DOI: 10.2352/ISSN.2470-1173.2016.13.IQSP-012

Optimized Tone Curve for In-Camera Image Processing

Praveen Cyriac, David Kane, Marcelo Bertalmío; Dept. of Information and Communication Technologies, Universitat Pompeu Fabra, Barcelona, Spain

Abstract

The sensor values captured by a digital camera are transformed in a non-linear manner prior to quantization in order to make the quantization rate approximately proportional to the sensitivity of the human visual system. We propose an image dependent non-linear transformation that can accurately reproduce the detail and contrast visible in the original scene. The principles underpinning the transform stem from an understanding of natural image statistics, as well as recent experimental and neurophysiological findings. To optimize the parameters of the model we collect user-feedback and develop a method that can predict the user defined parameters. The method we have developed has an extremely low computational complexity, therefore it operates almost instantaneously making it suitable for in-camera operations. The final image looks natural, without any halos, spurious colors or artifacts. It can also be applied to video sequences, after imposing temporal coherence on the parameter values by smoothing them over time. The proposed approach is validated through psychophysical tests that confirm that it outperforms other state of the art algorithms in terms of users' preference.

Introduction

The values captured by a digital camera sensor span between 3 and 4 orders of magnitude, and they must be transformed in-camera in a non-linear manner, prior to quantization, with two main purposes: to make the quantization rate approximately proportional to the sensitivity of the human visual system, and at the same time making the transformed signal suitable for the reduced dynamic range of standard displays.

The non-linearity applied in most digital cameras is well approximated by a simple power law, and while this may perform well on average, in general it is suboptimal. The literature provides a number of image-dependent non-linearities, some based on models of the human visual system, see [9][6] and references therein. Additionally, there are more complex, local tone mapping algorithms that could, in theory, transform the sensor data, but these tend to be computationally more expensive, sensitive to image fluctuations, and thus harder to incorporate in the image-processing pipeline of a camera. The aim of this work is to introduce a tone-curve non-linearity with the following properties:

- It is automatic (no need for user-selected parameters) and fast, such that it can work in-camera.
- Accurately reproduces the detail and contrast visible in the original scene.
- Produces no visible artifacts nor color distortions.
- Works on video sequences and results do not show temporal artifacts.

In order to fulfill this goal we propose a two stage non-linear transform that extends the method of Cyriac et al. [7], which is well suited to the statistics of natural scenes, and is in keeping with new psychophysical findings and neurophysical data. Our two main contributions are:

1. The modification of the model, incorporating a novel term performing adaptive clipping. In the original model one percentile of the highest input values are clipped prior to the transform to improve the global contrast of the image; however, this simple clipping is not suitable for all situations. The proposed modification provides a better global contrast and works reliably in all tested scenarios.
2. Changing the way the model parameters are estimated: the new functions that automatically select the parameter values have been designed so as to correlate well with user choices performed in psychophysical experiments.

The method we have developed has an extremely low computational complexity, therefore it operates almost instantaneously making it suitable for in-camera operations. The final image looks natural, without any halos, spurious colors or artifacts. It can also be applied to video sequences, after imposing temporal coherence on the parameter values by smoothing them over time. The proposed approach is validated through psychophysical tests that confirm that it outperforms other state of the art algorithms in terms of users' preference. Aside for in-camera implementation, our proposed method can be used as an off-line tone mapping method for converting high dynamic range (HDR) images into low dynamic range (LDR) ones, with applications to cinema shoots (on-set use of LDR monitors with an HDR camera), cinema post-production (color grading), television broadcast (making HDR signals compatible with LDR equipment), and rendering in computer graphics (for videogames, 3D animation, the integration of CGI onto real footage, etc.)

Original approach

In this section we briefly explain the two-stage method proposed by Cyriac et al.[7] for the non-linear transformation step in the camera image processing pipeline.

