

Design of a Tone Mapping Operator for High Dynamic Range Images based upon Psychophysical Evaluation and Preference Mapping

Frédéric Drago^a

William L. Martens^b

Karol Myszkowski^c

Norishige Chiba^a

^aIwate University. Computer Graphics Laboratory, Ueda 4-3-5, 020-8551 Morioka, Iwate, Japan.

^bUniversity of Aizu, Tsuruga Ikki-machi, Aizuwakamatsu City, 965-8580 Fukushima, Japan.

^cMax-Planck-Institut für Informatik, Stuhlsatzenhausweg 85, 66123 Saarbrücken, Germany.

ABSTRACT

A tone mapping algorithm for displaying high contrast scenes was designed on the basis of the results of experimental tests using human subjects. Systematic perceptual evaluation of several existing tone mapping techniques revealed that the most “natural” appearance was determined by the presence in the output image of detailed scenery features often made visible by limiting contrast and by properly reproducing brightness. Taking these results into account, we developed a system to produce images close to the ideal preference point for high dynamic range input image data. Of the algorithms that we tested, only the Retinex algorithm was capable of retrieving detailed scene features hidden in high luminance areas while still preserving a good contrast level. This paper presents changes made to Retinex algorithm for processing high dynamic range images, and a further integration of the Retinex with specialized tone mapping algorithms that enables the production of images that appear as similar as possible to the viewer’s perception of actual scenes.

Keywords: Tone Mapping, Retinex, High Dynamic Range Images

1. INTRODUCTION

High dynamic range (HDR) imaging is a very attractive way of capturing real world appearances, since it permits the preservation of complete information on luminance (radiance) values in the scene for each pixel. Initially, HDR images were obtained as the result of lighting simulation and physically-based rendering of synthetically-modeled scenes.¹ Recently, methods for capturing HDR photographs of real world scenes have been developed; while HDR cameras exploiting new sensor technologies are still relatively expensive, virtually every digital camera with controllable exposure might be used to quickly produce HDR photographs of static scenes using multiple-exposure techniques.² There also exist software solutions which can automatically register differently exposed photographs taken by a hand-held camera, facilitating HDR photography even for consumer use. In the near future we can expect even further expansion of HDR technology to make possible the capture of high quality stills and video sequences in an inexpensive and convenient way. Also, in multimedia applications, it can be envisioned that video streams will be distributed in an HDR format to make possible tuning of the display parameters on the end-user side to accommodate for display hardware characteristics as well as external lighting conditions. Also, users might adjust the video reproduction according to their preferences, to achieve a desired balance between reproduced contrast, details, and brightness, whereas this information at present must be decided by the distributor and is fixed for the streamed video.

Further author information: F.D.: E-mail: frederic@cis.iwate-u.ac.jp. W.L.M.: E-mail: wlm@u-aizu.ac.jp. K.M.: E-mail: karol@mpi-sb.mpg.de. N.C.: E-mail: nchiba@cis.iwate-u.ac.jp

While the prospects for proliferation of HDR capturing technologies are good, an important problem to solve regards how to reproduce the precise information for scene appearance using media with very limited dynamic range such as hard-copy prints, CRT displays, and projectors. While there are some efforts made to build high dynamic range display devices, they deal mostly with static images, and are expensive and laborious to use.³ To solve the problem of transforming high dynamic range images to the limited range of most existing devices, so-called tone mapping operators are being developed.

The goals of tone mapping can be differently stated depending on the particular application. In some cases producing just “nice-looking” images is the main goal, while other applications might emphasize reproducing as many image details as possible, or might maximize the image contrast. The goal in realistic rendering applications might be to obtain a perceptual match between a real scene and a displayed image even though the display device is not able to reproduce the full range of luminance values.⁴ However, achieving such a match is feasible because the visual system of humans has limited capacity to detect differences in absolute luminance levels, and concentrates more on aspects of spatial patterns when comparing two images.⁵

Many tone mapping operators that appeared in the last decade have been designed to specifically address some of the issues discussed here. The goal of this work is to develop a general purpose tone mapping technique which leads to “natural” looking images. The design of our technique is based on perceptual evaluation of several existing tone mapping techniques. Our prior study revealed that “natural” appearance was determined by the presence in the output image of detailed scenery features often made visible by limiting contrast and by properly reproducing brightness.⁶ Taking these results into account, we developed a tone mapping operator based on the Retinex algorithm⁷⁻⁹ to enhance details in very high illumination areas of a scene, coupled with an adaptive tone mapping function, to reproduce the lighting conditions of the scene.

