Brightness Adjustment for HDR and Tone Mapped Images

Grzegorz Krawczyk
MPI Informatik
Stuhlsatzenhausweg 85
66123 Saarbrücken, Germany

Rafał Mantiuk
MPI Informatik
Stuhlsatzenhausweg 85
66123 Saarbrücken, Germany

Dorota Zdrojewska
Szczecin University of Technology
Żołnierska 49
71-210 Szczecin, Poland

Hans-Peter Seidel
MPI Informatik
Stuhlsatzenhausweg 85
66123 Saarbrücken, Germany

Abstract

Both High Dynamic Range images and their tone mapped correspondents contain relative luminance values which have to be mapped on a scale of available gray levels of a display. Such mapping includes brightness adjustment, which has a direct impact on the final image appearance and the observers' assessment of image quality. We conduct a psychophysical experiment in which subjects adjust image brightness to match their preference. We observe that the brightness choice is consistent across subjects and is primarily affected by image content. We investigate popular methods for automatic brightness adjustment and show a significant inaccuracy for a group of images. The incorrect brightness adjustment degrades in these cases perceived image quality. We identify characteristics of images that are highly correlated with the subjects' choice of brightness and develop an improved model for the brightness adjustment.

1. Introduction

Digital images are almost never displayed at the same brightness levels as the original scenes. This is because output media, regardless whether it is a display or a print, can produce only limited range of brightness levels. Therefore brightness adjustment is a necessary step when visualizing images. This is especially apparent for high dynamic range (HDR) images, whose range of luminance values often exceeds that of a display. However, the brightness adjustment is equally important when visualizing standard low dynamic range (LDR) contents, particularly when their contrast range is lower than offered by a display device. Furthermore, the adjustment of brightness is in fact not only a matter of choosing how bright an image should look, but often involves decisions about whether at the cost of better visibility of certain areas, other areas should loose details by clipping the brightest and the darkest pixels. Consequently, the brightness adjustment has a high impact on the image appearance and communicated information, and directly influences observers in their judgment of image quality.

Brightness adjustment usually requires mapping a particular image luminance (HDR) or luma (LDR) value to a selected brightness of a display, while the rest of pixel values are mapped according to a predetermined tone function. In this paper we do not consider the problem of selecting a proper shape of tone curve, which is a domain of tone mapping. Instead, we investigate how people manually adjust image brightness and whether this process can be made automatic. This work validates and attempts to improve popular brightness estimation algorithms, which are important part of every tone mapping operator.

Although, it may seem that the choice of image brightness is a matter of individual subjective preferences and style, previous study shows that brightness adjustment is mostly determined by image content [16]. In this paper, we further investigate subjective choices of preferred brightness over a set of images with varied contents and dynamic range. The purpose of this study is twofold: to determine how consistent people are when selecting preferred brightness levels, and to verify whether it is possible to automatically determine the best brightness adjustment based on image characteristic.

2. Related Work

The problem of image brightness adjustments has its origins in photography, printing, television, and has recently evoked new interest in tone mapping. The major purpose of brightness (exposure) adjustment in photography is map-

0-7695-3009-5/07 $25.00 © 2007 IEEE
DOI 10.1109/PG.2007.38

ping the most important part of the scene to the middle densities of a film or the middle code value of a digital sensor, which offer the best preservation of contrast [7]. The most important brightness adjustment is in fact performed by a photographer, when he or she selects the point of interest, which the camera uses for exposure control. Analog photography offers limited possibilities for image processing inside a camera and therefore the final brightness adjustment is performed when a print is developed [1]. Digital cameras must produce acceptable images immediately after an image is taken, therefore they often employ a sophisticated image enhancement algorithms. However, the details of these algorithms are usually trade secrets and very scarce information can be found from publicly available publications. Such algorithms often involve the computation of an *image key* value, which is used to preserve the overall dark or bright appearance of the original scene and to produce better distinction between night and day-light scenes [6, 14]. The key value is usually computed from image histogram and it indicates whether an image is overall dark, bright or medium gray.

