Extending Quality Metrics to Full Luminance Range Images

Tunç Ozan Aydın [a] Rafał Mantiuk[a] Hans-Peter Seidel[a]

[a]MPI Informatik, Stuhlsatzenhausweg 85 , Saarbrücken, Germany;

ABSTRACT

Many quality metrics take as input gamma corrected images and assume that pixel code values are scaled perceptually uniform. Although this is a valid assumption for darker displays operating in the luminance range typical for CRT displays (from 0.1 to 80 cd/m^2), it is no longer true for much brighter LCD displays (typically up to 500 cd/m^2), plasma displays (small regions up to 1000 cd/m^2) and HDR displays (up to 3000 cd/m^2). The distortions that are barely visible on dark displays become clearly noticeable when shown on much brighter displays. To estimate quality of images shown on bright displays, we propose a straightforward extension to the popular quality metrics, such as PSNR and SSIM, that makes them capable of handling all luminance levels visible to the human eye without altering their results for typical CRT display luminance levels. Such extended quality metrics can be used to estimate quality of high dynamic range (HDR) images as well as account for display brightness.

Keywords: Image Quality Metric, JND, CSF, HDR, SSIM, PSNR, luminance masking

1. INTRODUCTION

Most of the commonly used quality metrics do not take into account the brightness of display devices. Such metrics take as input 8-bit code values (luma* or gamma corrected pixel values) and assume that they are perceptually uniform, regardless of how bright or dark the display is. However, the visibility of distortion can increase significantly as the display gets brighter. Taking into account the effect of display brightness is especially important for the new LCD TVs, whose peak brightness (over 500 cd/m^2) exceeds five or more times the typical peak brightness of a CRT display.

Accounting for luminance effects is also important for high dynamic range (HDR) images. They store linear radiance or luminance maps, instead of 8-bit gamma-corrected code values. The difference between luminance or radiance values has little correspondence with the actual visible difference, since the eye is sensitive to luminance ratios rather than absolute luminance values, the property sometimes referred as the luminance masking. Therefore, the PSNR measure computed on luminance or radiance maps have little correspondence with the actual image quality. In this paper we explain how absolute luminance values can be converted to an approximately perceptually uniform encoding, which in turn can give meaningful quality predictions when used with the image quality metrics that operate on pixel values.

In this paper we discuss how the perceived image quality is affected by the actual luminance levels. We propose an extension to a pair of well-known quality metrics in the form of a transfer function, referred as perceptually uniform (PU) encoding. The PU encoding transforms luminance values in the range from 10^{-5} to 10^8 cd/m^2 into approximately perceptually uniform code values. The resulting code values are passed to the quality metric instead of gamma corrected RGB or luma values. The proposed PU encoding is derived from the contrast sensitivity function (CSF) that predicts detection thresholds of the human visual system for a broad range of luminance adaptation conditions.

Further author information: (Send correspondence to Tunç Ozan Aydın)
Tunç Ozan Aydın: E-mail: tunc@mpi-inf.mpg.de, Telephone: +49 681 932 5426
Rafał Mantiuk: mantiuk@mpi-inf.mpg.de, Telephone: +49 681 932 5427
Hans-Peter Seidel: hpseidel@mpi-inf.mpg.de, Telephone: +49 681 932 5400

*Luma is a new word proposed by the NTSC in 1953 to prevent confusion between the Y' component of a color signal and the traditional meaning of luminance. While luminance is the weighted sum of the linear RGB components of a color video signal, proportional to intensity, luma is the weighted sum of the non-linear $R'G'B'$ components after gamma correction has been applied, and thus is not the same as either intensity or luminance. Source: http://www.wikipedia.org/

The PU encoding is designed so that it is backward-compatible with the sRGB non-linearity within the dynamic range of a CRT display. Consequently, the quality metrics using PU encoding show similar behaviour as the original metrics for CRT displays. We test the proposed PU encoding with two widely used visual quality measures: the Peak Signal to Noise Ratio (PSNR)[1] and the more sophisticated Structured Similarity Index Metric (SSIM).[1]

