Tone Mapping for Single-shot HDR Imaging

Johannes Herwig, Matthias Sobczyk and Josef Pauli
Intelligent Systems Group, University of Duisburg-Essen, Bismarckstr. 90, 47057 Duisburg, Germany

Keywords: High Dynamic Range Imaging, Tone-reproduction Operators, Noise Reduction, Image Segmentation.

Abstract: The problem of tone mapping for HDR (high dynamic range) to LDR (low dynamic range) conversion is introduced by a unified framework considering all the usual processing steps. Then the specific problem of single-shot HDR is outlined where special emphasis is taken on the effect of the greater noise floor of those images when compared to the usual exposure bracketing approach to HDR. We herein tailor the popular tone mapping operators proposed by Reinhard for single-shot HDR. A region-based approach for preprocessing any HDR image in order to increase SNR and perceptual sharpness is introduced as an extension to our initial tone mapping framework. The results are compared with respect to specially developed baseline tone mappers and an extensive subjective evaluation is performed.

1 INTRODUCTION

Tone mapping operators are used in a high dynamic range (HDR) image acquisition and processing chain (Reinhard et al., 2010) as the final completion. Tone mapping allows displaying or printing a HDR image on LDR (low dynamic range) media by compressing the wide tonal range of the HDR image into an image with lower tonal sampling. Thereby the bit-depth of the HDR image pixel commonly is 32-bit floating point but LDR images are only 8-bit unsigned integers. Tonal compression should be able to preserve the overall contrast, textural details and color fidelity of the original image (Frazor and Geisler, 2006).

Often only a HDR image is capable of capturing the real dynamic range of any natural scene, but both consumer digital displays and printing technologies are only capable of dealing with low dynamic range images (DiCarlo and Wandell, 2000). The crux thereby is that photographers usually want to create a ”true” reflection of their visual experiences which they want to convey to their viewers, but due to limited capturing and displaying technologies any photograph can never be as visually rich as the real scene. Therefore, the tone mapper is crucial in delivering an image that ”feels” as naturalistic as possible when viewed at low dynamic range.

Although, tone mapping operators are developed with the (in-)capabilities of the human visual system (HVS) in mind, their design can be considered more an art than engineering. This is also because of the vast amount of factors that influence the sensation of an image where lots of assumptions are to be made. For example, apart from the tonal richness of the particular HDR image additional properties of the viewing conditions and the audience are to be considered: specific display technology, viewing distance, ambient lighting, emotional state, cultural background, etc (Bodrogi and Khanh, 2012).

In this paper, we present several extensions and enhancements of the tone mapper for photographic tone reproduction originally introduced by Reinhard (Reinhard et al., 2002). Reinhard has developed different tone mapping operators in the past which are commonly acknowledged for their naturalistic results thereby advancing this field of research (Reinhard et al., 2010). His and other operators usually assume that the HDR image was created by fusing a series of differently exposed LDR images of the same scene.

1.1 Tone Mapping for Single-shot HDR

In our application scenario only one image is taken with a pixel depth of 12-bit unsigned integer, which means there are 4096 different values that we want to tone map to 8-bit or 256 different values per color channel. We use the so-called RAW imaging mode of a digital camera which directly stores the raw but color balanced pixels from the digital imaging sensor.

For our tone mapping we exploit the fact that modern digital consumer cameras internally have higher analog-to-digital conversion (ADC) capabilities than

Herwig J., Sobczyk M. and Pauli J..
Tone Mapping for Single-shot HDR Imaging.
DOI: 10.5220/0004695401450152
In *Proceedings of the 9th International Conference on Computer Vision Theory and Applications* (VISAPP-2014), pages 145-152
ISBN: 978-989-758-003-1

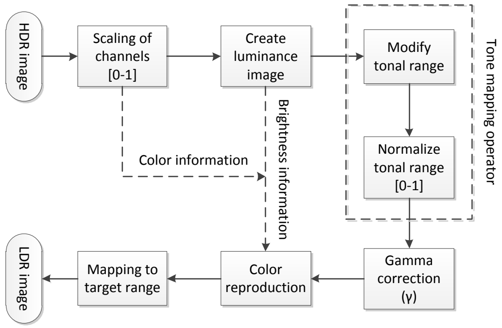

Figure 1: General framework for tone mapping.

