Display HDR Image using a Gain Map

Jian Guan and Guoping Qiu

School of Computer Science, University of Nottingham

Fig. 1. Direct display of an HDR map in the logarithmic scale (left); the gain map image of the HDR map (middle); the sum of the left and middle images (right). Note the gain map image is only computed for the luminance hence it is shown as grayscale image. Radiance map courtesy of Raanan Fattal, Dani Lischinski and Michael Werman

Abstract - **In this paper, we present a novel method for the display of high dynamic range images. The new method first computes a gain map image using a computational approach inspired by a machine learning algorithm and sums the gain map and the original image together; it then linearly scales the sum image to fit the dynamic range of the display devices. Results are presented to demonstrate the effectiveness of this new method and it is also shown that the new approach is an effective method for enhancing standard (8bits/pixl) images.**

Index Terms – Tone mapping, high dynamic range imaging, display, machine learning

I. INTRODUCTION

Today, digital cameras are ubiquitous. However, when imaging scenes containing wide variations of illumination intensities, the picture quality often turns out to be less than satisfactory. In such cases, the main cause of the poor image quality is the mismatch between the dynamic range of irradiance of the real-world scene and the number of binary bits used to represent pixel values in the standard image formats. Whilst the real world irradiance can have dynamic ranges exceed four to five orders of magnitude, typical standard image formats using 8 bits per pixel can only represent part of the real world dynamic range. The situations can be remedied by using high dynamic range (HDR) imaging technology [1-4] where the so called HDR radiance maps (>32bits/pixel) can record the actual dynamic range of the real world scenes; however, there is still the problem of faithfully reproducing the image in conventional low dynamic range (LDR) reproduction media such as print paper and monitors which normally have a useful contrast about 2 orders of magnitude.

Processing HDR maps for reproduction in LDR media is often called tone mapping or dynamic range compression. Even though high dynamic range display devices have started to emerge, they are very expensive. In the foreseeable future, conventional LDR devices will still be the dominant reproduction media for HDR pictures and hence effective processing techniques for the display of HDR images in LDR devices are still important.

In this paper, we present a novel method for the processing of HDR images for display in LDR media. Our idea is to "invent" a high dynamic range gain map image (GMI) the same size as the original image, which when summed with the original image will produce a contrast enhanced version of the original high dynamic range image. The enhancement is controlled in such a way that weak local contrasts are enhanced more whilst strong local contrasts are enhanced less. By linearly scaling down the new high dynamic range image to fit the dynamic range of the reproduction devices results in a LDR version of the HDR image in which low contrasts are boosted and high contrasts are suppressed thus achieving dynamic range reduction without causing heavy loss of visual details. We burrow ideas from a manifold learning technique [6] and formulate the problem of computing the GMI as a linearly constrained optimization problem and compute the GMI by solving a linear system of equations. Furthermore, we show that the GMI idea can be used to enhance standard (8 bits/pixel) images.

II. DISPLAY HDR IMAGE USING A GAIN MAP

The display pipeline of our new scheme is illustrated in Fig. 2. From the HDR map, $\boldsymbol{I}$, we compute a gain map image (GMI), $\boldsymbol{G}$, and the sum of the two, $\boldsymbol{I} + \boldsymbol{G}$, is then linearly scaled to fit the dynamic range of the display device for display. Note that $\boldsymbol{I}$ is in logarithmic scale, $\boldsymbol{G}$ is derived from $\boldsymbol{I}$, hence also in logarithmic scale, so is $\boldsymbol{I} + \boldsymbol{G}$. Unlike several other methods in the literature, e.g. [2, 3], we do not take the antilogarithm of the processed signal but display it directly because the HDR map data is linear and the display devices are nonlinear having an antilogarithmic-like curve.

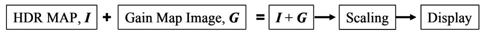

Fig. 2: Schematic of HDR map using gain map image, where the HDR map is in logarithm scale.