The first stage is a global transform based on the psychophysical study [12] demonstrating that subjects tend to prefer the images with a flat lightness histogram. Complete histogram equalization achieved by an intensity transform based on the cumulative histogram may produce images with an unnatural appearance due to sharp changes in contrast. Thus, Cyriac et al. [7] perform a constrained histogram equalization by modeling the cumulative histogram as a smooth function based on natural image statistics. Studies in natural image statistics [16, 11] reported that, on average, natural images have a triangular shaped histogram in log-log coordinates. Thus, the cumulative histogram (H) can be modeled as a piecewise linear function with two slopes (γ_L and

DOI: 10.2352/ISSN.2470-1173.2016.13.IQSP-012

γ_H) in log-log coordinates; therefore, in linear-linear coordinates H has the form

$$H(I) = I^{\gamma(I)}. \tag{1}$$

where I is the normalized image and γ is a slope function with $\gamma \simeq \gamma_L$ for small intensities, $\gamma \simeq \gamma_H$ for large intensities and a smooth transition at M_{lin} with a slope n as follows:

$$\gamma(I) = \gamma_H + (\gamma_L - \gamma_H)\Big(1 - \frac{I^n}{I^n + M_{lin}^n}\Big), \tag{2}$$

where the parameter values γ_L, γ_H and M_{lin} are computed from the cumulative histogram of the input intensity image.

So the first stage, that performs constrained histogram equalization based on natural image statistics, is given by:

$$I_1(x) = I(x)^{\gamma(I(x))} \tag{3}$$

The second stage performs contrast normalization based on the neurophysiological evidence [4, 5] which explains that the visual system performs normalization of the contrast by a factor depending on the standard deviation of the light intensity. The second stage is given by:

$$O(x) = \mu(x) + (I_1(x) - \mu(x)) * k/\sigma, \tag{4}$$

where x is the pixel position, I_1 is the output of the previous stage, μ and σ are the local mean and global standard deviation of I_1 respectively, k is a constant and $O(x)$ is the final output of our method.

Proposed modification

We propose two modifications to the first stage of Cyriac et al. [7]: incorporating an adaptive clipping term, and designing the automated parameter estimation method so as to correlate well with the results of psychophysical experiments.

We added an adaptive clipping term C to the first stage of [7] to overcome the limitation of one percentile clipping of highest input values and to preserve the global contrast of the image. The new first stage is given by:

$$I_1 = (I(x))^{\gamma(I(x))} C(I(x)) \tag{5}$$

where C is also a smooth curve with $C \simeq C_L$ for small intensities, $C \simeq C_H$ for large intensities with a smooth transition at M_{Lin} with a slope m as follows:

$$C(I) = C_L + (C_H - C_L)\Big(\frac{I^m}{I^m + M_{lin}^m}\Big), \tag{6}$$

where $C_L$and C_H are also computed from the cumulative histogram of the input intensity image and the approach is explained in the implementation section.

The parameter values of the function γ in Eq. 2 are also computed from the cumulative histogram of the intensity image. But we modify the automated parameter estimation method such that the new parameter values match as good as possible with the users' choice in the psychophysical experiment one described in Appendix B. The detailed explained of the approach is in the following section and in Appendix A.

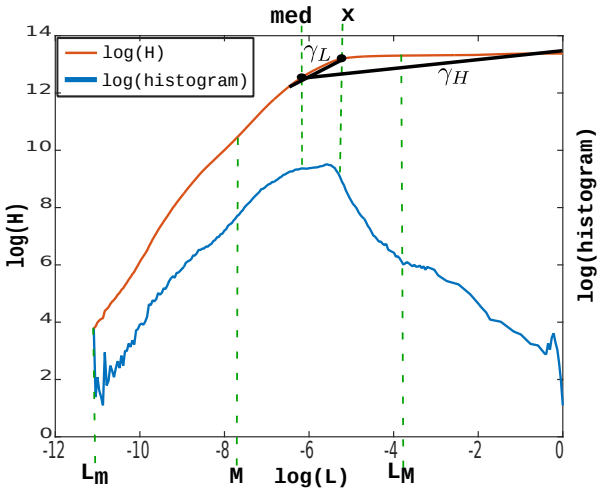

Figure 1. *Example of a cumulative histogram and histogram for a single natural image (in log-log axes) and our estimated parameters γ_L, γ_H and M.*

Implementation

In this section we present the implementation details. Our method consists of two stages, one global operation Eq. 5 and a local operation Eq. 4, and are applied separately to each of the red, green and blue color channels of the image. The input image is initially normalized to the [0,1] range and the luminance component L is computed to estimate the parameters of the first stage.