their low dynamic range JPEG output images. This ultimately means that these cameras already have implemented proprietary tone compression algorithms. These are however restricted to the light processing power of current imaging processors and therefore cannot perform advanced computations. In our experience the qualitative results of these global tone mappers generally come close to some sort of gamma correction which is one of the simplest tone mappers and does not sufficiently lighten up shadowed parts and tends to wash out textural details in brighter parts of the image. Our globally and locally adaptive approaches are designed to overcome these issues.

The single-shot HDR approach provides a lower dynamic range than real HDR capture using multiple exposures. It however is easier to apply to dynamic scenes. On the other hand, a single-shot HDR image is noisier because there is no averaging of multiple exposures and since tone mappers are designed for compressing wider dynamic ranges than a single shot provides the noise tends to be intensified because its is assumed to represent textural detail.

Although there are lots of different tone mapping algorithms these nevertheless can be described by the framework depicted in figure 1. Note that we always perform an additional gamma correction after tonal compression in order to comply with the usual $ITU-RBT.709$ HDTV-standard for displaying devices.

2 PREVIOUS WORK

2.1 Linear Mapping

With a linear mapping a given range of values is mapped to some target range by scaling with a constant factor. HDR images that are downscaled this way generally appear too dark and textural details in shadowed regions get lost. This can be accommodated by pre-scaling the HDR image like this:

$$I'(y,x) = \frac{I(y,x) - I_{mincut}}{I_{maxcut} - I_{mincut}}$$

with $I_{maxcut}, I_{mincut} \in \{v \mid v \in \mathbb{R} \ \wedge \ 0 < v \leq 1\}$, $I_{maxcut} > I_{mincut}$, $I(y,x) \in \{0.0, \ldots, 1.0\}$ and intensities $I'(y,x) > 1.0$ and $I'(y,x) < 0.0$ will get clipped.

I_{maxcut} could be set in such a way that e.g. the upper 5% of intensities of the HDR image get clipped, i.e. are collectively set to the maximum value of the target range. This will result in a brightened LDR image because fewer higher intensity outliers have less impact on the overall scaling factor. Although the resulting effect of burning pixels within highly lit image regions causes lost textural details, it is positive for perception (Reinhard et al., 2010) and enhances the overall image contrast. We experimentally found that clipping the higher 1% of intensities does not result in any visually perceivable loss of information.

I_{mincut} can be either set to zero or it can be similarly used in order to clip the noise within darker regions of the image which is especially useful for single-shot HDR and further contributes to preserving visually important textural details in the resulting LDR image. Thus I_{mincut} could be set to the estimated noise level (Immerkær, 1996) of the HDR image.

2.2 Global Photographic Mapping

Reinhard has successfully introduced an operator (Reinhard et al., 2002) that is inspired by the zone system of the famous photographer Ansel Adams.

A zone system splits the tonal range of the HDR image into 11 different tonal zones from pure black (zone 0) to pure white (zone 10). Zones $1-9$ have pre-set brightnesses which linearly blend into each other. Thus zone 5 is of medium brightness and the so-called scene key α maps to this zone. If α is of relatively low brightness, then the overall image is brightened up so that shadowed regions are better visible but lighter regions loose richness at the same time. If α is of relatively high brightness, then shadowed regions fully loose detail but brighter image regions feature more textural richness.

