A Triage Metric for Determining the Extent of DCT-Based Compression Artifacts in a Digital Image

Qing Yu,* Jiebo Luo,† and Michael E. Miller

Eastman Kodak Company, Electronic Imaging Product, Research & Development, Imaging Science Division, Rochester, New York

A triage metric is proposed to determine the extent of blocking and contouring artifacts in an image due to JPEG compression. It is found that this metric has a linear correlation with subjective judgements of print quality. This triage metric can be used to determine the amount of needed image enhancement or achieve a good trade-off between image quality and system throughput.

Journal of Imaging Science and Technology 46: 443–452 (2002)

Introduction

Digital images contain enormous amounts of data. Storage of this type of data on digital media is generally expensive and transmission of digital images requires either a large bandwidth or a long period of time. Many algorithms have been developed to compress image data by removing visually redundant information from the image. Discrete-cosine-transform-based (DCT-based) compression has been the most popular among existing techniques. DCT has very good energy compaction and data decorrelation properties. Moreover, DCT can be computed using fast algorithms and efficiently implemented using very large scale integration (VLSI) techniques. In DCT-based coding, an image is partitioned into small square blocks (typically 8×8) and DCT is computed over these blocks to remove the local spatial correlation. To achieve high compression, quantization of the DCT coefficients is performed. Quantization is an irreversible process that causes loss of information and distortions in the decompressed image. After quantization, the redundancy in the data is further reduced using entropy coding. At the decoder end, the received data is decoded, dequantized, and reconstructed by the inverse DCT. In general, a typical 8-bit gray-level image can be coded with compression ratio up to 10:1 without noticeable artifacts. Since JPEG is the international standard for implementing DCT-based image compression, we adopted JPEG in all the experiments in this study without loss of generality.[1]

Original manuscript received November 7, 2001

E-mail: yu@vimicro.com; jiebo.luo@kodak.com;
E-mail: michael.e.miller@kodak.com

* Qing Yu is currently with Vimicro Corporation

† Corresponding author

However, at low bit rates the reconstructed images generally suffer from visually annoying artifacts because of very coarse quantization. One major such artifact is the blocking effect, which appears as artificial block boundaries between adjacent blocks. At a low bit rate, each block is represented mainly by the first few low frequency coefficients and, because each block is processed independently, no inter-block correlation is accounted for in standard block DCT-based coding schemes. Therefore, discontinuity across the block boundary becomes noticeable.

Another common artifact resulting from coarse quantization, in particular the DC coefficients by JPEG compression, is false contours in smoothly varying image regions. This is actually an extreme case of the blocking artifacts where block boundaries only appear in one direction. This artifact will be referred to as contouring artifacts in this article. Examples of this artifact are often found in images that contain a smooth sky region with a gradient of transition in the vertical direction. In such an image, false contours are often visible in the horizontal direction and appear as horizontal contours when the image is highly compressed with coarse quantization.

A straightforward method for determining the extent of the blocking artifacts involves taking the ratio between the size of the compressed JPEG file and the size of the original image, which corresponds to the amount of uncompressed image data. This ratio is commonly referred to as the compression ratio. However, the compression ratio is not necessarily a good measure of the quality of the compressed image, nor the visibility or objectionability of blocking artifacts. In general, at the same compression ratio, a busy image would look worse than a less busy image because busy images are harder to compress (therefore a busy image has lost more details than the less busy image).

Another viable method determines the extent of the blocking artifacts by extracting the quantization table from the image header based on the assumption that the aggressiveness of the quantization table may determine the level of blocking/contouring artifacts within an image. There are two potential drawbacks with this approach. First, the quantization table may not be available. For example, an image can be highly compressed, then modified and saved into another image file format that does not support the quantization table in its metadata structures. Even if the image is saved as a JPEG file, it may have already been compressed one or more times before. In this instance, the quantization table in the image header will not reflect the true compression level of the image. Second, even if the right quantization table is available, wrong conclusions may be drawn for certain types of images because this approach is not image content dependent. For example, the most aggressive quantization table may be used to compress an image with little activity, i.e., a uniform patch, and the resulting image can still be of high quality, thus the image is not degraded at all. Conversely, a picture containing significant high frequency information may be severely degraded.

To achieve better prediction of the JPEG blocking/contouring artifacts, researchers have gone beyond image-content independent approaches as discussed above, and employed different levels of analysis of image content and structure. Several metrics have been proposed based on analysis of discontinuities at JPEG block boundaries. Ho[2] developed a parameter called NNND as a mathematical tool for measuring the magnitude of block artifacts. NNND is the ratio of averaged square difference between adjacent pixels immediately next to block borders and away from block borders. Jeon and co-workers[3] also took a similar approach. Liu and co-workers[4] and Coudoux and co-workers[5] went a step further by including the psychophysical properties of the human visual system during block discontinuity evaluation to achieve predictions that correlate with human judgement. Tan and co-workers[6] studied the blocking artifacts in the frequency domain. Using this approach, the image is partitioned into smaller blocks of pixels and converted into the frequency domain using the Fourier transform. The blocking artifacts are then translated into harmonics whose amplitude and phase provide vital information for quantifying the blocking artifacts. Also working in the frequency domain, Wang and co-workers[7] modeled a JPEG image as a non-blocky image interfered with a pure blocky signal. Consequently, the task of the blocking effect measurement algorithm is to detect and evaluate the power of the blocky signal. Datcu[8] studied the histogram of a JPEG image and calculated a blocking factor corresponding to the level of spikes in the histogram that are due to quantization during image compression and amplitude clustering during image decompression.

