

An Optimal Tone Reproduction Curve Operator for the Display of High Dynamic Range Images

Guoping Qiu and Jiang Duan

{qiu | jxd}@cs.nott.ac.uk

School of Computer Science, The University of Nottingham

Jubilee Campus, Nottingham NG8 1BB, United Kingdom

Thursday, 26 February 2004

Please send all correspondences to

Dr Guoping Qiu
School of Computer Science
The University of Nottingham
Nottingham NG8 1BB
United Kingdom

Email: qiu@cs.nott.ac.uk
Fax +44 115 951 4254

Abstract

We present a new tone mapping method for the display of high dynamic range images in low dynamic range devices. We formulate high dynamic range image tone mapping as an optimisation problem. We introduce a two-term cost function, the first term favours linear scaling mapping, the second term favours histogram equalisation mapping, and jointly optimising the two terms optimally maps a high dynamic range image to a low dynamic range image. We control the mapping results by adjusting the relative weightings of the two terms in the objective function. We also present a fast and simple implementation for solving the optimisation problem. We will present results to demonstrate that our method works very effectively.

1. Introduction

The ultimate goals of computer graphics and digital imaging are to develop systems that match or maybe even exceed the capabilities of the human visual system (HVS). One of the remarkable abilities of the HVS is that it can perceive real world scenes with brightness dynamic ranges of over 5 orders of magnitude and can distinguish even higher contrast through adaptation. Ordinary imaging sensors and reproduction media typically have dynamic ranges spanning a few orders of magnitude, much lower than those of the real world scenes.

High dynamic range (HDR) imaging technologies are designed to produce images that faithfully depict the full visual dynamics of real world scenes. There are two major technical challenges in HDR imaging, image capture and image reproduction. In the literature, several methods have been developed for creating

high dynamic range still images [1, 2] and videos [3]. In order to display HDR images on monitors or print them on paper, we must compress the dynamic range of the HDR images to reproduce low dynamic range (LDR) images supported by these media. In recent years, several techniques have been developed by various researchers for mapping HDR images to LDR images for display. Please see the Related Work box for a brief review of these methods available in the open literature.

In this article, we present a new optimization method for mapping HDR images to LDR images. For a given HDR image, our method first defines an objective function. An optimal solution to the objective function corresponds to an optimal mapping of the HDR image to an LDR image for display. We will present a fast computational solution to solve the optimization problem. Our new method only manipulates the pixel distribution of an image and is a tone reproduction curve (TRC) based high dynamic range compression method. Our final algorithm has one single variable, which the users can change to control a reproduction to suit individual images and users' subjective preference. Compared with existing HDR image compression methods, our technique is more flexible and computational much more efficient. We will show that our new method has an excellent performance in mapping a variety of high dynamic range scenes.

2. Optimal Tone Reproduction Curve Tone Mapping for HDR Images

For high dynamic range image tone mapping, there are at least two requirements. Firstly, it has to ensure that all features, from the darkest to the brightest, to be visible simultaneously. Secondly, it has to

preserve the original scene's visual contrast impression to produce a visual sensation matching that of the original scene.

To preserve the original scene's relative visual contrast impression, the simplest approach is to linearly map the pixels from a high dynamic range to a low dynamic range. However, since the dynamic range in the display devices is much narrower than that of the original scene, visibility will be lost due to compression. Also, linear scaling maps all values in the same way, some displayable values in the low dynamic range device may be empty or have too few pixels mapped onto them thus resulting in an under utilization of all displayable values. On the other extreme, one can render the low dynamic range image to have a maximum contrast, i.e., histogram equalized by distributing equal number of pixels to each display levels. However, this will alter the original scene's visual impression, because it exaggerates contrast in densely populated pixel value intervals while compresses too aggressively sparsely populated pixel value intervals.