This zone system is realized by the following tonal compression of the normalized luminance image I:

$$I'(y,x) = \frac{L_\alpha(y,x)}{1 + L_\alpha(y,x)} \cdot \left(1 + \frac{L(y,x)}{I_{maxcut}(y,x)^2}\right)$$

whereby $L_\alpha(y,x)$ denotes the intensity $I(y,x)$ which is scaled by the scene key α:

$$L_\alpha(y,x) = \frac{\alpha}{I_{avg}} \cdot I(y,x)$$

with $\alpha \in \{x \mid x \in \mathbb{R} \ \wedge \ 0 \leq x \leq 1\}$ that can automatically be computed (Reinhard, 2002) from the logarithmic average I_{avg} of the HDR luminance image as

$$\alpha = 0.18 \cdot 4^{\frac{2 \cdot log_2(I_{avg}) - log_2(I_{mincut}) - log_2(I_{maxcut})}{log_2(I_{maxcut}) - log_2(I_{mincut})}}.$$

2.3 Local Photographic Mapping

Based on the global tone mapping operator presented in the previous paragraph, a localized version of this operator is given in (Reinhard, 2002). This operator is inspired by the technique of ”dodging and burning” that has also been originally developed by Ansel Adams when making photographic prints on paper that were not straightforwardly capable of the full tonal range of his film negatives. Here, every pixel is selectively brightened up or darkened based on an adaptively computed local neighborhood of small enough intensity variance. Darker pixels that are comprised by similar but somewhat lighter pixels are damped more strongly than bright pixels that are surrounded by relatively darker pixels, thereby creating higher local contrast in the tone mapped result. This outlined the overall idea, but for more details the reader is referred to Reinhard's original publication.

3 PROPOSED APPROACHES

Based on the well-accepted algorithms of Reinhard we propose the following extensions and enhancements for single-shot HDR imaging.

3.1 Dynamic Scene Key

We present here the global photographic mapping with a dynamic scene key $\alpha(y,x)$ that is different for every image pixel. In the original algorithm α can only be varied for the whole image, whereby an increased α brightens the mid-tones but a decreased α darkens the the whole image.

For improving the overall perception of image contrast and simultaneously preserving or enhancing local textural details it is only necessary to selectively brighten up shadowed image regions but largely preserve the local brightnesses of already bright image regions in order to avoid the loss of visual detail. This desired qualitative behavior of our proposed function for computing $\alpha(y,x)$ is depicted by figure 2. Only very bright intensities are starkly dampened and also very dark intensities are smoothly cut-off since these often tend to represent image noise. We experimentally found the following adaptive $\alpha(y,x)$ which varies with pixel intensity $I(y,x)$:

$$\alpha(y,x) = 1 - \exp(-(I(y,x) \cdot (1 + d \cdot I(y,x)))^{1-(I_{avg})^{\frac{1}{c}}})$$

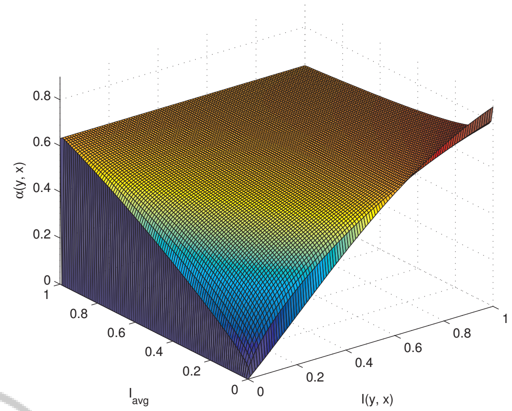

Figure 2: The qualitative behavior of our varying scene key $\alpha(y,x)$ with recommended $c = 4$. For better visualization variable d is constantly set to 1 here.

with the user defined $c \in \{x \mid x \in \mathbb{N} \wedge x > 0\}$ and

$$d = \begin{cases} -1 & \text{if } I(y,x) < I_{avg}, \\ 1 & \text{if } I(y,x) \geq I_{avg}. \end{cases}$$

The constant c determines the overall strength of the contrast enhancement. Thereby smaller values create greater contrast as is depicted in figure 3. Note that only the shadowed image regions are starkly brightened up with increasing c but that already bright image regions are only moderately brightened which works out as expected. This property makes the only user parameter for computing $\alpha(y,x)$ to be robust for a large amount of different scenes.

In our experiments a value $c = 4$ has been proven to produce pleasing results for most scenes. For different scenes with an overall dark appearance the resulting α is always between ≈ 0.63 and ≈ 0.86, but for scenes with higher average brightnesses it usually lies between ≈ 0.63 and ≈ 0.65.