2. PREVIOUS WORK

Objective visual quality metrics either model luminance masking (effect of luminance of the detection threshold) explicitly and include it in their processing, or implicitly, assuming that input code-values are "gamma-corrected" and thus perceptually linearized. The former group includes Sarnoff VDM,[2] PDM,[3] DVQ,[4] VDP,[5] HDR-VDP[6] and many other metrics that model Human Visual System (HVS). These metrics, however, due to their complexity, difficult calibration, on-going standardization effort or lack of freely available implementation, are not as popular as the latter group of metrics, which includes arithmetical and structural metrics. Two such popular metrics are peak signal-to-noise ratio (PSNR):

$$PSNR(x,y) = 20\,log_{10}\,\tfrac{D}{MSE(x,y)} \quad MSE(x,y) = \tfrac{1}{N}\sum_{i=1}^{N}(x_i - y_i)^2 \qquad (1)$$

and structural similarity index metric (SSIM):[1]

$$SSIM(x,y) = l(\mu_x,\mu_y)^\alpha\,c(\sigma_x,\sigma_y)^\beta\,s(\sigma_x,\sigma_y)^\gamma \qquad (2)$$

where x and y are pixel values in reference and distorted images, D is the dynamic range, μ and σ are the mean and standard deviations of the corresponding input images. The final quality measure $SSIM$ is a weighted combination of the luminance comparison function l, contrast comparison function c and structure comparison function s. These metrics rely on the perceptual linearity of input pixel values x_i and y_i, which should account for luminance masking. In the following sections we show that this is reasonable assumption for CRT displays, but it is less accurate for much brighter LCD displays. This is especially the case when the same "gamma" function is used for both a bright and a regular display. Finally, such metrics cannot be applied directly to HDR images.

The proposed PU encoding in conceptually similar to the DICOM Grayscale Standard Display Function,[7] but is intended to handle a larger dynamic range. The proposed encoding is an adaptation of the color space used for HDR image and video encoding[8] for quality metrics that ensures backward-compatibility with the sRGB color space.

3. DISTORTION VISIBILITY ON REGULAR AND BRIGHT DISPLAYS

The effect of luminance level on the sensitivity of the human visual system is often referred as luminance masking. Figure 1 shows the Campbell-Robson contrast sensitivity chart for two different background luminance levels. For the best viewing, the figure should be viewed on an LCD display of about $200\ cd/m^2$ and the display function close to the sRGB non-linearity. The solid lines denote the contrast sensitivity of the HVS, which is the contrast level at which the sinusoidal contrast patterns become invisible. Even though the same scales were used for both left and right plots, the CSF is shifted upwards (higher sensitivity) and right (towards higher spatial frequencies) for the brighter pattern. This shows that we are more likely to notice contrast changes, if the stimuli is brighter, as is the case of a brighter display.

But it is not clear if this observation for simple sinusoidal pattern can be assumed valid for complex images. Consequently, we cannot assume that a difference in sensitivity due to image brightness results in a difference in quality assessment. To verify this, we performed a subjective quality evaluation of distorted images shown on the displays of different brightness.

Our 16 test subjects were within the ages 23–48, all with near perfect or corrected vision. Each subject was presented a reference and distorted image side by side for 10 seconds. After that interval, a blank screen was displayed and the subjects were asked to assess the quality of the distorted image with respect to the reference on a 5 point scale, where higher values indicate better quality. Subjects were given the opportunity to view the image
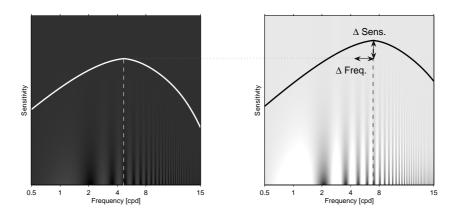
Figure 1.
Contrast sensitivity function (CSF) of the human eye in dark (left) and bright (right) viewing conditions. Arrows labelled as $\Delta Sens.$ and $\Delta Freq.$ denote the amount of difference in magnitude and frequency of the peak sensitivity between the dark and bright cases.
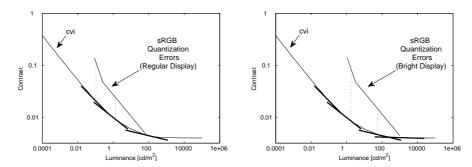