The TRC based HDR image tone mapping problem can be conveniently explained by referencing to Figure 1. Essentially, there are more (discrete) levels in the input HDR axis than in the output LDR axis. If the HDR data is represented by 32-bit floating point number, then there are 4,294,967,296 possible levels in the HDR input. For the LDR output, typically, there will be 256 levels (can be any other numbers, but we will use this number for convenience). Clearly and straightforwardly, some HDR values will have to be merged and represented by the same LDR value. To do this, we cut the HDR dynamic range axis into 256 intervals at positions c_1, c_2, \dots, c_{255} . HDR values falling into the same interval are then assigned the same LDR display value. For convenience, we assume both the HDR range and the total number of pixels in the image are unity. If we divide the HDR input axis in such a way that pixel populations falling into each interval are identical and equal to $1/256$, then this is histogram equalization

mapping. On the other hand, if we make each interval to have the same length and equal to $1/256$, then this is linear scaling mapping.

A TRC based tone mapping algorithm should assign relatively more display values to densely populated luminance intervals and relatively fewer display values to sparsely populated luminance intervals. A TRC based tone mapping algorithm should also maintain the relative visual contrast impressions of the original scene while simultaneously trying to fully utilize all available display levels.

Therefore, with reference to Figure 1, c_1, c_2, \dots, c_{255} must fall in between linear scaling and histogram equalization. For illustration purpose, Figures 3 and 4 show examples of mapping 8bit/pixel images to 1bit/pixel images for display. In Figure 3, the majority of pixels have values greater than mid grey (128). The overall impression of this image is bright. The display image should also ensure that there are more bright pixels than dark pixels, which means that the cut should be below c_e . It should not be below c_l either for obvious reasons. In Figure 4, there are more pixels have values below mid-grey, the image appears dark. The display should also have more dark pixels than bright pixels and this means that the cut should be above c_e . In this case, the cut should be below c_l . In these examples, we clearly see that if the mappings are outside the linear scaling or histogram equalization lines, then the mapped images will either have too many dark pixels or too many bright pixels. As a consequence, the overall visual impressions of the displays differ drastically from those of the originals. By cutting the image in between the linear scaling and histogram equalization lines, both the details and visual impressions of the original images are better preserved.

To formalize this mapping principle, again, with reference to Figure 1, we can define the following objective function

$$E = \sum_{i=1}^{255} \left(c_i - \frac{i}{256} \right)^2 + \lambda \sum_{i=1}^{255} \left(\int_0^{c_i} h(x) dx - \frac{i}{256} \right)^2 \quad (1)$$

where λ is the Lagrange multiplier. Setting $\lambda = \infty$, optimizing E becomes histogram equalization mapping, and $\lambda = 0$, optimizing E becomes linear scaling mapping. By choosing an appropriate λ , we can strike a balance between the two extreme forms of mapping to suit individual images. An optimal solution to (1) can be found by solving following linear equations

$$\frac{\partial E}{\partial c_i} = 0 \quad i = 1, 2, \dots, 255 \quad (2)$$

3. A Fast Implementation

Equation (2) may be seriously ill conditioned, and a straightforward numerical solution to optimize E in (1) may be difficult to obtain. We here present an approximated fast solution. Instead of trying to find all cuts in one go, we use a recursive binary cut approach. We first divide the dynamic range of the HDR image into two intervals based on a modified objective function of (1). These resultant two intervals are then subsequently and independently divided into two intervals in a similar way. The schematic is illustrated in Figure 2. We first find c_0 , which divides the full dynamic range into two intervals, we then find $c_{1,0}$ and $c_{1,1}$, which divide the resultant two intervals from previous level into 4 intervals. The process continues, until the desired numbers of intervals are found.