In television brightness adjustment primarily compensates for the glare on a TV screen, which is caused by ambient light [12, p. 25]. Brightness setting of a TV screen changes the display black level, which is the luminance emitted for black color. High-end displays offer automatic brightness compensation for ambient light, which is controlled by the luminance sensor installed in a TV display.

Hard copy printing usually involves mapping a *reference white*, which is the color that should be perceived as a white diffuse surface, to the maximum lightness (relative brightness) of the output medium, which is usually white paper. However, there is no reliable way to estimate which image color corresponds to the *reference white*, especially for the scenes that contain color brighter than that reference, for example specular highlights or light sources.

The tone mapping in computer graphics originally set different goals than tone mapping in photography. Tumblin and Rushmeier [15] attempted to design an algorithm producing images that would match the appearance of the original scenes. Such appearance match is not only difficult to realize because of the limited color reproduction capabilities of displays, but is also not always desirable. Sometimes the images that differ in the appearance from the original scenes are preferred to images that give a perfect match. In this respect, our goal is closer to the photographic image reproduction, as we try to find the most preferred images, rather than those that are the closest to the original scenes.

The problem of brightness adjustment for computed generated scenes is more difficult than for photographs, as there is no photographer that could choose the appropriate points of interest (focus and light measurement point) or adjust exposure settings. In practice, tone mapping operators used in computer graphics involve three different strategies of mapping scene luminance to the display brightness: with respect to a diffuse surface perceived as white, or to a middle gray surface, or to a black surface. These strategies are often grounded in perceptual theories assuming that such mapping is performed by the human visual system. Not only the psychophysical findings are often contradictory as to which of them is actually the correct one, but also a good estimation of which luminance level corresponds to a white diffuse surface is not trivial. We review the actual algorithms used in tone mapping in Section 5 and compare their performance with our experimental data.

Yoshida et al. [16] studied how people adjust brightness, contrast and color saturation to produce the most pleasing images on a display. Our study further refines the results of this work, by limiting the considerations to brightness and employing a larger number of images and subjects.

This work as well as [16] demonstrates an alternative approach to design of tone mapping operators. We clearly state our goal, which is the most preferred brightness adjustment, we collect experimental data and finally we fit mathematical models to them. We do not make any assumptions about correctness of a particular visual model and work only on the data collected in the experiment.

3. Experiment Design

We conducted an experiment in order to investigate how people manually adjust brightness in digital images. The subjective adjustments are later compared with automatic methods of brightness adjustments.

3.1. Subjects

A group of 30 people participated in the experiment. Their age was between 21 and 36, with the average of 24 years old. There were 8 females and 22 males in the group. They had normal or corrected to normal visual acuity. Most of the participants were students who completed a basic computer graphics course. None of them was aware of either the purpose or technical details of the experiment.

3.2. Stimuli

We used a set of 33 HDR images (see Figure 1), including 22 outdoor pictures (16 taken in a daytime, and 6 in the evening or at night) and 9 indoor pictures.

The images were displayed on the Miro TD490 19" LCD display. Its minimum and maximum luminance levels were 2 and 217.8 cd/m^2 respectively. Calibration of the monitor involved measuring its luminance response curve with the Minolta LS-100 luminance meter. The luminance was measured in a well illuminated room (300 lux), at the monitor

Figure 1. Set of HDR images used in the experiment. Their dynamic range was between 1.68 and 3.85 in log_{10} units. Images were tone mapped using the photographic tone reproduction [14].

brightness and contrast settings of 80%, and color temperature set to warm. Series of measurements were taken for increasing gray levels and at five different points on the screen - in the center and at four points positioned a quarter of the screen size from the corners. The luminance responses from all the points were averaged, and together with the corresponding gray level values formed an inverse lookup table used to display the images on the monitor. The luminance correction was performed on-the-fly by a fragment program running in the graphics hardware.

Each displayed image had its brightness adjusted using the formula:

$$\log R_{new} = \log R + bri \qquad (1)$$

where bri is the brightness adjustment parameter, R is the original linear value of the red tristimulus value (not gamma corrected). Similar formulas were used for the green and blue tristimulus values. Since the brightness adjustment operation is performed in the logarithmic domain, the changes of the bri parameter are approximately proportional to observable changes in image brightness.

