

Measurement of Texture Loss for JPEG 2000 Compression

Peter D. Burns and Don Williams*

Burns Digital Imaging and *Image Science Associates

ABSTRACT

The capture and retention of image detail are important characteristics for system design and subsystem selection. An established imaging performance measure that is well suited to certain sources of detail loss, such as optical focus and motion blur, is the Modulation Transfer Function (MTF). Recently we have seen the development of image quality methods aimed at more adaptive operations, such as noise cleaning and adaptive digital filtering. An example of this is the measure of texture (image detail) loss using sets of overlapping small objects, known as dead leaves targets. In this paper we investigate the application of the above method to image compression. We apply several levels of JPEG and JPEG 2000 compression to digital images that include scene content that is amenable to the texture loss measure. A modified form of the method was used. This allowed direct target compensation without data smoothing. Following a camera simulation, the texture MTF and acutance were computed. The standard deviation of the acutance measure was 0.014 (relative error of 1.63%), found by replicate measurements. Structured similarity index (SSIM) values, used for still and video image quality evaluation, were also computed for the image sets. The acutance and SSI results were similar; however the relationship between the two showed an offset between the JPEG and JPEG 2000 image sets.

Keywords: image quality, MTF, JPEG 2000, image compression, dead leaves, texture blur, SSIM

1. INTRODUCTION

The design and optimization of digital imaging systems is often guided by a statistical analysis of the capture and retention of image detail. An established imaging performance metric that is well suited to certain sources of detail loss, such as optical focus and motion blur, is the Modulation Transfer Function (MTF). As performance standards have developed for digital imaging systems, the MTF concept has been adapted and applied as the spatial frequency response (SFR). Measurement of the SFR is generally done using particular test target features such as edges,¹ repeating patterns of square or sign-waves.²

The use of special image features to derive quality measures is challenged when the effective system characteristics vary with local image (scene) content. For example, if the processing of an artificial test image, such as an edge or sinusoid, results in digital spatial processing that is different from that for natural scenes, the computed image quality measure may yield non-representative results. This has led to the development of image quality methods that rely on computed test image content that in some ways resembles natural scenes. Image texture is the term given to the information-bearing fluctuations such as those for skin, grass and fabrics. Since image processing aimed at reducing unwanted fluctuations (noise or other artifacts) can also remove important texture, good product design requires a balance. To aid in the image quality evaluation of digital and mobile-telephone cameras a method is being developed as part of an international standards effort. The method addresses the retention of image texture.^{3,4}

We investigate the application of the above methods to another common adaptive image processing operation – image compression. The texture-loss measure being developed by the Camera Phone Image Quality (CPIQ) Initiative is based on the capture and retention of image detail, as expressed by a signal power spectrum. This spectrum provides a statistical description of image fluctuations as a function of spatial frequency. We adapt this method to evaluate signal loss during image compression. The transfer of image detail is not from object to image, as for the CPIQ applications, but rather from original to compressed digital image data.

2. TEXTURE-LOSS MTF MEASUREMENT

A proposed method to evaluate digital cameras and cell-phone cameras uses a computed image field comprising randomly arranged, overlapping features.⁴ Methods using computed pseudo-images have been used for some time. For example, Fig. 1 shows a field of random polygons from reference 5. This was developed during the analysis of image

contact first author at: pdburns@ieee.org

detector sampling, reconstruction and optimized image processing for improved information capture.⁶ One advantage of such test image fields is the ability to generate large numbers of unique image fields to test, e.g., the stability of the components of the imaging chain.

