D. Baxter, F. Cao, H. Eliasson, J. Phillips, Development of the I3A CPIQ spatial metrics, Image Quality and System Performance IX, Electronic Imaging 2012.

http://dx.doi.org/10.1117/12.905752

Development of the I3A CPIQ spatial metrics

Donald Baxter[a], Frédéric Cao[b], Henrik Eliasson[c] and Jonathan Phillips[d]

[a]STMicroelectronics Ltd., 33 Pinkhill, EH12 7BF Edinburgh, Great Britain;
[b]DxO Labs, 3 Rue Nationale, 92100 Boulogne-Billancourt, France;
[c]Sony Ericsson Mobile Communications, Mobilvägen 10, SE-221 88 Lund, Sweden;
[d]Eastman Kodak Company, 343 State Street, Rochester, NY, USA 14650

ABSTRACT

The I3A Camera Phone Image Quality (CPIQ) initiative aims to provide a consumer-oriented overall image quality metric for mobile phone cameras. In order to achieve this goal, a set of subjectively correlated image quality metrics has been developed. This paper describes the development of a specific group within this set of metrics, the spatial metrics. Contained in this group are the edge acutance, visual noise and texture acutance metrics. A common feature is that they are all dependent on the spatial content of the specific scene being analyzed. Therefore, the measurement results of the metrics are weighted by a contrast sensitivity function (CSF) and, thus, the conditions under which a particular image is viewed must be specified. This leads to the establishment of a common framework consisting of three components shared by all spatial metrics. First, the RGB image is transformed to a color opponent space, separating the luminance channel from two chrominance channels. Second, associated with this color space are three contrast sensitivity functions for each individual opponent channel. Finally, the specific viewing conditions, comprising both digital displays as well as printouts, are supported through two distinct MTFs.

Keywords: Image quality, sharpness, noise, texture blur, MTF, noise power spectrum, contrast sensitivity function

1. INTRODUCTION

Surprisingly, the megapixel count of the mobile phone imager is still the only image quality metric widely available and marketed to the general public. Obviously, this measure of image quality is seriously flawed since it correlates very badly with most important aspects of image quality such as color rendition, sharpness, signal to noise ratio (SNR), and so on. Furthermore, even the most commonly used metrics today, e.g., sharpness[1] and SNR,[2] do not in many cases correlate very well with the perceived image quality. The main reason for this sometimes poor correspondence between measured and experienced image quality is not due to poor metrics, but rather the fact that important aspects of the human visual system are not taken into account. Furthermore, the properties of the medium used for watching the resulting images, such as a computer display or paper printout, need to be handled in an appropriate way.

The assessment of the perceived image quality is made even more difficult by the complex, often non-linear, processing performed in the camera image signal processor (ISP). The effect of such algorithms is indeed reduced noise and the sharpness is maintained, but a serious side-effect is the smearing out of low-contrast texture areas, leaving the impression of an "oil painting" in the worst cases. Unfortunately, neither the sharpness nor SNR measurement methods used today are able to pick up this effect in an effective manner. For this reason, a texture blur metric was developed which uses a test target known as the "dead leaves" target.[4] The effectiveness of this approach has been demonstrated in a range of investigations.[4–8]

The Camera Phone Image Quality (CPIQ) initiative has the goal to produce a consumer-oriented image quality rating system for mobile phone cameras. As such, it relies on having access to perceptually well-correlated image quality metrics. This paper describes a subset of these metrics, referred to as the spatial metrics. This encompasses metrics for sharpness,[9] SNR,[10] and texture blur.[11] The common feature of these metrics is that they are functions of spatial frequency and as such are dependent on viewing conditions, including the distance between image and observer, and the type of medium (print or display). This suggests a common framework incorporating a set of contrast sensitivity functions, spatial models of the printer/paper as well as display, and color space in which the measurement results should be analyzed. This paper describes the development of such a framework as implemented in CPIQ.