The non-linear function $\gamma(I)$ in Eq. 2 is computed using the parameter values of γ_H, γ_L and M_{lin}. The parameter values are estimated from the cumulative histogram (H) in log-log coordinates(see Figure 1). We compute M_{lin} as the exponential of M which is the average value of L_m and L_M. Where L_m and L_M are the values on the horizontal axis (log luminance) corresponding to 1 and 90% in the vertical axis (log cumulative histogram).

To compute γ_H and γ_L we distinguish 3 cases depending on the shape of the histogram.

Case 1: The histogram of the input image has roughly a triangular shape in log-log coordinates(see Figure 1). This is the general case and the majority of images fall in this category. The values of γ_H and γ_L are estimated with respect to $med = log(median(L))$ and $x = log(\sqrt{median(L) * trm(L)})$ respectively, where $trm(L)$ is the $mean(L)$ after clipping 1% of the extreme values.

Case 2: The histogram of the input image has a bi-modal distribution in log-log coordinates(see Figure 2). Our estimate of γ_H and γ_L are with respect to the log of *median* of the upper half of L (denoted by y) and $temp = log(mean(L))$ respectively.

Case 3: The bin of the histogram near to the *median* is overpopulated, resulting in a spike in the histogram (see Figure 3). we estimate γ_H and γ_L with respect to v, the intensity value just lower than those high frequent intensities.

More detailed explanations and an algorithm for the parameter estimation are given in the Appendix A. In all the cases, we set the slope value n equal to γ_L.

DOI: 10.2352/ISSN.2470-1173.2016.13.IQSP-012

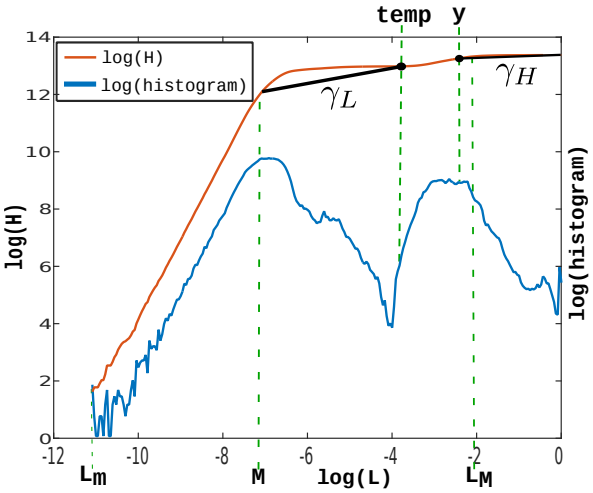

Figure 2. *Example of a cumulative histogram and histogram for a single natural image (in log-log axes) with bi-modal distribution and our estimated parameters γ_L, γ_H and M.*

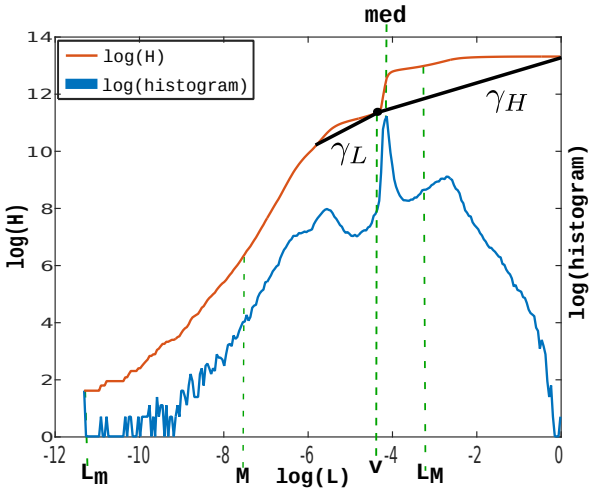

Figure 3. *Example of a cumulative histogram and histogram with a spike for a single natural image (in log-log axes) and our estimated parameters γ_L, γ_H and M.*

The clipping term $C(I)$ in Eq. 6 is computed by estimating its parameters (C_L and C_H) as follows:

$$\begin{aligned} C_L &= e^{a-\gamma_L Low_x}; \qquad a = log\Big(\frac{1}{255}\Big) \\ C_H &= e^{b-\gamma_H High_x}; \qquad b = log\Big(1-\frac{1}{255}\Big) \end{aligned} \tag{7}$$

where Low_x and $High_x$ correspond to 0.39 and 99.6 percentile of L in log domain.