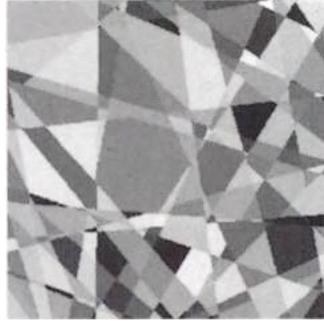


Figure 1: Example of a computed test image array used for imaging performance analysis from Ref. 5

2.1 Dead-Leaves Method

The current CPIQ method under development uses a computed image field of overlapping filled circles or rectangles, often called a ‘dead leaves target’. An example of such a target is shown in the lower left of Fig. 2. The texture blur MTF method relies on several underlying statistical characteristics of this image field. The variability of the method, when viewed as an estimation task was recently addressed.⁷

The texture loss MTF, also referred to as texture blur MTF, is the ratio of output (processed image) measurement with the corresponding, modeled or measured, characteristics for the input target image. The imaging characteristic used is the signal spectrum.* The basic steps of the proposed method are:

1. Transform the captured image array of the target field to one encoded as proportional to luminance. This requires measurement of the camera Opto-Electronic Conversion Function (OECF).
2. Compute the power-spectral density as the square of the amplitude of the two-dimensional DFT of the array.

$$U(m, n) = \left| \sum_{x=N/2+1}^{N/2} \sum_{y=N/2+1}^{N/2} I(x, y) e^{-2i\pi(mx+ny)} \right|^2, \quad (1)$$

where the (N x N) luminance image array data are, $I(x, y)$, corresponding to the dead leaves region.

3. This computed spectrum is corrected for image noise that is estimated from an image region corresponding to a uniform 50% (reflectance) target step,

$$U'(m, n) = U(m, n) - H(m, n). \quad (2)$$

$H(m, n)$ is the spectrum measured from this step, computed in the same way as Eq.(1) and scaled to remove the effect of different data array sizes.

4. Divide this array, frequency-by-frequency, by the modeled spectrum for the specific target to yield a two-dimensional array as the square of the effective MTF.

$$S(m, n) = \frac{U'(m, n)}{T(m, n)}, \quad (3)$$

where $T(m, n)$ is the modeled spectrum of the input target

5. Compute the square-root, frequency-by-frequency, of this array,

* Here we refer to the *signal spectrum* to indicate the intent of the measurement. However it is common to call this the *noise power spectrum* due to the use of the power spectral density to describe unwanted image noise.

$$MTF_{2D}(m, n) = \sqrt{S(m, n)}. \quad (4)$$

6. Compute the one-dimensional MTF vector by a radial-average of this array, $MTF(v)$, where

$$v = (m^2 + n^2)^{0.5} \text{ the radial frequency index } v = 1, 2, \dots, v_{\max}.$$

7. A summary measure, acutance, is computed by weighting the texture MTF by a visual contrast sensitivity function (CSF). The texture acutance is computed as,

$$A_{ref} = \sum_{v=1}^{v_{\max}} CSF(v), \quad A = \sum_{v=1}^{v_{\max}} MTF(v)M(v)CSF(v) \quad (5)$$

where M is the modeled display MTF. The acutance is,

$$\text{acutance} = \frac{A}{A_{ref}}. \quad (6)$$

We should note that this method is currently under development and refinement by members of the Camera Phone Image Quality (CPIQ) Initiative, and it is likely that a variant of this method will be adopted. McElvain *et al.*³ refer to a second method for estimating the effective texture MTF by reversing steps 4 and 5.

2.2 Test Target Image and Modified Method

The test object used for texture blur analysis is shown in Fig. 2. The element in the lower left is the texture blur target, containing random disks element surrounded by the step tablet used for OECF and noise measurements. The edge Spatial Frequency Response (SFR) target is shown on the right-hand side. This is helpful for comparison of the edge-SFR¹ with the texture blur results. The picture included (*Giverny*) is clearly the type of content to which image most image compression is aimed.

The dead leaves spectrum was computed using a modified procedure for this study.[†] One difference from the previously published approach was the scaling of the result so as to yield a conventional power spectral density, rather than the simple two-dimensional Discrete Fourier Transform (DFT) of Eq. 1.

$$U(m, n) = \frac{1}{(dxN)^2} \left| \sum_{x=N/2+1}^{N/2} \sum_{y=N/2+1}^{N/2} I(x, y) e^{-2i\pi(mx+ny)} \right|^2, \quad (7)$$

where dx is the data sampling interval, which we expressed in pixels, i.e. $dx = 1$ pixel. The frequency sampling of the power spectrum estimate is therefore, $df = 1/(dxN)$ cycles/pixel. An advantage of this form of the power spectral density is that the result does not scale with the size of the data array used. This is helpful when combining spectrum measurements from different sized regions (step 3, Eq. 2). In addition this two-dimensional spectrum, when integrated correctly over spatial frequency, yields the signal variance as it should. The units of the power spectrum are, signal variance/(cy/pixel)². The computed power spectrum for the dead leaves target is shown in Fig. 3, as is a fit to a simple equation. Note that the most common presentation of these results is in Fig. 3a, a log-log plot.