The paper is organized as follows. The "backbone" metrics for sharpness, texture blur and noise are first described in Section 2. In Section 3, the framework employed to transform the raw measurement data into numbers representing visually correlated metrics is introduced. The mapping to Just Noticeable Differences (JND), allowing the multivariate combination of the metrics into a summation value describing the overall subjective image quality degradation, is described in Section 4. Concluding the paper is a discussion on how to incorporate the proposed metrics into an overall image quality metric together with suggestions for future improvements.

2. SPATIAL METRICS

2.1 Edge acutance

The ISO 12233 standard[1] describes several methods to measure and calculate the spatial frequency response (SFR) of an imaging system. For the CPIQ sharpness metric, the edge SFR was found to be most appropriate. One reason for this choice is that the edge SFR provides a localized measurement, while in other methods the measurements at different spatial frequencies are spatially separated. For a mobile phone camera lens, where the sharpness can vary considerably across the field, this might lead to large errors in the SFR calculations.

The edge SFR as implemented in CPIQ has also been modified to allow for measurements in the tangential and sagittal directions, as this is typically what is being measured in an optical system. The CPIQ SFR test chart for camera arrays larger than VGA is shown in Figure 1. MTF measurements assume a linear system. Since the transfer func-

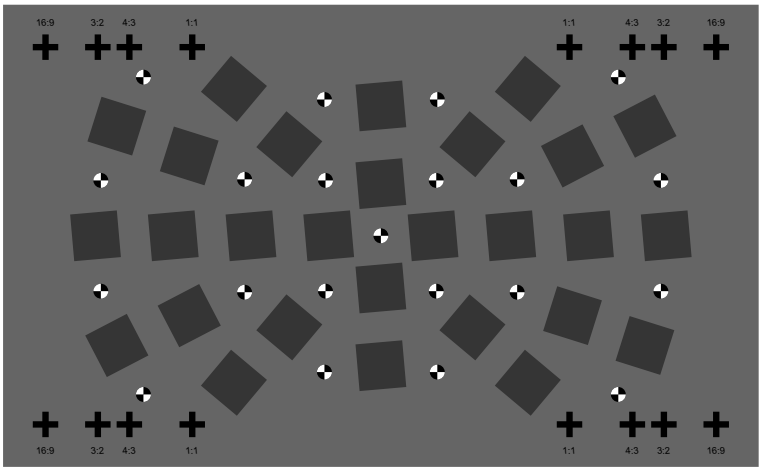

Figure 1. CPIQ SFR test chart for image sizes larger than VGA. Reproduced with permission from the International Imaging Industry Association (I3A).

tion of a digital camera is non-linear, the image for SFR analysis has to be linearized before measurement. In ISO 12233, this is performed through the inversion of the opto-electronic conversion function (OECF). However, as has been shown previously,[12] using a low-contrast chart may eliminate the need for this inversion. Furthermore, since the largest contribution to the non-linearity arises from the gamma curve, which is known, the combination of a low contrast test chart and inversion of the gamma curve will yield sufficient linearity. This is implemented in the CPIQ acutance metric by transforming the RGB image into CIEXYZ(D65) space and performing the analysis on the Y channel. As described below, this also fits well into the common analysis framework for the CPIQ spatial metrics.

2.2 Texture blur

Texture blur is the most recently developed metric of the spatial metrics presented in this paper. It addresses a problem which is specific to digital cameras and related to the content adaptive digital processing applied to digital images. The adaptation can be quite simple (e.g., based on the local gradient magnitude) or more complex (e.g., based on prior statistical learning). As a consequence, a camera can reduce noise in homogeneous areas and enhance the sharpness on isolated edges, yielding good noise and sharpness performance. However, images may still not look natural because the camera is unable to suitably render fine details with low contrast. A new test chart and protocol was therefore necessary to specifically evaluate the performance of cameras on texture. The test chart