3.3. Experimental Procedure

The participants adjusted the brightness level of the displayed images to match their preference. For each randomly selected image from Figure 1, a pair of its renderings of different brightness was presented side-by-side. Each participant chose the image that looked better in their opinion. After that, a next pair of images was presented. Brightness levels for each pair of images were found using the PEST method (Parameter Estimation by Sequential Testing) [5]. Initial brightness estimate was random and the PEST procedure was stopped when the brightness change between iterations was below visible difference. The images were displayed at a viewing distance of approximately 0.5 meters.

The participants were asked to make the choice within a few seconds, after comparing the overall image appearance, without focusing on details. Each participant ran through the experiment once. Single session took approximately 25 minutes and was preceded by a short training session including 2 images.

After the session each participant was asked to fill in a questionnaire. The questionnaire included the question about age, gender, experience in photography and experience in digital image editing (e.g. editing in Adobe Photoshop).

4. Experiment Results

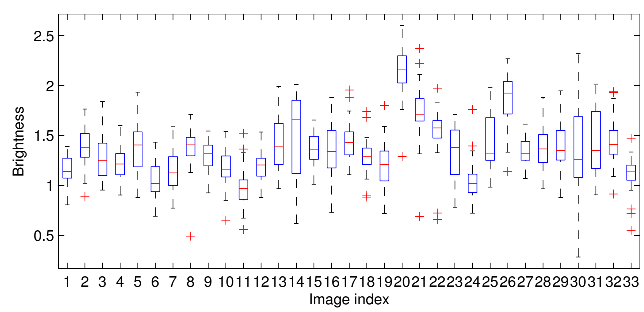

Figure 2. Brightness adjustment parameters, bri, for all subjects and for each image. The upper, center and lower line of the boxes are at 75th, 50th and 25th percentile. The whiskers show maximum and minimum value. The crosses are outliers.

The results of the experiment for each image separately are shown in Figure 2. The variances of the brightness adjustment across subjects (upper and lower quartiles on the plot) differ significantly between images; there are images such as IMAGE 4, for which subjects chose consistent brightness levels, as well as images for which preference of brightness differed significantly, such as IMAGE 14. Both images are shown in Figure 3 to illustrate the visible difference between the 25^{th} and 75^{th} percentile of subjective choices. Although brightness levels for IMAGE 14 are visibly different, all three renderings shown in Figure 3 are perfectly acceptable. Therefore the brightness adjustment cannot be understood as a single number, but rather a range of values (or random variable) that lead to plausible renderings.

Figure 3. Image 4 (top) and Image 14 (bottom) for the brightness adjustment corresponding to the 25th (left), 50th (center) and 75th (right) percentile of the experiment result (see Figure 2).

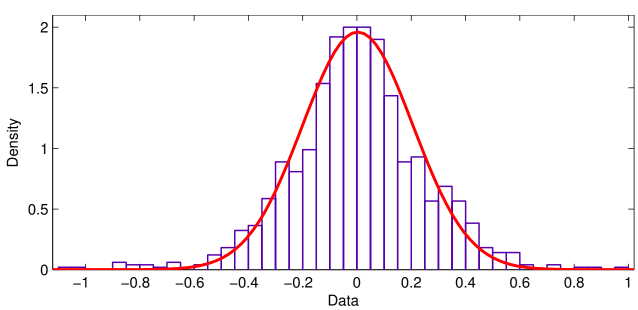

Figure 4. Distribution across subjects of the brightness parameter relative to the image mean brightness parameter.

It is interesting to further investigate how the choice of brightness adjustment parameter is distributed across subjects. Figure 4 illustrates the distribution of the brightness parameter that is corrected for the variation across images. The values are corrected by subtracting a mean brightness adjustment for the corresponding image. The data is well approximated with the normal distribution of $\sigma = 0.2$. Note that the difference in bri parameter of 0.2 corresponds to absolute luminance difference of 2/3 f-stop, which shows that the brightness visibly differs between subjects.

