$. Each element of the random vector X_i is uniformly quantized at rate R_i . Under these assumptions, an interval in the joint granular region of the quantizers can be indexed by a vector $\mathbf{k} = [k_0, k_1, \dots, k_{N-1}]$, where k_i varies between 0 and $2^{R_i} - 1$. The vector $\mathbf{X} (\in S_{\mathbf{k}})$ is quantized by a point $Q(\mathbf{X}) = [Y_{k_0} \ Y_{k_1} \ \dots \ Y_{k_{N-1}}]^T$. We ignore the contribution of the overload region to the average SSIM between \mathbf{X} and $Q(\mathbf{X})$, and consider only the granular region. The average SSIM index between \mathbf{X} and $Q(\mathbf{X})$ is computed as shown in (8) at the bottom of the page. In practice, the most common DCT block size used in image and video coding applications is 8×8 . The expression in (8) however, is quite formidable to evaluate and implement even for DCT block sizes as small as 2×2 . Therefore, directly using (8) in a practical scenario appears extremely difficult, if not impossible. To make this problem tractable, we develop upper and lower bounds on (8). These bounds are shown to be accurate in estimating the range of the average SSIM index between the reference and quantized versions of a variety of sources. Further, it is also shown that the bounds are easier to implement and evaluate than an explicit solution to (8).

III. BOUNDS ON THE SSIM INDEX

In this section, we present upper and lower bounds on the average SSIM index as a function of quantization rate, evaluate these bounds for uniform, Gaussian, and Laplacian sources, and discuss several properties of these bounds. We assume that the DCT coefficients X_i are independent, and each coefficient is quantized separately at step size Δ_i . The high-resolution assumption is made, and only the contribution of the granular region is considered.

Theorem 3.1: For a random vector \mathbf{X} with independent components, the average SSIM index [as defined in (7) and (8)] between \mathbf{X} and its uniformly quantized version

$Q(\mathbf{X}) = [Q_0(X_0), Q_1(X_1), \dots, Q_{N-1}(X_{N-1})]^T$ is bounded with probability p by

$$\begin{aligned} \bar{M} \left(1 - \frac{\sum_{i=1}^{N-1} \Delta_i^2/12}{U + C_2} \right) &\leq \mathbb{E}[\text{SSIM}(\mathbf{X}, Q(\mathbf{X}))] \\ &\leq \bar{M} \left(1 - \frac{\sum_{i=1}^{N-1} \Delta_i^2/12}{V + C_2} \right) \end{aligned} \quad (9)$$

where Δ_i is the step size assigned to quantizer Q_i to quantize random variable X_i , $\bar{M} = \mathbb{E}[(2X_0Q(X_0) + NC_1)/(X_0^2 + Q(X_0)^2 + NC_1)]$ is the average value of the contribution from the mean term, U, V are quantities defined below, $2L$ is the range of the granular region of the quantizer with the largest span, and C_2 is a stabilizing constant [from (7)].

The terms V and U for a given probability p are

$$\begin{aligned} V &= \sqrt{2}\sigma_v \text{erf}^{-1}(2p - 1) + \mu_v \\ U &= \sqrt{2}\sigma_u \text{erf}^{-1}(1 - 2p) + \mu_u \end{aligned} \quad (10)$$

where $\sigma_v, \sigma_u, \mu_v, \mu_u$ are dependent on the source distribution. For the case of uniform sources, $U = (\min(X^2) + \min(Q(X)^2))$, $V = (\max(X^2) + \max(Q(X)^2))$, which makes the bounds hold with probability $p = 1$.

The main idea used to arrive at the bounds is to relate the SSIM index in (7) to MSE. Once this relation is established, the average SSIM index can be bounded with terms that are a function of the MSE, and the standard high resolution MSE result for fixed rate uniform quantization can be applied. The detailed proof can be found in the Appendix. The use of the standard MSE result gives the bounds several useful properties that are discussed in the following, and in Section IV.

The term \bar{M} is evaluated next for uniform, Gaussian, and Laplacian sources. The uniform source is considered as it is best suited to uniform quantization. Its upper bound provides an estimate of the highest average SSIM index that is achievable at a given quantization rate. Gaussian and Laplacian sources are considered as they are commonly used to model DCT coefficients [19]. The expressions for these bounds can be very easily implemented for these sources for any DCT block size. Most

$$\begin{aligned} &\mathbb{E}[\text{SSIM}(\mathbf{X}, Q(\mathbf{X}))] \\ &\approx \mathbb{E}[\text{SSIM}_{\text{gran}}(\mathbf{X}, Q(\mathbf{X}))] \\ &= \sum_{\mathbf{k}} \int_{S_{\mathbf{k}}} \text{SSIM}(\mathbf{X}, Q(\mathbf{X})) f_{\mathbf{X}}(\mathbf{x}) d\mathbf{x} \\ &= \sum_{\mathbf{k}} \int_{L_{k_0}}^{U_{k_0}} \frac{2x_0 Y_{k_0} + NC_1}{x_0^2 + Y_{k_0}^2 + NC_1} f(x_0) dx_0 \\ &\quad \times \int_{L_{k_1}}^{U_{k_1}} \int_{L_{k_2}}^{U_{k_2}} \dots \int_{L_{k_{N-1}}}^{U_{k_{N-1}}} \left(\frac{2\sum_{i=1}^{N-1} x_i Y_{k_i} + (N-1)C_2}{\sum_{i=1}^{N-1} x_i^2 + Y_{k_i}^2 + (N-1)C_2} \right) \\ &\quad \times f(x_1) f(x_2) \dots f(x_{N-1}) dx_1 dx_2 \dots dx_{N-1} \end{aligned} \quad (8)$$

importantly, we show that the bounds are indeed accurate not only for these sources, but also for a large set of natural images.

Before presenting expressions for \bar{M} , we revisit the notation used. These expressions correspond to the DC coefficient X_0 , quantized by quantizer Q_0 . The rate assigned to this quantizer is R_0 . The quantization levels of Q_0 are indexed using k_0 . The upper and lower limits of an interval S_{k_0} are notated by U_{k_0} and L_{k_0} respectively. Finally, a quantization level is denoted by Y_{k_0} . The number of DCT coefficients is denoted by N , and C_1 is a constant from (7). The steps in arriving at the expressions for \bar{M} are presented in the Appendix.

A. Uniform Source

Suppose that the DC coefficient X_0 is uniformly distributed over $[-A_0/2, A_0/2]$. For this case, the expression for \bar{M} is given by

$$\begin{aligned} \bar{M} = & \frac{1}{A_0} \sum_{k_0=0}^{2^{R_0}-1} Y_{k_0} \log \frac{U_{k_0}^2 + Y_{k_0}^2 + NC_1}{L_{k_0}^2 + Y_{k_0}^2 + NC_1} \\ & + \frac{NC_1}{\sqrt{Y_{k_0}^2 + NC_1}} \left[\tan^{-1} \frac{U_{k_0}}{\sqrt{Y_{k_0}^2 + NC_1}} \right. \\ & \left. - \tan^{-1} \frac{L_{k_0}}{\sqrt{Y_{k_0}^2 + NC_1}} \right]. \end{aligned} \quad (11)$$

B. Gaussian Source

If the DC coefficient X_0 is Gaussian distributed with zero mean and variance σ_0^2 , the expression for \bar{M} is given by

$$\begin{aligned} \bar{M} \approx & \frac{1}{\sqrt{2\pi}\sigma_0} \sum_{k_0=0}^{2^{R_0}-1} Y_{k_0} e^{\frac{Y_{k_0}^2 + 2\sigma_0^2 NC_1}{2\sigma_0^2}} \\ & \times \left[E_1 \left(\frac{L_{k_0}^2 + Y_{k_0}^2 + 2\sigma_0^2 NC_1}{2\sigma_0^2} \right) \right. \\ & \left. - E_1 \left(\frac{U_{k_0}^2 + Y_{k_0}^2 + 2\sigma_0^2 NC_1}{2\sigma_0^2} \right) \right] \end{aligned} \quad (12)$$

where $E_1(x) = \int_x^\infty (e^{-t})/(t)dt$ ($x > 0$), is the exponential integral. The expression for \bar{M} is an approximation in this case since we consider the contribution of only one term in the numerator (see Appendix for details).

If the AC coefficients X_1, X_2, \dots, X_{N-1} are independent and Gaussian distributed with zero mean and variance $\sigma_1^2, \sigma_2^2, \dots, \sigma_{N-1}^2$, respectively

$$\begin{aligned} \sigma_v^2 &= \frac{\sum_{i=1}^{N-1} 2\sigma_i^4}{(N-1)^2}, \mu_v = \frac{\sum_{i=1}^{N-1} \sigma_i^2}{(N-1)} + L^2 \\ \sigma_u^2 &= \frac{\sum_{i=1}^{N-1} 2\sigma_i^4}{(N-1)^2}, \mu_u = \frac{\sum_{i=1}^{N-1} \sigma_i^2}{(N-1)}. \end{aligned} \quad (13)$$

C. Laplacian Source

If the DC coefficient X_0 is Laplacian distributed with zero mean and variance σ_0^2 , the expression for \bar{M} is a combination of three terms depending on the values of the upper and lower limits of the interval S_{k_0} . Suppose that there are N_1 intervals

corresponding to Case 1 ($L_{k_0} > 0, U_{k_0} > 0$), N_2 intervals in Case 2 ($L_{k_0} \leq 0, U_{k_0} \leq 0$), and N_3 intervals in Case 3 ($L_{k_0} < 0, U_{k_0} > 0$), with $N_1 + N_2 + N_3 = 2^{R_0}$. Each case is evaluated as follows.