For the second stage, Eq. 4, the parameters are the same as in Cyriac et al [7]. The local mean $\mu(x)$ is computed by the convolution of a kernel W with image I_1, where W is generated by the linear combination of two Gaussian kernels with standard deviations $\sigma = 5$ and $\sigma = 25$ and the kernels are weighted by 0.9 and 0.1 respectively. The constant k determines the contrast level of the image, larger k values give images with higher contrast. We set the value of $k = 0.33$ that produces final images that have a natural appearance and good contrast.

The proposed approach can also be applied to video if we incorporate temporal coherence in the parameter estimation. Thus, for the first stage we follow a two pass approach. In the first pass, parameters are estimated separately for each frame and then a temporal low pass filter is applied to get new parameters. In the second pass, we apply Eq. 5 using the new parameters.

Results and Discussion

In this section we first show the potential of the proposed approach to be used as a method for the in-camera non-linear processing and also to be used as a tone mapping operator. Then we validate our approach through psychophysical tests in 2 experiments, whose setup is detailed in Appendix B.

In Figure 5 we illustrate the advantage of our method over the conventional non-linearity applied in a camera imaging pipeline. Three sample images each from a consumer camera, smart phone and cinema camera, are shown along with the results of applying the proposed method on the corresponding RAW sensor values. Our results look natural in appearance, with enhanced overall contrast and without any visual artifacts.

In Figure 6 we show some results to illustrate that the proposed approach can also be used as a tone mapping operator and applied offline to HDR images. Our method when applied to HDR video sequences from the ARRI dataset [10] produces results that are natural looking with no visible flicker and without any sort of spatiotemporal artifacts.

Now we discuss the results of experiment one, in which subjects adjust the parameter values γ_L and γ_H using sliding bars in order to get a pleasing image. Figure 7 shows that for both parameters γ_H and γ_L there is a strong correlation, correlation coefficient $R = 0.83$ and $R = 0.88$, between the automatically estimated and user chosen parameters. This result indicates that the parameters estimated by the proposed approach match well with the users' choice.

In Figure 4 we show the result of experiment two, in which the subjects evaluated three tone mapping operators: Mantiuk et al. [13], Cyriac et al. [7] and the proposed approach. We choose these tone mapping operators as they were the ones performing best in [7]. The experiment shows that the subjects have a preference for the results of the proposed approach over the other operators. This contradict the prediction of the tone mapping metrics

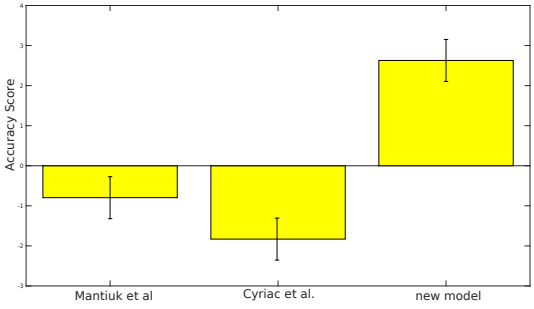

Figure 4. *Result of experiment two. Pairwise comparison of 3 tone mapping operators: Mantiuk et al. [13], Cyriac et al. [7] and proposed approach.*