With the scheme in Figure 2, we can now reformulate (1) for each interval we want to divide into two. Assuming that the total length of the interval is L , the minimum is L_{\min} and the maximum is L_{\max} , the number of pixels falling within the interval is N , we find c that cuts the interval into two by optimizing the following objective function

$$E_b(c) = \frac{(c - 0.5(L_{\max} + L_{\min}))^2}{L^2} + \lambda \frac{\left(\left(\sum_{x=0}^c h(x) \right) - 0.5N \right)^2}{N^2} \quad (3)$$

Now to find a c that minimizes $E_b(c)$ in (3) is simple. We simply compute all possible values of E_b , $E_b(L_{\min})$, $E_b(L_{\min} + \Delta L)$, $E_b(L_{\min} + 2\Delta L)$, ..., $E_b(L_{\min} + n\Delta L)$, ..., $E_b(L_{\max})$ and find the minimum:

$$c = \arg \left(\min_x (E_b(x)) \right) \quad (4)$$

ΔL controls the precision of the division. We used the same ΔL for all the segments in the hierarchy of Figure 2 and we found that setting ΔL between 0.01% ~ 0.1% of the full dynamic range of the HDR image worked very well. Obviously, the smaller ΔL is, the more accurate is the division and the longer it will take to compute the optimization. Using a PC with a 2.66 GHz Pentium 4 Processor, calculating a 256-interval division according to (3) and (4) takes 0.016 to 0.672 second for ΔL ranging between 0.1% and 0.01% (codes written in C and not optimized for speed). For each intermediate segment, we apply (3) and (4) to divide it into two segments. Although it is possible to vary λ for different levels of the hierarchy in Figure 2, for simplicity, we use the same λ for all the segments.

4. Results

Our technique has been tested on a variety of high dynamic range images. In our experiments, the luminance signal is calculated as: $L = 0.299*R+0.587*G+0.114*B$. $\text{Log}(L)$ is computed to compile a histogram. After mapping, the LDR images are rendered for display using following formula

$$R_{out} = \left(\frac{R_{in}}{L_{in}} \right)^{\gamma} L_{out}, G_{out} = \left(\frac{G_{in}}{L_{in}} \right)^{\gamma} L_{out}, B_{out} = \left(\frac{B_{in}}{L_{in}} \right)^{\gamma} L_{out} \quad (5)$$

where L_{in} and L_{out} are luminance values before and after compression, γ controls display color (setting it between 0.4 and 0.6 worked well). How to compute the mapped luminance for display devices is a well-studied problem [4]. In our experiments, we simply give all pixels mapped to the first (lowest) interval a display value of 0, those mapped to the second interval a display value of 1, and so on. We did not carry out any color correction for the mapped images. Notice that for the best viewing effects, it is necessarily to carry out the correct calibration of the display devices. Because compression will inevitably loose some fine details, we found that using a standard sharpening filter can improve the spatial sharpness of the mapped images.

Our new method is very flexible. By controlling the value of λ , we can control the final appearance of the image in a simple and elegant way. Figures 5 shows results of mapping HDR images using different values of λ . A smaller λ renders the image with lower contrast whilst a larger λ renders the image with higher contrast. By choosing an appropriate λ , we can achieve a balanced reproduction. Figure 6 shows more examples of HDR images mapped by our new method.

Figure 7 shows the Memorial Church image mapped by our new method and several other methods in the literature. It is seen that our results compared very favorably with these methods (It is to be noted that this comparison is informal. Different images may be best viewed with different display gamma. Other factors, such as viewing environment also affect the appearances greatly).

It is interesting to note that apart from Ward Larson's histogram adjustment method, all others involve spatial processing and are therefore computationally more demanding than our method. Furthermore, these spatial operators have more manually adjusted parameters. Compared with Ward Larson's histogram adjustment method, our method is much more flexible and more general. For example, the histogram adjustment method only prevents the display contrast (produced by simple histogram equalization) from exceeding that of the real scene, but this will not prevent the display from having too low contrast. Whilst the histogram adjustment method is restrictive and inflexible, our method can adjust the contrast of the display by simply changing a single parameter.