3.2 Modified Local Photographic

For single-shot HDR we experienced that Reinhard's local photographic tone mapper can be safely used with its default parameters because the comparatively low dynamic range of RAW camera images does not take advantage of those. Also shadowed regions are adequately brightened up.

But we found that there is severe burning-in among brighter parts of the scenes which causes superfluous loss of detail and local contrast. Here we did not change the original operator itself but the method for normalizing the tonal range which is the first post-processing step as depicted in figure 1. Here the standard method is to simply cut-off values $I'(y,x) > 1.0$.

After tonal compression, we propose to project all the values $I'(y,x)$ which are out of range into the valid

Figure 3: This image series shows the effect of parameter c (= 1, 4, 6 and 11) of the dynamic scene key $\alpha(y,x)$.

$0.0,\dots,1.0$ range:

$$I'_{normalized}(y,x) = \frac{I'(y,x) - I'_{min}}{I'_{max} - I'_{min}}$$

Thereby I'_{min} denotes the smallest and I'_{max} is the largest value of $I'(y,x)$. In order to avoid that large values for I'_{max} result into high loss of information in darker image regions, we additionally dampened these by using the square root. Therefore very large brightness values ($>> 1.0$) are more starkly dampened than smaller out of range values (> 1.0). This also avoids that the linear scaling by $I'_{normalized}(y,x)$ produces overly dark images.

This minor modification has only an effect if the original tone mapper produced values $I'(y,x) > 1.0$. Then the overall contrast of the result images gets reduced but at the same time local detail in bright regions is enhanced as intended. The result is comparable to the global photographic mapping without burning-in, but shadowed regions are more intense here which was the original benefit of the local photographic mapping approach.

3.3 Region-based Preprocessing

The previously presented methods already result into visually pleasing results. The problem with tone mapping single-shot HDR images is however the extended noise floor when compared to bracketed HDR exposure series. Therefore we extend the usual tone mapping framework of figure 1 with a universally applicable pre-processing step. Here we want to selectively increase the signal-to-noise ratio (SNR) of low-lit and therefore noise-prone image regions by using a smoothing filter. At the same time we want to increase perceived image sharpness in areas with already originally satisfactory SNR in order to enhance textural detail. The image segmentation occurs on the HDR luminance image as depicted in figure 4.

For both smoothening an sharpening we use the same "unsharp masking" algorithm, which is a common tool in image processing software. Thereby we use the filter mask I_G which is the HDR input image I that is convolved with a 3×3 Gaussian smoothing

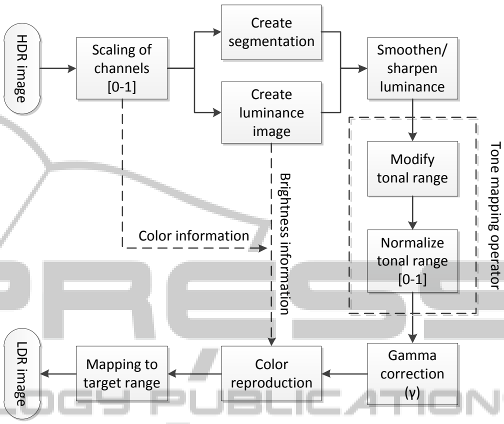

Figure 4: Extended framework for tone mapping.

kernel with $\sigma = 0.8$:

$$SMOOTH(I,a) = I \cdot (1+a) + I_G \cdot (-a) \text{ with } a < 0$$
$$SHARPEN(I,a) = I \cdot (1+a) + I_G \cdot (-a) \text{ with } a > 0$$

and $a \in \mathbb{R}$. When smoothening the SNR is increased but local contrast is necessarily reduced, we however tackle single-shot noise in order to improve the overall quality in darker image regions. Sharpening on the other hand always reduced SNR but increases the visual perception quality of textural detail in the image. Sharpening is always applicable to global tone mapping operators, whereas most local tone mapping operators already perform some inherent sharpening (Mantiuk et al., 2009), so that additional sharpening may overdo the intended perceptual effect.