DOI: 10.2352/ISSN.2470-1173.2016.13.IQSP-012

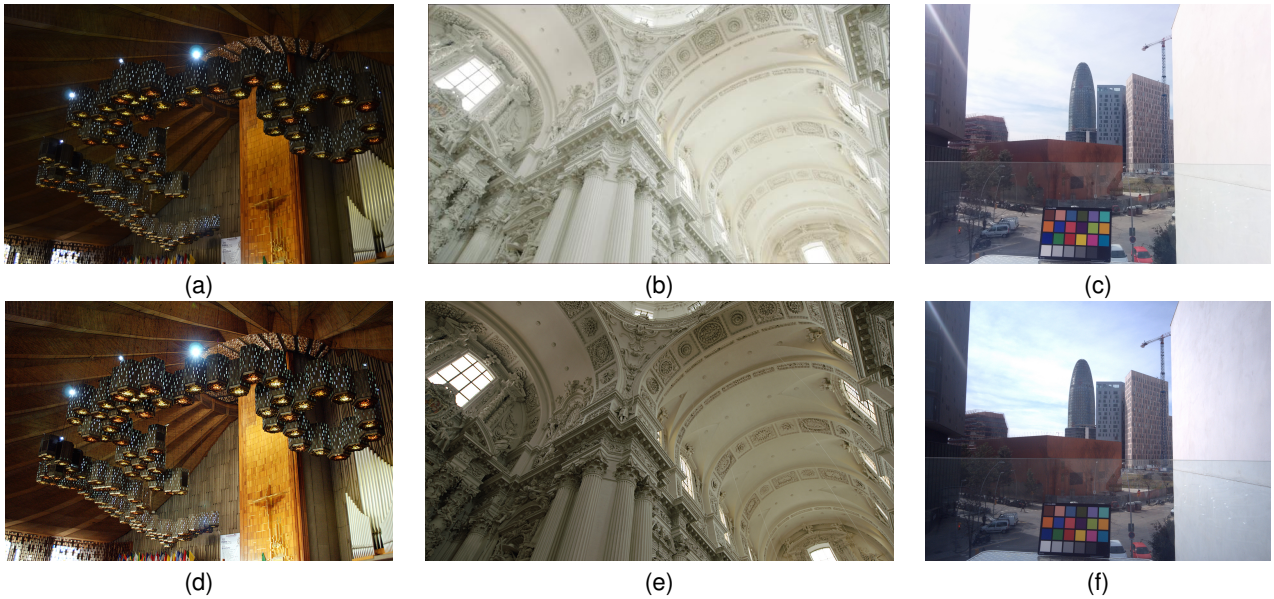

Figure 5. *Top row: original JPEG images as recorded by the camera, with the exception of image (b) which is generated by applying a S-shaped curve to a LogC image. Bottom row: results of applying our method to the corresponding RAW images. Camera models: left column, Nikon D3100 consumer photography camera; middle column, ARRI Alexa digital cinema camera [1]; right column, Nexus 5 smartphone camera.*

Figure 6. *Results of our method applied to HDR images from the Fairchild dataset (top row) [8] and to video frames from the ARRI dataset (bottom row)[10].*

DOI: 10.2352/ISSN.2470-1173.2016.13.IQSP-012

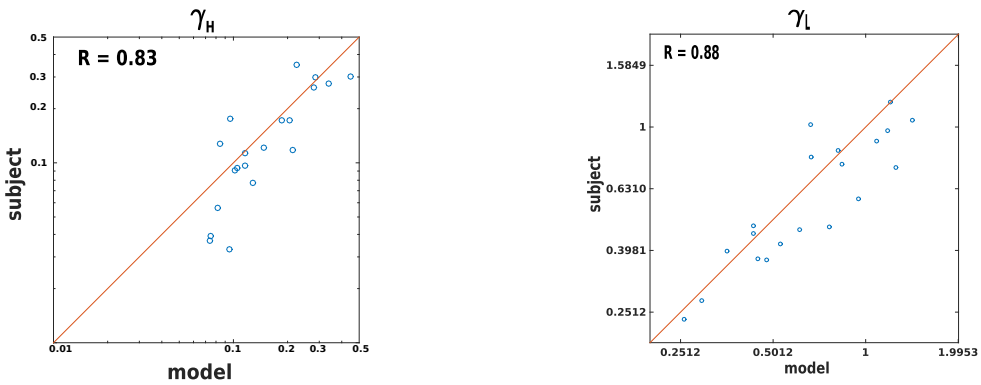

Figure 7. *Average subject selected values from 7 observers plotted against model estimated parameter values. Left: plot for parameter γ_H. Right: plot for parameter γ_L.*

[2] and [18], and this disparity puts into question the validity of visibility metrics to estimate subjective image quality.