A second modification to the texture blur method was to compensate for the target-spectrum in Eq. (3) by direct measurement of the input spectrum, rather than a model spectrum. This was needed because our input image was modified by an image simulation (see below) so as to present more natural image microstructure characteristics to the image compression operations under evaluation. To accomplish this, the radial averaging (step 6), which results in a smooth spectral estimate, was performed before step 4. This had the added benefit of reducing the variability in the estimated power spectra prior to the frequency-by-frequency division operation. The texture MTF was computed as,

$$S'(v) = \frac{U'(v)}{U'_{target}(v)}. \quad (8)$$

[†] Information on a Matlab implementation of the modified method is available from PDB.



Figure 2: Test target used in the texture blur evaluation of image compression

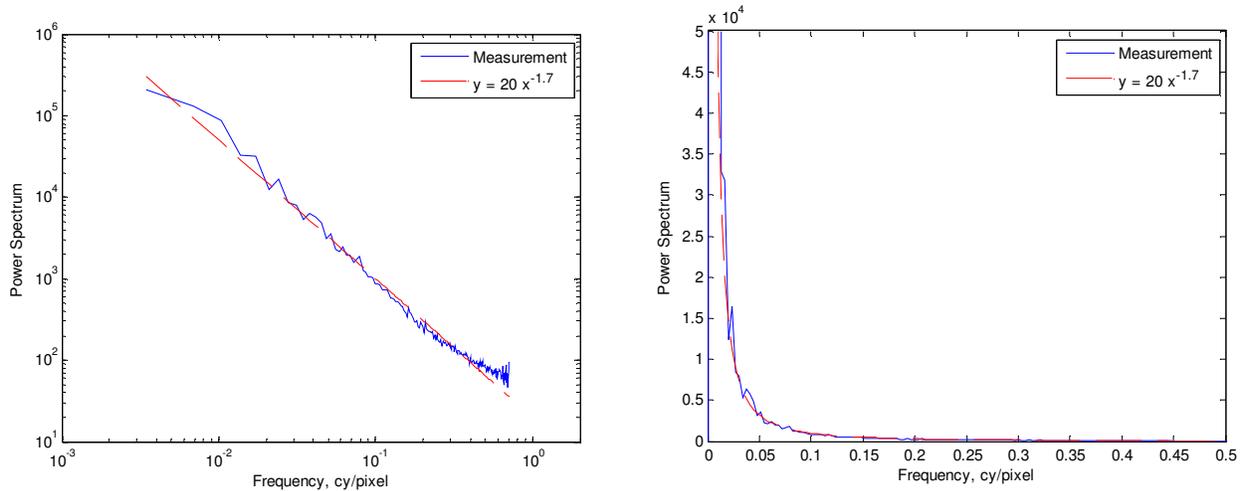


Figure 3: Computed spectral density for the dead leaves patch of the texture blur target. Note log and linear axes.

3. APPLICATION TO IMAGE COMPRESSION

For camera evaluation the test target, as shown Fig. 2, would be printed at a size appropriate for the optical field of view. For our purposes this is not necessary, so we can avoid the printing and any problems with optical distortion, etc. However we would like the test images used to represent captured images, rather than the computed noise-free content of the texture blur and edge SFR targets. The solution adopted was to introduce optical and signal processing characteristics via a simple functional image simulation. Since we are primarily interested in the image microstructure for this study we included; an optical MTF, detector image noise, and color filter array interpolation.