Here we use a popular graph-based segmentation algorithm (Felzenszwalb and Huttenlocher, 2004) that was extended for coping with HDR images. This greedy algorithm has three parameters: (1) the similarity measure K controls which neighboring pixels will belong to the same region, (2) the constant min denotes the minimum number of pixels of every region, (3) and σ is the smoothing parameter of a Gaussian pre-filtering step. We found experimentally that parameters $K = 0.36$, $min = 10$ and $\sigma = 1.6$ provide good enough segmentation results. Note that for our purpose no exact segmentation is needed. However, an over-segmentation is preferable over under-segmentation. Some segmentation results with different parameters are exemplarily shown in figure 5.

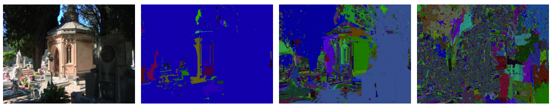

Figure 5: The graph-based segmentation leads to under-segmented, acceptable, and over-segmented results, resp.

In the following we describe the targeted processing that is chosen for every segmented image region s based on different quality measures:

1. For small standard deviation: replace pixel values within this segment with its average intensity.
$$0 \leq I^s_\sigma \leq \sigma_{noise} \rightsquigarrow I^s(y,x) = I^s_{avg}$$
2. For small entropy: replace pixel values within this segment with its average intensity.
$$0 \leq I^s_H \leq \sigma_{noise} \rightsquigarrow I^s(y,x) = I^s_{avg}$$
3. For small SNR: smooth this segment.
$$0 \leq I^s_{SNR} \leq 30 \rightsquigarrow SMOOTH(I^s, -1)$$
4. For mean SNR: moderately sharpen this segment.
$$30 \leq I^s_{SNR} \leq 39 \rightsquigarrow SHARPEN(I^s, 1)$$
5. For high SNR: starkly sharpen this segment.
$$I^s_{SNR} \geq 39 \rightsquigarrow SHARPEN_+(I^s, 1.5)$$

The decision parameters have been found by excessive testing with a wide variety of images. These were chosen to avoid negative perceptual effects around the borders of image regions where a smoothed region is adjacent to a sharpened image region at normal viewing distances. The trained eye however can experience minor visual artifacts when the image would be enlarged. These are not disturbing but could however be further dampened by applying additional blending techniques.

In this paper, we evaluate the effect of this region-based preprocessing by the global photographic and the simplistic linear tone mapping operator. Thereby the normalization of the compressed tonal range (see figure 4) always cuts off out of range values. When comparing the results of the tone mapping with and without region-based preprocessing there visually occur only minor differences between the results and also the global contrast does not change which is as intended. The local contrast is however slightly increased. Depending on the amount of smoothed versus sharpened image regions the SNR either increases or decreases, respectively. This is as intended, because sharpening adds more edges as textural detail but is falsely interpreted as noise by the SNR measure, and on the other hand smoothing increases SNR.

3.4 Alternating Global and Local Photographic Mapping

As has been mentioned previously, local tone mapping operators often inherently perform some form of sharpening the image (Mantiuk et al., 2009). However, sharpening is not desired for noise image regions because noise will be unnecessarily enhanced. Since the local photographic tone mapping operator can be reduced to Reinhard's original global version, we propose to alternatingly use both of these depending on the local SNR of an image region.

Thereby we make use of our previously introduced graph based segmentation as follows:

1. For small SNR: process with the global operator.
$$0 \leq I^s_{SNR} \leq 30 \rightsquigarrow R(I^s, V(s_{min}), 0)$$
2. For mean SNR: process with the local operator with moderate sharpening.
$$30 \leq I^s_{SNR} \leq 39 \rightsquigarrow R(I^s, V(s_{max}), 8)$$
3. For high SNR: process with the local operator with starker sharpening.
$$I^s_{SNR} \geq 39 \rightsquigarrow R(I^s, V(s_{max}), 16)$$

Here, function R encapsulates the parameterization of the local tone mapping operator (Reinhard, 2002) as $R(segment, neighborhood, sharpening)$. Hence, the first action uses the smallest Gaussian neighborhood $V(s_{min})$ of size 1, which effectively transforms the local mapping operator into its global counterpart. Whereas the two remaining actions select the largest possible Gaussian neighborhood $V(s_{max})$ but use different amounts of sharpening 8 and 16, respectively.