To compress the dynamic range of the HDR map to fit the dynamic range of the LDR reproduction media, linear scaling should have been the simplest and the correct way because linear scaling preserves the relation of the pixels, i.e., a relatively bright HDR pixel will be displayed as relatively bright and a relatively dark HDR pixel will be displayed as relatively dark, which ensures that the relative brightness of the display matches that of the original scene. However, when the dynamic range between the scene and that of the reproduction devices differ greatly, a large down scaling factor will have to be used to make the scene dynamic range fit within that of the display device. The consequence of which is that, for large contrasts, they would have been suppressed to within the display's dynamic range and will still be visible in the display; however, for small contrasts, which would have been visible on the original scene, would become invisible because of too aggressive compression. The end effect is that the linearly scaled image appears blurry and lack of local details. See the left image in Fig. 1 for a visual example of such effect.

The introduction of the gain map image in Fig.2 is to engineer a solution such that when we scale down the HDR map to fit the dynamic range of the reproduction device, not only the relative brightness of the pixels is well preserved, but local details will also be protected. To find such a gain map image, our idea is very simple: we find an image (the gain map) with the same edge directions as the original HDR map but with the edge strengths as a function of the edge magnitude of the corresponding edge in the original HDR map. We use Fig. 3 to illustrate how the idea works.

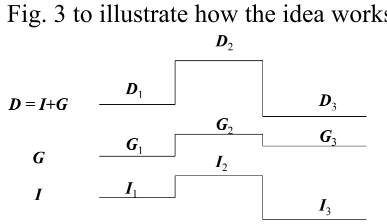

Fig. 3: Illustration of how the idea of using a gain map image will work in rendering a HDR map for display

From the input HDR signal ***I***, we find a signal which has the same edge directions, i.e., in the Figure, ***I*** changes from low to high and then to low, ***G*** follows ***I*** and changes from low to high and then to low at the same changing points. The edges in the display image ***D*** (sum of ***I*** and ***G***) will have exactly the same direction as those in the original image ***I***. In other words, the display image ***D*** will not have edge direction reversed thus is free from halo artifact.

Only having the correct edge orientations in the display image is not enough, we also have to reduce the dynamic range of the display image. One way to achieve this is to boost relatively weak contrasts and compress high contrasts as suggested by several authors [2]. We adopt a similar approach and our task is to design the GMI such that when summing the original image and the GMI together to produce a display image in which, when compared with the original image, weak contrasts are enhanced and strong contrasts are suppressed. Such a GMI can be obtained by ensuring that, for a larger edge magnitude in ***I***, the corresponding edge magnitude in ***G*** will be smaller; conversely, for a smaller edge magnitude in ***I***, the edge magnitude in ***G*** will be larger. To see why such a GMI can achieve the goal of boosting weaker contrasts and suppressing strong contrasts in the display image, lets consider two edges at location u and v, in the original image ***I*** with magnitudes $\Delta(u)$ and $\Delta(v)$, and the corresponding edge magnitudes in the GMI are $\delta(u)$ and $\delta(v)$. Assuming that $\Delta(u) > \Delta(v)$, we will have $\delta(u) < \delta(v)$, thus $\Delta(u)\delta(v) > \Delta(v)\delta(u)$. The edge magnitudes of the display signal at these locations are $\Delta(u) + \delta(u)$ and $\Delta(v) + \delta(v)$. We have

$$\frac{\Delta(u)}{\Delta(v)} - \frac{\Delta(u)+\delta(u)}{\Delta(v)+\delta(v)} = \frac{\Delta(u)\delta(v)-\Delta(v)\delta(u)}{\Delta(v)\left(\Delta(v)+\delta(v)\right)} > 0 \qquad (1)$$

$$\Rightarrow \quad \frac{\Delta(u)}{\Delta(v)} > \frac{\Delta(u)+\delta(u)}{\Delta(v)+\delta(v)}$$

What equation (1) means is that the ratio of a larger magnitude over a smaller magnitude is reduced in the display image, which can only mean one of three things: (i) the larger edge is reduced: $\Delta(u)>\Delta(u)+\delta(u)$, (ii) the smaller edge is boosted: $\Delta(v) < \Delta(v) + \delta(v)$ and (iii) the larger edge is reduced and the smaller edge is enhanced: $\Delta(u)>\Delta(u)+\delta(u)$ and $\Delta(v) < \Delta(v) + \delta(v)$.