Figure 2.
Quantization errors of sRGB encoding for maximum luminance $80 cd/m^2$ (left) and $1000 cd/m^2$ (right), in comparison to contrast versus intensity (cvi) function of the human visual system. The discrepancy between the slopes of both functions is large, especially for the bright case.

pair again for additional 10 second intervals until deciding on the image quality. A set of distorted test images was generated by applying 3 types of distortions (random pixel noise, gaussian blur and JPEG compression) at 2 levels (high and low) to 3 images. Each image pair was shown on a Brightside DR-37P HDR display, which simulated either a regular (1–100 cd/m^2) or a bright display (10–1000 cd/m^2). The simulated displays had the same response as an actual LCD display (measured with a Minolta LS-100 luminance meter), only the absolute luminance levels were shifted for the bright display. The order of trials were entirely randomized and each image was shown 2 times to ensure subject reliability

Our experimental setup and grading scale is adopted from ITU-T Rec. P.910 standard.[9] We determined the mean quality value for the regular display as 3.15, and for the bright display as 2.85, indicating that subjects tend to perceive the quality of distorted images to be lower on the bright display. In other words, distortions of the same type and with the same magnitude are more annoying when the overall brightness of the image is higher. An evaluation of the data with the ANalysis Of VAriance (ANOVA) method resulted in an F-value of 20.57 and the corresponding p-value $\ll 0.05$ for the display brightness parameter, showing that the effect of display brightness to perceived quality is statistically significant.

4. WEBER-FECHNER LAW AND LUMINANCE MASKING

Figure 1 reveals that the threshold contrast $\Delta y/y$ is different for dark and bright stimuli. This is contrary to the commonly assumed Weber-Fechner law, which would require that the ratio $\Delta y/y$ stays constant. This observation is better illustrated on the contrast versus intensity (cvi) plot shown in Figure 2. The cvi function indicates the threshold contrast (y-axis) at particular luminance adaptation level (x-axis). The region where such contrast is constant, and the Weber-Fechner law holds ($\Delta y/y = const.$), can be found for luminance values greater than approximately 500 cd/m^2. For lower luminance levels the detection threshold rises significantly. This indicates that the Weber-Fechner law is in fact very inaccurate model of luminance masking for the range of luminance shown on typical displays (from about 0.1 cd/m^2 to 100-1000 cd/m^2).

5. SRGB NON-LINEARITY AND DETECTION THRESHOLDS

The compressive non-linearity (transfer function) used in the sRGB color space[10] accounts not only for the response of a typical CRT, but also partly for the drop of the HVS sensitivity for dark luminance levels. The sRGB non-linearity has the form:

$$l(L') = \begin{cases} \left(\frac{L'+0.055}{1.055}\right)^{2.4} & \text{if } L' > 0.04045 \\ \frac{L'}{12.92} & \text{otherwise} \end{cases} \tag{3}$$

where L is the trichromatic value (for simplicity we assume luminance) normalized by the peak display luminance and L' is the "gamma-corrected" luma value. In Figure 2 we plot the peak quantization errors due to 8-bit coding of L', assuming the peak display luminance of 80 cd/m^2 on the left (CRT) and 1000 cd/m^2 on the right (bright LCD or Plasma). We compute the peak quantization errors as

$$e(L') = \frac{1}{2}\frac{\max |l(L'{\pm}1) - l(L')|}{l(L')} \tag{4}$$

but plot them in the luminance domain (L), instead of luma domain (L'), to compare different displays. The slopes of the error quantization functions give closer match to the cvi function for the darker display (80 cd/m^2), suggesting that the sRGB has better perceptual uniformity for CRT displays. The slopes start to deviate much stronger for brighter displays, making perceptual uniformity of the sRGB non-linearity for LCD and Plasma displays questionable.

Another observation that we can make in Figure 2 is that the quantization errors of 8-bit code value encoding are actually larger than the detection threshold of the human eye. This means that when we display a smooth gradient on a display driven by 8-bit input, we can see contouring artifacts. This is true even for darker displays, but is more noticeable for bright displays, where the discrepancy between encoding quantization errors and the cvi gets larger. Such contouring artifacts could be easily hidden by adding random noise to the gradient (spatial or temporal dithering). For the same reason, medical displays are usually driven by signals of 10- or more bits to reduce the quantization errors to an undetectable level.