Figure 4 shows the steps implemented for the processing of the input image shown in Fig. 2. The optical MTF was introduced by convolution, with no field variation. The color filter array (CFA) sampling operation resulted in a single image array to which signal-dependent noise was added. The subsequent spatial interpolation to a color image resulted in spatial and color (pixel-to-pixel) correlation being introduced into the final color image. The combined resulting MTF

was evaluated using the slanted-edge feature of the SFR target, and is shown in Fig. 5a. The image noise characteristics that were introduced are indicated as the pixel standard deviation in Fig. 5b.

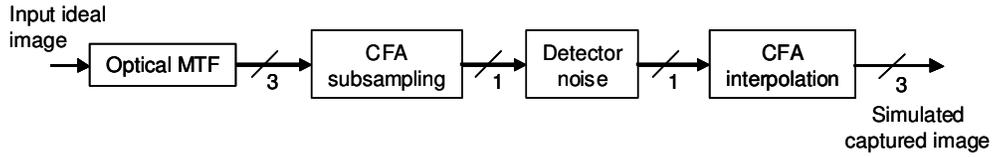


Figure 4: Image simulation used to prepare example image set for compression study.

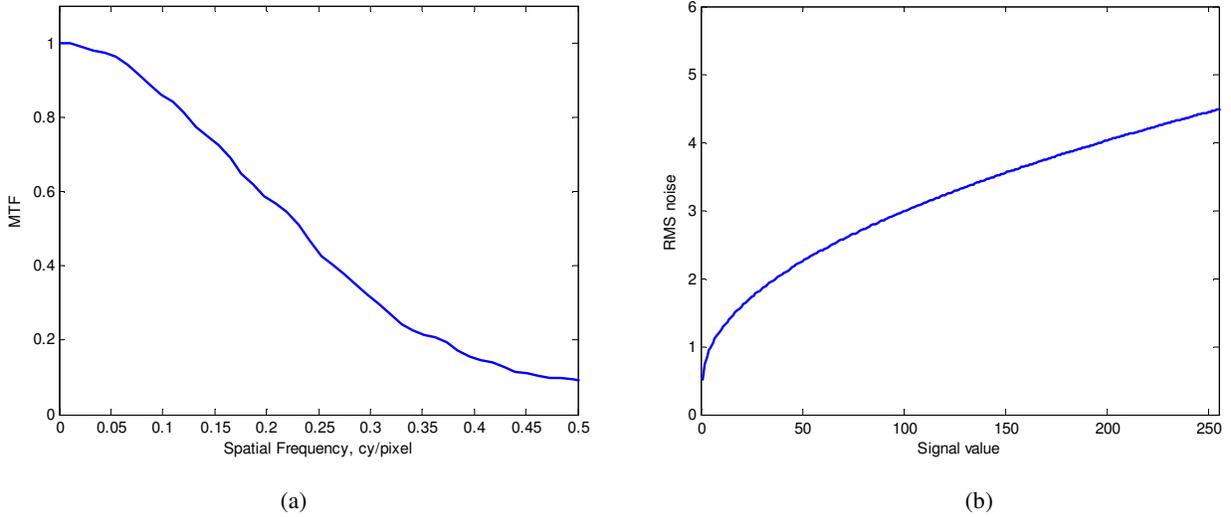


Figure 5: Simulated image capture. (a) Luminance MTF including optics, sensor and CFA interpolation, (b) image noise amplitude introduced as a function of signal level.

The power spectrum for the dead leaves simulation target was computed and corrected for the corresponding noise spectrum. From the results in Fig. 6 we see that the correction due to the noise spectrum is minor in this case.

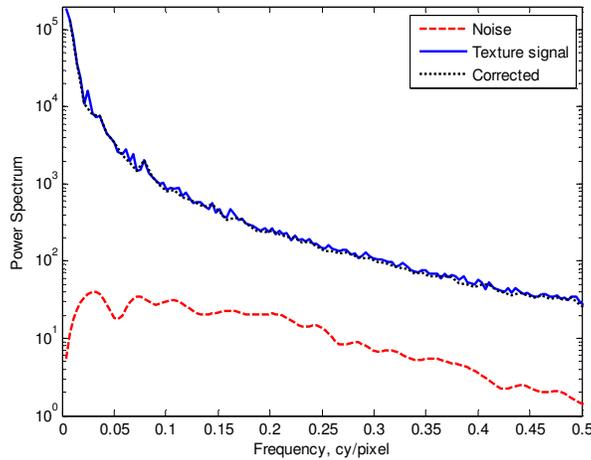


Figure 6: Computed dead leaves (texture signal) radial averaged power spectrum for original simulated image, with corresponding noise spectrum from uniform 50% target.