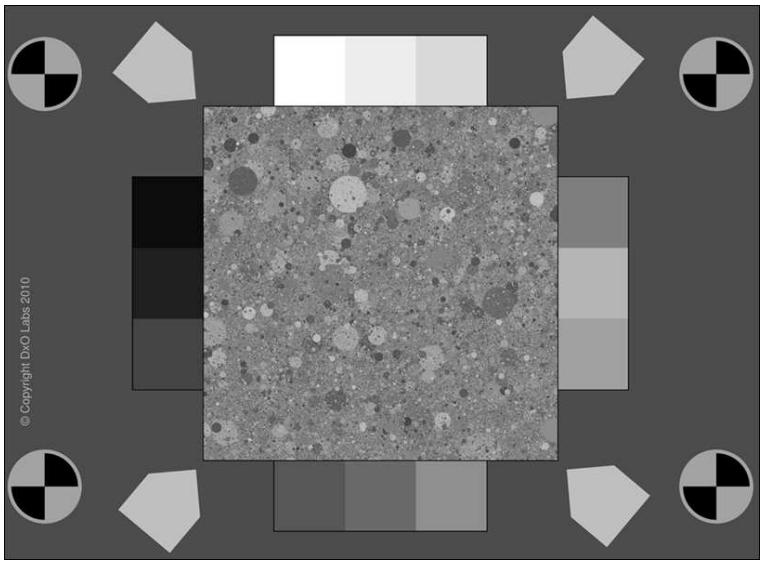

Figure 2. Texture blur chart.

that was eventually chosen is composed of disks with random radii following a specific distribution. The theory about this target was developed in previous work.[4,5] The key properties of this test chart are: low contrast (reflectance between 0.25 and 0.75), isotropic, scale invariant, occluding objects, and known ground truth statistics. The power spectrum of the ideal target is known and follows a power function, which is a property shared with many natural textures as proved in some statistical studies.[13]

The square root of the ratio of the power spectrum of the photograph and the theoretical chart defines a texture MTF. The interpretation is exactly the same as for the usual MTF computed on an edge, via the SFR algorithm. The value at a given frequency is the attenuation (or sometimes amplification) due to the camera. It is worth noting that in good illumination conditions, the attenuation is mainly related to optical blur. Therefore it is expected that edge SFR and texture MTF are very close in this case. When the illumination level decreases, the noise level increases and digital processing (noise reduction) will usually degrade low contrast texture faster for low end cameras, such as camera phones. As proposed by McElvain et al.,[6] the power spectrum of noise is calculated on a uniform patch and subtracted from the power spectrum of the texture part. The aim of this operation is to make the texture MTF insensitive to noise. We conducted tests with different noise levels (ISO settings from 100 to 3200 on a camera with RAW outputs) with a simple image pipe with no noise reduction at all, showing that texture acutance was indeed independent from noise when using this refinement. The texture power spectrum is computed in the linear gray level scale. For this purpose, gray level patches are used around the texture area in order to invert the camera OECF. Then, the matrix transforming linear sRGB values