Case 1: $L_{k_0} > 0, U_{k_0} > 0$

$$\begin{aligned} T_1 = & \frac{1}{\sqrt{2}\sigma_0} \sum_{k_0=0}^{N_1-1} Y_{k_0} \{ e^{iab} [E_1(a(L_{k_0} + ib)) \\ & - E_1(a(U_{k_0} + ib))] + e^{-iab} [E_1(a(L_{k_0} - ib)) \\ & - E_1(a(U_{k_0} - ib))] \} \\ & + i \frac{NC_1}{\sqrt{b}} \{ e^{iab} [E_1(a(L_{k_0} + ib)) \\ & - E_1(a(U_{k_0} + ib))] \\ & - e^{-iab} [E_1(a(L_{k_0} - ib)) - E_1(a(U_{k_0} - ib))] \} \end{aligned} \quad (14)$$

where $a = (\sqrt{2})/(\sigma_0)$, $b = \sqrt{Y_{k_0}^2 + NC_1}$, $i = \sqrt{-1}$, $E_1(x)$ is the exponential integral.

Case 2: $L_{k_0} \leq 0, U_{k_0} \leq 0$

$$\begin{aligned} T_2 = & \frac{1}{\sqrt{2}\sigma_0} \sum_{k_0=0}^{N_2-1} Y_{k_0} \{ e^{-iab} [E_1(-a(L_{k_0} + ib)) \\ & - E_1(-a(U_{k_0} + ib))] + e^{iab} [E_1(-a(L_{k_0} - ib)) \\ & - E_1(-a(U_{k_0} - ib))] \} \\ & + i \frac{NC_1}{\sqrt{b}} \{ e^{-iab} [E_1(-a(L_{k_0} + ib)) \\ & - E_1(-a(U_{k_0} + ib))] \\ & - e^{iab} [E_1(-a(L_{k_0} - ib)) - E_1(-a(U_{k_0} - ib))] \} \end{aligned} \quad (15)$$

where $a = (\sqrt{2})/(\sigma_0)$, $b = \sqrt{Y_{k_0}^2 + NC_1}$, $i = \sqrt{-1}$, $E_1(x)$ is the exponential integral.

Case 3: $L_{k_0} < 0, U_{k_0} > 0$

$$T_3 = T_1 + T_2 \quad (16)$$

with T_1 as above evaluated over the interval $(0, U_{k_0}]$ and T_2 also as above, evaluated over $(L_{k_0}, 0]$

$$\bar{M} = T_1 + T_2 + T_3. \quad (17)$$

If the AC coefficients X_1, X_2, \dots, X_{N-1} are independent and Laplacian distributed with zero mean and variance $\sigma_1^2, \sigma_2^2, \dots, \sigma_{N-1}^2$, respectively

$$\begin{aligned} \sigma_v^2 &= \frac{\sum_{i=1}^{N-1} 5\sigma_i^4}{(N-1)^2}, \mu_v = \frac{\sum_{i=1}^{N-1} \sigma_i^2}{(N-1)} + L^2 \\ \sigma_u^2 &= \frac{\sum_{i=1}^{N-1} 5\sigma_i^4}{(N-1)^2}, \mu_u = \frac{\sum_{i=1}^{N-1} \sigma_i^2}{(N-1)} \end{aligned} \quad (18)$$

are a conservative set of parameters that satisfy the bounds.

The distribution of X_0 is assumed to have zero mean mainly to simplify notation. The essence of these results is the same irrespective of the mean.

D. Properties of the Bounds

The bounds in (9) possess several useful properties. (a) The terms $(\bar{M}, \sigma_v, \sigma_u, \mu_v, \mu_u)$ can be easily evaluated for several




Fig. 2. *Uniform source*: For i.i.d. uniform sources (over $[-0.5, 0.5]$) of different lengths, that have been uniformly quantized at the *same* rate, shown are the upper and lower bounds on SSIM index. Also shown is the true SSIM index. (a) Source size 4×4 . (b) Source size 8×8 .




Fig. 3. *Gaussian source*: For i.i.d. Gaussian sources (zero mean, unit variance) of different lengths, that have been uniformly quantized at the *same* rate, shown are the upper and lower bounds on SSIM index with $p = 0.9$. Also shown is the true SSIM index. (a) Source size 4×4 . (b) Source size 8×8 .

commonly used unbounded source types (as shown in the previous subsections). (b) The second term in the product is easy to evaluate. These two properties make the bounds tractable when compared to (8). (c) In practice, different DCT coefficients are quantized at different rates in order to optimally allocate bits. The bounds hold for any combination of rates, thereby making them attractive in practical rate allocation problems. (d) From the expression for the bounds, we see that they can be implemented efficiently and easily (even for the complex looking Laplacian case). Note that the second term in the bound involves only N summation and division operations. This property could be very useful if these bounds were to be used in real-time codec implementations. This property also allows for fast computation at any practical DCT block size. (e) These bounds can easily be extended to SSIM index's predecessor—the universal image quality index (UQI).

A point to note is that though the analysis considers a 1-D DCT, it is easy to show that the results carry over to the 2-D DCT case. The 2-D DCT obeys the Parseval's theorem, and the relation between the space domain and DCT domain means and inner products also hold. In the following section, we present several results that illustrate the useful properties of the proposed bounds.

IV. RESULTS

In this section, simulation results for a variety of examples are presented. Results for uniform, Gaussian, and Laplacian are presented first, followed by results for natural images. Finally, to illustrate the usefulness of these bounds in a practical image coding scenario, a bit-allocation problem is solved using the proposed bounds.

A. Uniform, Gaussian, and Laplacian Sources

The results are classified into two parts: equal rate and unequal rate allocation. As the names suggest, in the equal rate case, all the elements of the random vector are quantized at the same rate, while in the unequal rate case, different rates are assigned to different elements of the random vector.

1) *Equal Rate*: The first set of results are presented in Figs. 2–4. These results illustrate several properties of the bounds. The true SSIM index lies within the bounds over a range of rates for all three source types, over a range of commonly used DCT block sizes. In these results, zero mean i.i.d source have been used with the uniform source in the range $[-0.5, 0.5]$, and the Gaussian and Laplacian sources both having unit variance. The source sizes correspond to random




Fig. 4. *Laplacian source*: For i.i.d. Laplacian sources (zero mean, unit variance) of different lengths, that have been uniformly quantized at the *same* rate, shown are the upper and lower bounds on SSIM index with $p = 0.9$. Also shown is the true SSIM index. (a) Source size 4×4 . (b) Source size 8×8 .

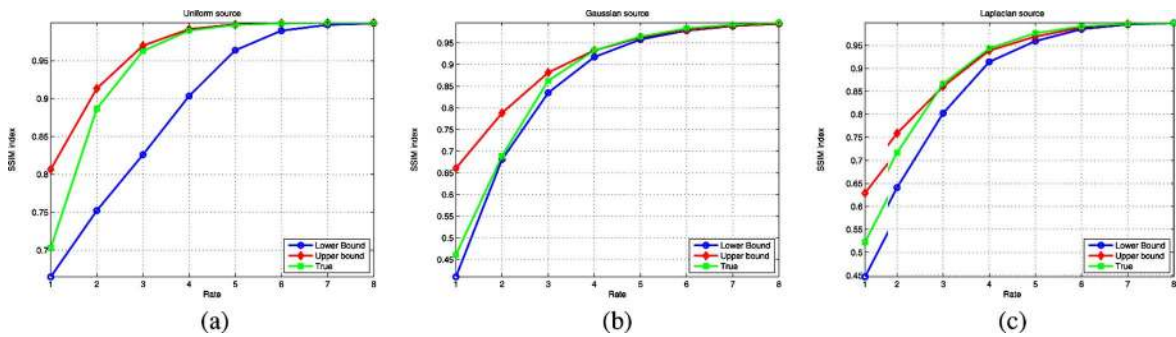


Fig. 5. *General variance example*: Upper and lower bounds on SSIM index for sources whose variance is different from unity, along with true SSIM Indices. (a) 64 independent uniform random variables (RVs) that have been divided into four groups. The first group is uniformly distributed over $[-2, 2]$, the second group over $[-1.5, 1.5]$, the third group over $[-1, 1]$, and the last group over $[-0.5, 0.5]$. (b) 64 zero mean independent Gaussian RVs divided into four groups, with the first group having standard deviation of 4, the second group having a standard deviation of 3, the third having a standard deviation of 2, and the last group having a standard deviation of 1. (c) 64 zero mean independent Laplacian RVs also divided into four groups, and assigned variances identical to the Gaussian case. $p = 0.9$. (a) Uniform source. (b) Gaussian source. (c) Laplacian source.

vectors of size 4, 16, and 64, respectively. It is important to note that these results are general in the sense that they can be used for comparison with any normalized data set.

The bounds presented in the previous section hold for any source variance. Fig. 5 shows the results for the case of equal rate allocation to the three sources with zero mean, and variance different from one. In this case, the source vector contains 64 zero mean independent random variables (RVs). For simplicity, the component RVs are divided into four groups, and each group is i.i.d. For the uniform source, the first group is distributed over $[-2, 2]$, the second group over $[-1.5, 1.5]$, the third group over $[-1, 1]$, and the last group over $[-0.5, 0.5]$. For the Gaussian source, the first group has a standard deviation of 4, the second group has a standard deviation of 3, the third has a standard deviation of 2, and the last group has a standard deviation of 1. Finally, the four groups in the Laplacian case have variances identical to the Gaussian case. From the figure, we see that the true SSIM index lies within these bounds over a range of rates. These plots point to the fact that the proposed bounds could be used as-is in a practical image coding scenario where there is no restriction on the variance of the DCT coefficients. This example is still not general enough since we assume that equal rate is assigned to all the RVs. We present more general results in the following.