Conclusion and Future work

We have presented an image dependent tone curve for in-camera image processing that operates in real time and can accurately reproduce the detail and contrast visible in the original scene. We optimized the automated parameter estimation method based on the users' parameter choice. The proposed approach can also be used off-line as a tone mapping operator to convert HDR images and video sequences into LDR ones. The results look natural, without halos, flicker or spatiotemporal artifacts of any kind. We validate our method through psychophysical tests that confirm that proposed approach outperforms other state of the art algorithms in terms of users' preference.

In future work we would take into account the effect of surround and viewing conditions and modify the proposed method to produce reliable results in all scenarios.

Acknowledgment

A patent application based on the research in this article has been filed at the European patent office, Application no 15154172.9-1906. We would like to thank Syed Waqas Zamir for sharing the code of the pairwise comparison experiment. This work was supported by the European Research Council, Starting Grant ref. 306337, by the Spanish government, grant ref. TIN2012-38112, and by the Icrea Academia Award.

Appendix A

The automated parameter estimation method is given in Algorithm 1 and explained below.

The approach relies on the cumulative histogram (H) of the normalized luminance channel (L) in log-log coordinates. Initially, γ_H and γ_L are computed by step 1 of the Algorithm 1, as illustrated in Figure 1. γ_H is the slope of the line joining the point in H that corresponds to $med = log(median(L))$ and the top end of H and γ_L is the slope of the line joining the point in H corresponding to $x = log((sqrt(median(L) * trm(L))))$ and the point in H that is 1 unit lower than the above point in vertical axis.

If the condition in step 2 of the Algorithm 1 is satisfied, i.e, γ_H is less than the slope of the line joining the point in H corresponding to $temp = log(mean(L))$ and the top end of H, then γ_H and γ_L are computed as in step 2.1 and 2.2 respectively. This condition is an indication of bi-modal distribution of the luminance histogram. Then a new *median* is computed by excluding all the luminance value below original *median* and γ_H is recomputed. γ_L is recomputed as the slope of line from point in H corresponding to $log(mean(L))$ and 1 unit lower than the above point in vertical axis.

But if the condition in step 3 of the Algorithm 1 is satisfied, i.e, the slope of H is very high in some region around the $median(L)$ due to high concentration of some luminance values, then γ_H and γ_L are computed as in step 3.1 and 3.2 respectively and explained in what follows. Select a value $v \in log(L)$ that satisfies the condition in step 3 and is within unit distance from *med*, then γ_H is the slope of the line joining the point in H corresponding to $v - \delta$ and the top end of H and γ_L is the slope of the line joining the same point and the point in H that is 1 unit lower than the above point in vertical axis. We used $Te = 4$ and $\delta = 0.1$.

Algorithm 1 Parameter estimation

Input: Image I
Result: γ_L and γ_H
L: Luminance channel of I normalized to [0,1].
$H(L)$: Cumulative histogram of L in log-log domain.
$\mathbf{S}(a,b)$: Slope of line joining point on H corresponding to a and b; where a, b $\in log(L)$
$trm(L)$: Mean of L after clipping 1% of extreme values.

1. $\gamma_H = \mathbf{S}\Big(med, 0\Big)$; $\quad med = log(median(L))$
$\gamma_L = \mathbf{S}\Big(x, H^{-1}\big(H(x) - 1\big)\Big)$; $\quad x = log(\sqrt{median(L) * trm(L)})$

2. If $\gamma_H < \mathbf{S}\Big(temp, 0\Big)$ then $\quad temp = log(mean(L))$

2.1. $\gamma_H = \mathbf{S}\Big(y, 0\Big); y = log\Big(median\big(L > median(L)\big)\Big)$

2.2. $\gamma_L = \mathbf{S}\Big(temp, H^{-1}\big(H(temp) - 1\big)\Big)$

3. If $\frac{dH}{d(log(L))} > Te$, for some $v \in log(L)$ and within unit distance from *med*, then

3.1. $\gamma_H = \mathbf{S}\Big(v - \delta, 0\Big)$

3.2. $\gamma_L = \mathbf{S}\Big(v - \delta, H^{-1}\big(H(v - \delta) - 1\big)\Big)$

DOI: 10.2352/ISSN.2470-1173.2016.13.IQSP-012

Appendix B

We explain the details of the psychophysical experiments in this section.