Using the data collected in our questionnaire we examine the influence of the factors, such as image content, gender, experience in photography and experience in digital image editing. Since the variances of bri parameter for images are significantly different, we cannot use the ANOVA analysis and instead we use the nonparametric Kruskal-Wallis test. As expected, brightness adjustment is affected by image content ($\chi^2 = 399.16, p = 0$)[1], which is also the most significant factor. More interestingly, female participants choose images that are significantly darker than male participants ($\chi^2 = 7.57, p = 0.0059$), and the difference of mean brightness adjustment is 0.07 in log_{10} units or 0.23 f-stops. There is also a significant difference between the groups of participants that are or are not experienced in photography ($\chi^2 = 35.81, p < 0.3 \cdot 10^{-9}$), with photographers choosing images 0.15 log_{10} units or 1/2 f-stops darker. Experience in digital imaging was not statistically significant at the 1% confidence level ($\chi^2 = 4.87, p = 0.0274$).

In the further analysis we will consider only the effect due to the most significant factor, which is image content. The other factors, though statistically significant, lead to rather moderate visible differences (up to 1/2 f-stops). It is also not clear that the results for the group of participants can be generalized for a larger population.

5. Brightness Adjustment Algorithms

We compare several commonly used models for automatic brightness adjustment with the results of our experiment. We assume that an image is adjusted using the formula:

$$\log R_{new} = \log R - f(x; Y) + c \tag{2}$$

where $f(x; Y)$ is one of the models from Table 1, x are model parameters, Y are image luminance values, and c is a target logarithm luminance on a display. The intuitive interpretation of the above formula is that we map the log-luminance in an image equal $f(x; Y)$ (where $f(x; Y)$ could be for example the mean of Y) to the log-luminance of the display equal c.

The models from Table 1 can be categorized according to their assumptions into the following groups (model name given in italic):

[1]The larger χ^2 value indicates higher influence of the factor on the brightness adjustment. The p value is the probability that the clustered groups (e.g. by gender) belong to the same population.

Gray-world assumption. The most common algorithms in tone mapping fall into the gray-world assumption category. An average pixel luminance in an image is assumed to be perceived as a gray tone of a medium brightness level. Such match is also often assumed to be the goal of the adaptation processes in the human visual system. The luminance of middle gray tone is either calculated as the *mean* of all luminance in the image, or the *logarithmic average* (geometric mean) [2, 14], or the *median* of log-luminance values.

Reference white. The human eye can easily identify white diffuse surfaces in the context of a scene, thus such a surface should also be depicted as white on a target display. The major difficulty of the algorithmic approach is how to reliably estimate the luminance of a white diffuse surface in an image. There are several heuristics that give such estimation, from which the most common one assumes that a fixed percentage of the brightest pixels creates specular reflections. Therefore, the white reference is estimated as the maximum luminance after cutting off the given percent of the highest luminance values (*high percentile*) [9]. Alternatively, one can assume that specular reflections exist only in high frequency contents of an image [11], therefore a luminance estimate of a white diffuse surface is the maximum luminance of a low-pass filtered image (*maximum after blur*). Also, the *minimum derivative* of the luminance histogram can be used as an estimator of the threshold luminance above which only highlights exist [8].

Reference black. We can also assume that the eye perceives the tone scale of a scene with the reference to a black diffuse surface. The reference black can be estimated as the minimum luminance in the image, often after ignoring a group of outliers using *low percentile* or after blurring to remove the influence of noise (*minimum after blur*).

Hybrid methods. The brightness adjustment based on the *image key* [6, 14, 13] involves also the gray-world assumption, but is additionally corrected towards darker shades for predominantly bright scenes and towards brighter colors for predominantly dark scenes. The *image key* is usually computed from low-percentile, high-percentile and the mean or median value of log-luminance (refer to Table 1).