2) *Unequal Rate*: We now present results for the more practically relevant case of unequal rates being assigned to the ele-

ments of the random vectors. As before, both unit-variance, and a general variance source set is considered.

The results for the unit-variance (unit step for the uniform source) are shown in Fig. 6. In this example, the source vector is composed of 64 zero mean i.i.d RVs. For the uniform source, the RVs are distributed over $[-0.5, 0.5]$. The Gaussian and Laplacian sources have zero mean and unit variance. The 64 RVs are divided into four groups, and each group is assigned the same rate. The division is motivated by the practice of grouping DCT coefficients based on their perceptual importance [43]. The combinations of rates considered are (8, 6, 4, 2), (5, 5, 3, 3), (4, 3, 2, 1), and (3, 3, 1, 1). These rates were chosen as they are a representative set of practical rate combinations. The figure reveals that the true SSIM index lies within the bounds, and provide a good estimate of the range of the average SSIM index for all the combinations considered. This example further strengthens the case for the applicability of these bounds in a practical setting.

The results for the most general case of unequal rates being assigned to a source composed of independent RVs is shown in Fig. 7. As with the unit variance case above, the source is composed of 64 zero mean independent RVs that are divided into four groups. Each group has the same variance, and is assigned the same quantization rate. The sources considered here have the same distributions as in the equal-rate general variance example. The same rate combinations as in the unit variance case are used. The figure shows that the true SSIM index again




Fig. 6. Upper and lower bounds on the SSIM index for a set of 64 i.i.d zero-mean sources that have been divided into four groups, and each group quantized at a different rate. The rate profiles include (8, 6, 4, 2), (5, 5, 3, 3), (4, 3, 2, 1), and (3, 3, 1, 1). (a) Uniform source distributed over $[-0.5, 0.5]$. (b) Zero mean, unit variance Gaussian source. (c) Zero mean, unit variance Laplacian source. $p = 0.9$. (a) Uniform. (b) Gaussian. (c) Laplacian.

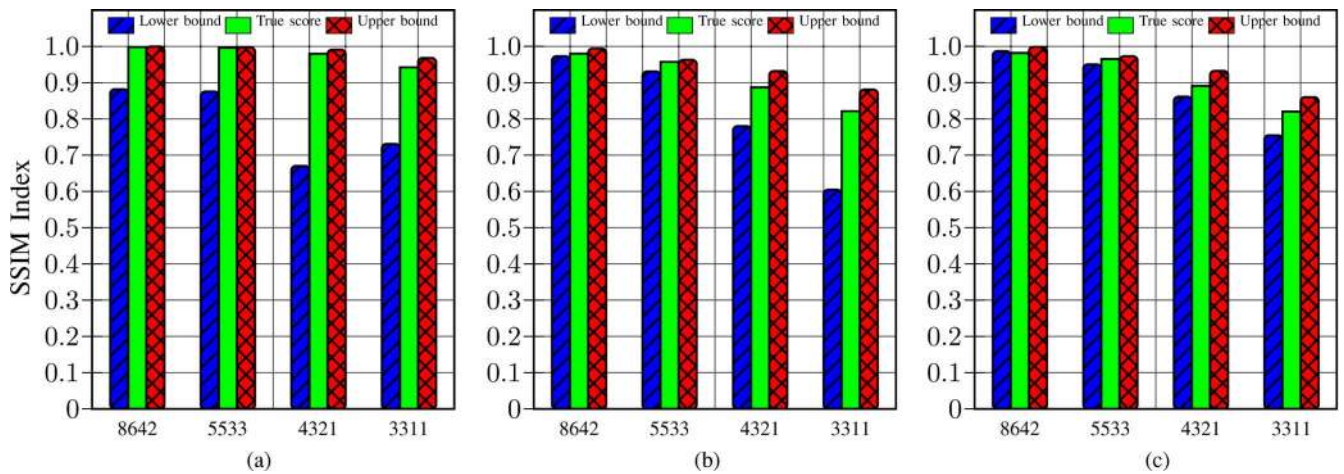


Fig. 7. Upper and lower bounds on the SSIM index for a set of 64 independent zero-mean sources that have been divided into four groups, and each group quantized at a different rate. The rate profiles include (8, 6, 4, 2), (5, 5, 3, 3), (4, 3, 2, 1), and (3, 3, 1, 1). (a) 64 independent uniform RVs that have been divided into four groups. The first group is uniformly distributed over $[-2, 2]$, the second group over $[-1.5, 1.5]$, the third group over $[-1, 1]$, and the last group over $[-0.5, 0.5]$. (b) 64 zero mean independent Gaussian RVs divided into four groups, with the first group having standard deviation of 4, the second group having a standard deviation of 3, the third having a standard deviation of 2, and the last group having a standard deviation of 1. (c) 64 zero mean independent Laplacian RVs also divided into four groups, and assigned variances identical to the Gaussian case. $p = 0.9$. (a) Uniform. (b) Gaussian. (c) Laplacian.

lies within the bounds for all the combinations considered. This example demonstrates that the bounds can indeed be used in a practical image coding scenario where there is no restriction on the variance of the source or on the rate that is assigned to each DCT coefficient. We highlight this with examples of quantization of DCT coefficient of natural images in the following.

All the simulation results presented so far have the following experimental setup. For each distribution, a realization consisting of 100 000 samples is used. Each point in the plot is the average of ten iterations. The value L for the granular region for the Gaussian and Laplacian source have been chosen based on the optimal values presented in Hui and Neuhoff [17].

B. Natural Images

In case of natural images, the DCT coefficients are quantized at various rates, and the SSIM index between the reference and quantized images is computed. Since DCT coefficients of image patches are well approximated either by Gaussian or Lapla-

cian probability density functions (pdf's), the bounds for these sources are compared with the true SSIM Indices. The mean and variance of the DCT coefficients are used to estimate bounds for Gaussian and Laplacian sources. In all our examples, an 8×8 DCT block size is used. As with the distributions above, we consider two sets of results for natural images as well, one where equal rate is assigned to all the coefficients, and the other where there is no such restriction.

1) *Equal Rate*: Two examples are presented, one where the DCT coefficients of the image are normalized (subtract mean and divide by variance), and the other where no normalization is performed. In both cases, like frequency DCT coefficients are grouped together, and quantized at the same rate.

The bounds for the zero-mean unit variance case for Gaussian and Laplacian sources are general in the sense they could be compared with the quantization results for any normalized data set (and saved as a lookup table). The results for equal rate allocation to normalized DCT coefficients is shown in Fig. 8. It pro-

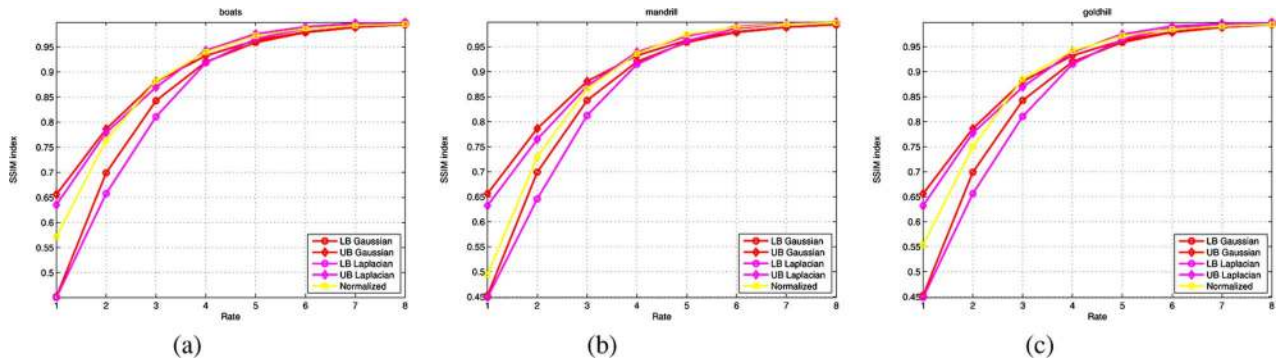


Fig. 8. Equal rate normalized coding. DCT coefficients are first normalized, and then quantized at the rates shown. The yellow plot (squares) shows the SSIM index between the reference and quantized normalized coefficients. Note that the true SSIM Indices fall nicely within the the bounds for all three images. $p = 0.9$. (a) Boats. (b) Mandrill. (c) Goldhill.