In this article, two novel predictors are proposed to measure the blocking and contouring artifacts in an image caused by JPEG compression. An image quality metric based on these two predictors is further derived based on experiment data collected from a psychophysical evaluation of photographic prints from JPEG compressed images. We first propose a new algorithm to extract the locations of block boundary and block size in a JPEG compressed image. We further propose two predictors for quantifying JPEG blocking and contouring artifacts, respectively. Next, we report a psychophysical experiment that has been conducted to gather information on human perception of JPEG artifacts. Based on the above findings, we propose a JPEG image quality metric that correlates with the psychophysical data. Finally, the potential application for this metric is discussed.

Block Boundary Detection

Most prior approaches for evaluating JPEG artifacts assume prior knowledge of JPEG block boundary locations. However, because digital images compressed with the DCT-based technique may have been further modified, for example, through cropping and zooming, the block boundary locations may have changed and become unknown at the time of evaluation. Therefore, there is a need for an automatic technique for detecting block boundary locations in digital images having blocking artifacts.

In this section, we propose a new algorithm to extract the locations of block boundary and block size in a JPEG compressed image.

A digital image may be a gray scale image containing the intensity channel, or a color image containing RGB channels. For color images, a color transform is typically performed before image compression to take advantage of the redundancy in the color perception of the human visual system (HVS). For example, in JPEG compression, a color image is first converted from RGB to the YCbCr color space with the following equations:

$$Y = 16+65.481*R+128.553*G+24.966*B;$$

$$Cb = 128-37.797*R-74.203*G+112*B;$$

$$Cr = 128+112*R-93.786*G-18.214*B;$$

where Y is the luminance channel, and Cb and Cr are the two chrominance channels. R, G and B are all normalized to 1. Thus, an input RGB image is first converted to YCbCr space before the block boundary detection is performed. Since most of the blocking artifacts are perceived in the luminance channel (Y) due to the lower sensitivity of the human visual system to changes in chrominance channels Cb and Cr, the block boundary detection is carried out for that channel only. In fact, the lower chrominance sensitivity is the basis for more aggressive subsampling in Cb and Cr by JPEG.[1]

Referring to Fig. 1, for the Y image channel, a column difference image is first generated by calculating the absolute difference between two adjacent columns. For example, subtracting the pixel values of the second column of the image from the pixel values of the first column of the image to generate a column of difference values, and set the pixel values of the first column of the column difference image as the absolute value of the column of difference values. The same procedure is repeated to set the remaining columns of the column difference image except the last columns where all the values of that column are set to zero.

The column difference image is further averaged in the vertical direction to generate a one-dimensional column difference array, (VA). Therefore, assuming that the original image has M rows and N columns of image data, VA should have N entries.

To prevent image edges from confusing the block boundary detection, the contribution of a pixel in the column difference image is discarded if the corresponding pixel in the original Y image channel is near an edge. The classification of an edge pixel is achieved with the following steps:

1. Generate an intensity gradient image from the original Y image channel. Sobel operators are used as the intensity gradient operators with the intensity gradient equal to the sum of the absolute values from the horizontal and vertical Sobel operators;