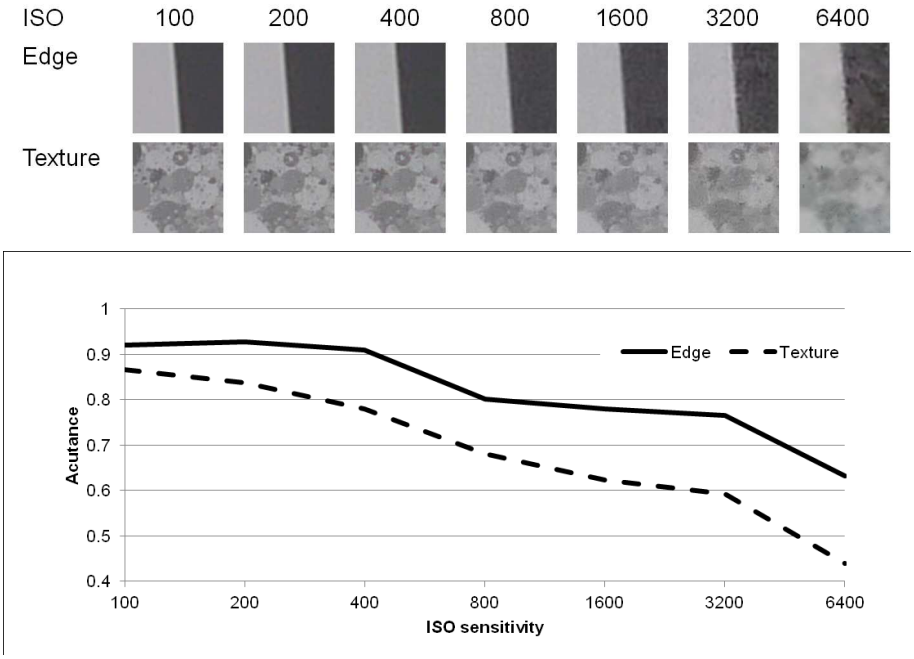

Figure 3. Crop of an edge and texture chart for a same camera at different ISO settings. The measurement shows a faster degradation on texture than on edges.

into CIEXYZ values is used. An acutance is computed on the luminance channel Y. The CSF is the same as for the edge acutance.

2.3 Visual noise

Two classes of noise metrics exist, namely, those with and those without spatial filtering. Visually aligned spatial filtering is mandatory for the I3A CPIQ visual noise metric for two reasons. The first is noise assessment under different viewing conditions. The second is the increasing intelligence of in-camera processing which is leading to an increasing disparity between the two classes of noise metrics. This disparity is due to the changes in the noise power spectrum by in-camera processing. Example of visually aligned noise metrics include ISO 15739 visual noise,[2] S-CIELAB[14,15] and vSNR.[16]

The ISO 15739 visual noise standard was chosen as the starting point as it is a pre-existing standard and the frequency-based spatial filtering allows multiple frequency filters to be easily cascaded. However, the optical non-uniformities prevalent in cell phone cameras and the potentially higher noise level pose significant challenges to the current ISO 15739 visual noise protocol.

The high peak of the ISO 15739 luminance channel contrast sensitivity function (CSF) results in significant amplification of luminance noise. For cell phone images captured in

low light, this amplification results in clipping of the noise. This is the primary reason for changing to the luminance CSF which is expressed by Eq. 2 in Section 3.1.

The ISO 12232 high pass filter[17] is able to remove non-uniformities, but the relatively small kernel size results in valid information near the luminance CSF peak response also being removed. The proposed frequency-based high pass filter enables better control over the cut-off frequency. In the proposed visual noise metric, three frequency-based filters are cascaded: the CSF, the display MTF, and the high pass filter.

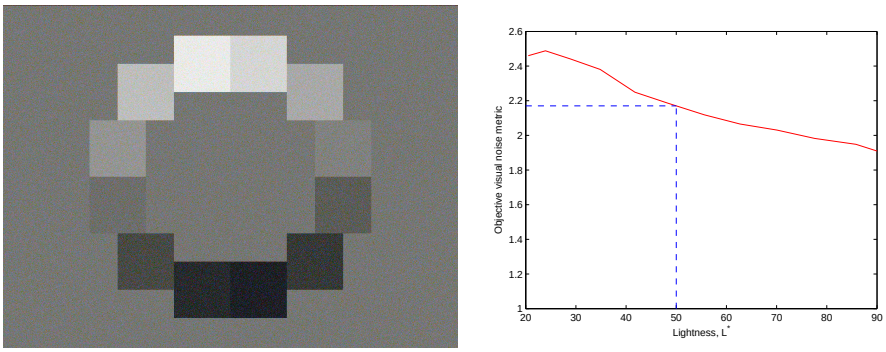

Figure 4. Left: Simulated OECF chart under 30 lux, 3200 K illumination. On the right is the patch mean L^* lightness versus patch objective noise metric plot. Also illustrated is the interpolated 50% brightness objective visual noise metric value.