Computationally, our method is very fast, using non optimized C code running on a PC with a 2.66 GHz Pentium 4 Processor, mapping a 512 x 768 pixel image takes 0.75 second, a 768 x 1024 image takes 1.28 seconds, and a 2048 x 1536 image takes 5.375 seconds (note that these times do not include file I/O operations but include all other necessary computations).

5. Concluding Remarks

In this article, we have introduced an optimization method for HDR image tone mapping. We have introduced an objective function and developed a fast method to solve such an optimization problem. Experimental results demonstrated that our new method worked very well on a variety of HDR images. The merits of our new method includes that it is intuitive, simple, flexible and effective.

References

1. P. E. Debevec and J. Malik, “Recovering high dynamic range radiance maps from photographs”, Proc. ACM SIGGRAPH’97, pp. 369 – 378, 1997
2. T. Mitsunaga and S. K. Nayar, “High dynamic range imaging: Spatially varying pixel exposures”, Proc. CVPR’2000, vol. 1, pp. 472-479, 2000
3. S. B. Kang, M. Uyttendale, S. Winder and R. Szeliski, “High dynamic range video”, ACM Transactions on Graphics, vol.22, no. 3, Pages: 319 – 325, July 2003
4. R. Hall, Illumination and color in computer generated imagery, Springer-Verlag, 1989

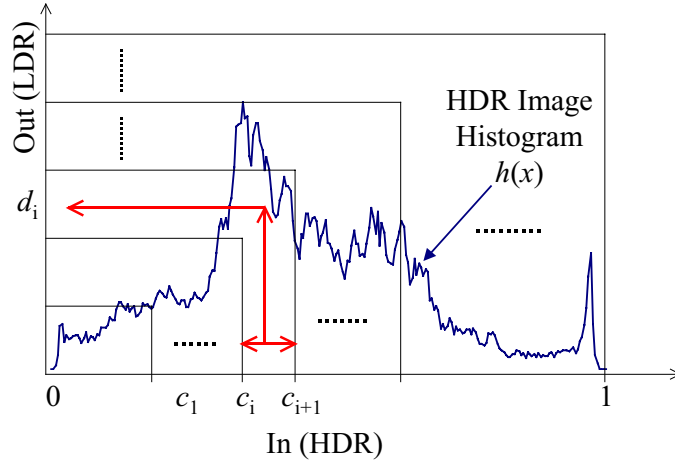


Figure 1, Tone reproduction curve based HDR image tone mapping problem can be solved by dividing the HDR range axis into segments, HDR pixels falling into the same segment are mapped to the same display value in the LDR image.




Figure 2, The full dynamic range of the HDR image is first divided into two intervals, each of which is then again independently divided into two intervals. The process is applied recursively until the desired numbers of intervals are created.