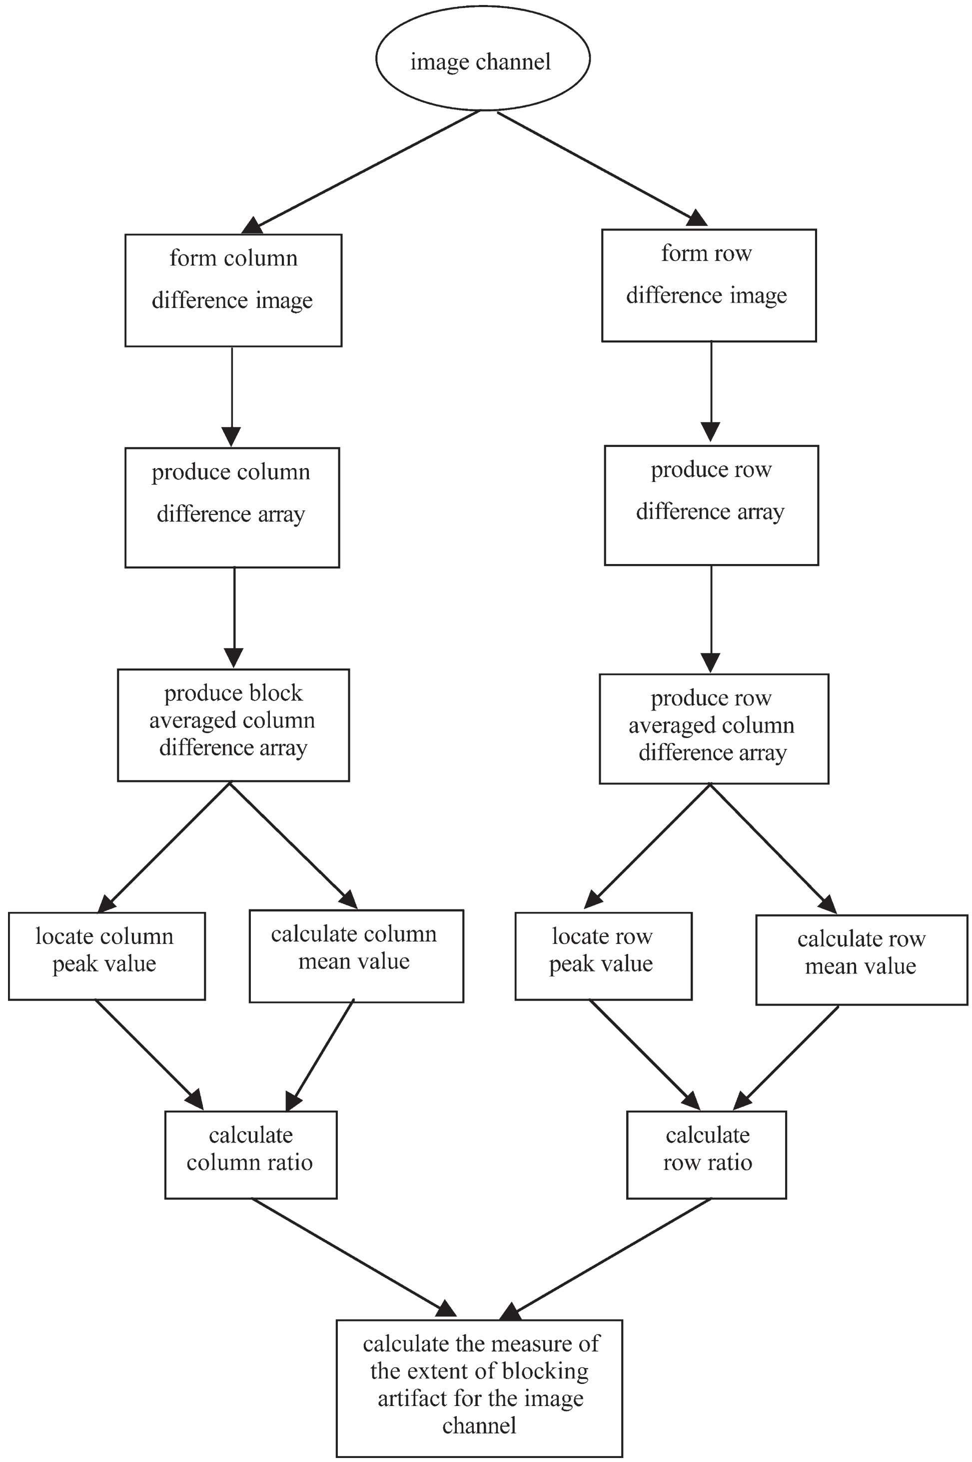

Figure 1. Flow chart for block boundary detection and column/row ratio calculation

2. Calculate the standard deviation (SD) of the intensity gradient image;
3. Set a threshold value T that equals to twice SD;
4. If the magnitude of the intensity gradient of a pixel in the original image channel is larger than T, it is flagged as an edge pixel.

The threshold selection that equals twice SD is common practice in image processing because it is well known that the distribution of the gradient magnitude can be modeled by a Laplacian distribution, and the strong edges fall in the tails of the distribution beyond twice SD.

VA is further averaged using a periodicity of the JPEG block width, i.e., 8, to generate a block averaged column difference array, VAA, which has eight entries. In other words, every eighth entry of VA will be averaged, and the result will be used to set the eight entries of VAA. Clearly, image crop without zoom can be handled properly by the proposed method. In the case where a

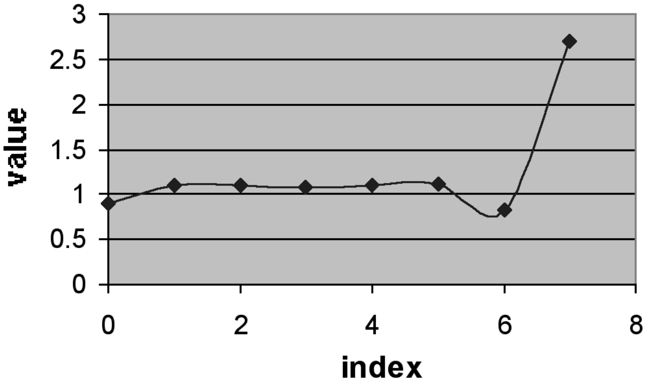

Figure 2. A typical block averaged difference

zoom operation has been applied to the image, as long as the zoom factor or the resizing factor results in an integer block size, different periodicities, i.e., 2, 4, 6, 8, 10, 12, 14, 16, can be tried and the periodicity that results in the highest peak to base ratio (see the definitions of peak and base below) is selected to be the effective block size. It would be difficult to detect the periodicity if the zoom factor results in non-integer or irregular block size. From VAA, the maximum value is first located and defined as the column peak value, then the mean value of VAA is calculated excluding the peak value. The mean value is defined as the column base or floor value.