6. DETECTION THRESHOLDS IN COMPLEX IMAGES

Before we can derive a perceptually uniform encoding, we need to estimate contrast detection thresholds as a function of pixel luminance. Many aspects of complex images, such as spatial frequency, orientation and masking pattern, can significantly rise the detection threshold. Figure 3 illustrates how the sensitivity (inverse of the contrast detection thresholds) changes with spatial frequency and adapting luminance. Since the perceptually uniform encoding is a function of pixel value, we need to reduce all these factors except adapting luminance L_a, and assume that they will be taken into account by the actual quality metric. To ensure that the estimated detection threshold is always conservative, we choose the value that corresponds to the maximum sensitivity for each factor we want to reduce. Therefore, we define our cvi function as:

$$cvi(L, L_a) = \left(\max_{\mathbf{x}} [CSF(L_a, \mathbf{x})\, MA(|L - L_a|)]\right)^{-1} \tag{5}$$
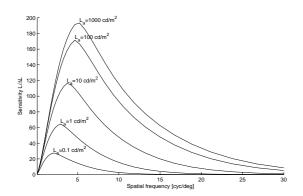
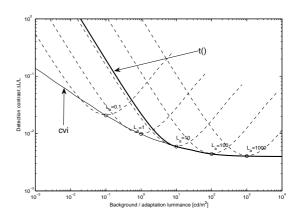Figure 3. Contrast sensitivity function variation with adaptation luminance.

Figure 4. Continous line - cvi function for different adaptation levels; Dashed lines - contrast detection thresholds for fixed adaptation and varying background luminance. Refer to the text for the description of function t.
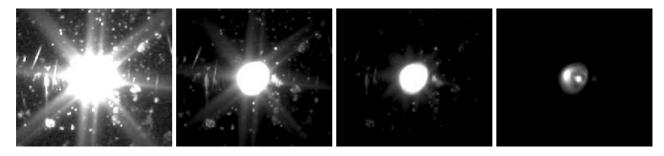

Figure 5. A specular highlight on a piece of metal captured using multi-exposure technique. The exposure time decreases from left to right. The rightmost image reveals the reflection of a lamp, which is not visible to the human eye in the actual setup.

where the CSF is the contrast sensitivity function and $\mathbf{x}$ corresponds to all the parameters (spatial frequency, orientation, stimuli size, etc.) except adapting luminance L_a and the background luminance L. The $MA()$ function estimates the loss of sensitivity due to maladaptation, as explained below. We use the CSF function from Daly's VDP,[5] as it is valid for a large range of luminance values (both photopic and scotopic viewing).

To properly utilize the cvi function, it is important to distinguish between the adapting luminance, L_a, and the background luminance, L. When viewing a complex scene the human eye can adapt locally to small regions. For example our eyes are in one state of luminance adaptation when looking outside a window on a sunny day, and in a different state when looking at the interior of a room. However, the eye is hardly ever perfectly adapted for each tiny luminance patch in a scene. For example, when looking at bright specular reflections, we usually cannot see the reflected features of a light bulb or the sun, since we are adapted to the diffuse light reflected from an object, rather than the tiny specular spot. Figure 5 shows that such tiny features are in fact reflected, but we usually don't see them. The situation when the eye is maladapted has been studied in so-called *probe-on-flash* experiments,[11] in which a threshold stimuli on a background was briefly flashed, thus bypassing the adaptation process. The typical characteristics measured in such experiments are shown in Figure 4. The plots were derived by combining the typical cvi function with an S-shaped photoreceptor response curve, as done by Irawan, et al.[12]

To make our extension spatially independent and possibly compatible with the sRGB non-linearity, we make two simplifying assumptions about the luminance adaptation process. Firstly, we assume that there is a minimum luminance level to which the eye can adapt, L_{a-min}. When viewing complex images, the darkest areas are usually affected by the glare (light scattering in the eye's optics), therefore the minimum luminance level that reaches the retina and to which the eye can adapt is elevated. Secondly, we assume that the eye is perfectly adapted for

all luminance levels above L_{a-min}, that is the adapting luminance is equal the luminance of the pixel ($L_a = L$). The second assumption results in the most conservative estimates of the contrast detection thresholds (refer to Figure 4). Our final estimates of the detection thresholds are:

$$t(L) = cvi(L, max(L, L_{a-min})). \tag{6}$$

7. PERCEPTUALLY UNIFORM ENCODING

The goal of perceptually uniform encoding is to ensure that the distortion visibility is approximately uniform along all encoded values. This is achieved when the differentials of such encoding are proportional to the detection thresholds. The easiest way to find such mapping from the detection threshold estimates t (Equation 6) is to use the following recursive formula:

$$f_i = f_{i-1}\left(1 + t(f_{i-1})\right) \quad where \quad f : L' \longrightarrow L,\ i \in [2 \cdots N] \tag{7}$$

where f_1 is the minimum luminance we want to encode ($10^{-5}\ cd/m^2$ in our case) and N is selected so that f_N is larger than the maximum luminance to be encoded ($10^{10}\ cd/m^2$). Note that $cvi(L) \cdot L$ gives an absolute detection threshold in cd/m^2. The values of f_i give the luminance value associated with particular luma value i, that is the inverse mapping from luma to luminance. To find a forward mapping function, which we denote with $PU : L \longrightarrow L'$, we use the values of f as a lookup table and find the nearest (or interpolated) index i for a given luminance value L. For a more complete mathematical formulation of this problem, refer to[8] or.[13]

Ideally, we would like our PU encoding to be backward-compatible with the sRGB non-linearity (Equation 3), meaning that it should result in similar luma values within the dynamic range of a CRT display, while still retaining perceptual uniformity. We achieve this by minimizing the squared difference between both encodings within the range $0.1 - 80\ cd/m^2$ with respect to three parameters m, s and L_{a-min}:

$$\sum_{L=0.1}^{80} \left((s\ PU(L, L_{a-min}) + m) - l^{-1}(L)\right)^2 \tag{8}$$

where the summation is performed for 256 logarithmically distributed luminance values L, $l^{-1}(L)$ is the inverse of Equation 3, and $PU(L)$ is the inverse of Equation 7. The result of such fitting together with the sRGB non-linearity is shown in Figure 7. The fit is not perfect, as the sRGB non-linearity does not fully agree with the cvi function. Note that neither of the parameters m and s affect our initial assumption since the differentials of the PU encoding are still proportional to the detection thresholds. The parameter s can be understood as the absolute sensitivity factor, which in fact varies among observers. By performing the optimization we implicitly assume the same sensitivity as the sRGB encoding. The other parameter m adjusts the absolute encoding response to fit to sRGB.

We store the resulting PU encoding as a look-up table, rather than trying to fit an analytic function. A look-up table offers better accuracy and is usually faster to compute than power or logarithmic functions approximating such encodings. We provide the look-up table and the matlab code for the PU encoding at `http://www.mpii.mpg.de/resources/hdr/fulldr_extension/`.

The data flow diagram of the extended metrics is given in Figure 6. Similar to non-extended metrics, the input is a pair of reference and distorted images. Both images are converted to display luminance values using the response function of the display on which the images are viewed. Next, the PU encoding transforms the luminance values into perceptually uniform pixel values. At the final quality assessment step, no modification on the metric part is necessary since the PU encoding merely provides perceptually uniform pixel values, which was the metric's assumption in the first place (Section 2).

8. VALIDATION: BACKWARD COMPATIBILITY

We validate the compatibility of the PU encoding with the sRGB non-linearity by comparing the extended and non-extended metric responses for a set of images viewed on a CRT display ($0.1 - 80 cd/m^2$). The test set of distorted images is generated by converting the reference images to display luminance values and applying a
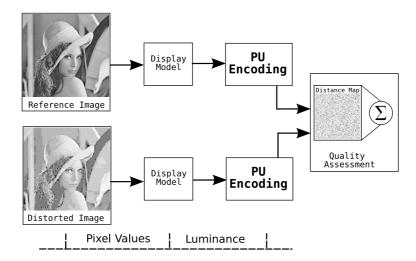
Figure 6.
Data flow diagram of the extended metrics for typical 8-bit images. Pixel values are converted to luminance and re-encoded with PU encoding before quality assessment.
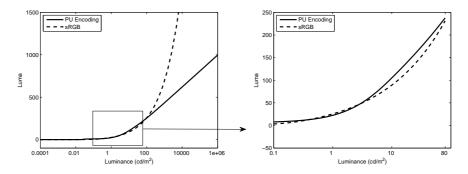

Figure 7.
The best fit of PU encoding to sRGB within the range $0.1 - 80 \ cd/m^2$ in a least squares sense. Resulting curve is shown along the entire dynamic range (left), and only within the range that is considered for optimization (right).