Subjects

Seven subjects completed both experiments. All had corrected to normal vision. Two are authors of the paper. Ethics was approved by the Comité Etico de Investigacion Clinica, Parc de Salut MAR, Barcelona, Spain and all procedures complied with the declaration of Helsinki.

Apparatus

Both experiments were conducted on an ASUS VS197D LCD monitor set to 'sRGB' mode with a luminance range from $0.1 cdm^{-2}$ to $106 cdm^{-2}$, with spatial and temporal resolutions of 1366 by 768 pixels and $50 \sim 75$ Hz. The display was viewed at a distance of approximately 70 cm so that 40 pixels subtended 1 degree of visual angle. The full display subtended 33 by 18 degrees. The decoding nonlinearity of the monitor was recorded using a Konica Minolta LS 100 photometer and was found to be closely approximated by a gamma function with an exponent of 2.2. Stimuli were generated under Ubuntu 12.04 LTS running MATLAB (MathWorks) with functions from the Psychtoolbox[3, 15].

The experiment was conducted in an office environment and the ambient luminance levels was recorded with a Sinometer LX1010B which could record the incoming light from at 180 degree angle. The results indicated the average ambient illumination was 147 lux. The surround luminance of the display as measured by the photometer was 65 cdm^{-2}.

Stimuli

Experiment one: 20 base images were taken from the high dynamic range survey by Mark Fairchild [8] including indoor and outdoor scenes, night time images and landscapes. Images were resized to a quarter of the original areas using Matlabs `imresize` and the setting `nearest` which performs a simple subsampling of pixel value. Each image then covered approximately 80% of the viewing area. The remaining area was presented with a surround luminance of 65 cdm^{-2} corresponding to the average luminance of the surrounding area.

Experiment two: 30 base images were taken from the high dynamic range survey by Mark Fairchild [8], excluding the 20 used in experiment one. Tone mapped versions of the original HDR radiance maps were produced according to the original [7] and the proposed approach and the approach of Mantiuk et al. [13] created using `pfstools` [14] using the parameters of 'lcd_office' display type except for ambient illumination which is set to 147 lux. The images were viewed side be side and each image was presented over 14 and 10 degrees.

In both experiments the images were presented without correcting for the decoding nonlinearity of the monitor.

Procedure

Experiment one: Subjects manipulated the parameter γ_H and γ_L of Eq. 2 via two scroll bars. The subject interacted with the scroll bar via a mouse and a press of the space key initiated the next image. The range of values corresponding to the scroll bar location was fixed at 0 to 1 for γ_H and 0 to 3 for γ_L. The initial position of the scroll bar was the midway point and the scaling was linear. The parameter values dynamically and in real time, updated the displayed image by applying Eq. 5 to the original HDR image. The subject was asked to manipulate the scroll bars until the most pleasing image was achieved. No further instruction was provided. Aside from the two authors the participants did not know what the action of the scroll bars was. Subjects had unlimited time to choose the appropriate values.

Experiment two: Subjects were asked to select which of the two simultaneously presented images they preferred. Given the 30 base images and 3 tone-mapped versions for each image, the total number of comparisons was 30*3=90. Subjects had unlimited time to make the comparison. Thurstone's law of comparative judgments [17] was used to obtain the preferred image.

References

[1] Stefano Andriani, Harald Brendel, Tamara Seybold, and Joseph Goldstone. Beyond the kodak image set: A new reference set of color image sequences. In *ICIP*, pages 2289–2293, 2013.

[2] Tunç Ozan Aydin, RafałMantiuk, Karol Myszkowski, and Hans-Peter Seidel. Dynamic range independent image quality assessment. In *ACM SIGGRAPH 2008 Papers*, SIGGRAPH '08, pages 69:1–69:10, 2008.

[3] David H Brainard. The psychophysics toolbox. *Spatial vision*, 10:433–436, 1997.