The visual noise objective metric is the base 10 logarithm of the weighted sum of the variance and covariance values for the spatially filtered CIELAB image. This is the same as the equation proposed by Keelan et al.[18] The assumption is that the perception of noise above threshold is approximately logarithmic:[19]

$$\Omega = \log_{10}\left[1 + w_1\sigma^2(L^*) + w_2\sigma^2(a^*) + w_3\sigma^2(b^*) + w_4\sigma^2(L^*a^*)\right] \quad (1)$$

The test target is an ISO 14524:2009 compliant OECF chart[20] (Figure 4). The ISO 15739 protocol requires the chart background to be around 118 codes (8-bit sRGB). Practically this is problematic due to the wide variation in and increasing intelligence of auto exposure correction(AEC) algorithms and the lack of exposure compensation options in many cell phone cameras. Two complementary methods are provided to manage this. The first is to use either a changeable patch or neutral density filter in the centre of the chart to "fool" the AEC to making the background approximately 118 codes. Second, the 50% brightness ($L^* = 50$) objective noise value is interpolated from the set of L^* mean versus objective noise data values for each patch.

3. COMMON FRAMEWORK

In order to distill the objective metrics described above into subjectively correlated quantities, a model for the human visual system as well as the output medium has to be

established. For this purpose, we define a set of contrast sensitivity functions (CSF) as well as display and printer/paper MTFs.

3.1 Contrast sensitivity function

The CSFs used in the CPIQ spatial metrics come from the work of Johnson and Fairchild.[15] The functional form of the CSFs can be expressed as follows

$$\mathrm{CSF}(\nu) = \frac{a_1 \nu^{c_1} \exp\left(-b_1 \nu^{c_2}\right) + a_2 \exp\left(-b_2 \nu^{c_3}\right) - S}{K} \tag{2}$$

and the coefficients for the luminance CSF, CSF_A, and the two chrominance CSFs, CSF_{C1} and CSF_{C2}, are shown in Table 1. The spatial frequency, ν, is in this case expressed in cycles per degree. The bandpass nature of the luminance CSF implies that when using this

Table 1. Coefficients defining the luminance and chrominance CSFs.

Coefficient	CSF_A	CSF_{C1}	CSF_{C2}
a_1	1	109.1413	7.0328
a_2	0	93.5971	40.691
b_1	0.2	0.0004	0
b_2	0	0.0037	0.1039
c_1	0.8	0	0
c_2	1	3.4244	4.2582
c_3	0	2.1677	1.6487
K	1	202.7384	40.691
S	0	0	7.0328

function as a spatial filter for the visual noise metric, there will be no signal at zero spatial frequency. In the ISO 15739 standard, this is handled by normalizing the luminance CSF to 1 at zero spatial frequency. However, this has a side-effect of amplifying the signal, leading to clipping problems as described above. Another approach, which is the one adopted in CPIQ, is to subtract the average value prior to filtering and adding this value back afterwards. This makes more sense intuitively as well, since the CSF describes *contrast transfer*, where zero contrast implies that only the average value survives.

3.2 Color space

The CSF filtering is performed in a color opponent space AC_1C_2, as described in ISO 15739.[2] Assuming that the source image has first been transformed into XYZ with a D65 white point, the transformation into this space is performed via XYZ with a CIE illuminant

E white point, according to

$$\begin{bmatrix} X_E \\ Y_E \\ Z_E \end{bmatrix} = \begin{bmatrix} 1.0503 & 0.0271 & -0.0233 \\ 0.0391 & 0.973 & -0.00927 \\ -0.00241 & 0.00266 & 0.918 \end{bmatrix} \begin{bmatrix} X_{\mathrm{D65}} \\ Y_{\mathrm{D65}} \\ Z_{\mathrm{D65}} \end{bmatrix} \tag{3}$$

and

$$\begin{bmatrix} A \\ C_1 \\ C_2 \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 \\ 1 & -1 & 0 \\ 0 & 0.4 & -0.4 \end{bmatrix} \begin{bmatrix} X_E \\ Y_E \\ Z_E \end{bmatrix} \tag{4}$$

It should be observed that $A = Y_E \approx 0.973 Y_{\mathrm{D65}} \approx Y_{\mathrm{D65}}$. Therefore, since the edge acutance and texture acutance metrics both work on the luminance channel only, the analysis is performed on the Y_{D65} channel in these cases.