Finally, the ratio between the peak value and the base value is calculated to generate a column ratio.

Figure 2 shows a typical block averaged difference array with 8 elements (index starts at 0) from a typical JPEG compressed image. The x-axis is the array index and the y-axis is the magnitude. In this case, the peak value (2.62) occurs at index 7, the mean value is 1.03, and the ratio is about 2.6.

Row peak value, row base value, and row ratio are also calculated analogously.

The indexes of peak value in VAA and HAA (block averaged row difference array) are the JPEG block boundary locations in the horizontal and vertical directions for the image channel. For example, if the index for the peak value in VAA is 4, it is assumed that JPEG block boundary will appear at column numbers 4, 12, …, in the image channel. Note that the indexes for the columns of the image channel as well as for VAA and HAA start at zero in this embodiment.

In the case where digital images are not only cropped but also zoomed (shrunk or magnified), the block size is no longer the original JPEG block size of 8. For potentially magnified digital images, the above mentioned process for determining the extent of the JPEG blocking artifacts can be repeated for a predetermined series of hypothetical block width values to produce a series of ratios between the peak value and the base value. The hypothetical block width that produces the maximum artifact measure is chosen to be the effective block size after zooming. The zooming factor is, in turn, determined to be the ratio between the effective block size and the original block size, i.e., 8.

A New Predictor for JPEG Blocking Artifacts

The column and row ratios are good indicators of the presence of the blocking artifacts in the column and row directions, respectively. In general, it can be argued that the peak values represent the magnitude of the discontinuity at block boundaries while the floor values are related to the busyness of the image. It is observed that, for the same scene, peak value normally increases while floor value decreases when the compression becomes more aggressive. Therefore, the column and row ratios could be potentially used as a metric to measure the extent of blocking artifacts of an entire image.

The visibility or objectionability of the blocking artifacts may depend on the actual structures in an image. For example, textured areas tend to hide the blocking artifacts better than flat areas due to spatial frequency masking within the human visual system.[9] The column and row base values are good indicators of the amount of texture. To achieve a higher degree of adaptivity to the image content, a measure of the extent of the blocking artifacts can be defined as a function of both the column/row ratios and the column/row base values. In general, the higher the base values, meaning a higher degree of texture activity (capable of hiding a larger extent of blocking artifacts), the higher threshold on the artifact measure should be used. In practice, a look-up table can be built empirically to quantitatively characterize the relationship between the base values and the threshold.

A New Predictor for JPEG Contouring Artifacts

As mentioned in the introduction session, JPEG compression can produce contouring artifacts in smoothly varying image regions and the contouring is primarily due to coarse quantization, in particular of the DC coefficients, of the image. This conjecture is supported by the fact that in smooth regions the energy of each block is concentrated at the DC coefficient. The contouring artifact is actually an extreme case of the blocking artifacts where block boundaries are primarily visible in one dimension. For example, if an image contains a smoothly varying sky region and the gradient of transition is in the vertical direction, then false contours will show up in the horizontal direction as horizontal contours when the image is highly compressed with coarse quantization. In general, the coarser the quantization, the lower the frequency of the contours, and the lower the image quality. This artifact may be difficult for the new MaxRatio metric to detect. For example, the distance between two contours in Fig. 3 is about 80 pixels, which is about 10 blocks. Thus, block boundaries actually show up only every 10 blocks, posing a problem for the peak/floor calculation where periodic averaging is carried out on a block by block basis. Therefore images containing this artifact have a relatively low MaxRatio value.

Because the contouring artifacts in an image are directly related to the degree of quantization of the DC value of each block in the image, an efficient approach to measuring the contouring artifacts in the image is to determine the DC quantization step used by JPEG compression. A new approach is, therefore, proposed to estimate the DC quantization step (S), and a modified image quality metric will be introduced that is based on both MaxRatio and S.

Referring to Fig. 4 for the luminance channel of the image, we first determine block size and positions based on the locations of peak values of the column and row difference arrays, as discussed previously. Next, the DC value of each block is calculated by taking the arithmetic mean of the pixel values within each block. Then, a histogram of these block DC values is generated. A histogram from a typical highly compressed image is shown in Fig. 5.

Figure 3. An image with a distinct contouring artifact.

Because of the quantization of the DC coefficients introduced by JPEG, the histogram of block DC values will be sparse so that it contains only values of multiples of the DC quantization step size. In other words, the histogram is periodic, as can be seen clearly in Fig. 5. To estimate the periodicity, i.e., the quantization step size, the Fourier transform of the histogram is further calculated. Fig. 6 illustrates the Fourier transform of the histogram in Fig. 5.

Next, the first non-DC peak of the histogram in the Fourier transform domain is identified. The first non-DC peak should be at least half of the DC peak in amplitude. Other non-DC peaks correspond to higher frequency harmonics (multiples of the basic periodicity). Finally, the size of the DC quantization step is calculated based on the frequency of the first non-DC peak using the following equation:

$$S = D/f_p$$

where f_p is the frequency of the first non-DC peak in the Fourier transform domain, D is the length of the histogram, i.e., 256, and S is the size of the DC quantization step, which will serve as a measure of the extent of the contouring artifacts in the digital image.