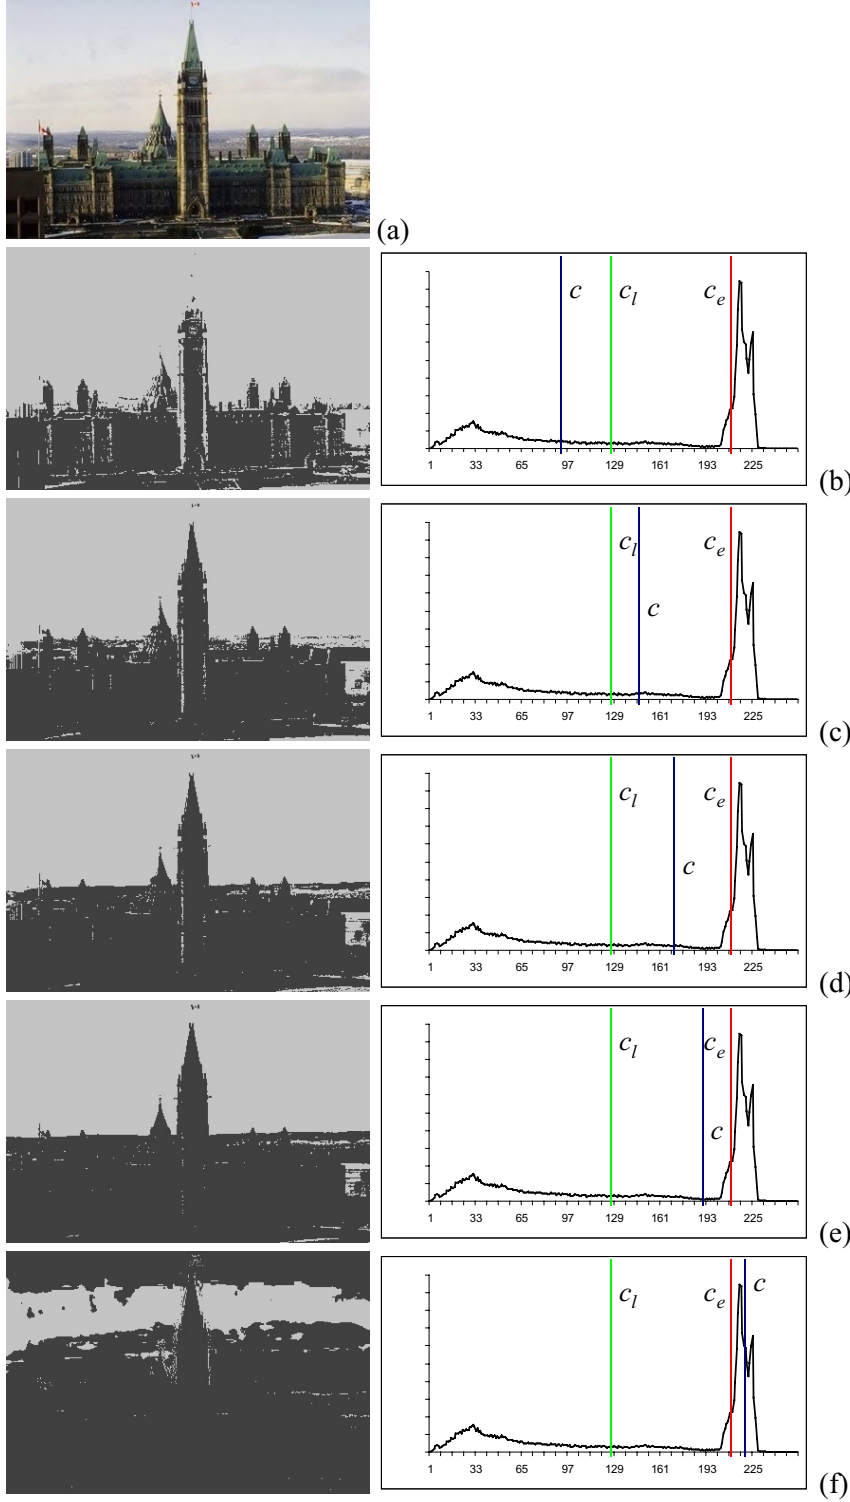


Figure 3, Mapping an 8bit/pixel image to 1bit/pixel image for display. Linear mapping cut the HDR range at $c_l = 128$ and histogram equalization mapping cut the HDR range at $c_e = 208$. Pixels above c are displayed as bright and those below c are displayed as dark pixels. (a) Original image, the overall impression of this image is bright. (b) Cut below linear mapping. (c), (d) and (e), Cuts between linear scaling and histogram equalization. (f) Cut outside histogram equalization. Outside c_l and c_e , there are either too many bright or too many dark pixels. Cuts in between c_l and c_e , visual impressions are more faithful to that of the original and also have more details.