This approach results into minor enhancements of the SNR when compared to the original approaches, and also the correlation coefficient between the original HDR and the resulting LDR image is increased. Furthermore, artifacts that occurred at borders of different regions as with the previous region-based approach are not seen here.

3.5 "Comic" Algorithm

This is an artistic mapping that we developed as a negative baseline that helps to better interpret the results in our evaluation section. Its output is not naturalistic and does not adhere to the goals of this paper.

First, we compute the cosinus-weighted RMS contrast as in (Frazor and Geisler, 2006). This local contrast measure is computed over a 3×3 neighborhood for every pixel and the resulting image is denoted I_K. The resulting pixels are normalized between 0.0 and 1.0 and inverted, so that originally small local contrasts result into larger values than originally high local contrasts. This inverted I_K^{-1} is then transformed by a common histogram equalization resulting into $I_K^{-1,eq}$ where the probabilities of occurrence of all the available intensities are approximately equal. Since we perform histogram equalization on floating point HDR luminances, we have scaled intensities by 10^6, so that the resulting bins are meaningful and not mostly empty. In the last processing step, we scaled the pixels of the here computed contrast image $I_K^{-1,eq}$ with the unprocessed HDR luminance I:

$$I'(y,x) = \sqrt{I(y,x)} \cdot I_K^{-1,eq}(y,x)$$

Since noise is greatly enhanced due to the contrast inversion, we convolved the original luminance HDR with a 11×11 box-filter before the color reproduction step in the processing framework of figure 1.

The resulting tone mapping traces textural boarders between homogeneous image regions and therefore is termed the "comic" operator. The resulting LDR image features high local contrast and very low SNR. This properties make the algorithm suitable as a baseline for the evaluation and comparison of other tone mapping operators.

4 EVALUATION

The evaluation of tone mapping operators often takes place in photometrically calibrated environments (Kuhna et al., 2011), (Ledda et al., 2005). This is very complex to set up and the comparison results are practically questionable because ordinary users do not have calibrated monitors and also the environmental effects on image perception can never be canceled out. Therefore, we propose a quantitative evaluation with baseline operators in the evaluation set in order to arrange the results. Here we use our "comic" operator and the adaptive logarithmic (Drago et al., 2003) operator as tone mappers producing extreme results at both ends of the evaluation scale. Whereas the "comic" operator produces extremely high contrast and noisy results, the operator by Drago produces very dull but highly detailed results. We think that both results are not perceptually preferable and hence a good algorithm should produce results in-between.

There are 27 single-shot HDR images in our evaluation set featuring a broad range of scenes with higher and lower overall contrast and more or less scenic details. Due to space constraints we present only cropped results from two different scenes in figure 6 obtained for every algorithm.

(a) Linear with 5% cut-off

(b) Global photographic without burning

(c) Global photographic with burning

(d) Dynamic scene key

(e) Local photographic

(f) Local photographic with modified normalization

(g) Region-based global photographic without burning

(h) Region-based with adaptive unsharp masking

(i) Region-based with linear mapping

(j) Comic algorithm

(k) Adaptive logarithmic (by Drago)

Figure 6: Cropped images from LDR results of various tone mapping algorithms as described in this paper.

4.1 Quantitative Evaluation

For the quantitative evaluation we chose the following criteria: correlation ratio, signal-to-noise ratio, global contrast, and local contrast. The correlation ratio measures the correlation between an LDR im-

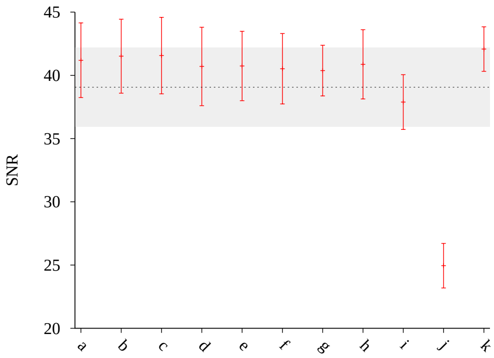

Figure 7: Average SNR values and standard deviations of LDR images. The shaded area denotes the SNR and standard deviation of the original HDR images for comparison.