6. Analysis

We confront the models introduced in the previous section with the results of our subjective study. We fit each of these models to the experimental data, by minimizing the difference between the mean brightness adjustments from our experiment, $\hat{bri}_k$, and each model f (refer to Table 1):

$$\arg\min_{c,x} \sum_{k=1..K} ||\hat{bri}_k - f(x;Y_k) - c||^2 \qquad (3)$$

method	equation
mean	$f = \log\left(\frac{1}{N}\sum_{n=1}^{N} Y_n\right)$
log average	$f = \frac{1}{N}\sum_{n=1}^{N} L_n$
median	$f = L^{[50]}$
high percentile	$f = L^{[p]},\ p \in (50, 100]$
low percentile	$f = L^{[p]},\ p \in [0, 50)$
min after blur	$f = \min(L * K_\sigma)$
max after blur	$f = \max(L * K_\sigma)$
image key	$f = \frac{L^{[1]} + L^{[99]} - 2\frac{1}{N}\sum_{n=1}^{N} L_n}{L^{[99]} - L^{[1]}}$
min derivative	see [8] for details

Table 1. Equations for brightness adjustment. Y_n denotes luminance of the n-th pixel in an image, L_n equals $\log Y_n$ and N is the total amount of pixels. The $*$ denotes convolution with the Gaussian kernel K_σ with standard deviation σ. $L^{[p]}$ is the p-th percentile of the logarithmic luminance values in an image.

The minimization is performed for logarithm of the target display luminance c (refer to the Equation 2) and the model parameter x, which can be the percentile number, σ of the Gaussian blur or no parameter, depending on the model. Y_k denotes luminance values of the k-th image and K is the total number of images used in the experiment. We judge the accuracy of the fit by measuring the χ^2 error. This measure tests whether the predicted brightness lies within the standard deviation range of the preferred brightness adjustment ($\chi^2 < 1$). Additionally, we measure correlation between data to identify increased probability that a given assumption indeed influences the choice of brightness.

The goodness of fit and the correlation coefficients for a range of percentile numbers and Gaussian extends are illustrated in Figure 5(a)–5(c). The low percentile values result in the best fit to the experiment data (9^{th} percentile), while the high values are considerably worse with the best result for 99^{th} percentile. Interestingly, high 100^{th}, middle 44^{th}, and low 2^{nd} percentile values show increased correlation with the subject preferences. Figure 5(b) shows analysis for maximum luminance in a blurred image. The small blur gives better predictions than no blur or larger blur, and is also better correlated with the choices of subjects. The best prediction is obtained for $\sigma = 2$. In case of the minimum luminance in a blurred image, Figure 5(c), the best prediction and the strongest correlation is obtained for $\sigma = 16$.

The results of analysis are summarized in Table 2. To interpret the target luminance on the display c from equation (2), which is optimized for each model, we map it to a uniform scale of gray shades G where 1 denotes white, 0.5 middle gray and 0 is black. The G value in Table 2 distinctively identifies the assumption that each model fol-

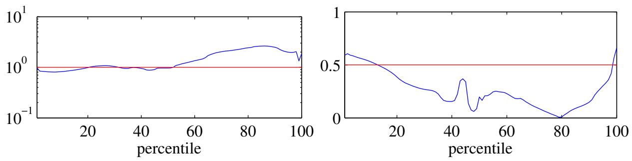

(a) percentile

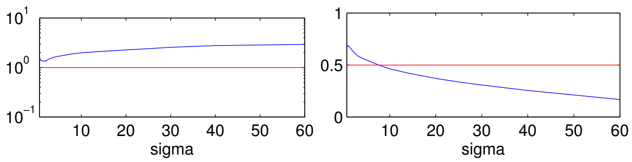

(b) maximum of a blurred image

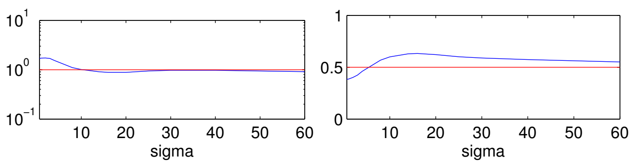

(c) minimum of a blurred image

Figure 5. Goodness of fit χ^2 (left) and correlation coefficients r (right) to estimate best parameters for percentile and blur based brightness adjustment. The horizontal lines indicate value 1 for χ^2 and 0.5 for r.