[4] Naama Brenner, William Bialek, and Rob de Ruyter Van Steveninck. Adaptive rescaling maximizes information transmission. *Neuron*, 26(3):695–702, 2000.

[5] Matteo Carandini and David J Heeger. Normalization as a canonical neural computation. *Nature Reviews Neuroscience*, 13(1):51–62, 2012.

[6] Praveen Cyriac, Marcelo Bertalmio, David Kane, and Javier Vazquez-Corral. A tone mapping operator based on neural and psychophysical models of visual perception. In *SPIE/IS&T 2015 Human vision and electronic imaging XX*, 2015.

[7] Praveen Cyriac, David Kane, and Marcelo Bertalmio. Perceptual dynamic range for in-camera image processing. In *BMVC15*, 2015.

[8] Mark D Fairchild. The hdr photographic survey. In *Color and Imaging Conference*, volume 2007, pages 233–238. Society for Imaging Science and Technology, 2007.

[9] Sira Ferradans, Marcelo Bertalmio, Edoardo Provenzi, and Vicent Caselles. An analysis of visual adaptation and contrast perception for tone mapping. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 33(10):2002–2012, 2011.

[10] Jan Froehlich, Stefan Grandinetti, Bernd Eberhardt, Simon Walter, Andreas Schilling, and Harald Brendel. Creating cinematic wide gamut hdr-video for the evaluation of tone mapping operators and hdr-displays. 2014.

[11] Jinggang Huang and David Mumford. Statistics of natural images and models. In *Computer Vision and Pattern Recognition, 1999. IEEE Computer Society Conference on.*, volume 1. IEEE, 1999.

[12] David Kane and Marcelo Bertalmío. Is there a preference for linearity when viewing natural images? In *SPIE/IS&T 2015 Image Quality and System Performance XX*, 2015.

[13] RafałMantiuk, Scott Daly, and Louis Kerofsky. Display

DOI: 10.2352/ISSN.2470-1173.2016.13.IQSP-012

adaptive tone mapping. *ACM Trans. Graph.*, 27(3):68:1–68:10, August 2008.

[14] Rafał Mantiuk, Grzegorz Krawczyk, Radosław Mantiuk, and Hans-Peter Seidel. High dynamic range imaging pipeline: Perception-motivated representation of visual content. In *Human Vision and Electronic Imaging XII*, volume 6492. SPIE, February 2007.

[15] Denis G Pelli. The videotoolbox software for visual psychophysics: Transforming numbers into movies. *Spatial vision*, 10(4):437–442, 1997.

[16] Daniel L Ruderman. The statistics of natural images. *Network: computation in neural systems*, 5(4):517–548, 1994.

[17] Louis L Thurstone. A law of comparative judgment. *Psychological review*, 34(4):273, 1927.

[18] H. Yeganeh and Zhou Wang. Objective quality assessment of tone-mapped images. *Image Processing, IEEE Transactions on*, 22(2):657–667, Feb 2013.

Author Biography

Praveen Cyriac *received the Bachelors in computer engineering form Kannur University, India (2009) and the Masters in image processing from the University of Kerala, India (2011). Currently, he is a Ph.D.student in the image processing for enhanced cinematography (IP4EC) group at the Universitat Pompeu Fabra, Barcelona, Spain. His research focus in developing algorithms to map HDR images into lower dynamic range displays by taking into account the properties of human visual system.*

David Kane *received his PhD from University Collage London where he studied under Steven Dakin and Peter Bex and learnt the basics of low-level psychophysics. From there David chose to apply his skills at the intersection of perception and technology and took a position working under Professor Marty Banks at The University of California Berkeley investigating issues surrounding the development of stereo 3D displays. Currently, David is working as in-house experimental psychologist in the image processing lab Marcelo Bertamlío at Universitat Pompeu Fabra.*

Marcelo Bertalmío *received the Ph.D. degree in electrical and computer engineering from the University of Minnesota in 2001. He is an Associate Professor at University Pompeu Fabra in Barcelona, Spain. His interests are Image Processing and Computer Vision for digital cinema applications, although he prefers the (analog) films of Joseph L. Mankiewicz and Ernest Lubitsch.*