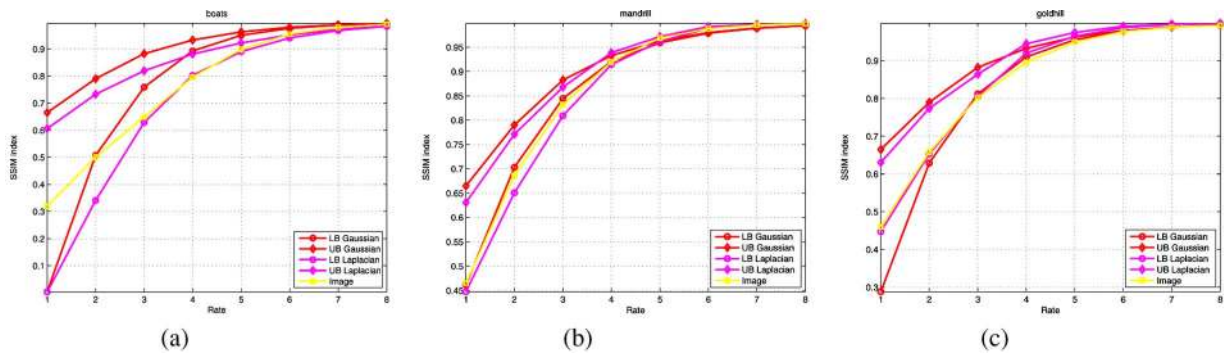


Fig. 9. Equal rate general coding. The DCT coefficients are quantized at the rates shown. The yellow plot (squares) shows the SSIM index between the reference and quantized *images*. $p = 0.9$. (a) Boats. (b) Mandrill. (c) Goldhill.

vides initial validation that the true SSIM index of normalized DCT coefficients of natural images also lies within the bounds. It also shows that the Gaussian and Laplacian models for DCT coefficients are indeed a good fit. The SSIM index values reported in these plots are measured between the normalized DCT coefficients, and their quantized versions.

In a practical setting, DCT coefficients are not usually normalized. The proposed bounds are general, and do not impose any normalization restrictions. In our next example, the DCT coefficients are quantized after their mean is removed (for simplicity). The variance of the DCT coefficients (grouped according to their frequency) is used to determine the granular region of the quantizer, based on the results from Hui and Neuhoff [17]. The results for this case are shown in Fig. 9. The SSIM Indices for the images are true values measured between the reference and quantized images. These plots reiterate that the bounds provide a good estimate of the range of the expected SSIM index for natural images. Observe that the Laplacian lower bound and Gaussian upper bound form good lower and upper bounds for natural images. Note that this example is still restricted by the equal rate assumption. We present more general results in the following.

2) *Unequal Rate*: We now present results for unequal rate quantization of DCT coefficients of natural images. The results are presented for both normalized and regular DCT coefficients. To present the results succinctly, yet retain the flavor of a practical setting, the following experimental setup is used. The 64 DCT coefficients are lexicographically ordered, and divided into

four groups. The first group contains the DC coefficient, and is given the highest precedence in rate allocation. The remaining three groups are given progressively decreasing importance. All the coefficients in a group are assigned the same rate. Two cases are considered—one at high rate, and the other at low rate.

The results are presented first for the normalized case. As before, the normalized case serves as a good first test to validate the proposed bounds. Fig. 10 presents the results for the high rate case where the four DCT coefficient groups are assigned 8, 6, 4, and 2 bits/coefficient respectively, according to their importance. The results for the low rate case where the four DCT coefficients are assigned 3, 2, 1, and 1 bit/coefficient respectively are presented in Fig. 11. From these results, it is seen that the true SSIM Indices lie well within the Gaussian and Laplacian bounds for all the images. One main reason for the accuracy of these bounds is the goodness of Gaussian and Laplacian fits to the DCT coefficients. The constants that go into the bounds have been chosen empirically so that the bounds are as tight as possible for as many images as possible. This is discussed in detail in the following subsections.

Finally, we present the most general example. In this example, the DCT coefficients are quantized (after mean subtraction), at two rates, as before. The variance of the DCT coefficients is used to determine the extent of the granular regions of the quantizer, and to form the bounds. The results for the high rate case is shown in Fig. 12, and for the low rate case in Fig. 13. From the figures, we can conclude that the bounds are indeed useful even in the most practical case. The true SSIM Indices

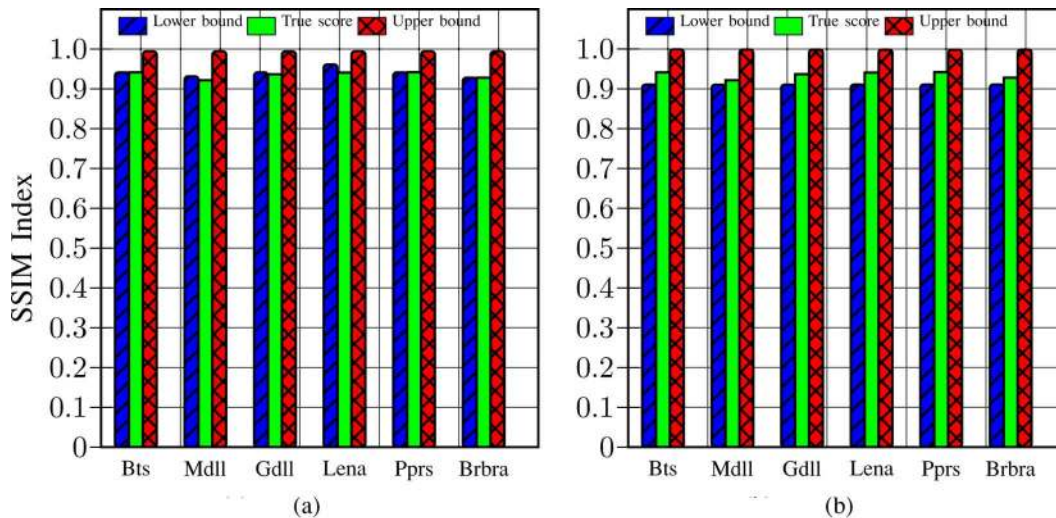


Fig. 10. Upper and lower bounds on the SSIM index of quantized normalized DCT coefficients of six natural images—Boats, Mandrill, Goldhill, Lena, Peppers, and Barbara. The DCT coefficients have been lexicographically ordered, and divided into four groups. The first group is assigned 8 bits/coefficient, the next group 6 bits/coefficient, the third group is allocated 4 bits/coefficient, and the last group is assigned 2 bits/coefficient. (a) Upper and lower bounds for a zero mean unit variance i.i.d Gaussian source consisting of 64 components. (b) Upper and lower bounds for a zero mean unit variance i.i.d Laplacian source consisting of 64 components. The bounds vary across images since the constants C_1, C_2 are image dependent. $p = 0.9$. (a) Gaussian. (b) Laplacian.

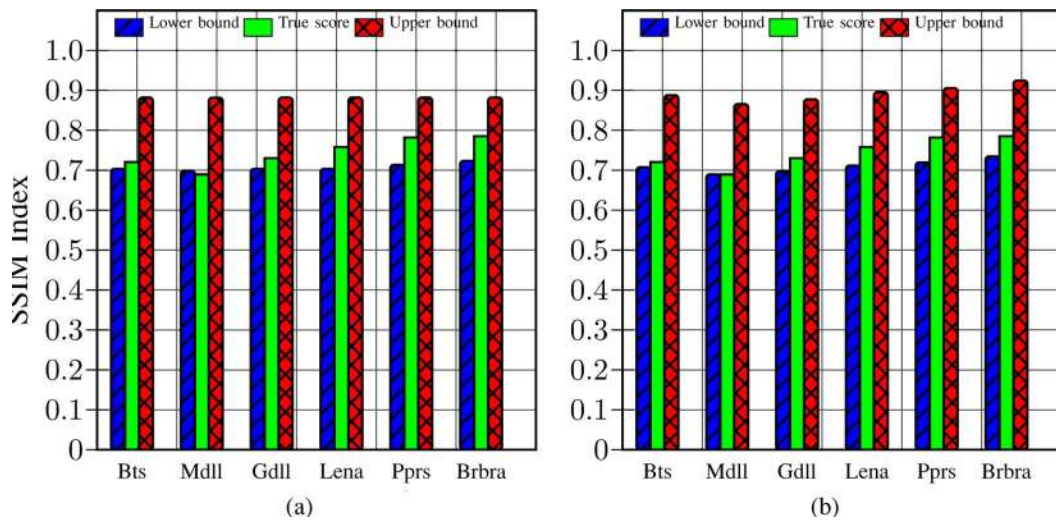


Fig. 11. Upper and lower bounds on the SSIM index of quantized normalized DCT coefficients of six natural images—Boats, Mandrill, Goldhill, Lena, Peppers, and Barbara. The DCT coefficients have been lexicographically ordered, and divided into four groups. The first group is assigned 3 bits/coefficient, the next group 2 bits/coefficient, the third group is allocated 1 bit/coefficient, and the last group is assigned 1 bit/coefficient. (a) Upper and lower bounds for a zero mean unit variance i.i.d Gaussian source consisting of 64 components. (b) Upper and lower bounds for a zero mean unit variance i.i.d Laplacian source consisting of 64 components. The bounds vary across images since the constants C_1, C_2 are image dependent. $p = 0.9$. (a) Gaussian. (b) Laplacian.

lie well within the bounds for both the high and low rate cases. The bounds are tighter for the high rate case as compared to the low rate case, which can be attributed to the underlying high resolution result. As with the normalized case, our aim is to use U and V in the bounds that are as general as possible, across rates, and images. This aspect is discussed further in the following.