Note that the sum signal ***D*** will have to be linearly scaled to fit the dynamic range of the display devices. However, linear scaling does not change the relative values of the pixels hence the display image will have weak contrasts relatively boosted and high contrasts relatively suppressed. From a human perception's point of view, the visual system is sensitive to relative intensities rather than the absolute intensities. The task now is to compute such a GMI.

III. COMPUTING THE GAIN MAP

For a given image $I(x, y)$, we seek a gain map image $G(x, y)$ to produce a display image $D(x, y) = I(x, y) + G(x, y)$. Based on the discussion in the previous Section, we know that the GMI should have the same edge orientations as the original image and should have edge magnitudes inversely proportional to those of the original image. To compute such GMI, our basic idea is illustrated in Fig. 4, which consists of two steps; we first compute a linear relation for each pixel and its local neighbors in the original image and then embed these linear relations in the gain map image. These local neighborhood relationships can be thought of as constraints to ensure that the GMI changes with the original image; in other words, to ensure that ***G*** to have the same edge orientations as the original image. We sparsely constrain the gain map image locally, i.e., seed some initial GMI pixel values such that the contrasts of these seeds follow the requirements of GMI as discussed in previous section. These initial seeds are then propagated to the whole image by adhering to the local neighbor relations of the original image. Such propagation of the initial seeds is achieved by solving a constrained optimization problem. We compute the locally linear relations and solve the global embedding problem by borrowing the computational techniques of the "think globally, fit locally" manifold leaning framework [6]. Note all processing is done in logarithm scale.

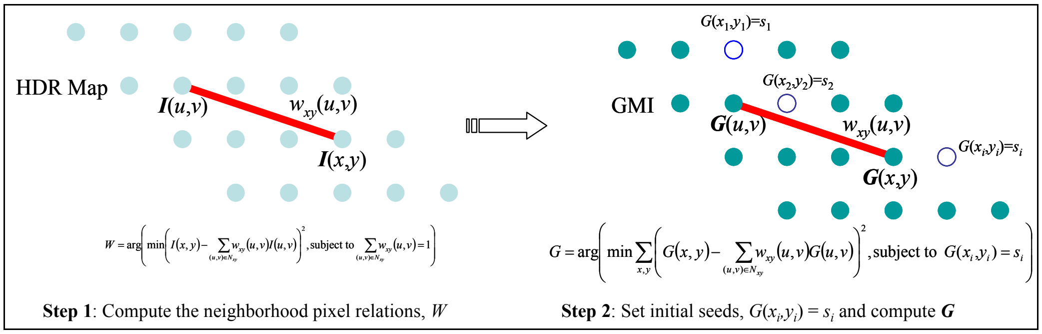

Fig. 4: Schematic of linear neighborhood embedding (LNE) and its application to computing the gain map image. The process has two steps. First, we compute the local neighborhood pixel relations by solving a quadratic optimization problem. Secondly, we seed some initial GMI pixels and use the neighborhood pixel relations computed from step 1 to compute the GMI by solving a constrained optimization problem.

To compute the local neighborhood pixel relations, we can perform following constrained optimization:

Minimizing

$$E(W) = \sum_{x,y}\left(I(x,y) - \sum_{u,v} w_{xy}(u,v) I(u,v)\right)^2 \qquad (2)$$

Subject to

$$\sum_{u,v} w_{xy}(u,v) = 1 \quad and \quad w_{xy}(u,v) = 0 \quad if \quad (u,v) \notin N_{xy}$$

where N_{xy} denotes a local neighborhood surrounds the pixel at location (x, y), $w_{xy}(u, v)$ is the weight which quantifies the contribution of the neighborhood pixel at location (u, v) to reconstructing the pixel at (x, y). Note that the relation is made local by setting the weights of pixels outside a local neighborhood of the pixel to zero. All weights summed to 1 to be invariance to the absolute intensity of the image.