In Section 2 we discuss previous work on tone mapping techniques, which have been developed mostly for computer graphics applications. Then, in Section 3 we summarize the results obtained in our psychophysical experiments⁶ which had as their goal the evaluation of six popular tone mapping algorithms. In Section 4 we provide some background information on Retinex which is the main focus in this work. Our extensions of Retinex to accommodate for HDR images are presented in Section 5. In Section 6 is a discussion of the further computations necessary to convert an image processed by Retinex into an image that should be preferred according to perceptual criteria. Finally we conclude this work and suggest some directions for future research.

2. PREVIOUS WORK ON TONE MAPPING OPERATOR

Although the tone mapping problem is still not well studied, a number of surveys summarizing the state-of-the-art solutions have been recently published.¹⁰⁻¹³ In general, tone mapping can be classified as time dependent for handling animated sequences and time independent designed specifically for static images. In this work we deal exclusively with the latter category.

Another classification takes into account whether the same mapping function is used for all pixels (*spatially uniform* or *global* operators) or if the mapping function varies depending on a given pixel neighborhood which is often achieved through the modeling of spatial adaptation (*spatially varying* or *local* operators).

Yet another classification takes into account whether a given tone mapping algorithm tries to model some characteristics of the Human Visual System or perhaps is a pure image processing technique which essentially transforms one image into another.

One of the best known global and perception-motivated operator was proposed by Tumblin and Rushmeier.¹⁴ It is based on the idea of brightness preservation in order to get matching between the view of real and displayed scenes. The relation between scene/display luminance and perceived brightness is modeled using a psychophysical model of brightness perception developed by Stevens and Stevens.¹⁵ More recently Tumblin and Rushmeier proposed a revised method, which improves the performance of the original for scotopic viewing conditions.¹⁶ To the same category can be included operators developed by Ward¹⁷ and Ferwerda et al.¹⁸ which emphasize on preserving the perceived contrast. The Visual Adjustment operator proposed by Ferwerda et al. models also the eye sensitivity and color appearance for photopic, mesopic, and scotopic viewing conditions. The method is very efficient in terms of computation because it boils down to a linear scaling coefficient, but this means that its compressive power is somewhat limited. Ward et al.¹⁹ proposed the Histogram Adjustment technique which allocates dynamic range space in proportion to the percentage of pixels with similar brightness. This means that the brightness levels which are poorly represented in the image are strongly compressed, while

pixels featuring even small differences in brightness but well represented in the image can obtain significant portion of the available dynamic range. As a result the mapping function may have a complicated shape but remains monotonic. It is applied globally to all pixels and results in high contrast images. Local contrast can be additionally modified by glare and visual acuity modeling.

An example of global method which relies purely on image processing without considering perception is the Uniform Scaling Quantization method proposed by Schlick.²⁰ He uses rational functions which are efficient to compute, have controllable shape, and when required may compress luminance strongly.

Methods based on local image contents proved to be more efficient in high dynamic range compression, because of considering the neighborhood of each pixel before applying locally customized compressive function. Early local methods^{21,22} lead to inverse gradients and halo effects at the proximity of high luminance gradients. The same problems occur for another local method Retinex.^{8,23} We discuss it in detail in Section 4 and apply solutions to avoid those image artifacts. The Low Curvature Image Simplifier method proposed by Tumblin and Turk²⁴ is not affected by halos, and represents well high frequency details, however, it is slow in terms of computation and often leads to images with an excessively low contrast.