C. Bit-Allocation Example

So far, examples that highlight several useful properties of the bounds have been presented. Now, an example that illustrates the practical applicability of the bounds is presented. Consider

the following rate allocation problem, and the associated constraints. Suppose that a bit budget of 128 bits is to be allocated to the 64 DCT coefficients. To make this problem tractable, the following constraints are introduced. The DCT coefficients are divided into four groups, each containing 16 coefficients. Further, the first group is assumed to contain the most important coefficients, the next group to contain the next most important coefficients, and so on. Finally, the more important group is always assigned bits greater than or equal to the number of bits assigned to the group immediately lower in importance. Though this setup is simple, it is a fair reflection of a true coding scenario. Under these assumptions, four combinations are possible—(5, 1, 1, 1),

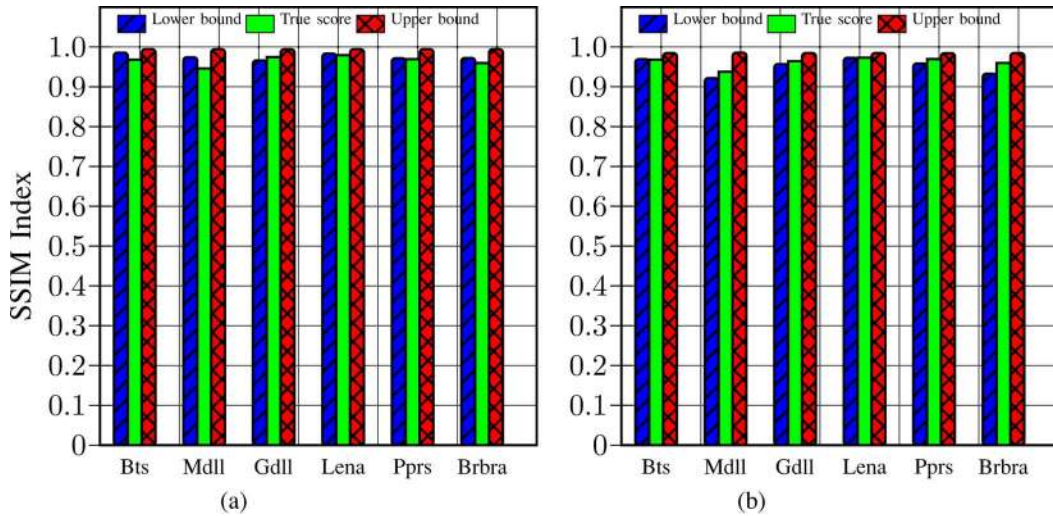


Fig. 12. Upper and lower bounds on the SSIM index of six quantized natural images—Boats, Mandrill, Goldhill, Lena, Peppers, and Barbara. The DCT coefficients have been lexicographically ordered, and divided into four groups. The first group is assigned 8 bits/coefficient, the next group 6 bits/coefficient, the third group is allocated 4 bits/coefficient, and the last group is assigned 2 bits/coefficient. (a) Upper and lower bounds for a zero mean Gaussian source consisting of 64 independent components. The variance of the components is equal to the variance of the 64 DCT coefficients, respectively. (b) Upper and lower bounds for a zero mean Laplacian source consisting of 64 independent components. Again, the variance of the components is equal to the variance of the 64 DCT coefficients, respectively. Also shown is the true SSIM index. $p = 0.9$. (a) Gaussian. (b) Laplacian.

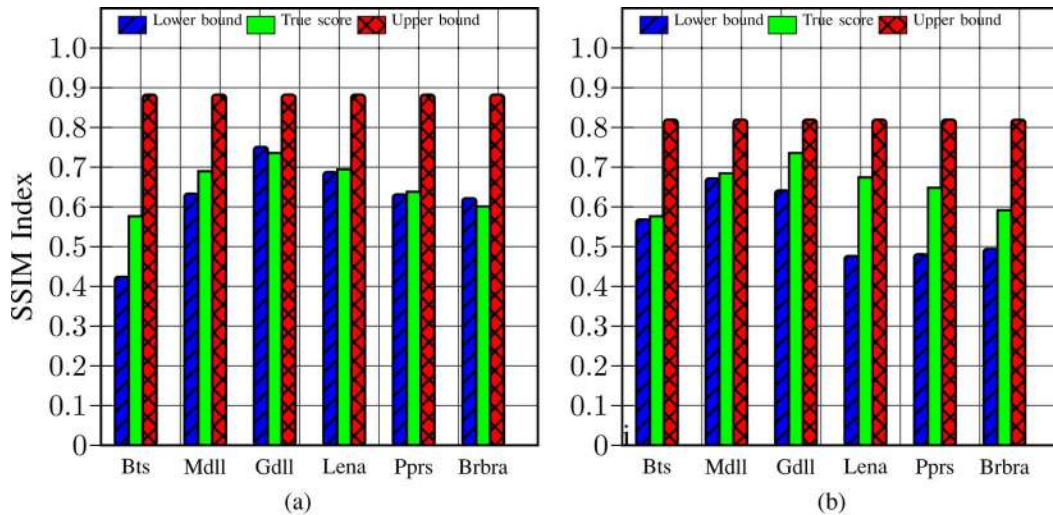


Fig. 13. Upper and lower bounds on the SSIM index of six quantized natural images—Boats, Mandrill, Goldhill, Lena, Peppers, and Barbara. The DCT coefficients have been lexicographically ordered, and divided into four groups. The first group is assigned 3 bits/coefficient, the next group 2 bits/coefficient, the third group is allocated 1 bit/coefficient, and the last group is assigned 1 bit/coefficient. (a) Upper and lower bounds for a zero mean Gaussian source consisting of 64 independent components. The variance of the components is equal to the variance of the 64 DCT coefficients, respectively. (b) Upper and lower bounds for a zero mean Laplacian source consisting of 64 independent components. Again, the variance of the components is equal to the variance of the 64 DCT coefficients, respectively. Also shown is the true SSIM index. $p = 0.9$. (a) Gaussian. (b) Laplacian.

(4, 2, 1, 1), (3, 3, 1, 1), and (2, 2, 2, 2). The problem is to find the rate combination that results in the highest SSIM index of the quantized image.

The proposed bounds give a range over which the average SSIM index can be expected to lie. This however, is not directly useful in a bit-allocation problem, where a single score for the expected SSIM index is desired. In this example, the average of Laplacian upper and lower bounds is used as a coarse estimate of the expected SSIM index. The decision rule is to choose the combination of rates that gives the highest estimate of the SSIM index from the Laplacian bounds. This decision rule is empirical, based on the results for over 50 training images from the

‘Austin and its Vicinity’ database. The results are reported for six test images (popularly used in the literature) in Table I.

From the table we see that the correct combination is chosen (based on true SSIM Indices), five of the six times. Even the erroneous choice is not very expensive in terms of the reduction of the SSIM index. The most important test for the bit-allocation technique is the visual quality check. The main motivation to derive these bounds is the fact that the SSIM index is a powerful IQA algorithm. It means that high SSIM Indices correspond to high visual quality of the images in question (and low scores correspond to poor visual quality). The bit-allocation results for the Boats image is shown in Fig. 14. We see that the image

TABLE I
ESTIMATE OF THE SSIM INDEX OF SIX QUANTIZED NATURAL IMAGES BASED ON THE LAPLACIAN BOUNDS. THE DCT COEFFICIENTS OF THESE IMAGES HAVE BEEN DIVIDED INTO FOUR GROUPS, AND QUANTIZED AT DIFFERENT RATE COMBINATIONS [(5, 1, 1, 1), (4, 2, 1, 1), (3, 3, 1, 1), AND (2, 2, 2, 2)] SUCH THAT A BIT BUDGET OF 128 BITS IS SATISFIED. THE COMBINATION THAT GIVES THE HIGHEST SSIM INDEX ESTIMATE IS CHOSEN TO BE SSIM-OPTIMAL AT THE GIVEN BIT BUDGET $p = 0.9$

Image	Bit allocation	Laplacian estimate	True score	MSE	Choice
Boats	5, 1, 1, 1	0.786540	0.774372	130.28	✓
	4, 2, 1, 1	0.767669	0.755105	191.62	×
	3, 3, 1, 1	0.670477	0.668960	301.39	×
	2, 2, 2, 2	0.444552	0.533513	796.13	×
Mandrill	5, 1, 1, 1	0.825605	0.688594	401.07	✓
	4, 2, 1, 1	0.806932	0.728961	367.72	×
	3, 3, 1, 1	0.755143	0.726923	398.63	×
	2, 2, 2, 2	0.663860	0.662344	614.71	×
Goldhill	5, 1, 1, 1	0.792467	0.806561	107.24	✓
	4, 2, 1, 1	0.776348	0.795592	138.31	×
	3, 3, 1, 1	0.702920	0.734362	216.86	×
	2, 2, 2, 2	0.544282	0.609219	1086.51	×
Lena	5, 1, 1, 1	0.851016	0.821189	109.27	✓
	4, 2, 1, 1	0.816963	0.792384	156.26	×
	3, 3, 1, 1	0.730336	0.709415	252.00	×
	2, 2, 2, 2	0.565082	0.588451	401.44	×
Peppers	5, 1, 1, 1	0.845451	0.783734	152.64	✓
	4, 2, 1, 1	0.811786	0.732363	226.99	×
	3, 3, 1, 1	0.727653	0.651803	347.59	×
	2, 2, 2, 2	0.571540	0.536406	572.70	×
Barbara	5, 1, 1, 1	0.645786	0.664828	287.84	×
	4, 2, 1, 1	0.685298	0.678023	323.80	✓
	3, 3, 1, 1	0.632041	0.639972	411.19	×
	2, 2, 2, 2	0.411487	0.516585	623.97	×

quantized using the rate combination (5, 1, 1, 1) in Fig. 14(b) has the highest visual quality and SSIM index. The fact that our bit-allocation technique does indeed pick this image provides the strongest validation of the usefulness of the proposed bounds. Note that the true score is reported only to validate the choice based on the estimates. Through this example, we demonstrate that the proposed bounds could be used in a practical image coding scenario.