In a recent independent study, Fan and co-workers proposed a Bayesian method for estimating the entire quantization table from a compressed image.[10] Their approach involves taking the DCT of the compressed image, forming histograms of individual DCT coefficients, and performing Bayesian estimation of the quantization step sizes.

A Psychophysical Experiment

To study the effectiveness of the two predictors, a psychophysical experiment was conducted. The primary goal of this experiment was to determine the quality degradation produced by DCT-based compression algorithms and to understand how well the two predictors explained the overall quality degradation induced by the DCT-based compression algorithm. It is recognized that the predictors that are discussed do not attempt to detect unsharpness and certain other artifacts that can be created by applying DCT-based compression algorithms and therefore these predictors are not likely to account for these additional artifacts. However, our end goal was to find a method for predicting quality degradation of compressed images and therefore, additional predictors could have been considered if the results from these two predictors did not adequately predict image degradation from the compression algorithms.

A secondary goal was to determine how the quality degradation mapped to consumer print acceptability. As described later, this knowledge is important since one would not want to apply an algorithm to remove artifacts that are not detectable to a typical observer, especially when this enhancement algorithm requires significant processing time and may, in some cases, degrade the quality, e.g., sharpness and detail, of certain images.

The details of the psychophysical experiment are presented in this section.

Participants

A total of 22 individuals (7 female and 15 male) took part in this experiment. All participants were naïve to the purpose of the experiment. All of them were screened to insure they had normal color vision and near field visual acuity of at least 20/40. Participants with congenital color vision defects were screened from the study using Ichikawa standard pseudoisochromatic plates.[11] Visual acuity was verified using a standard Snellen eye chart[12] mounted at a viewing distance of 16 inches.

Stimulus

Prints were used as stimulus for this experiment. Eight scenes (four from color negative scans and four from a high-end digital camera) were selected as originals. Shown in Figs. 7 and 8, these eight images are all uncompressed images of size 1536 × 1024, and should represent the most common types of consumer images. These eight originals were compressed with five different quantization tables (four for the digital camera images) to generate 36 original JPEG compressed images. Finally, the JPEG images were printed at 4 × 6 inch size with a high quality, calibrated continuous tone thermal printer (Kodak Professional 8650 thermal printer).