In 2002, a number of successful local methods have been proposed which avoid the typical problems of their predecessors. Reinhard et al. proposed the Photographic Tone Reproduction,²⁵ which is inspired by “*odging and burning*” used in photography and makes possible local contrast adjustment. A recent version of this method operates automatically freeing the user from setting parameters that are not particularly intuitive.²⁶ Ashikhmin²⁷ proposes a perception-motivated algorithm which uses local adaptation level for each pixel to decide upon the non-linear compression applied to the luminance of the pixel. Since this suppresses high frequency details in the image, those details must be re-introduced at a final step. Fattal et al.²⁸ proposed an image processing method which identifies gradients at various scales and attenuates more strongly drastic luminance changes, while preserving smaller changes which are responsible for high frequency details in the image. A multi-scale decomposition based on the Gaussian pyramid approach is applied to separate low frequency details (to be strongly compressed) from high frequency details (to be even slightly amplified). Although the gradient attenuation function is computed for each pyramid level separately the actual scaling is performed at the finer image resolution to avoid halo artifacts. Durand and Dorsey²⁹ also decompose image, but they consider only two different spatial frequency layers: base and detail. For the decomposition they use a non-linear, edge preserving filter. They suppress high contrast in the base layer and combine it with the detail layer which essentially remains unchanged. All four recent methods are efficient in terms of computations and lead to nicely looking images.

The problem of color reproduction on the limited gamut of a display device has attracted much less attention in computer graphics. However, it is well addressed in the more specialized literature of imaging and color science. An example of an advanced color reproduction model, which predicts a wide range of visual phenomena, is the Hunt model³⁰ developed in the Kodak Research Laboratories. A comprehensive survey of the color appearance models can be found in.³¹

3. PERCEPTUAL EXPERIMENTS

As we discussed in the previous section the problem of mapping full range luminance data into the limited luminance range of display devices has been addressed through the development of many different tone mapping algorithms, but evaluating the success of these solutions is rather complex. It is in fact quite difficult to judge whether putative improvements in existing algorithms actually produce images that would be preferred to those produced by the algorithms before their “*improvement*.” Changes should be made to achieve specific, well-defined goals. One such goal might be to obtain images that look as “*good*” as possible and are always preferred by the most viewers. In this work we focus specifically on the goal of faithful reproduction of real world appearance, which we have termed “*naturalness*.”

Our approach to the design of a tone mapping algorithm to display HDR images is based upon both psychophysical evaluation and preference mapping. In order to best generalize the results of this evaluation, a number of synthetic images and photographs were processed using six popular tone mapping operators: Photographic Tone Reproduction (P),²⁵ Uniform Scaling Quantization (Q),²⁰ Retinex (X),^{8,23} Visual Adjustment (V),¹⁸ Revised Tumblin-Rushmeier (T),^{11,14} and Histogram Adjustment (H),¹⁹ which we briefly characterize in Section 2. In a preliminary, exploratory stage of this study, global dissimilarity judgments were made for all pairwise comparisons of the resulting images, and these judgments were submitted to INDividual Differences SCALing (INDSCAL) analysis. Then, in order to aid in the interpretation of the INDSCAL results, psychophysical scale values were obtained for the same images with respect to three perceptual

attributes: apparent contrast, level of detail, and naturalness. Multiple regression analysis revealed that the first perceptual dimension describing the differences between stimuli was associated linearly with apparent level of detail ($r = .85$), and the second INDSCAL-derived dimension was most highly associated with apparent naturalness ($r = .67$). A more complex relationship was found for apparent contrast, the fit regression equation having significant quadratic components on both dimensions of the INDSCAL-derived “Stimulus Space.”

Next, the stimuli were ordered with respect to viewer preference via pairwise forced-choice comparisons, and the resulting ordinal-scaled preference values were submitted to PReFeRence MAPping (PREFMAP) analysis. This analysis related preference values to stimulus coordinates on the INDSCAL-derived dimensions. The results identified an ideal preference point in INDSCAL-derived “Stimulus Space,” the location of which suggested that the ideal tone mapping operator should produce higher-than-average detail, but slightly lower-than-average naturalness (with respect to the stimuli examined). Furthermore, higher contrast images were not preferred nor regarded as more natural; rather, a moderate contrast coupled with relatively higher detail (in geometry and textures) was most preferred. In addition, several unacceptable artifacts were identified, including improper reproduction of overall brightness, cropped luminance around light sources, and loss of detail in the images darker regions.