D. Discussion

The results have shown that the proposed bounds work well over a range of sources. It must be noted that the bounds were not accurate for a very small percentage of examples, especially when working with natural images. The bounds are directly impacted by the choice of the numbers U and V . It is important that these are as general as possible, while remaining as tight as possible. A straightforward choice for U is 0, and for V is $2(2L)^2$ (for the case of bounded sources), where $2L$ is the extent of the granular region of the quantizer. This choice was found to give loose bounds, especially the lower bound. For the case of uniform sources, it is easy to see that $U = (\min(X^2) + \min(Q(X)^2))$, $V = (\max(X^2) + \max(Q(X)^2))$ not only make the bounds hold with $p = 1$, but also give good results over a large set of examples. The minimum and maximum values are determined from the data that is actually quantized,

and the quantization levels of the quantizers used. The maxima and minima are global over the set of random variable realizations or DCT coefficients.

For the case of unbounded inputs (Gaussian and Laplacian sources), it was found that choosing $p = 0.9$ in (10) gives bounds that are tight over a large set of images. Further, the expressions in (10) are conservative in the sense that letting $p = 0.9$ actually gives bounds that hold with a probability close to 1.

The results have demonstrated several useful properties of the bounds. The key to arriving at the bounds is the relation between the SSIM index and the MSE (see Appendix for details). Once the relation between the SSIM index and MSE is established, the well-known result for fixed rate uniform quantization of a general random variable ($MSE(\Delta) \approx \Delta^2/12$) is used to formulate the bounds. This MSE result is very strong over the range of quantization rates generally used in image and video coding applications. This is one of the main reasons for the strength of the proposed bounds. It also explains the relatively poorer performance of the bounds at lower rates. Further, the expression for \bar{M} is accurate for all the three distributions considered in this analysis. Finally, since the distribution of DCT coefficients of natural images is well approximated by either a Gaussian or Laplacian pdf, the bounds for Gaussian and Laplacian sources are well-behaved even for natural images. One other useful property that the proposed bounds possess is their



Fig. 14. *Rate allocation example.* (a) The original Boats image. (b) Original quantized using the 5, 1, 1, 1 rate profile. SSIM index = 0.7743, Laplacian estimate of SSIM index = 0.7865, MSE = 130.28. (c) Original quantized using the 4, 2, 1, 1, rate profile. SSIM index = 0.7551, Laplacian estimate of SSIM index = 0.7676, MSE = 191.62. (d) Original quantized using the 3, 3, 1, 1 rate profile. SSIM index = 0.6689, Laplacian estimate of SSIM index = 0.6704, MSE = 301.39. $p = 0.9$. (a) Original. (b) 5111. (c) 4211. (d) 3311.

ease of implementation, even for large DCT block sizes [as compared to (8)]. This is again a byproduct of the underlying MSE result.

V. CONCLUSION AND FUTURE WORK

In this paper, we presented bounds on the SSIM index as a function of quantization rate for fixed-rate uniform quantization. The proposed bounds make use of a well-known relation between MSE and quantization rate for fixed rate uniform quantization under the high rate assumption. We have demonstrated the strength of the proposed bounds using a wide variety of examples, and their usefulness in a practical scenario. Through

these results, we have taken a step in the path of designing perceptually optimal image coding algorithms, and more generally in designing perceptually optimal image processing algorithms that leverage the strength of the emerging IQA algorithms.

Several improvements to the proposed bounds could be made for future research. It is well known that one of the main reasons for the compression efficiency of transform based coding is the unequal perceptual importance of the transform coefficients (DCT or wavelet). Our results currently do not incorporate this feature of DCT coefficients. The SSIM index would have to be modified to do so, and we hope to address this soon. Similar analysis for the wavelet transform would also be very useful. We believe that such an analysis would be similar in

flavor to the current analysis. An extension to video would involve applying the current results on a frame-by-frame basis. Finally, the bounds could be improved by a more careful choice of constants.

APPENDIX

In the following, we present the Proof of Theorem 3.1 and the derivation of \bar{M} for uniform, Gaussian, and Laplacian sources.

Let $\mathbf{X} = [X_0, X_1, \dots, X_{N-1}]$ and $Q(\mathbf{X}) = [Y_{k_0}, Y_{k_1}, \dots, Y_{k_{N-1}}]$ represent a vector and its quantized version in the DCT domain, respectively. The subscripts k_0, k_1, \dots, k_{N-1} index the appropriate code points in the joint granular region of the quantizer denoted by $S_{\mathbf{k}}$. We assume that X_0 is the DC coefficient. Further, let \mathbf{x} and $\hat{\mathbf{x}}$ represent the space domain versions of \mathbf{X} and $Q(\mathbf{X})$ respectively.

Proof: Let vector $Q(\mathbf{X}) = [Y_{k_0}, Y_{k_1}, \dots, Y_{k_{N-1}}]$ be a set of quantization levels corresponding to $\mathbf{X} = [X_0, X_1, \dots, X_{N-1}]$

SSIM($\mathbf{X}, Q(\mathbf{X})$) = MS , with

$$M = \left(\frac{2 \frac{X_0 Y_{k_0}}{N} + C_1}{\frac{X_0^2 + Y_{k_0}^2}{N} + C_1} \right)$$

$$S = \left(\frac{2 \frac{\sum_{i=1}^{N-1} X_i Y_{k_i}}{N-1} + C_2}{\frac{\sum_{i=1}^{N-1} X_i^2 + Y_{k_i}^2}{N-1} + C_2} \right) \quad (19)$$

where M corresponds to the mean term, and S corresponds to the structure term. It is easy to show that $M > 0$ (for natural images), and $-1 \leq S \leq 1$. Now

$$1 - S = 1 - \frac{2 \frac{\sum_{i=1}^{N-1} X_i Y_{k_i}}{N-1} + C_2}{\frac{\sum_{i=1}^{N-1} X_i^2 + Y_{k_i}^2}{N-1} + C_2}$$

$$= \frac{\sum_{i=1}^{N-1} (X_i - Y_{k_i})^2}{\frac{\sum_{i=1}^{N-1} X_i^2 + Y_{k_i}^2}{N-1} + C_2}. \quad (20)$$

To simplify the denominator in the above equation, a variable $Z = (\sum_{i=1}^{N-1} X_i^2)/(N-1)$ is introduced. Since N is large for the most interesting case of an 8×8 DCT, the distribution of Z can be approximated well by a Gaussian distribution due to the central limit theorem. For a specified probability p , choose V and U such that $P(Z \leq V - L^2) = p$, $P(Z > U) = p$. The term L^2 is subtracted from V since it is larger than the highest value that any of quantizer levels Y_{k_i} can take. Since Z is Gaussian, it follows that the expression for V, U that make the bounds hold with probability p is given by (10), where the first and second moments of Z are computed based on the source distribution. Since $M > 0$, the following bound holds with probability p

$$M \frac{\sum_{i=1}^{N-1} (X_i - Y_{k_i})^2}{V + C_2} \leq M(1 - S)$$

$$\leq M \frac{\sum_{i=1}^{N-1} (X_i - Y_{k_i})^2}{U + C_2}. \quad (21)$$

Applying the expectation operator, using the high-rate uniform quantization result $MSE = (\Delta^2)/(12)$ [1], [15], and the independence assumption

$$\mathbb{E}[M] \left[\frac{\sum_{i=1}^{N-1} \Delta_i^2/12}{V + C_2} \right] \leq \mathbb{E}[M] - \mathbb{E}[MS]$$

$$\leq \mathbb{E}[M] \left[\frac{\sum_{i=1}^{N-1} \Delta_i^2/12}{U + C_2} \right]. \quad (22)$$

Replacing $\mathbb{E}[M]$ with \bar{M} and rearranging terms

$$\bar{M} \left(1 - \frac{\sum_{i=1}^{N-1} \Delta_i^2/12}{U + C_2} \right) \leq \mathbb{E}[\text{SSIM}(\mathbf{X}, Q(\mathbf{X}))]$$

$$\leq \bar{M} \left(1 - \frac{\sum_{i=1}^{N-1} \Delta_i^2/12}{V + C_2} \right). \quad (23)$$

□

The derivation of the term \bar{M} for different source types is given below. For simplicity, we assume that X_0 is zero-mean. The results hold irrespective of this assumption.