Figure 8.
Sample images from our validation test set. We consider random pixel noise (left), gaussian blur (center) and JPEG compression (right) as distortion types.
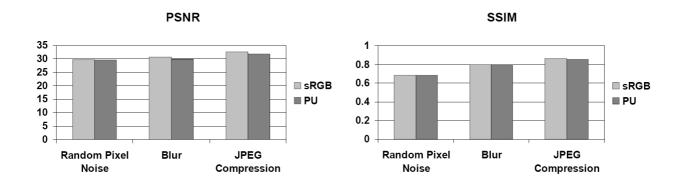
Figure 9.
Backward-compatibility with sRGB encoding. The average PSNR (left) and SSIM (right) responses of PU encoded images for different distortion types provides a good match to corresponding sRGB encoded images.
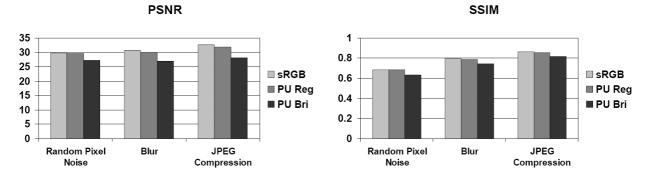

Figure 10.
Image quality on bright display. The pixel values of sRGB encoded images are the same for both regular $(1 - 100 cd/m^2)$ and bright $(10 - 1000 cd/m^2)$ displays. PU encoding successfully accounts for the effect of display brightness.

distortion which can be either of the following types: random pixel noise, gaussian blur or JPEG compression (Figure 8). Each type of distortion is applied to 3 reference images at 2 different levels. The image luminance is converted to pixel values using sRGB and PU encodings, and the quality of the distorted images in both cases are assessed by PSNR and SSIM. Figure 9 shows the average responses for both extended and non-extended metrics separately for each type of distortion. We observe that the match between the responses is not exact, since our optimization procedure does not result in a perfect fit of PU encoding to sRGB non-linearity (Figure 7). Still, the difference between extended and non-extended metric responses are quite low (< 1 dB for PSNR and < 0.01 for SSIM), indicating that they can be used interchangeably for typical CRT dynamic range if small deviations in metric responses are acceptable.

9. QUALITY ASSESSMENT FOR BRIGHT DISPLAYS

The subjective experiment in Section 3 revealed that distortions of the same type and magnitude appear more annoying on a bright display than a regular one. In this section we show that the extended metrics can correctly predict this effect, while non-extended metrics fail to do so. In parallel with the subjective study, we simulate the brightness of an LDR image on two hypothetical displays: a regular $(1 - 100 \ cd/m^2)$ and a bright display $(10 - 1000 \ cd/m^2)$, both with the same dynamic range $(1 : 100)$. The resulting luminance values from both display models are transformed to perceptually uniform pixel values with the proposed encoding.

In Figure 10, we compare the metric predictions for sRGB encoded images side by side with the extended metric responses for both display models. Note that the pixel values generated by sRGB non-linearity are exactly
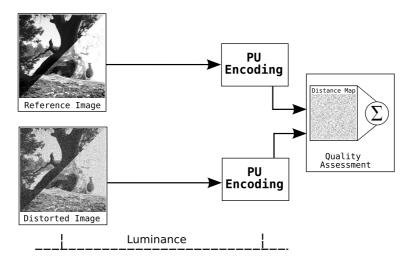
Figure 11.
Data flow diagram of the extended metrics for HDR images. HDR images are scene referred and store luminance, which is directly converted to percepually uniform pixel values by the PU encoding

the same for both displays, and consequently the quality estimates are also the same. On the other hand, quality of the PU encoded images viewed on the bright display are noticeably lower than the quality of the same images viewed on the regular display, in agreement with the outcome of our subjective experiment.