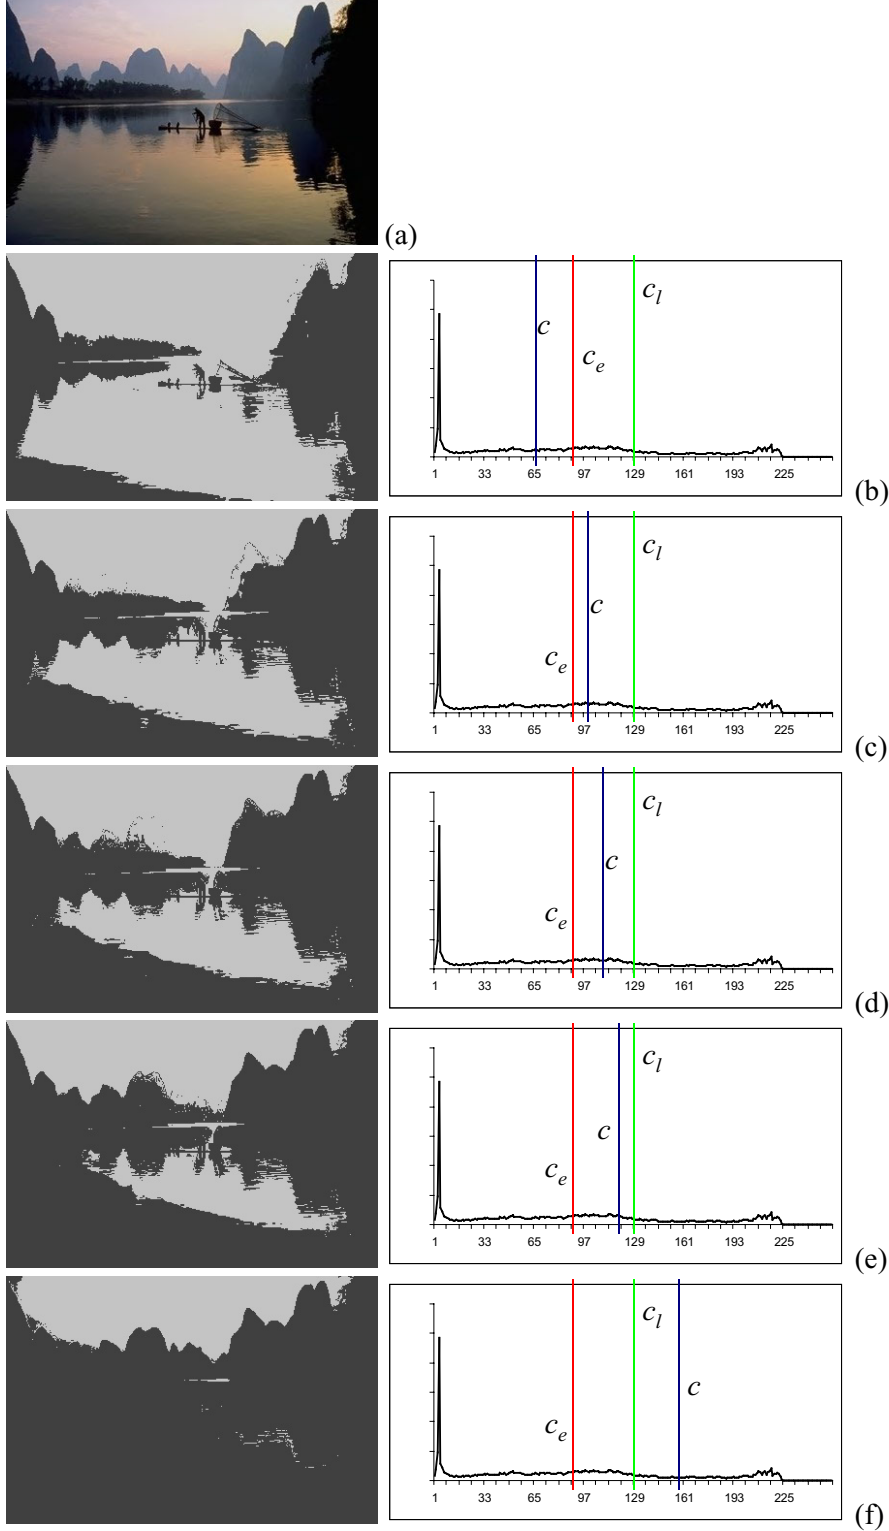


Figure 4, Mapping an 8bit/pixel image to 1bit/pixel image for display. Linear mapping cut the HDR range at $c_l = 128$ and histogram equalization mapping cut the HDR range at $c_e = 88$. Pixels above c are displayed as bright and those below c are displayed as dark pixels. (a) Original image, the overall impression of this image is dark. (b) Cut below linear mapping. (c), (d) and (e), Cuts between linear scaling and histogram equalization. (f) Cut outside histogram equalization. Outside c_l and c_e , there are either too many bright or too many dark pixels. Cuts in between c_l and c_e , visual impressions are more faithful to that of the original and also have more details.




Figure 5, Mapping HDR radiance maps to LDR images with different contrasts and visibilities by controlling the parameter λ . The HDR radiance maps were created with multiple photos of the same scene using Paul Debevec's HDRShop software. The radiance map of the left column image has a dynamic range of 488,582:1 and the radiance map of the right column image has a dynamic range of 602,055:1. The two images in the top row were mapped by setting $\lambda = 0$ (linear scaling). The images in the middle row were mapped by setting $\lambda = 0.5$, and the images in the bottom row were mapped by setting $\lambda = 1$. It is seen that using a smaller λ , the operator maps the radiance map to an image with less contrast whilst using a larger λ , the operator maps the radiance map to an image with increased contrast. By simply changing the value of λ , we can obtain a reproduction with the right amount of visibility and contrast.



Figure 6, Results of various HDR images mapped by our new method. Notice that no effort has been spent on selecting the best λ for individual images. Data courtesy of Paul Debevec and Sumant Pattanaik.




Figure 7, Memorial Church image. HDR data courtesy of Paul Debevec. (a) Result of Ward Larson’s histogram adjustment technique, image courtesy of Dani Lischinski. (b) Result of Fredo Durand’s bilateral filtering. (c) Results of our new optimal method, from left to right, $\lambda = 0.3$ and 0.7 . (d) Erik Reinhard’s photographic tone mapping method. (e) Result of Dani Lischinski’s gradient domain method. Images printed with the kind permission of the authors. Note that this comparison is informal. Different images may be best viewed with different gamma.