A. Uniform Source

The source pdf is given by

$$f_{X_0}(x_0) = \begin{cases} \frac{1}{A_0}, & \text{if } -A_0/2 \leq x_0 < A_0/2 \\ 0, & \text{elsewhere} \end{cases}$$

and the expression

$$\bar{M} = \sum_{k_0=0}^{2^{R_0}-1} \int_{L_{k_0}}^{U_{k_0}} \left(\frac{2 \frac{x_0 Y_{k_0}}{N} + C_1}{\frac{x_0^2 + Y_{k_0}^2}{N} + C_1} \right) f_{X_0}(x_0) dx_0$$

$$= \sum_{k_0=0}^{2^{R_0}-1} \int_{L_{k_0}}^{U_{k_0}} \left(\frac{2 \frac{x_0 Y_{k_0}}{N}}{\frac{x_0^2 + Y_{k_0}^2}{N} + C_1} \right) \frac{1}{A_0} dx_0$$

$$+ \int_{L_{k_0}}^{U_{k_0}} \left(\frac{C_1}{\frac{x_0^2 + Y_{k_0}^2}{N} + C_1} \right) \frac{1}{A_0} dx_0$$

(substituting $t = x_0^2 + Y_{k_0}^2$ in the first integral and using the standard result [14] in the second)

$$= \frac{1}{A_0} \sum_{k_0=0}^{2^{R_0}-1} Y_{k_0} \log \frac{U_{k_0}^2 + Y_{k_0}^2 + NC_1}{L_{k_0}^2 + Y_{k_0}^2 + NC_1}$$

$$+ \frac{NC_1}{\sqrt{Y_{k_0}^2 + NC_1}} \left[\tan^{-1} \frac{U_{k_0}}{\sqrt{Y_{k_0}^2 + NC_1}} \right.$$

$$\left. - \tan^{-1} \frac{L_{k_0}}{\sqrt{Y_{k_0}^2 + NC_1}} \right]. \quad (24)$$

B. Gaussian Source

For this case, we present the expression for the contribution to \bar{M} only from the first term in the numerator of M . The contribution from C_1 in the numerator cannot be evaluated in closed form. Note that the presence of C_1 in the denominator helps M remain stable. Here, $f_{X_0}(x_0) = (1)/(\sqrt{2\pi}\sigma_0)e^{-(x_0^2)/(2\sigma_0^2)}$, and

$$\begin{aligned} \bar{M} &\approx \sum_{k_0=0}^{2^{R_0}-1} \int_{L_{k_0}}^{U_{k_0}} \left(\frac{2 \frac{x_0 Y_{k_0}}{N}}{x_0^2 + Y_{k_0}^2 + NC_1} \right) \frac{1}{\sqrt{2\pi}\sigma_0} e^{-\frac{x_0^2}{2\sigma_0^2}} dx_0 \\ &\quad \left(\text{substituting } t = \frac{x_0^2 + Y_{k_0}^2 + NC_1}{2\sigma_0^2} \right. \\ &\quad \left. \text{and simplifying} \right) \\ &= \sum_{k_0=0}^{2^{R_0}-1} Y_{k_0} e^{\frac{Y_{k_0}^2}{2\sigma_0^2}} \int_{\frac{L_{k_0}^2 + Y_{k_0}^2}{2\sigma_0^2}}^{\frac{U_{k_0}^2 + Y_{k_0}^2}{2\sigma_0^2}} \frac{e^{-t}}{t} dt \\ &= \sum_{k_0=0}^{2^{R_0}-1} Y_{k_0} e^{\frac{Y_{k_0}^2 + NC_1}{2\sigma_0^2}} \left[E_1 \left(\frac{L_{k_0}^2 + Y_{k_0}^2 + NC_1}{2\sigma_0^2} \right) \right. \\ &\quad \left. - E_1 \left(\frac{U_{k_0}^2 + Y_{k_0}^2 + NC_1}{2\sigma_0^2} \right) \right] \end{aligned} \quad (25)$$

where $E_1(x) = \int_x^\infty (e^{-t})/t dt$ ($x > 0$), is the exponential integral.

C. Laplacian Source

Here, $f_{X_0}(x_0) = (1)/(\sqrt{2}\sigma_0)e^{-(\sqrt{2}|x_0|)/(\sigma_0)}$. We present the steps for the case where $L_{k_0} > 0$, and $U_{k_0} > 0$, and assume that there are N_1 intervals that satisfy this case. The other cases follow similar steps

$$\begin{aligned} T_1 &= \sum_{k_0=0}^{N_1-1} \int_{L_{k_0}}^{U_{k_0}} \left(\frac{2 \frac{x_0 Y_{k_0}}{N} + C_1}{\frac{x_0^2 + Y_{k_0}^2}{N} + C_1} \right) \frac{1}{\sqrt{2}\sigma_0} e^{-\frac{\sqrt{2}|x_0|}{\sigma_0}} dx_0 \\ &= \sum_{k_0=0}^{N_1-1} \int_{L_{k_0}}^{U_{k_0}} \left(\frac{2 \frac{x_0 Y_{k_0}}{N}}{\frac{x_0^2 + Y_{k_0}^2}{N} + C_1} \right) \frac{1}{\sqrt{2}\sigma_0} e^{-\frac{\sqrt{2}|x_0|}{\sigma_0}} dx_0 \\ &\quad + \int_{L_{k_0}}^{U_{k_0}} \left(\frac{C_1}{\frac{x_0^2 + Y_{k_0}^2}{N} + C_1} \right) \frac{1}{\sqrt{2}\sigma_0} e^{-\frac{\sqrt{2}|x_0|}{\sigma_0}} dx_0 \end{aligned}$$

using partial fraction expansions for both terms

$$\begin{aligned} \frac{2 \frac{x_0 Y_{k_0}}{N}}{\frac{x_0^2 + Y_{k_0}^2}{N} + C_1} &= \frac{2Y_{k_0}}{2(x_0 + ib)} + \frac{2Y_{k_0}}{2(x_0 - ib)} \\ \frac{C_1}{\frac{x_0^2 + Y_{k_0}^2}{N} + C_1} &= i \frac{C_1}{2b(x_0 + ib)} - \frac{C_1}{2b(x_0 - ib)} \end{aligned}$$

with $b = \sqrt{Y_{k_0}^2 + NC_1}$, $a = \frac{\sqrt{2}}{\sigma_0}$, $i = \sqrt{-1}$

and simplifying

$$\begin{aligned} &= \frac{1}{\sqrt{2}\sigma_0} \sum_{k_0=0}^{N_1-1} Y_{k_0} \{ e^{iab} [E_1(a(L_{k_0} + ib)) \\ &\quad - E_1(a(U_{k_0} + ib))] + e^{-iab} [E_1(a(L_{k_0} - ib)) \\ &\quad - E_1(a(U_{k_0} - ib))] \} + i \frac{NC_1}{\sqrt{b}} \{ e^{iab} [E_1(a(L_{k_0} + ib)) \\ &\quad - E_1(a(U_{k_0} + ib))] - e^{-iab} [E_1(a(L_{k_0} - ib)) \\ &\quad - E_1(a(U_{k_0} - ib))] \}. \end{aligned} \quad (26)$$

ACKNOWLEDGMENT

The authors would like to thank Prof. C. Caramanis for useful discussions on the probabilistic bounds.

REFERENCES

- [1] W. R. Bennett, "Spectra of quantized signals," *Bell Syst. Tech. J.*, vol. 27, pp. 446–472, Jul. 1948.
- [2] R. W. Buccigrossi and E. P. Simoncelli, "Image compression via joint statistical characterization in the wavelet domain," *IEEE Trans. Image Process.*, vol. 8, no. 12, pp. 1688–1701, Dec. 1999.
- [3] D. M. Chandler and S. S. Hemami, "Dynamic contrast-based quantization for lossy wavelet image compression," *IEEE Trans. Image Process.*, vol. 14, no. 4, pp. 2284–2298, Apr. 2005.
- [4] D. M. Chandler and S. S. Hemami, "VSNR: A wavelet-based visual signal-to-noise ratio for natural images," *IEEE Trans. Image Process.*, vol. 16, pp. 397–410, 2007.
- [5] S. S. Channappayya, A. C. Bovik, C. Caramanis, and R. W. Heath, Jr., "Design of linear equalizers optimized for the structural similarity index," *IEEE Trans. Image Process.*, vol. 16, no. 7, pp. 857–872, Jul. 2007.
- [6] S. S. Channappayya, A. C. Bovik, and R. W. Heath, Jr., "A linear estimator optimized for the structural similarity index and its application to image denoising," *Proc. IEEE Int. Conf. Image Processing*, 2006.
- [7] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York: Wiley, 1991.
- [8] S. Daly, "The visible difference predictor: An algorithm for the assessment of image fidelity," in *Digital Images and Human Vision*, A. B. Watson, Ed. Cambridge, MA: MIT Press, 1993, pp. 179–206.
- [9] M. P. Eckert and A. P. Bradley, "Perceptual quality metrics applied to still image compression," *Signal Process.*, vol. 70, pp. 177–200, Nov. 1998.
- [10] J. D. Eggerton and M. D. Srinath, "A visually weighted quantization scheme for image bandwidth compression at low data rates," *IEEE Trans. Commun.*, vol. COM-34, no. 12, pp. 840–847, Aug. 1986.
- [11] C. Fogg, D. J. LeGall, J. L. Mitchell, and W. B. Pennebaker, *MPEG Video Compression Standard*. New York: Springer, 1996.
- [12] A. Gersho and R. M. Gray, *Vector Quantization and Signal Compression*. New York: Springer, 1991.
- [13] B. Girod, A. B. Watson, Ed., "What's Wrong with Mean-squared Error?," in *Digital Images and Human Vision*. Cambridge, MA: MIT Press, 1993, pp. 207–220.
- [14] I. S. Gradshteyn and I. M. Ryzhik, *Table of Integrals, Series, and Products*. San Diego, CA: Academic, 2000.
- [15] R. M. Gray and D. L. Neuhoff, "Quantization," *IEEE Trans. Inf. Theory*, vol. 44, no. 6, pp. 2325–2383, Oct. 1998.
- [16] I. Hontsch and L. J. Karam, "Adaptive image coding with perceptual distortion control," *IEEE Trans. Image Process.*, vol. 11, no. 3, pp. 213–222, Mar. 2002.
- [17] D. Hui and D. L. Neuhoff, "Asymptotic analysis of optimal fixed-rate uniform scalar quantization," *IEEE Trans. Inf. Theory*, vol. 47, no. 3, pp. 957–977, Mar. 2001.
- [18] A. K. Jain, *Fundamentals of Digital Image Processing*. Englewood Cliffs, NJ: Prentice-Hall, 1989.
- [19] E. Y. Lam and J. W. Goodman, "A mathematical analysis of the DCT coefficient distributions for images," *IEEE Trans. Image Process.*, vol. 9, no. 10, pp. 1661–1666, Oct. 2000.
- [20] Z. Liu, L. J. Karam, and A. B. Watson, "JPEG2000 encoding with perceptual distortion control," *IEEE Trans. Image Process.*, vol. 15, no. 7, pp. 1763–1778, Jul. 2006.