	method	χ^2	P	r	G
1.	model, eq.5	0.37	100%	0.73	0.6
2.	model, eq.4	0.37	94%	0.76	0.5
3.	image key	0.48	87%	0.75	0.5
4.	log average	0.67	79%	0.44	0.5
5.	low percentile (9^{th})	0.81	76%	0.54	0.2
6.	min after blur $\sigma = 16$	0.90	64%	0.63	0.1
7.	median	0.96	76%	0.10	0.5
8.	mean	1.30	61%	0.25	0.7
9.	high percentile (99^{th})	1.33	70%	0.56	1.0
10.	max after blur $\sigma = 2$	1.35	58%	0.63	1.1
11.	min derivative	5.13	34%	0.01	0.7

Table 2. Ranking of brightness adjustment methods. $\overline{\chi}^2$ is the mean error with which the model describes the experiment, $\overline{\chi}^2 < 1$ denotes error within the standard deviation. P is the percentage of adjustments within the standard deviation of experiment data, r is the correlation coefficient, and G is the target brightness on the display. Refer to Section 6.

lows. The *max after blur* model maps to a G value above white what suggests that this model predicts certain level of highlights luminance rather than a diffuse white surface.

Out of the methods we analyzed, the best model of preferred brightness in the experiment can be obtained with the *image key*. The reference black and gray-world methods generally perform similarly and superior to the reference white methods. On the other hand, reference white and reference black methods have higher correlation with the experiment data than gray-world methods, although again the *image key* has the highest correlation. Still several images remained unpredicted by any of the methods, although the subjects have been consistent about their preference. Below, we discuss the results in detail.

6.1. Gray-world Assumption

The algorithms based on the gray-world assumption adjust brightness close to the preferred setting in the experiment, with logarithmic average giving the smallest error $\overline{\chi}^2 = 0.67$. However, they fail ($\chi^2 \geq 1$) on about 21% of images while choices of the subjects are consistent. This includes images containing details in the bright areas, which have been adjusted too bright, thus saturating interesting parts, and the images with blurred bright backgrounds, as in many portraits, which have been adjusted too dark. Also low contrast images in 50% of cases have been preferred much brighter than adjusted by the algorithm.

The adjustment based on the *image key* improves the accuracy for several images by maintaining the expected subjective scene appearance. However, the images with blurred background and containing night scenes are still not well adjusted. Yoshida et al. [16] also reports high correlation of the *image key* with the percentage of clipped pixels, which is related to the image brightness adjustment.

6.2. Reference White

The reference white assumption, a competing theory to the gray-world assumption, performed the least accurately. Our analysis does not undermine the theory, but rather imply that current heuristics to identify diffuse white surface are not robust enough. Particularly, the *minimum derivative* approach successfully accounts for spiky specular reflections, but fails in the presence of broad reflections or direct lights commonly encountered in natural images. However, an increased correlation of *high percentile* and *maximum after blur* deserves attention when developing an improved brightness adjustment model.

6.3. Reference Black

The models based on reference black assumption, which is generally not supported by psychophysical evidence, perform surprisingly well on the observations from the experiment in case of both *low percentile* and *minimum after blur*. Again, rather than favoring the reference black over the reference white assumption, the stability of the results suggests that the estimation of black diffuse surface is far more accurate than of the white one. However, since this group of adjustments consistently failed on images with detailed bright areas, one can conclude that it is the information in the bright areas that people prefer to see.

7. Improved Brightness Adjustment

While there are a few images that cannot be accurately predicted by any of the models from Table 1, we observe that some images are well corrected by one model, while the other images by the other model. This motivates us to check if there exists a linear combination of models which results in an improved goodness of fit.