3.1 JPEG and JPEG2000 results

The simulated digital target file was compressed using a reference JPEG 2000 implementation⁸ for several levels of compression ranging from 30 to 140:1. The image distortion introduced ranged from the imperceptible to severe, as expected. Our input test file, with simulated camera characteristics was 2000 x 1300 pixels and required 7.5 Mb of disk space. Lossless compression resulted in a compression ratio of approximately 2:1, as is usually observed for natural scenes. Figure 7 shows a range of performance observed for the input image, after 40:1 and 100:1 compression.



Figure 7: Results of JPEG 2000 compression: (a) input (b) 40:1, and (c) 100:1 compression ratio

The texture blur MTF was computed for these all cases, and two are shown in Figs. 8 and 9. The measured texture MTF is based in direct measurement of the signal spectra (as explained above), and no smoothing has been applied. The fluctuations in the texture MTF estimates indicate the variability in estimates.

From the measured texture MTF results we computed the texture acutance. This requires the selection of a viewing distance and display/printer MTF. Following the current CPIQ recommendations, our results are computed for a computer display with sampling at 100 pixels/inch, and a 60 cm viewing distance. Table 1 shows the results for the sets of JPEG and JPEG 2000 compressed images for a wide range of image quality.

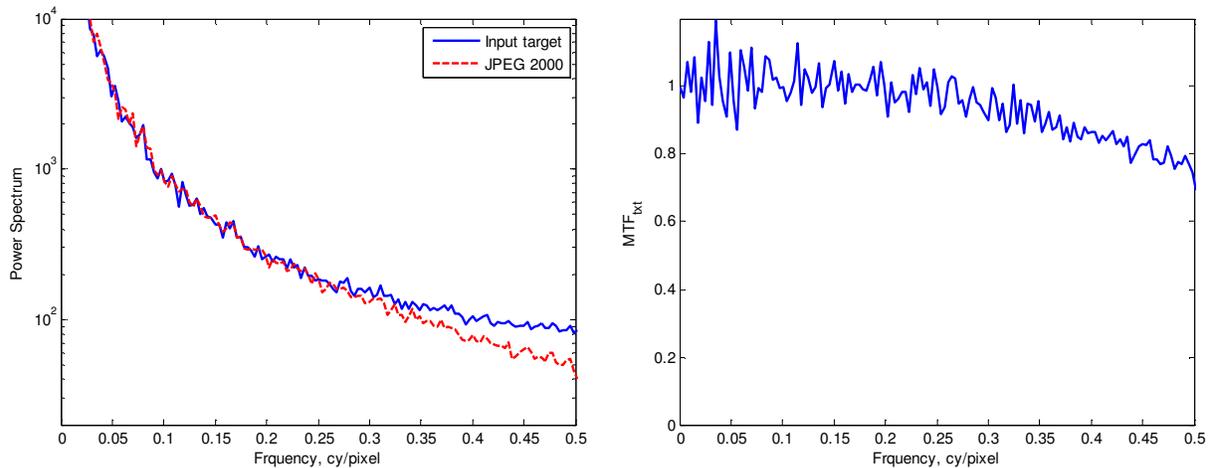


Figure 8: The dead leaves power spectra before and after 40:1 compression, and the corresponding texture MTF. The acutance is 0.95.