More details on this research can be found in our report.⁶ For the reader’s convenience we include Figure 1 which summarizes the obtained results presented in that report.




Figure 1. The 2D “group” *Stimulus Space* solution based upon INDSCAL analysis of dissimilarity ratings for six tone mappers applied to four scenes (24 images total). The letter in each circular symbol identifies the tone mapper applied: Revised Tumbin and Rushmeier (T), Photographic Tone Reproduction (P), Uniform Rational Quantization (Q), Histogram Adjustment (H), RetineX (X), and Visual Adaptation (V). The concentric circles correspond to iso-contrast contours, and the diamond symbols plots the “ideal” preference point that resulted from PREFMAP ideal-point unfolding. The x axis of the graph is labeled “Detail” because the coordinates on this dimension correlated most highly with obtained “detail” ratings. The y axis of the graph is labeled “Naturalness” by the same rationale.

The subsequent research reported here is based upon these results, but is mostly focused on an attempt to extend the Retinex algorithm to handle HDR images and produce naturally looking images. Essentially we use Retinex to extract the detailed features of the scene, and then follow this with a global tone mapping function to further restore the brightness and adapt the final image to the displaying method. We describe all the involved steps in detail in the following sections. However, before moving on to the description of the extensions, we provide some background information on standard Retinex algorithms.

4. RETINEX BACKGROUND

The Retinex theory of human vision presented by Edwin Land⁷ was designed to explain perceptual phenomena, and as such seems an ideal basis for a computational method to predict human adaptation to light intensity. Lightness transfer among pixels in a scene has already been applied.^{21,22} In an effort to solve the tone mapping problem Jobson et al.⁹

implemented a Retinex variation based on a Gaussian center-surround function to simultaneously achieve dynamic range compression, color constancy, and lightness rendition. The efficiency of their approach was proven with medium contrast 24-bit images, but high dynamic content caused some artifacts.

Although the principle and basic processing steps were always similar to the first implementation by Land and McCann,⁷ many variations of Retinex have been proposed, mainly differing in the choice and randomness of path in the image plane along which processed pixels are selected. For example Marini et al.³² implemented an approximation of Brownian motion to generate random paths.

We adapted the Retinex variation proposed by Frankle and McCann⁸ and later implemented by Funt et al.²³ This iterative scheme is based on interactions between luminance values for pairs of pixels whose image plane coordinates are denoted as (x, y) and (x_s, y_s) . The number of Retinex iterations at each step of the calculation is decided by the user. The first step processes pixels horizontally separated by a predefined distance D , then the direction for comparison is rotated 90° the computation is repeated a number of time. The distance D is halved at each step until unit pixel distance is reached. In each iteration the ratio of luminance values $R(x, y)$ and $R(x_s, y_s)$ in the original image is multiplied by the luminance value $OP(x_s, y_s)$ computed for the previous Retinex iteration. The resulting value is then averaged with the luminance value $OP(x, y)$. As a result a new luminance value $NP(x, y)$ is obtained:

$$\log NP(x, y) = \frac{(\log OP(x_s, y_s) + \log R(x, y) - \log R(x_s, y_s)) + \log OP(x, y)}{2} \quad (1)$$

For the next iteration, the current value $NP(x, y)$ will replace $OP(x, y)$ and so on. In the first iteration OP is set to the maximum luminance found in the whole image for each spectral color band (in this research we consider only the maximum luminance since we do not perform the Retinex computation for the chroma channels). During Retinex processing, each time the product $L = \log OP(x_s, y_s) + \log R(x, y) - \log R(x_s, y_s)$ in equation (1) is found bigger than maximum luminance, L is reset to maximum luminance. The latter step of the Retinex iteration is called the reset operation.

The initial distance D depends on the size of the image and might be adjusted for particular applications. A small starting distance results in local processing while an infinite path would theoretically generate an image identical to the original. To simulate the HVS, an ideal initial distance D needs to be determined. Following Funt et al.²³ we used the largest power of 2 which is smaller than the maximum image resolution in the horizontal and vertical directions. The optimum number of iterations is dependent on image content and must be determined by the user. Since the overall result does not depend on the size of the image, fast experimentation can be carried using down-sampled images.