10. QUALITY ASSESSMENT OF HDR IMAGES

Unlike 8-bit images that store gamma corrected code values tailored towards particular display devices, the content of an HDR image is related to the actual photometric characteristics of the scene it depicts, which in turn directly correspond to physical luminance. In order to get meaningful responses from PSNR and SSIM when comparing a pair of HDR images, physical luminance of both images need to be converted to perceptually uniform pixel values. The use of sRGB encoding for HDR images brings in an ambiguity in the choice of the white point value. The straightforward approach of setting the white point to the maximum luminance of the image generally results in suppression of details in dark image regions. Instead, the logarithmic function is a simple and often used approximation of the HVS response along the entire visible luminance range. Although logarithmic encoding adheres to the Weber-Fechner law (Section 4), it provides a very coarse approximation and does not predict the loss of sensitivity for the low light conditions. These shortcomings can be avoided by employing the PU encoding to generate perceptually uniform pixel values for HDR images. The data flow of the extended "HDR metrics" is shown in Figure 11. Since HDR images already contain physical luminance information, the use of a diplay model is not necessary.

11. CONCLUSION

We proposed an extension to two popular image quality metrics, namely PSNR and SSIM, that makes them capable of handling all luminance levels visible to the human eye, without altering their response at the dynamic range of a typical CRT display. Our extension consists of transforming image luminance to perceptually uniform pixel values, that are optimized to fit gamma correction non-linearity within the range from 0.1 to 80 cd/m^2 in a least squares sense. The proposed extension does not impose any changes on the quality metric part. Another consequence of this modularity is that it can potentially be applied to any quality metric that takes gamma corrected pixel values as input.

In the future, we would like to validate the metric responses for HDR images through subjective experiments. We are also interested in exploring the application of our extension to other quality metrics.

REFERENCES

1. Z. Wang and A. C. Bovik, *Modern Image Quality Assessment*, Morgan & Claypool Publishers, 2006.
2. J. Lubin, "A visual discrimination model for imaging system design and evaluation," *Vision Models for Target Detection and Recognition* , pp. 245–283, 1995.
3. S. Winkler, *Digital Video Quality: Vision Models and Metrics*, John Wiley & Sons, Ltd, West Sussex, England, 2005.
4. A. Watson, J. Hu, and J. McGowan III, "Digital video quality metric based on human vision," *Journal of Electronic Imaging* **10**, p. 20, 2001.
5. S. Daly, "The visible differences predictor: An algorithm for the assessment of image fidelity," in *Digital Images and Human Vision*, A. B. Watson, ed., pp. 179–206, MIT Press, 1993.
6. R. Mantiuk, S. Daly, K. Myszkowski, and H.-P. Seidel, "Predicting visible differences in high dynamic range images - model and its calibration," in *Human Vision and Electronic Imaging X, Proc. of SPIE*, **5666**, pp. 204–214, 2005.
7. "Part 14: Grayscale standard display function," in *Digital Imaging and Communications in Medicine (DI-COM)*, 2001.
8. R. Mantiuk, K. Myszkowski, and H.-P. Seidel, "Lossy compression of high dynamic range images and video," in *Proc. of SPIE Human Vision and Electronic Imaging XI*, p. 60570V, 2006.
9. ITU-T, "Subjective video quality assessment methods for multimedia applications," 1999.
10. IEC 61966-2-1:1999, *Multimedia systems and equipment - Colour measurement and management - Part 2-1: Colour management - Default RGB colour space - sRGB*, International Electrotechnical Commission, 1999.
11. J. Walraven, C. Enroth-Cugell, D. Hood, D. MacLeod, and J. Schnapf, "The control of visual sensitivity: receptoral and postreceptoral processes," *Visual Perception: The Neurophysiological Foundations* , pp. 53–101, 1990.
12. P. Irawan, J. A. Ferwerda, and S. R. Marschner, "Perceptually based tone mapping of high dynamic range image streams," in *Proceedings of the Eurographics Symposium on Rendering*, pp. 231–242, 2005.
13. R. Mantiuk, G. Krawczyk, K. Myszkowski, and H.-P. Seidel, "Perception-motivated high dynamic range video encoding," *ACM Transactions on Graphics* **23**(3), pp. 730–738, 2004.