We experiment with several linear combinations of the models listed in Table 1 and find the best fit by optimizing Equation 3. Parameters x include both the linear coefficients and the parameters of the tested methods. An improvement in brightness adjustment accuracy for the combination of two models is observed for the two percentiles:

$$f = 0.55 \cdot L^{[2]} + 0.45 \cdot L^{[99]}. \tag{4}$$

Equation (4) predicts 94% of images with the fit quality $\overline{\chi}^2 = 0.37$ and outperforms any method analyzed so far. While these particular percentiles, 2^{nd} and 99^{th}, give the best prediction, in general a combination of high and low percentiles is significantly more accurate than other possibilities as illustrated in Figure 6. In these cases the brightness adjustment is driven by the amount of information clipped in both dark and bright areas of the image. The 100% prediction rate for images in the experiment is achieved with the combination of three percentiles:

$$f = 0.28 \cdot L^{[9]} + 0.37 \cdot L^{[42]} + 0.35 \cdot L^{[100]}, \tag{5}$$

with the fit quality $\overline{\chi}^2 = 0.37$. Interestingly these three estimates correspond to the local maxima of correlation coefficient for percentile, Figure 5(a). Therefore one can argue that a robust brightness adjustment has to be based on the relative distribution of low, high and mid-tones in the image. Interestingly, this is confirmed by the equation of empirically derived *image key* [13] which also performed well in our analysis.

Linear combinations of other brightness adjustment methods performed less accurately. The model fitting results for the two best performing linear combinations are given in Table 2.

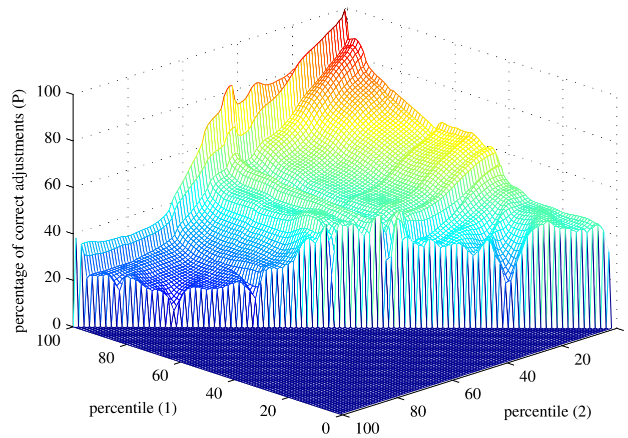

Figure 6. Percentage of adjustments within the standard deviation of experiment data for linear combination of two percentiles. The linear coefficients are optimized for each pair such that the best accuracy is obtained.

8. Applications

There are three major applications where the presented brightness models can improve the results: gradient domain image enhancements, tone mapping operators, and conversion of camera RAW files to 8bit images. We discuss them in detail below.

The derivation and integration steps used in the gradient domain algorithms produce images with an unknown constant offset, which is responsible for brightness adjustment. In case of local processing an appropriate offset is chosen to match the brightness at the borders of a local region. However, when the full image is processed, no clues are available and in most cases such final adjustment is left to the user for manual correction. Instead, applying the presented brightness adjustment can assure good appearance of images processed with such methods as gradient domain tone mapping [4, 10] or bilateral filtering [3], as shown in Figure 7.

Several tone mapping operators, such as [2, 14], include brightness adjustment methods based on the gray-world assumption, which is usually implemented as the logarithmic average. Since the model from equation (4) maps to the same display gray level, it can replace the logarithmic average, thus achieving more preferable brightness and higher overall image quality. Alternatively, the model can be applied as post-processing of a tone mapping operator. Once the images have been tone-mapped to the dynamic range

Figure 7. The *image key* outputs a dark image (top) which perhaps resembles the scene appearance, while the model from Equation 5 (bottom) gives bright image which better depicts strength of detail preserving tone mapping [10].

that matches the target display, the model can be used to fine-tune their brightness.

A critical brightness adjustment happens during the automatic exposure in digital cameras. Unfortunately, the light measuring devices in most digital cameras calculate the exposure from incomplete information about the scene and in certain conditions interesting parts of photograph become overexposed. Photographers often intentionally underexpose images in such conditions and capture in RAW format which has a larger dynamic range. Thus a preferable exposure can be adjusted using the presented models at a later stage, while viewing or converting the RAW files to 8bit images.