Tone mapping processing does not in principle attempt to correct the chromatic values of a scene; so we separated luminance from colors by a conversion from RGB to Yxy. In this paper, further operations will be carried only on the luminance channel (Y). It is of course possible to develop a more involved model to adapt the tone mapping process to each channel.

5. CHANGES MADE TO RETINEX TO PROCESS HIGH DYNAMIC RANGE IMAGES

In our efforts towards developing a perceptually-based tone mapping operator, we noticed that Retinex has the unique capabilities of bringing out detailed features while preserving optimum contrast throughout the resulting image. For this reason, we decided to make changes to the base Retinex algorithm to extend it for use in our tone mapping framework. The biggest concern was the removal of black halos around light sources caused by abrupt changes in contrast. Though Retinex is most often used to enhance contrast and make visible details in dark image regions, in this work we paradoxically use Retinex to do the opposite; i.e. to reduce contrast and to show details hidden in high luminance areas of the scene.

We first attempted to use Retinex as a tone mapping tool in a straightforward manner, but faced the black halo phenomenon described by Chiu.²¹ This is in fact a reverse intensity gradient, which is due to the extreme difference of luminance between close pixels. For example, Figures 2 and 4 show a computer rendered scene with a luminance ratio of 29207:1, the light bulb has a luminance of 500 cd/m^2 while the pixels of its nearest neighborhood average to values between $1-2 \text{ cd/m}^2$, resulting in a change ratio of 500:1 in an image region that is just a few pixels. This is not a particular case, each time a light source is present in a scene such ratios or greater are encountered. To reduce black halos common to local tone mapping operators, we first tried to scale down the high luminance values to the maximum displayable luminance of the output display,⁶ this effectively reduced the problem but we lost most details in high luminance areas, thus canceling the possible benefits of Retinex.




Figure 2. The gradient reversal phenomenon caused by the high luminance ratio between the light bulb and its proximity. Image A was obtained after four iterations of a straightforward Retinex processing. Image B shows the corresponding results after 20 iterations. Increasing the number of iterations leads to the expansion of halo effect and the deformation of its shape. These negative effects can be significantly attenuated by applying a counterclockwise rotation, which is demonstrated in image C for 20 iterations.

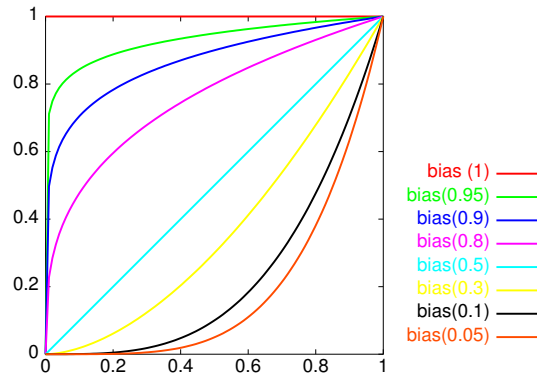


Figure 3. The bias function, with varying parameters from 0.05 to 1. The shape of the curve is somewhat similar to the Gamma function used to correct displays. Using a parameter of 0.5 has no action on the input. We used the bias power function to smoothly control contrast clipping with distance.

Our C-language implementation of Retinex is based on the “Frankle-McCann Retinex” Matlab code presented by Funt et al.²³ In this algorithm, there is no real concern for a path; but as a result of halving the distance for comparison after each series of iterations and using a clockwise rotation, the pixels (x_s, y_s) compared with pixel (x, y) are naturally placed at the corners of a squared spiral centered at (x, y) . Coopers³³ suggested to use a complementary counterclockwise rotation to evenly distribute compared pixels. This proved to be an effective solution to balance dark halos around light sources and prevent their uncontrolled growth with each iteration of the algorithm, masking most details in the scene. Figure 2 shows the growth of the gradient reversal phenomenon with the increasing number of Retinex iterations (Figure 2B) and the benefits of applying a counterclockwise path (Figure 2C).