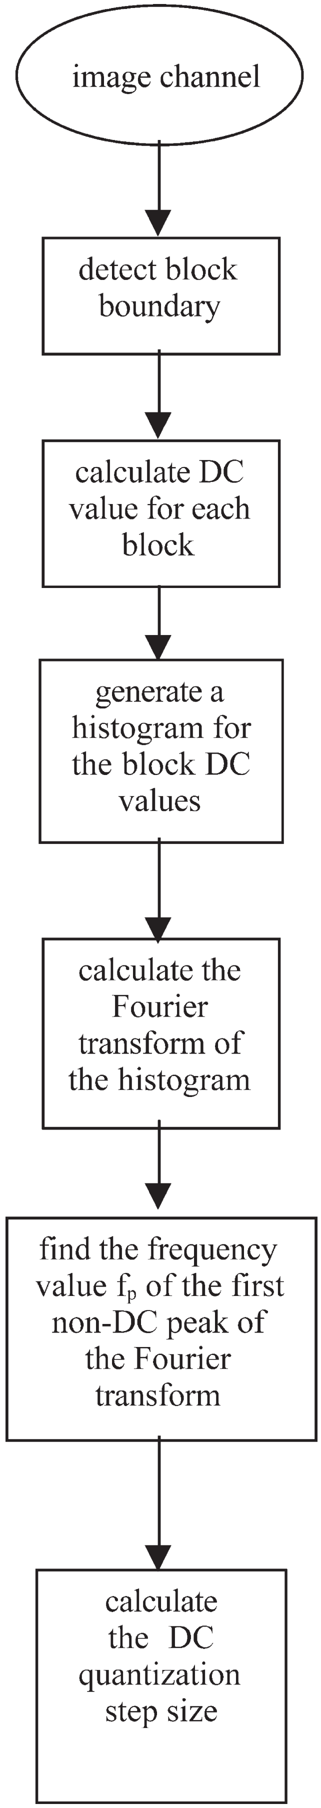

Figure 4. Flow chart of DC quantization step estimation

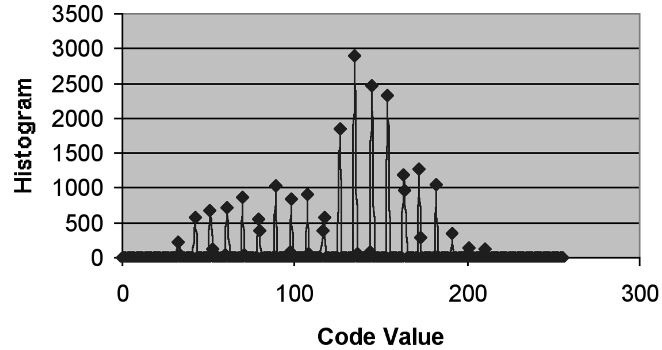

Figure 5. A typical histogram of block DC values

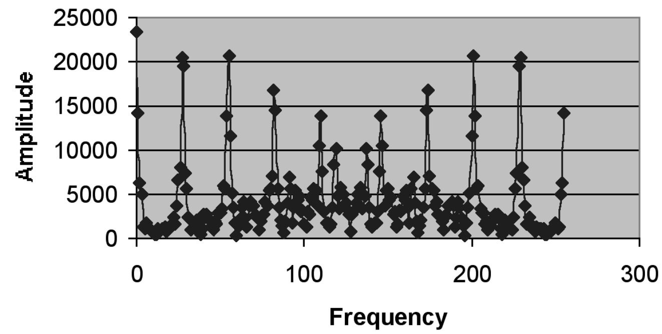

Figure 6. shows a plot of a typical Fourier transform of the histogram of block DC values

It was confirmed through visual assessment that these prints demonstrated a wide range of DCT-based compression blocking artifacts.

It should be noted that all of the prints included in this study were pictures that were not familiar to the participants. That is, the participants had little or no emotional attachment to the people or objects represented in the photographs and were considered third-party participants.[13]

All prints were viewed under GE Chroma 50 bulbs that were suspended and adjusted to provide an illumination level of 1000 lux at the desktop. The lighting fixtures were angled to minimize the visibility of specular reflections on the hand-held prints.

Procedure

The experimental procedure employed a category sort procedure.[14] Each observer was instructed to look for artifacts in the prints such as unsharpness, blocking, contouring, or unnatural color and to place the print into a category based on the severity of the artifact. Note that these instructions were phrased to encourage the participants to indicate the severity of any visible artifacts that are produced by DCT-based compression as the overall goal of this study was to create a metric that would correlate with the participants' impression of quality degradation.

We did not restrict the inspection to blocking and contouring for several reasons. First, the original images have been calibrated to have excellent ratings in quality and the printer is well calibrated and has compensation for MTF loss, therefore the only cause of quality degradation is block-DCT compression. Second, it is unnatural for consumers to evaluate image quality in

Figure 7. Four consumer images from color negative scans

Figure 8. Four consumer images from a digital camera

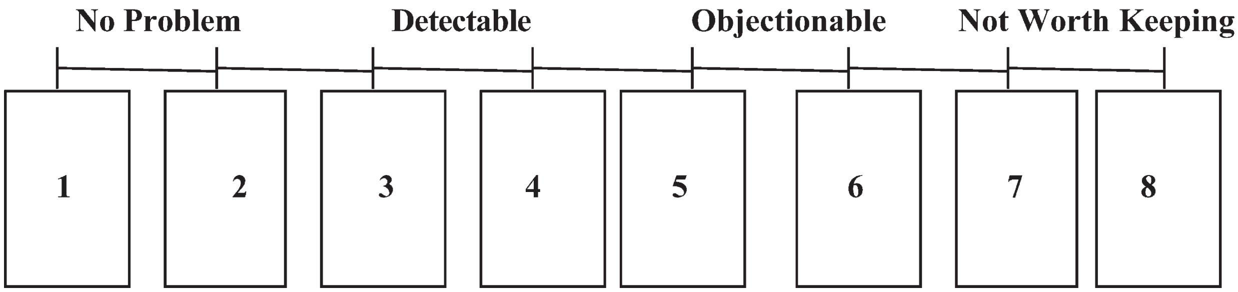

Figure 9. Image quality category label

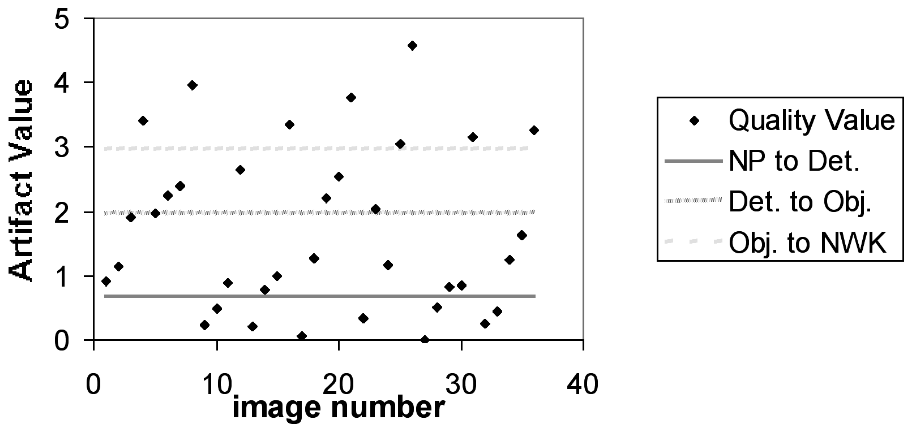

Figure 10. Artifact values for the prints gathered from the experiment (solid, colored lines are quality boundaries)

specific terms such as "blocking" and "contouring", and only limited to these two terms. Third, the ultimate source of the all these artifacts is quantization. Unsharpness and color artifacts are simply different aspects of the compression artifacts, all of which contribute to the overall quality degradation. Last but not least, it is our intention to derive a metric that correlates to the overall quality degradation and the metric proposed in this article was derived based on statistics related to quantization.