3.3 Printer and display MTF

The use of a contrast sensitivity function, with the spatial frequency expressed in cycles per degree, implies that specific viewing conditions have to be defined. As part of these viewing conditions, the display medium used should be defined as well. In CPIQ, two kinds of viewing devices are employed. The display device has an MTF described by a sinc function:

$$M_{\mathrm{disp}}(\nu) = \left| \frac{\sin \pi k_{\mathrm{disp}} \nu}{\pi k_{\mathrm{disp}} \nu} \right| \tag{5}$$

The factor k_{disp} is dependent on viewing condition, as summarized in Table 2. The second medium, the print, has an associated MTF given by

$$M_{\mathrm{print}}(\nu) = \exp\left(-\frac{\nu}{k_{\mathrm{print}}} \right) \tag{6}$$

where the factor k_{print} is also shown in Table 2. The values for the factor k_{print} were found in the literature[21,22] from two representative investigations on inkjet printers.

3.4 Viewing conditions

The CPIQ viewing conditions are shown in Table 2. These conditions are the same as those specified in ISO 15739,[2] with the exception of the photo frame condition, which is unique for CPIQ.

4. JND MAPPING

4.1 Edge acutance

In order to obtain a number describing the perceived sharpness for a particular viewing condition, the edge SFR is first integrated, weighted by the luminance CSF and medium MTF:

$$Q_{\mathrm{edge}} = \frac{\int_0^{\nu_c} S(\nu) C(\nu) M(\nu) d\nu}{16.88}, \tag{7}$$

Table 2. CPIQ viewing conditions.

Condition	k_{disp} **(degrees)**	k_{print} **(cy/degree)**
Small print 10×15 cm		21.8
Large print 40×60cm		65.4
Computer monitor viewing: 100% at 100 ppi	0.0243	
3" cell phone VGA display	0.0218	
42" 1080p HDTV	0.0159	
7" digital photo frame	0.0127	

where $S(\nu)$ is the edge SFR, $C(\nu)$ the CSF, $M(\nu)$ the medium MTF and Q_{edge} the integrated acutance value. This product is integrated up to a cutoff spatial frequency, ν_c. All quantities are expressed as functions of spatial frequency in units of cycles per degree. The cutoff frequency is the lowest frequency of the imager or display (or print) half-sampling frequency.

Once the acutance value has been calculated, it has to be mapped to JND values. ISO 20462-3[3] provides a method to accomplish this by mapping JND values from MTFs through the relationship

$$\text{JNDs} = \frac{17249 + 203792k - 114950k^2 - 3571075k^3}{578 - 1304k + 357372k^2} \tag{8}$$

where k is a parameter describing a set of MTFs

$$m(\nu) = \begin{cases} \frac{2}{\pi}\left(\cos^{-1}(k\nu) - k\nu\sqrt{1-(k\nu)^2}\right) & k\nu \le 1 \\ 0 & k\nu > 1 \end{cases} \tag{9}$$

A set of acutance values as functions of parameter k, $Q(k)$, are then generated by integrating the functions $m(\nu)$ multiplied by the CSF defined above. Next, an objective metric of blur, B, is defined as

$$B = \begin{cases} 0.8859 - Q_{\text{edge}} & Q_{\text{edge}} \le 0.8859 \\ 0 & Q_{\text{edge}} > 0.8859, \end{cases} \tag{10}$$

where the value 0.8859 is the acutance value above which increases in acutance are not accompanied by increases in perceived quality. In order to relate the objective blur metric to JND loss, we define

$$\text{JND quality loss} = \text{JND}_{\max} - \text{JND} \tag{11}$$

This quantity is then fitted to the blur values, yielding the relation between acutance and JND quality loss as

$$\text{Edge JND loss} = \frac{3.360 \times 10^{-3} - 2.330B + 164.1B^2 - 191.8B^3 + 16.32B^4}{1 - 0.08655B + 0.9680B^2 - 2.306B^3} \tag{12}$$

4.2 Texture blur

Similar to calculating edge acutance, the integrated texture acutance is calculated using the form in Equation 7. The $C(\nu)$ and $M(\nu)$ terms remain the same. Here, however, $S(\nu)$ is the noise-compensated texture MTF.