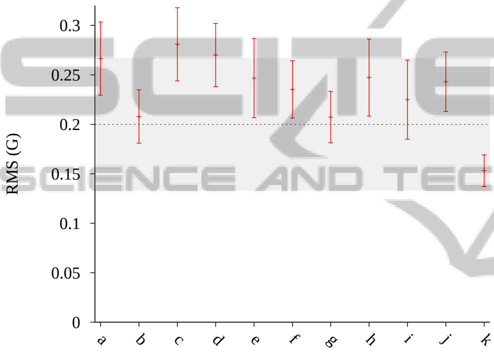

Figure 8: Average RMS contrast and standard deviations of LDR images. The shaded area denotes the RMS and standard deviation of the original HDR images for comparison.

age and its original HDR scene by using the Pearson coefficient. Global and local contrast were both measured using the RMS contrast approach (Frazor and Geisler, 2006), whereby the global contrast was measured for the whole image and the local contrast result is the average of multiple windowed RMS measurements. Results of the SNR and global contrast measures are depicted in figures 7 and 8 (please refer to figure 6 for identifying the different algorithms). It can be noted that the results greatly differ from each other, but that there is no clear winner, so trade-offs have to be made. In the figures, however, we have also indicated the variance of those quantitative measures within the original image set of 12bit data. We have experienced that a good tone mapping algorithm should have values whose mean lies well within this shaded stripe, and whose top values perform little better than on the original data although their range of variance should preferably be small.

4.2 Qualitative Evaluation

We chose a small control group of five students who are experienced in image processing. Every tone mapped LDR image was rated with values between 0 for poor and 5 for exceptional performance within the following qualitative categories: brightness, global contrast, local contrast, textural details, artifacts, fidelity, and naturalness. It is however very difficult to distinguish between some of these categories because participants sometimes have different ideas about those concepts like contrast vs. textural details.

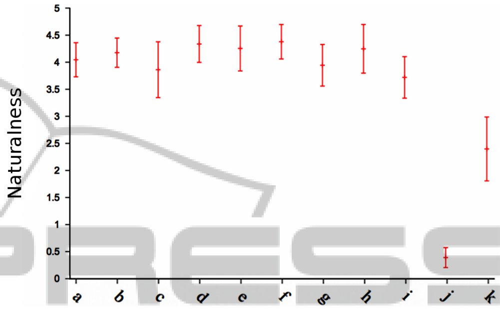

Figure 9: Average subjective naturalness and standard deviations of LDR images.

We exemplarily show the results for the subjectively perceived naturalness in figure 9. To summarize all subjective results we compiled a table of the final ranks. We chose to create composite measures where we evaluate the naturalness with respect to other desired image features like contrast (global and local) and textural detail. From table 1 it can be concluded that algorithms f and d (compare with figure 6) show good overall perceptual performance. It is interesting to note that algorithm f did not modify the original algorithm e but only the tonal normalization: as can be seen from the table, this had a great effect on its rank. Algorithm d is based on b with the intend to enhance the global contrast which greatly succeeded. At the same time the detail reproduction of d is worse then that of b which is a direct consequence of increasing global contrast (Smith et al., 2006), and the evaluation data exactly reflects that.

4.3 Discussion

According to figure 9 algorithms d, e and f perform best on average concerning the perceived naturalness, whereby d and f also show a reasonably small variance over the whole set of images. These two algorithms are also ranked best in our comparison table 1 whereby d creates higher perceived contrast but f is better balanced between image contrast and preserving textural detail. These characteristics are verified by our quantitative measurements where d is second best in terms of RMS contrast as shown in figure 8 and

Table 1: Rank of the subjective visual performance.