When performing this categorization, participants were ask to categorize the print based on the level of the problems present into the picture by placing the print into one of four high level categories labeled "No Problem", "Detectable", "Objectionable" or "Not Worth Keeping". Within each of these bins, there are two numbered categories, e.g., 1 and 2 for "No Problem", 3 and 4 for "Detectable", 5 and 6 for "Objectionable" and 7 and 8 for "Not Worth Keeping", as shown in Fig. 10. It was clear to the observers that these categories refer to image quality and not any semantic or emotional statements about the pictures. The observer was then asked to assign each print to a numbered category which was closest to the nearest labeled category boundary. For example, if a print contained an "Objectionable" level of an artifact and the nearest bin was the "Detectable" bin, then the print is assigned a value of 5. Otherwise, if the print contained a level of artifact that was closer to the "Not Worth Keeping" bin, then it was assigned a value of 6. To perform this categorization, participants placed each print upside down onto a board with labeled categories as shown in Fig. 9.

During the experiment, each participant was given a sample set of pictures that displayed the range of artifact levels present in the experiment and were asked to perform this task for a set of six sample prints. The participants were then provided the opportunity to look through the stack of 36 prints they would see during the experiment to further familiarize themselves with the stimuli. Finally, the participants were asked to lay the stack of prints upside down on the table in front of them, remove the first print from the stack and place the print upside down in the appropriate bin on the table in front of them. The participants were asked to categorize each of the 36 prints using the same procedure. The data were recorded after each participant completed the task and the image order was randomized for the next participant.

While viewing distance was not constrained during the experiment, each participant was instructed to hold the print at a natural viewing distance. The experimenter reminded the participants of this instruction if any of the participants began to overly scrutinize a print. Previous experience has shown that most individuals will hold a print at a viewing distance around 16 inches when provided this instruction.

Experimental Data

Once the data were collected from all the observers, a new interval scale was developed using the category scaling procedure described by Torgeson.[14] This procedure assumes that observer variability is due primarily to confusion between whether a print belongs in one category or another. As such, an image that is often placed into two neighboring categories by different participants are assumed to lie on the image quality scale near the boundary between the two categories while an image that is consistently placed into one of these categories by all participants is believed to be further from this boundary. Using this rationale, the scaling procedure converts the category data to an interval scale and provides an estimate of the boundary between each of the categories into which the prints were placed. Each category boundary therefore corresponds to the point on the scale where half of the participants would place a print into either of the neighboring categories. A linear transformation was then applied to the scale values such that a value of 0 corresponded to the print with the lowest problem level and prints with higher values were scaled as having a higher degree of one of the problems. For this interval scale, the boundaries between the four primary categories were as follows:

0.685	No Problem/Detectable
1.960	Detectable/Objectionable
2.958	Objectionable/Not Worth Keeping

Figure 10 illustrates the artifact values for all the 36 prints. The three problem category boundaries are also shown as solid lines.

A New Image Quality Metric for JPEG Blocking/ Contouring Artifacts

A new image quality metric for JPEG blocking and contouring artifacts is generated based on a statistical analysis using the JMP software. Various models based on these two predictors, i.e., first order, second order, cross terms, were evaluated, and a metric that linearly combines the two predictors produced the best fit. The new metric, defined as the potential Artifact Value (AV), can be expressed in the following equation:

$$AV = 0.752 \times MaxRatio + 0.281 \times S - 1.336$$

where *MaxRatio* is the maximum value of the column and row ratios, and *S* is the estimated DC quantization step size. Fig. 11 shows a plot of the predicted AV using the new metric vs the subjective artifact values from the experiment (in terms of severity of the JPEG artifacts: 0 – no artifacts; 5 – most severe artifacts). The R^2 value of the correlation is around 0.86. It is worth noting that when comparing the results of any two individuals while performing image quality rating tasks, it is rare for the ratings of any two individuals to correlate with an R^2 value of greater than 0.7. Therefore, this metric performs at least as well as any single observer when rating the quality of an image.

Potential Application of the New Metric

Many techniques[15,16] have been developed to reduce the blocking effect. Because most of these techniques employ some kind of image filtering technique and thus reduce image sharpness to some extent, it is imperative that these techniques not be used on a "good image," i.e., one that has not been compressed highly enough to exhibit the blocking artifacts. Most of these techniques are also computationally expensive in general, and it is not wise to waste a lot of resource on a "good image" either.

The image quality metric proposed in this article may be used as a first step toward a robust triage procedure to tackle the problem of tradeoff between print quality and processing throughput, and between artifact removal and image detail preservation. For example, whereras the quality boundary between "Objectionable" and "Detectable" is around 2.0, a conservative triage threshold can be set as 2.0. When the predicted AV exceeds 2.0 for an image, the deblocking algorithm should automatically be applied to the compressed image to reduce the blocking and contouring artifacts.