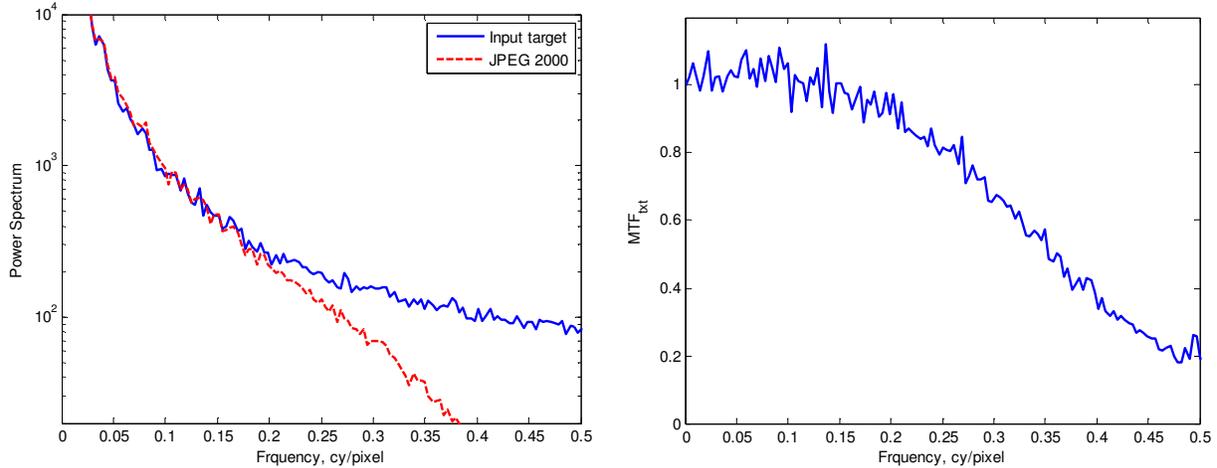


Figure 9: The dead leaves power spectra before and after 100:1 compression, and the corresponding texture MTF. The acutance is 0.819.

Table 1: Results from the texture blur evaluation for set of JPEG and JPEG 2000 compressed test images (24-bit original). The acutance was computed 60 cm viewing of a computer display.

Compression rate	Bits/pixel/color	JPEG 2000		JPEG	
		Texture acutance	SSI index	Texture acutance	SSIM index
30	0.80	0.986	0.939	1.02	0.939
40	0.60	0.950	0.924	0.991	0.922
50	0.48	0.951	0.920	0.961	0.904
60	0.40	0.859	0.908	0.930	0.886
80	0.30	0.826	0.884	0.890	0.844
100	0.24	0.819	0.865	0.841	0.790
120	0.2	0.796	0.835	0.777	0.742
140	0.17	0.731	0.799	0.667	0.667

To evaluate the general utility of the texture blur acutance to image compression, we also computed the Structured Similarity Index (SSIM).⁹ The SSIM is an objective measure developed to predict image quality when a reference image is available. In our case the input image provides the reference. The method computes a visual-difference map, based on a model of visually important information. The average value of the difference image is reported as the SSI. The precision of the acutance measure was investigated by computing repeated measurements for the same compressed file (0.4 bits/pixel/color). The standard deviation was found to be a remarkably low, 0.014, for a relative error of 1.63%.

Comparison of the two summary quality measures is also shown in Figs. 10. All results are reported as computed with no scaling or offset applied. We conclude that, for the image compression rates and scene content used, there is general agreement. Note, however, there is an offset between the acutance-SSIM characteristic between the JPEG and JPEG 2000 image sets in Fig. 10b.