- [21] J. Lubin, A. B. Watson, Ed., "The use of psychophysical data and models in the analysis of display system performance," in *Digital Images and Human Vision*. Cambridge, MA: MIT Press, 1993, pp. 163–178.
- [22] J. Lubin, *Visual Models for Target Detection and Recognition*. Singapore: World Scientific, 1995, ch. 10, pp. 245–283.
- [23] J. L. Mannos and D. J. Sakrison, "The effects of a visual fidelity criterion on the encoding of images," *IEEE Trans. Inf. Theory*, vol. 20, no. 7, pp. 525–536, Jul. 1974.
- [24] N. B. Nill, "A visual model weighted cosine transform for image compression and quality assessment," *IEEE Trans. Commun.*, vol. COM-33, pp. 551–557, Jul. 1985.
- [25] T. N. Pappas and R. J. Safranek, A. C. Bovik, Ed., "Perceptual Criteria for Image Quality Evaluation," in *Handbook of Image and Video Processing*. New York: Academic, 2000, pp. 669–684.
- [26] W. B. Pennebaker and J. L. Mitchell, *JPEG Still Image Data Compression Standard*. New York: Van Nostrand Reinhold, 1993.
- [27] M. F. Sabir, H. R. Sheikh, R. W. Heath, Jr., and A. C. Bovik, "A joint source-channel distortion model for JPEG compressed images," *IEEE Trans. Image Process.*, vol. 15, no. 6, pp. 1349–1364, Jun. 2006.
- [28] R. Safranek and J. Johnston, "A perceptually tuned sub-band image coder with image dependent quantization and post-quantization data compression," in *Proc. Int. Conf. Acoustics, Speech, and Signal Processing*, May 23–26, 1989, vol. 3, pp. 1945–1948.
- [29] O. Schwartz and E. P. Simoncelli, "Natural signal statistics and sensory gain control," *Nature Neurosci.*, vol. 4, no. 8, pp. 819–825, Aug. 2001.
- [30] H. R. Sheikh and A. C. Bovik, "Image information and visual quality," *IEEE Trans. Image Process.*, vol. 15, no. 2, pp. 430–444, Feb. 2006.
- [31] H. R. Sheikh, A. C. Bovik, and G. de Veciana, "An information fidelity criterion for image quality assessment using natural scene statistics," *IEEE Trans. Image Process.*, vol. 14, no. 12, pp. 2117–2128, Dec. 2005.
- [32] H. R. Sheikh, M. F. Sabir, and A. C. Bovik, "A statistical evaluation of recent full reference image quality assessment algorithms," *IEEE Trans. Image Process.*, vol. 15, no. 11, pp. 3440–3451, Nov. 2006.
- [33] E. Simoncelli, W. Freeman, E. Adelson, and D. Heeger, "Shiftable multiscale transforms," *IEEE Trans. Inf. Theory*, vol. 38, no. 2, pp. 587–607, Mar. 1992.
- [34] P. C. Teo and D. J. Heeger, "Perceptual image distortion," *Proc. SPIE*, vol. 2179, pp. 127–141, 1994.
- [35] M. Vetterli and J. Kovacevic, *Wavelets and Subband Coding*. Englewood Cliffs, NJ: Prentice-Hall, 1995.
- [36] M. J. Wainwright and E. P. Simoncelli, "Scale mixtures of gaussians and the statistics of natural scenes," *Adv. Neural Inf. Process. Syst.*, vol. 12, pp. 855–861, 2000.
- [37] Z. Wang and A. C. Bovik, "A universal image quality index," *IEEE Signal Process. Lett.*, vol. 9, no. 3, pp. 81–84, Mar. 2002.
- [38] Z. Wang and A. C. Bovik, *Modern Image Quality Assessment*. San Rafael, CA: Morgan and Claypool, 2006.
- [39] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, Apr. 2004.
- [40] Z. Wang, Q. Li, and X. Shang, "Perceptual image coding based on a maximum of minimal structural similarity criterion," presented at the IEEE Int. Conf. Image Processing, Sep. 2007.
- [41] Z. Wang and E. P. Simoncelli, "Translation insensitive image similarity in complex wavelet domain," in *Proc. IEEE Int. Conf. Acoustics, Speech, and Signal Processing*, 2005, vol. 2, pp. 573–576.
- [42] Z. Wang, E. P. Simoncelli, and A. C. Bovik, "Multi-scale structural similarity for image quality assessment," in *Proc. IEEE Asilomar Conf. Signals, Systems, Computers*, 2003, vol. 2, pp. 1398–1402.
- [43] A. Watson, "Visually optimal DCT quantization matrices for individual images," in *Proc. Data Compression Conf.*, Mar. 30–Apr. 2, 1993, pp. 178–187.
- [44] A. B. Watson, *Digital Images and Human Vision*. Cambridge, MA: MIT Press, 1993.
- [45] S. Winkler, "A perceptual distortion metric for digital color video," *Proc. SPIE*, vol. 3644, pp. 175–184, 1999.



Sumohana S. Channappayya (S'01–M'08) received the B. E. degree from the University of Mysore, India, in 1998, the M.S. degree in electrical engineering from the Arizona State University, Tempe, AZ, in 2000, and the Ph.D. degree in electrical and computer engineering from The University of Texas at Austin in 2007.

He worked for PacketVideo Corporation, San Diego, CA, from 2001 to 2003. He is currently with PacketVideo Corporation. His research interests include image restoration, image and video coding,

image and video quality assessment, and multimedia communication.



Alan Conrad Bovik (S'80–M'81–SM'89–F'96) received the B.S., M.S., and Ph.D. degrees in electrical and computer engineering from the University of Illinois at Urbana-Champaign, Urbana, in 1980, 1982, and 1984, respectively.

He is currently the Curry/Cullen Trust Endowed Professor at The University of Texas at Austin, where he is the Director of the Laboratory for Image and Video Engineering (LIVE) in the Center for Perceptual Systems. His research interests include image and video processing, computational vision, digital microscopy, and modeling of biological visual perception. He has published over 450 technical articles in these areas and holds two U.S. patents. He is also the author of *The Handbook of Image and Video Processing* (Elsevier, 2005, 2nd ed.) and *Modern Image Quality Assessment* (Morgan & Claypool, 2006).

Dr. Bovik has received a number of major awards from the IEEE Signal Processing Society, including: the Education Award (2007); the Technical Achievement Award (2005), the Distinguished Lecturer Award (2000); and the Meritorious Service Award (1998). He is also a recipient of the Distinguished Alumni Award from the University of Illinois at Urbana-Champaign (2008), the IEEE Third Millennium Medal (2000), and two journal paper awards from the international Pattern Recognition Society (1988 and 1993). He is a Fellow of the Optical Society of America and a Fellow of the Society of Photo-Optical and Instrumentation Engineers. He has been involved in numerous professional society activities, including: Board of Governors, IEEE Signal Processing Society, 1996–1998; Editor-in-Chief, IEEE TRANSACTIONS ON IMAGE PROCESSING, 1996–2002; Editorial Board, PROCEEDINGS OF THE IEEE, 1998–2004; Series Editor for Image, Video, and Multimedia Processing, Morgan and Claypool Publishing Company, 2003–present; and Founding General Chairman, First IEEE International Conference on Image Processing, Austin, TX, November 1994. He is a registered Professional Engineer in the State of Texas and is a frequent consultant to legal, industrial, and academic institutions.



Robert W. Heath, Jr. (S'96–M'01–SM'06) received the B.S. and M.S. degrees from the University of Virginia, Charlottesville, in 1996 and 1997, respectively, and the Ph.D. degree from Stanford University, Stanford, CA, in 2002, all in electrical engineering.

From 1998 to 2001, he was a Senior Member of the Technical Staff then a Senior Consultant at Iospan Wireless, Inc., San Jose, CA where he worked on the design and implementation of the physical and link layers of the first commercial MIMO-OFDM communication system. In 2003, he founded MIMO Wireless, Inc., a consulting company dedicated to the advancement of MIMO technology. Since January 2002, he has been with the Department of Electrical and Computer Engineering, The University of Texas at Austin, where he is currently an Associate Professor and member of the Wireless Networking and Communications Group. His research interests include several aspects of MIMO communication: limited feedback techniques, multihop networking, multiuser MIMO, antenna design, and scheduling algorithms, as well as 60-GHz communication techniques and multimedia signal processing.

Dr. Heath has been an Editor for the IEEE TRANSACTIONS ON COMMUNICATIONS and an Associate Editor for the IEEE TRANSACTIONS ON VEHICULAR TECHNOLOGY. He is a member of the Signal Processing for Communications Technical Committee in the IEEE Signal Processing Society. He was a Technical Co-Chair for the 2007 Fall Vehicular Technology Conference, the General Chair of the 2008 Communication Theory Workshop, and a co-organizer of the 2009 Signal Processing for Wireless Communications Workshop. He is the recipient of the David and Doris Lybarger Endowed Faculty Fellowship in Engineering and is a registered Professional Engineer in Texas.