9. Conclusions

The results of our experiment indicate that image content is the major factor affecting people's decision on the preferred image brightness. We fit several popular models for computing brightness adjustment to the data collected in the experiment. We find that a linear combination of three percentiles gives the best goodness of fit. This suggests that the most reliable method of brightness adjustment should take into account all three anchoring approaches: the gray-world assumption, reference white and reference black. We demonstrate the application of an automatic brightness adjustment in the context of tone mapping, gradient-domain processing and rendering of camera RAW files.

In the future work we would like to validate more complex models of brightness adjustment, which take into account spatial image information, such as presence of details. We would also like to extend our study to a larger set of images and use several displays of different brightness and contrast levels.

References

[1] A. Adams. *The Print, The Ansel Adams Photography Series 3*. New York Graphic Society, 1981.

[2] F. Drago, K. Myszkowski, T. Annen, and N. Chiba. Adaptive logarithmic mapping for displaying high contrast scenes. In P. Brunet and D. Fellner, editors, *Proc. of Eurographics*, pages 419–426, 2003.

[3] F. Durand and J. Dorsey. Fast bilateral filtering for the display of high-dynamic-range images. In *Proc. of ACM SIGGRAPH 2002*, Computer Graphics Proceedings, Annual Conference Series, 2002.

[4] R. Fattal, D. Lischinski, and M. Werman. Gradient domain high dynamic range compression. In *Proc. of ACM SIGGRAPH 2002*, pages 249–256, 2002.

[5] G. A. Gescheider. *Psychophysics: The Fundamentals*. Lawrence Erlbaum, 1997.

[6] J. Holm. Photographic tone and colour reproduction goals. *CIE Expert Symposium on Colour Standards for Image Technology*, pages 51–56, 1996.

[7] D. Kerr. Apex - the additive system of photographic exposure. In *http://doug.kerr.home.att.net/pumpkin/APEX.pdf*. 2005.

[8] E. A. Khan, E. Reinhard, R. W. Fleming, and H. H. Bülthoff. Image-based material editing. In *SIGGRAPH '06: ACM SIGGRAPH 2006 Papers*, pages 654–663, 2006.

[9] G. Krawczyk, K. Myszkowski, and H.-P. Seidel. Lightness perception in tone reproduction for high dynamic range images. In *The European Association for Computer Graphics 26th Annual Conference EUROGRAPHICS 2005*, volume 24 of *Computer Graphics Forum*. Blackwell, 2005.

[10] R. Mantiuk, K. Myszkowski, and H.-P. Seidel. A perceptual framework for contrast processing of high dynamic range images. *ACM Transactions on Applied Perception*, 3(3):pp. 286 – 308, 2006.

[11] L. Meylan, S. Daly, and S. Süsstrunk. Tone mapping for high dynamic range displays. In *Proc. of SPIE on Human Vision and Electronic Imaging XII*, volume 6492, page 649210, 2007.
[12] C. A. Poynton. *A Technical Introduction to Digital Video*. John Wiley & Sons, New York, 1996.
[13] E. Reinhard. Parameter estimation for photographic tone reproduction. *Journal of Graphics Tools*, 7(1):45–51, January 2003.
[14] E. Reinhard, M. Stark, P. Shirley, and J. Ferwerda. Photographic tone reproduction for digital images. In *SIGGRAPH 2002 Conference Proceedings*. ACM SIGGRAPH, Addison Wesley, Aug. 2002.
[15] J. Tumblin and H. E. Rushmeier. Tone reproduction for realistic images. *IEEE Computer Graphics and Applications*, 13(6):42–48, Nov. 1993.
[16] A. Yoshida, R. Mantiuk, K. Myszkowski, and H.-P. Seidel. Analysis of reproducing real-world appearance on displays of varying dynamic range. In *EUROGRAPHICS 2006 (EG'06)*, volume 25 of *Computer Graphics Forum*, pages 415–426. 2006.