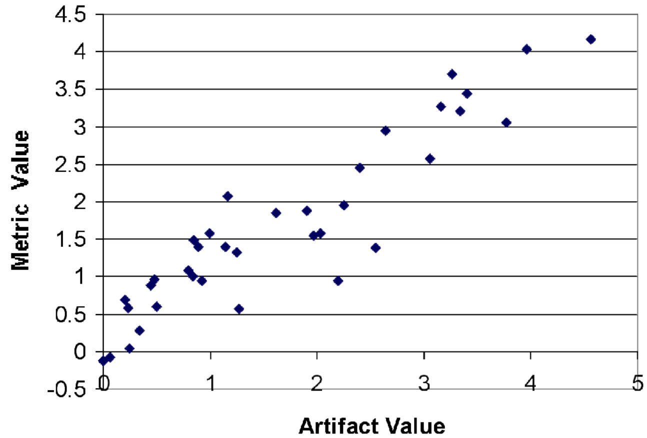

Figure 11. The composite artifact metric vs the subjective artifact value

To develop such a robust JPEG triage procedure, additional studies need to be carried out. For example, more prints should be generated and evaluated. These prints may come from originals of different size (VGA, base, 4base), and they themselves may be of different sizes (4 × 6, 8 × 10, 12 × 18). The metric itself can also be modified/enriched. For example, the floor and peak values of the row/column difference array may be used as predictors besides MaxRatio, and quantization steps for low frequency coefficients may also be included.

The amount of artifact removal also may be controlled, based on the new quality metric. In other words, the new quality metric provides more than a decision on whether or not to apply artifact removal; it also provides a lead to the optimal parameters used by the specific artifact removal algorithm.

Conclusion

In this article, two predictors are proposed to measure the blocking and contouring artifacts in an image caused by JPEG compression. An image quality metric based on these two predictors is further derived based on experimental data collected from a psychophysical evaluation of prints from JPEG compressed images. This new quality metric could potentially be used to make a decision on whether or not to apply JPEG artifact removal; it may also guide the specific artifact removal algorithm.

Acknowledgment. The authors would also like to thank their colleagues at Kodak, especially Dr. Rajan Joshi and Mr. Paul Jones, for the valuable discussion related to this research project. Comments and suggestions made by the anonymous reviewers helped improve the presentation of this article.

References

1. W. B. Pennebaker and J. L. Mitchell, *JPEG still image data compression standard*, Van Nostrand Reinhold, New York, 1993.
2. B. Ho, Mathematical model to quantify JPEG block artifacts, *Proc. SPIE* **1897**, 269–275 (1993).
3. B. W. Jeon and J. Jeong, Blocking artifacts reduction in image compression with block boundary discontinuity criterion, *IEEE Trans. Circuits Systems in Video Technology* **8**(3) 345–357 (1998).
4. C. M. Liu and J. Y Lin, Objective image quality measure for block-based DCT coding, *IEEE Trans. Consumer Electronics* **43**(3), 511–516 (1997).

5. Z. Wang and A. Bovik, Blind measurement of blocking artifacts in images, *Proceedings of International Conference on Image Processing,* IEEE Computer Society Press, Los Alamitos, CA, 2000.
6. F. X. Coudoux, M. Gazalet and P. Corlay, Reduction of blocking effect in DCT-based images based on a visual perception criterion, *Signal Processing: Image Communication* **11,** 179–186 (1998).
7. K. Tan and M. Ghanbari, Frequency Domain Measurement of Blockiness in MPEG-2 Coded Video, *Proceedings of International Conference Image Processing,* IEEE Computer Society Press, Los Alamitos, CA, 2000.
8. H. Datcu, Method for quality evaluation of the block transform compressed images, *Proc. SPIE* **2418**, 118–126 (1995).
9. N. Graham, *Visual Pattern Analyzers*, Oxford University Press, Oxford, 1989.
10. Z. Fan and R. de Queiroz, Maximum Likelihood Estimation of JPEG Quantization Table in the Identification of Bitmap Compression History, *Proceedings of International Conference on Image Processing,* IEEE Computer Society Press, Los Alamitos, CA, 2000.
11. H. Ichikawa, K. Hukami, S. Tanabe, and G. Kawakami, *Standard Pseudoisochromatic Plates: Part 1 for Congenital Color Vision Defects*, Igaku-Shon, New York, 1978.
12. H. Stein, B. Slatt and R. Stein, *Ophthalmic Terminology*, Mosby-Year Book, Inc., St. Louis, 1992.
13. B. W. Keelan, *Handbook of Image Quality: Characterization and Prediction*, Marcel Dekker, Inc., New York, 2002.
14. R. Torgeson, *Theory and Methods of Testing*, Elsevier, Robert E. Krieger Publishing Company, Malabar, Florida, 1985.
15. H. C. Reeve and J. S. Lim, Reduction of blocking effect in image coding, *Proceedings of International Conference on Acoustics, Speech, and Signal Processing,* IEEE Computer Society Press, Los Alamitos, CA, 1983.
16. J. Luo, C. W. Chen, K. J. Parker, and T. S. Huang, Artifacts Reduction in DCT-Based Low Bit Rate Image Compression, *IEEE Trans. Image Processing* **5**(9), 1363–1368 (1996).