To remove the black halos around light sources we further investigated Sobol’s³⁴ ratio modification operator to change the nature of the contrast mask. His aim was principally to compress large contrast ratio while keeping or enhancing small ratios to show the details in a photograph. While Sobol originally has limited the scope of his work to a fixed clipping value, he suggested that further modifications of the clipping value should result in smoother contrast masks. A fixed clipping value either removes large contrast ratios and equalizes the brightness of the image if the clipping value is small, or on the other hand does not eliminate gradient reversal. A possible solution is to use a smooth function to vary contrast clipping with the distance (D) used in Retinex computation. We used the bias function (2) proposed by Perlin and Hoffer.³⁵




Figure 4. A: Soft clipping ratio from ± 0.6 to ± 1 with a bias value of 0.8. B: Soft clipping ratio from ± 0.6 to ± 1 with a bias value of 0.95. C: Soft clipping ratio from ± 0.3 to ± 1 with a bias value of 0.95 and 4 iterations of Retinex. Image D uses the same clipping but shows the result of 30 Retinex iterations. Essentially the dark halos are greatly reduced and not further propagated.

This function is a power curve taking a parameter value a in the interval $[0,1]$ to remap a given function by boosting or lowering its value (refer to Figure 3):

$$f(x) = x^{\frac{\log a}{\log 0.5}} \quad (2)$$

The initial intent for the bias function was its use in three-dimensional texture synthesis, however, it proved to be useful for other applications in computer graphics as well. For our purpose the Cartesian distance D between two points of the square spiral path is the input value to the bias function:

$$ContrastClip = D^{\frac{\log a}{\log 0.5}} - 0.7 \quad (3)$$

The constant 0.7 was arbitrary decided to avoid values lower than 0.3 at which artifacts such as a blurred image appear. A bias parameter $a = 0.5$ will have essentially no effect on the halos while $a = 0.95$ greatly compresses differences of contrast and $a = 1$ will clip contrast to ± 0.3 throughout the image. A bias parameter between 0.85 and 0.9 seems appropriate in most cases. The value for $\log R(x,y) - \log R(x_s,y_s)$ in equation (1) is set to $\pm ContrastClip$ if greater or smaller than $ContrastClip$.

Our last modification of McCann’s ratio-product-reset-average Retinex processing is a different selection of the reset value. In standard Retinex, the reset value is the maximum luminance found in the scene. For HDR images where the maximum value is far greater than the display capabilities, setting the maximum value to the maximum luminance of the output medium helped in retrieving details in very high luminance areas. This is evident in Figures 5C and 5D, where 5C was computed using the maximum value equal to $\log_{10}(350000)$ (the maximum luminance in the scene), while for 5D the maximum was set to $\log_{10}(100)$ (which roughly corresponds to the maximum luminance of a typical CRT display).

6. TONE MAPPING

Retinex processing retrieved most of the details in the scene at every luminance level. Surprisingly, the resulting images might not look very realistic. In most of the cases, the semantic information of the scene has been lost, even if a good level of contrast is preserved, essential informations such as light source position and strength are not apparent anymore. Another concern is the brightness of the resulting image which depends on the number of Retinex iterations. Usually a trial and error approach is needed to decide upon the optimum number of iterations, and the choice about a satisfying brightness is arbitrary made. To address these two problems, we further processed the Retinex image with automatic tone mapping functions based on modeling of the HVS.

Most of the algorithms presented in Section 2 can be applied for this task, the resulting image will be a detail enhanced version of a straightforward approach. Since the maximum luminance of the scene is reduced to 100 cd/m^2 or less by Retinex processing, the Visual Adjustment¹⁸ operator based on contrast sensitivity and developed to display natural images is a good candidate for this post processing. We also applied our still in development tone mapping function. In this scheme, the luminance of the scene is sampled with varying logarithmic bases depending on scene content. In both cases




Figure 5. The Stanford Memorial Church, image courtesy of Paul Debevec. Luminance ratio $\approx 350000 : 1$. Image A is displayed with a linear luminance mapping, B is the result of 4 straightforward Retinex iterations. C and D are the results of our extension to the Retinex algorithm. In the case of image D the maximum luminance was set to 100 cd/m^2 . All the images have been further corrected with a gamma value of 2.2.