Related Methods

In the literature, a number of techniques have been developed for tone reproduction for high contrast images. There are two broad categories of technology [1]. Tone reproduction curve (TRC) based techniques manipulate the pixel distributions. Earlier pioneering work in this category include that of [2] which introduced a tone reproduction method that attempted to match display brightness with real world sensations. Recently, [3] presented a tone mapping method that modeled some aspects of human visual system. More recently, we have also developed a learning-based TRC tone mapping method for high dynamic range compression [12]. Perhaps the most comprehensive technique in this category is that of [4], which introduced a quite sophisticated tone reproduction curve technique that incorporated models of human contrast sensitivity, glare, spatial acuity and color sensitivity to exploit the limitations of human visual system.

Tone reproduction operator (TRO) based techniques involve the spatial manipulation of local neighboring pixel values, often at multiple scales. The scientific principle of this type of technique is based on the image formation model: $I(x, y) = L(x, y) R(x, y)$, which states that image intensity function $I(x, y)$ is the product of the luminance function $L(x, y)$ and the scene reflectance function $R(x, y)$. Because real world reflectance $R(x, y)$ has low dynamic range (normally not exceeding 100:1), reducing the dynamic range of $I(x, y)$ can be achieved by reducing the dynamic range of $L(x, y)$ if one could separate $L(x, y)$ from $R(x, y)$. Methods based on this principle include [5], [6] and [7]. They mainly differ in the way in which they attempted to separate the luminance component from the reflectance component. All TRO based methods can be regarded as related to the Retinex theory [10]. A direct use of the Retinex theory for high dynamic range compression was presented by Jobson and co-workers [11].