In order to establish the relationship to convert texture acutance into JNDs of quality, the texture acutance values of flat field patches with a series of known noise cleaning levels were compared to psychophysical ratings of photographic scenes with corresponding noise cleaning levels.[8] These scenes were rated using the ISO 20462-3 standard quality scale (SQS) JND ruler.[3] Updates to the relationship between texture acutance and JNDs of quality have been made in this paper due to upcoming post-experiment revisions of ISO 20462 SQS JND calibration values and revisions to the texture acutance calculation. A prediction model for the SQS JNDs for the texture acutance range tested is shown as

$$\text{Texture JND loss} = \begin{cases} 20.4 - 21.5Q_{\text{texture}} & Q_{\text{texture}} \leq 0.95 \\ 0 & Q_{\text{texture}} > 0.95, \end{cases} \tag{13}$$

where Q_{texture} is the texture acutance calculated using the described variant of Eq. 7.

4.3 Visual noise

The visual noise JND mapping was derived using the set of 11 quality loss calibrated ICCLAB noise images published by Keelan.[18,23] The best fit for the variance and covariance weighting factors in the objective metric versus the published quality loss values was obtained via regression analysis. An integrated hyperbolic increment function (IHIF)[24] is used to map the objective metric output to quality loss JND values:

$$\Delta Q(\Omega) = \begin{cases} \frac{\Omega - \Omega_r}{\Delta\Omega_\infty} - \frac{R_r}{\Delta\Omega_\infty^2} \ln\left(1 + \frac{\Delta\Omega_\infty(\Omega - \Omega_r)}{R_r}\right) & \Omega > \Omega_r \\ 0 & \Omega \leq \Omega_r, \end{cases} \tag{14}$$

where $\Omega_r = 0.4627$, $R_r = 0.1418$, and $\Delta\Omega_\infty = 0.02312$.

For the regression analysis a cost function was constructed from the objective metric in Eq. 1 and the IHIF.[18] The RMS errror of in terms of Quality Loss JND values was used to judge the goodness of the fit for the Levenberg-Marquardt regression algorithm. This produced the following expression for the objective visual noise metric:

$$\Omega = \log_{10}\left[1 + 23.0\sigma^2(L^*) + 4.24\sigma^2(a^*) - 5.47\sigma^2(b^*) + 4.77\sigma^2(L^*a^*)\right] \tag{15}$$

5. SUMMARY AND OUTLOOK

This paper describes the spatial metrics developed within CPIQ. As with all such metrics, there will be situations where the accuracy will be compromised to some extent, and it is certainly important to understand when those cases will occur. It should be observed that

the edge acutance metric is not an artifact metric, i.e., the effects of oversharpening such as halos around sharp edges is not taken into account. On the other hand, such effects will lead to the impression of overall higher sharpness levels which should be picked up by the acutance metric. However, the mapping to JNDs is not taking the characteristic overshoot of the MTF into account. This might produce acutance values larger than 1, the effect of which could motivate more investigation. For the texture blur metric, the subtraction of the power spectrum performed in order to suppress high-frequency artifacts might be less effective in the case of substantial noise reduction, in which case the frequency response in the uniform patch might be considerably different from the response in the dead leaves part.

The overall aim of CPIQ is to obtain a consumer-oriented image quality rating system. For this to be accomplished, all individual metrics for sharpness, texture acutance, noise, etc, must be combined into one general image quality loss metric. Any combination algorithm must take into account the fact that even if all other attributes give excellent scores, one excessively poor attribute should bring down the total score equally excessive. Keelan[24] describes one method for taking this into account through the following, simplified, relation:

$$\text{Total JND loss} = \left(\sum_n \text{JND}_n^\gamma \right)^{1/\gamma} \tag{16}$$

where JND_n is the JND loss of metric n (edge acutance, texture acutance, etc). By an appropriate choice of the exponent γ, any one excessively poor metric will dominate the total sum and thus provide an overall image quality score that better correlates with the perceived perception of image quality.

Another issue relating to the combination metric is the orthogonality of the individual metrics. The relation described in Eq. 16 requires orthogonal (independent) metrics. For metrics such as sharpness and color reproduction, orthogonality should be more or less trivial, but for metrics such as sharpness and texture blur, the situation gets more complicated. Because of the adaptive nature of most noise reduction algorithms, the relation between sharpness and texture blur may be different for different illumination levels. For high illumination, those metrics may provide very similar results, but could start to deviate considerably in low illumination situations.