Figure 6. Final result after applying two different tone mapping functions to the Retinex processed scene. Left image was corrected by Ferwerda et al. Visual Adjustment, on the right a varying logarithmic correction was applied.

the luminance adaptation to the scene was based on the logarithmic average of scene luminance. Figure 6 shows the result of the two tone mapping functions applied to Figure 5D. This final tone mapping operation is essential to preserve the appearance of natural lighting conditions and remove the subjective decision about the ideal number of Retinex iterations needed to match scene brightness.

7. CONCLUSIONS

We have demonstrated the usefulness of Retinex processing in the development of a perceptually-based tone mapping operator. Limitations in Retinex with regard to the handling of high contrast scenes were corrected essentially by a counterclockwise rotation of the path, spatially varying levels of pixel interaction based on luminance contrast information, and a reset ratio adjusted to the maximum luminance capability of the output medium instead of the maximum luminance of the scene. The tone mapping operation in itself was carried using global tone mapping functions, mainly to accurately adjust the brightness of the resulting image and retrieve the visual impression of light sources placed in the field of view. Further developments of a more specific tone mapping function adapted to Retinex output are currently planned.

ACKNOWLEDGMENTS

We would like to thank Peter Shirley and Paul Debevec for making their HDR images available. This work was supported in part by the European Community within the scope of the RealReflect project IST-2001-34744 “Realtime visualization of complex reflectance behavior in virtual prototyping”.

REFERENCES

1. G. J. Ward, "The RADIANCE Lighting Simulation and Rendering System," in *SIGGRAPH 94 Conference Proceedings*, pp. 459–472, 1994.
2. P. E. Debevec and J. Malik, "Recovering High Dynamic Range Radiance Maps from Photographs," in *SIGGRAPH 97 Conference Proceedings*, pp. 369–378, 1997.
3. P. Ledda, A. Chalmers, and G. Ward, "Perceptual Tone Mapping Operators for High Dynamic Range Displays," in *Siggraph 02: Conference Abstracts and Applications*, p. 229, 2002.
4. F. Drago and K. Myszkowski, "Validation Proposal for Global Illumination and Rendering Techniques," *Computers & Graphics* **25**, pp. 511–518, June 2001.
5. B. Wandell, *Foundations of Vision*, Sinauer Associates, Inc., Sunderland, Massachusetts, 1995.
6. F. Drago, W. Martens, K. Myszkowski, and H.-P. Seidel, "Perceptual Evaluation of Tone Mapping Operators with Regard to Similarity and Preference," Technical Report MPI-I-2002-4-002, Max-Planck-Institut fuer Informatik, Oct. 2002. <http://data.mpi-sb.mpg.de/internet/reports.nsf/AG4NumberView?OpenView>.
7. E. H. Land and J. J. McCann, "Lightness and the Retinex Theory," *Journal of the Optical Society of America* **61**(1), pp. 1–11, 1971.
8. J. Frankle and J. J. McCann, "Method and Apparatus for Lightness Imaging," *US Patent No: 4,384,336*, 1983.
9. D. J. Jobson, Z. Rahman, and G. A. Woodell, "Properties and Performance of a Center/Surround Retinex," *IEEE Transactions on Image Processing* **6**(3), pp. 451–462, 1997.
10. K. Matkovic, L. Neumann, and W. Purgathofer, "A Survey of Tone Mapping Techniques," in *13th Spring Conference on Computer Graphics*, pp. 163–170, Comenius University, Bratislava, Slovakia, June 1997.
11. J. Tumblin, *Three Methods For Detail-Preserving Contrast Reduction For Displayed Images*. PhD thesis, Georgia Institute of Technology, 1999.
12. A. McNamara, "Visual Perception in Realistic Image Synthesis," *Computer Graphics Forum* **20**(4), pp. 211–224, 2001.
13. K. Devlin, A. Chalmers, A. Wilkie, and W. Purgathofer, "Tone Reproduction and Physically Based Spectral Rendering," in *Eurographics 2002: State of the Art Reports*, pp. 101–123, Eurographics, 2002.
14. J. Tumblin and H. E. Rushmeier, "Tone Reproduction for Realistic Images," *IEEE Computer Graphics and Applications* **13**, pp. 42–48, Nov. 1993.
15. S. Stevens and J. Stevens, "Brightness Function: Parametric Effects of Adaptation and Contrast," *Journal of the Optical Society of America* **50**, p. 1139A, Nov. 1960.
16. J. Tumblin, J. K. Hodgins, and B. K. Guenter, "Two Methods for Display of High Contrast Images," *ACM Transactions on Graphics* **18**, pp. 56–94, January 1999.
17. G. Ward, "A Contrast-Based Scalefactor for Luminance Display," *Graphics Gems IV*, pp. 415–421, 1994.
18. J. A. Ferwerda, S. Pattanaik, P. Shirley, and D. P. Greenberg, "A Model of Visual Adaptation for Realistic Image Synthesis," in *SIGGRAPH 96 Conference Proceedings*, pp. 249–258, 1996.
19. G. W. Larson, H. Rushmeier, and C. Piatko, "A Visibility Matching Tone Reproduction Operator for High Dynamic Range Scenes," *IEEE Transactions on Visualization and Computer Graphics* **3**(4), pp. 291–306, 1997.
20. C. Schlick, "Quantization Techniques for the Visualization of High Dynamic Range Pictures," in *Photorealistic Rendering Techniques*, pp. 7–20, Springer-Verlag, 1994.
21. K. Chiu, M. Herf, P. Shirley, S. Swamy, C. Wang, and K. Zimmerman, "Spatially Nonuniform Scaling Functions for High-Contrast Images," in *In Graphics Interface '93*, pp. 245–253, 1993.
22. S. N. Pattanaik, J. A. Ferwerda, M. D. Fairchild, and D. P. Greenberg, "A Multiscale Model of Adaptation and Spatial Vision for Realistic Image Display," in *SIGGRAPH 98 Conference Proceedings*, pp. 287–298, 1998.
23. B. Funt, F. Ciuera, and J. McCann, "Retinex in Matlab," in *Proc. IS&T/SID Eighth Color Imaging Conference*, pp. 112–121, 2000.
24. J. Tumblin and G. Turk, "LCIS: A Boundary Hierarchy for Detail-Preserving Contrast Reduction," in *SIGGRAPH 99 Conference Proceedings*, pp. 83–90, 1999.
25. E. Reinhard, M. Stark, P. Shirley, and J. Ferwerda, "Photographic Tone Reproduction for Digital Images," in *SIGGRAPH 02 Conference Proceedings*, 2002.
26. E. Reinhard, "Parameter Estimation for Photographic Tone Reproduction," *Journal of Graphics Tools* **7**(1), pp. 45–52, 2002.