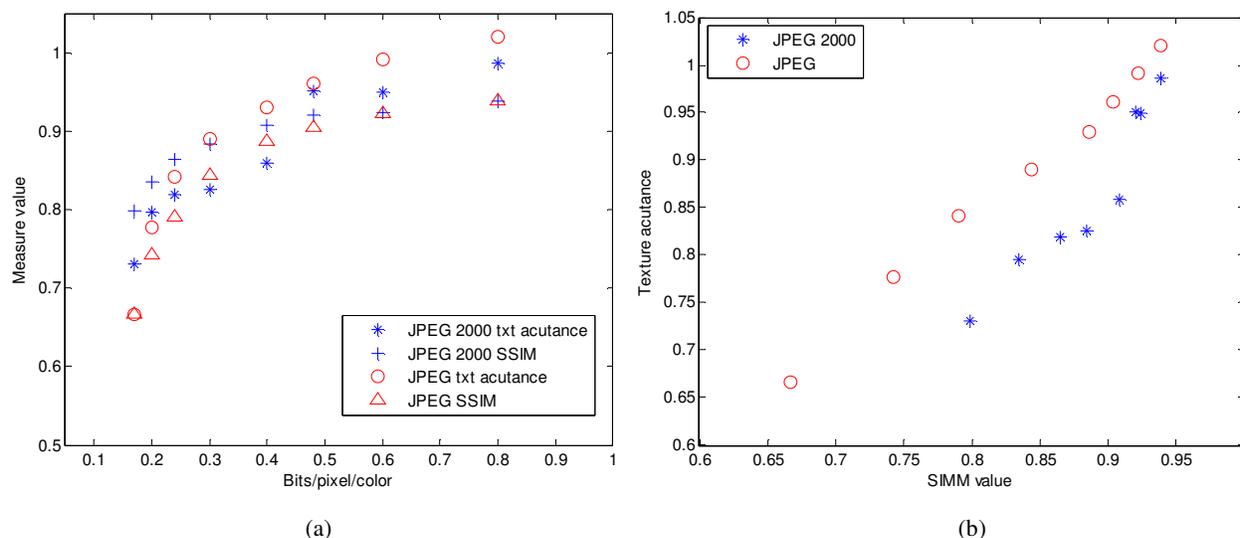


Figure 10: Results from the texture acutance and structured similarity index analysis for sets of JPEG 2000 and JPEG-compressed test images (24-bit original). The acutance was computed 60 cm viewing of a computer display

4. CONCLUSIONS

The texture blur measure that is based on the loss of important image information via power spectrum analysis has been applied to the JPEG and JPEG 2000 image compression. We adapted the method to include (1) ideal computed test images used in camera testing, (2) simulation of image capture and (3) direct calibration for the input target signal spectrum. When direct target calibration is desired, a simple reversal of two steps in the current texture-blur method facilitates this. For our implementation, without any data smoothing, the standard deviation of the measured acutance was found to be 0.014, and relative error of 1.63%. The texture acutance compared well with an established measure used for still and video image quality evaluation, the structured similarity index. The acutance and SSIM results were similar; however the relationship between the two showed an offset between the JPEG and JPEG 2000 images sets.

Acknowledgements

We are indebted to the members of the I3A CPIQ initiative and ISO/TC42 teams for their discussions, in particular Frédéric Cao, Herve Hornung, Dietmar Wueller, and Uwe Artmann.

REFERENCES

- [1] ISO/TC WG18, [Photography – Electronic still picture cameras – Resolution – measurements], ISO, 1998
- [2] Artmann, U. and Wueller, D., “Differences of digital Camera Resolution Metrology to Describe Noise Reduction Artifacts,” Proc. SPIE 7529, 7529L (2010)
- [3] McElvain, J., Campbell, S. P., Miller, J., and, Jin, E. W., “Texture-based measurement of spatial frequency response using the dead leaves target: extensions, and application to real camera systems,” Proc. SPIE 7573, 75370D (2010)
- [4] Cao, F., Guichard, F., and Hornung, H., “Measuring texture sharpen of a digital camera,” Proc. SPIE 7250, 72500H (2009)
- [5] Huck, F., Fales, C., Rahman, Z, [Visual Communication: communication: an information theory approach], Fig. 2.3 page 21, Kluwer Academic Publishers, Norwell MA, USA 1997
- [6] Fales, C. L., Huck, F. O., McCormick, J. A., and Park, S. K., “Wiener restoration of sampled image data: end-to-end analysis,” JOSA A, Vol. 5, Issue 3, pp. 300-314 (1988)
- [7] Burns, P. D., “Estimation error in image quality measurements,” Proc. SPIE 7867, 7867H (2011)
- [8] *kdu_compress*, available from <http://www.kakadusoftware.com>
- [9] Wang, Z., Bovik, A., Sheikh, H., and Simonelli, E., “Image Quality Assessment: From Error Visibility to Structural Similarity,” IEEE Trans. Image Processing, Vol. 13, pp.600-612, (2004)