Recent development has also attempted to incorporate traditional photographic technology to the digital domain for the reproduction of high dynamic range images [8].

An impressive latest development in high dynamic range compression is that of [9]. Based on the observation that human visual system is only sensitive to relative local contrast, the authors developed a multiresolution gradient domain technique. This is also a TRO type technique and the authors reported very good results that are free from halo effects.

TRO based methods involve multiresolution spatial processing and are therefore computationally very expensive. Because TRO methods can reverse local contrast, they can sometimes cause “halo” effects in the reproduction. Another difficulty of these techniques is that there are too many parameters the users have to set and this makes them quite difficult to use. In many cases, the setting of these controlling parameters is rather *ac hoc* and involves many trial and errors.

TRC based methods do not involve spatial processing, they are therefore computationally very simple. This is useful in real time applications such as high dynamic range video. TRC techniques also preserve the lightness orders of the original scenes and avoid artifacts such as halo that is often associated with TRO based methods.

One of the weaknesses of TRC approaches as compared with TRO methods is that it may cause noticeable lost of spatial sharpness in some images. In many cases, this is not a serious problem if the TRC method is well designed and a simple standard image sharpening operation often suffices to bring

back the sharpness. However, TRO methods sometimes could introduce too much (artificial) detail. Both types of techniques have their own merits and drawbacks in terms of computational complexity, easy implementation and practical application, and therefore are likely to co-exist for tone mapping in HDR imaging for the foreseeable future.

References

1. J. DiCarlo and B. Wandell, “Rendering high dynamic range images”, Proc. SPIE, vol.3965, pp. 392 – 401, 2001
2. J. Tumblin and H. Rushmeier, “Tone reproduction for realistic images”, IEEE Computer Graphics and Applications, vol. 13, pp. 42 – 48, 1993
3. M. Ashikhmin, “A tone mapping algorithm for high contrast images”, Proc. Eurographics Workshop on Rendering, P. Debevec and S. Gibson Eds., pp. 1 – 11, 2002
4. G. W. Larson, H. Rushmeier and C. Piatko, “A visibility matching tone reproduction operator for high dynamic range scenes”, IEEE Trans on Visualization and Computer Graphics, vol. 3, pp. 291 – 306, 1997
5. K. Chiu, M. Herf, P. Shirley, S. Swamy, C. Wang and K. Zimmerman, “Spatially nonuniform scaling functions for high contrast images”, Proc. graphics Interface’93, pp. 245 – 253, 1993
6. J. Tumblin and G. Turk, “LCIS: A boundary hierarchy for detail preserving contrast reduction”, ACM SIGGRAPH 1999

7. F. Durand and J. Dorsey, "Fast bilateral filtering for the display of high-dynamic-range images", Proc. ACM SIGGRAPH'2002
8. E. Reinhard, M. Stark, P. Shirley and J. Ferwerda, "Photographic tone reproduction for digital images", Proc. ACM SIGGRAPH'2002
9. R. Fattal, D. Lischinski and M. Werman, "Gradient domain high dynamic range compression", Proc. ACM SIGGRAPH'2002
10. E. H. Land and J. J. McCann, "Lightness and retinex theory", Journal of the Optical society of America, vol. 61, pp. 1-11, 1971
11. D. J. Jobson, Z. Rahman and G. A. Woodell, "A multiscale Retinex for bridging the gap between color images and the human observation of scenes", IEEE Transactions on Image processing, vol. 6, pp. 965-976, 1997
12. J. Duan, G. Qiu and G. D. Finlayson, "Learning to display high dynamic range images", CGIV'2004, IS&T's Second European Conference on Color in Graphics, Imaging and Vision, Aachen, Germany, April 5-8, 2004