27. M. Ashikhmin, "A Tone Mapping Algorithm for High Contrast Images," in *Rendering Techniques 2002*, pp. 145–156, Eurographics, 2002.
28. R. Fattal, D. Lischinski, and M. Werman, "Gradient Domain High Dynamic Range Compression," *ACM Transactions on Graphics* **21**, pp. 249–256, July 2002.
29. F. Durand and J. Dorsey, "Fast Bilateral Filtering for the Display of High-Dynamic-Range Images," *ACM Transactions on Graphics* **21**(3), pp. 257–266, 2002.
30. W. Hunt, *The Reproduction of Colour*, Fountain Press, England, 1995.
31. M. D. Fairchild, *Color Appearance Models*, Addison-Wesley, Reading, Massachusetts, 1998.
32. D. Marini, A. Rizzi, and C. Carati, "Color Constancy Effects Measurements of the Retinex Theory," in *IS&T/SPIE's Electronic Imaging 1999*, January 1999.
33. T. Cooper, "Modifications to Retinex to Relax RESET Nonlinearity and Implement Segmentation Constraints," in *Proc. IS&T/ SPIE Symposium on Electronic Imaging: Science and Technology*, San José January 2002.
34. R. Sobol, "Improving the Retinex Algorithm for Rendering Wide Dynamic Range Photographs," in *Proc. IS&T/ SPIE Symposium on Electronic Imaging: Science and Technology*, San José January 2002.
35. K. Perlin and E. M. Hoffert, "Hypertexture," in *Computer Graphics (Proceedings of SIGGRAPH 89)*, **23**, pp. 253–262, (Boston, Massachusetts), July 1989.