The above issues are being evaluated within CPIQ and will be addressed in forthcoming investigations.

ACKNOWLEDGMENTS

The authors would like to express their thanks to the other participants of CPIQ for many interesting and helpful discussions and also the International Imaging Industry Association (I3A) for managing the CPIQ initiative within which the metrics described in this paper were developed.

REFERENCES

1. ISO 12233:2000 – Photography – Electronic still-picture cameras – Resolution and spatial frequency response measurements.
2. ISO 15739:2003 – Photography – Electronic still-picture cameras – Noise measurements.
3. ISO 20462-3:2005 – Psychophysical experimental methods for estimating image quality – Part 3: Quality ruler method.
4. F. Cao, F. Guichard and H. Hornung, "Dead leaves model for measuring texture quality of a digital camera", *Proc. SPIE* **7537**, 75370E (2010).
5. F. Cao, F. Guichard and H. Hornung, "Measuring texture sharpness of a digital camera", *Proc. SPIE* **7250**, 72500H, (2009).
6. J. S. McElvain et al., "Texture-based measurement of spatial frequency response using the dead leaves target: extensions, and applications to real camera systems", *Proc. SPIE* **7537**, 75370D (2010).
7. J. B. Phillips et al., "Correlating objective and subjective evaluation of texture appearance with applications to camera phone imaging", *Proc. SPIE* **7242**, 724207 (2009).
8. J. B. Phillips and D. Christoffel, "Validating a texture metric for camera phone images using a texture-based softcopy ruler attribute", *Proc. SPIE* **7529**, 752904 (2010).
9. Camera Phone Image Quality – Phase 3 – Acutance - Spatial Frequency Response (I3A, 2011).
10. Camera Phone Image Quality – Phase 3 – Visual Noise (I3A, 2011).
11. Camera Phone Image Quality – Phase 3 – Texture Blur Metric (I3A, 2011).
12. P. D. Burns, "Tone transfer (OECF) characteristics and spatial frequency response measurements for digital cameras and scanners", *Proc. SPIE* **5668**, 123-8 (2005).
13. D.L. Ruderman, "The statistics of natural images", *Network* **5**, 517-48 (1994).
14. X. Zhang and B. A. Wandell, "A spatial extension to CIELAB for digital color image reproduction", in *Society for Information Display Symposium Technical Digest*, 731-34 (1996).
15. G. M. Johnson and M. Fairchild, "A top down description of S-CIELAB and CIEDE2000", *Color Res. Appl.* **28**, 425-35 (2003).
16. J. Farrell et al., "Using visible SNR (vSNR) to compare the image quality of pixel binning and digital resizing", *Proc SPIE* **7537**, 75370C (2010).
17. ISO 12232:2006 – Photography – Electronic still-picture cameras – Determination of exposure index, ISO speed ratings, standard output sensitivity, and recommended exposure index.
18. B. W. Keelan, E. W. Jin and S. Prokushkin, "Development of a perceptually calibrated objective metric of noise", *Proc. SPIE* **7867**, 786708 (2011).
19. C. J. Bartleson, "Predicting graininess from granularity", *J. Photogr. Sci.* **33**, 117-26 (1985).
20. ISO 14524:2009 – Photography – Electronic still-picture cameras – Methods for measuring opto-electronic conversion function (OECF).
21. C. Koopipat et al., "Effect of ink spread and optical dot gain on the MTF of ink jet image", *J. Imaging Sci. Technol.* **46**, 321-5 (2002).
22. N. Bonnier and A. J. Lindner, "Measurement and compensation of printer modulation transfer function", *J. Electronic Imaging* **19**, 011010 (2010).
23. ICCLab noise image set with known quality loss JND values, http://www.aptina.com/ImArch.
24. B. W. Keelan, *Handbook of Image Quality: Characterization and Prediction*, Marcel-Dekker, New York, NY, 2002.