C = Contrast, D = Details, N = Naturalness					
Rank	C	D	C/N	D/N	C/D/N
1	d	k	d	f	f
2	j	f	f	b	d
3	c	g	h	e	h
4	h	b	e	h	e
5	i	e	a	d	b
6	a	h	c	g	g
7	e	d	i	a	a
8	f	a	b	i	i
9	g	i	g	c	c
10	b	c	j	k	k
11	k	j	k	j	j

f is clearly worse but still within the upper range concerning the original 12bit RMS values. Therefore, we can recommend algorithm d as a tone mapper in industrial image processing applications where fast acquisition times and high-contrast images are needed.

5 CONCLUSIONS

We have presented a unified framework and modified tone mapping operators for the purpose of single-shot HDR imaging. The goal was to enhance the visually perceived contrast of tone mapped LDR images, thereby preserving most textural detail of the original HDR images in both bright and shadowed regions. The qualitative evaluation shows that this was successfully achieved with our newly introduced dynamic scene key approach. It has been shown that the implementation of tonal normalization after tonal compression should be taken care of because the clamping strategy for out of range intensities has a measurable effect on the subjective perception of the mapping result. Finally, we introduced a region-based noise reduction and selective sharpening approach that can be added to the general tone mapping framework in order to enhance the performance of already existing mapping operators. In our evaluation section we have outlined general criteria for subjective evaluation of tone mapping results.

REFERENCES

Bodrogi, P. and Khanh, T. Q. (2012). *Illumination, Color and Imaging: Evaluation and Optimization of Visual Displays*. Wiley Series in Display Technology. Wiley-VCH Verlag GmbH & Co. KGaA.

DiCarlo, J. M. and Wandell, B. A. (2000). Rendering high dynamic range images. *Sensors and Camera Systems for Scientific, Industrial, and Digital Photography Applications*, 3965:392–401.

Drago, F., Myszkowski, K., Annen, T., and Chiba, N. (2003). Adaptive logarithmic mapping for displaying high contrast scenes. In Brunet, P. and Fellner, D. W., editors, *Proc. of EUROGRAPHICS*, Computer Graphics Forum, pages 419–426. Blackwell Pub.

Felzenszwalb, P. F. and Huttenlocher, D. P. (2004). Efficient graph-based image segmentation. *International Journal of Computer Vision*, 59:167–181.

Frazor, R. A. and Geisler, W. S. (2006). Local luminance and contrast in natural images. *Vision Research*, 46:1585–1598.

Immerkær, J. (1996). Fast noise variance estimation. *Computer Vision and Image Understanding*, 64:300–302.

Kuhna, M., Nuutinen, M., and Oittinen, P. (2011). Method for evaluating tone mapping operators for natural high dynamic range images. In Imai, F. H. and Xiao, F., editors, *Digital Photography VII*, SPIE Proc., pages 78760O–78760O–12. SPIE.

Ledda, P., Chalmers, A., Troscianko, T., and Seetzen, H. (2005). Evaluation of tone mapping operators using a high dynamic range display. *ACM Transactions on Graphics*, 24:640–648.

Mantiuk, R., Tomaszewska, A., and Heidrich, W. (2009). Color correction for tone mapping. In Dutr, P. and Stamminger, M., editors, *Proc. of EUROGRAPHICS*, Computer Graphics Forum, pages 193–202. Blackwell Pub.

Reinhard, E. (2002). Parameter estimation for photographic tone reproduction. *J. of Graphics Tools*, 7:45–52.

Reinhard, E., Heidrich, W., Debevec, P., Pattanaik, S., Ward, G., and Myszkowski, K. (2010). *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting*. Morgan Kaufmann, 2nd edition.

Reinhard, E., Stark, M., Shirley, P., and Ferwerda, J. (2002). Photographic tone reproduction for digital images. *ACM Transactions on Graphics*, 21:267–276.

Smith, K., Krawczyk, G., Myszkowski, K., and Seidel, H.-P. (2006). Beyond tone mapping: Enhanced depiction of tone mapped HDR images. *Computer Graphics Forum*, 25:427–438.