The locally linear spatial relation at pixel location (x,y) in the original image is captured in the weight matrix $W_{xy} = \{w_{xy}(u, v)\}$. These weight matrices should also capture the spatial variations of the gain map G because the gain map G should follow the variations of I as discussed before. Therefore, we can construct G by embedding W_{xy}'s in the gain map by solving following constrained optimization problem:

Minimizing

$$E(G) = \sum_{x,y}\left(G(x,y) - \sum_{u,v} w_{xy}(u,v) G(u,v)\right)^2 \qquad (3)$$

Subject to

$G(x_i, y_i) = s_i$, $i = 1, 2, \ldots$

where s_i's are pre-se seed values of the gain map at location (x_i, y_i). The function is quadratic and the constraints are linear, and therefore the optimization problem leads to large, sparse linear system of equations which can be solved using a number of standard methods.

Note that although the reconstruction weight matrix for each pixel is computed from a local neighborhood in the original image and is independent of the weights of other pixels, the embedding is a ***global*** operation that couples all gain map pixels. Therefore, $\boldsymbol{G}$ should follow $\boldsymbol{I}$ locally and globally as well. Informally, we can view (3) as fitting the local constraints $G(x_i, y_i)$ to the whole image globally. Another way to view this solution is that $\boldsymbol{G}$ is a connected graph with each pixel corresponding to a vertex and the connection weights corresponding to the edges. Therefore the pixel values of $\boldsymbol{G}$ are affected by all the initial seeds and the connection weights.

Implementation of the gain map computation is relatively straightforward. To solve the constrained least squares fit problem of (2), we follow the computational method of LLE [6] by solving a linear system of equations. However, since in our case, the data is 1-d and there will always be more neighbors than input dimensions, the least squares problem for finding the weights does not have a unique solution. We follow the method of [5, 6] by adding a regularization term to the reconstruction cost function to solve the problem. The computational complexity of this step scales as $O(mn^3)$ where m is the number of pixels and n is the neighborhood pixels ($n = 8$ in all our results).

For the embedding problem of (3), since the cost function is quadratic and the constraints are linear, this optimization problem yields a large, sparse system of linear equations, which may be solved using a number of standard methods. The embedding step of LLE solve a similar optimization problem but under different constraints. Without special optimization, the complexity of this step scales as $O(m^3)$, where m is the number of pixels. To speed up the computation, there are several alternative methods for solving the embedding problem, such as multigrid solver which will lead to a complexity scales as $O(m)$.

To set the constraints, we divide the image into 17x17 (other sizes are also possible) blocks, and for each block we identify the largest and the smallest pixels and fix the gain map pixels at these two locations such that the difference between these two pixels in the sum image is enhanced if it is small and suppressed if it is large. For each block, the two constraint values are set as

$$G_{\min} = 0, \quad G_{\max} = \alpha\left(\frac{B_{\max} - B_{\min}}{\alpha}\right)^{\beta} - \left(B_{\max} - B_{\min}\right)$$

where $B_{\max}$ and $B_{\min}$ are the largest and the smallest pixel values in a 17x17 block, α is set to 0.3 times the average local patch contrast of the image and β between 0.6 – 0.8.

IV. EXAPERIMENTAL RESULTS

High dynamic range image compression. We have applied our technique to compress high dynamic range radiance maps for display in low dynamic range device. Fig. 1 shows an example of applying our method. To inspect the effect of the method more closely, we extract one line of pixels from the images in Fig.1 and show the plot of this line of pixels in Fig.5. Fig. 6 shows more examples of our results and comparison with other methods in the literature. It is seen that our method is quite competitive.

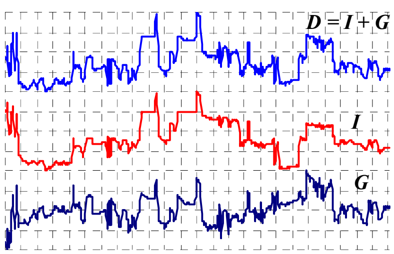

Fig. 5. A scan line from the images in Fig. 1, all three signals are scaled to 0 ~ 255. It is seen that the gain map image (bottom line) strictly follows the changes of the original image (middle line). The result image (top line) and the original image (middle line) have exactly the same edge directions. It is also seen that the sum image *D* (top line) has more local details than the logarithm of the original radiance map (middle line).

Fig. 6 Left column: our results. Top right result of [2] (note the halo artifacts in this image) and bottom right result of [3]. Images are courtesy of Raanan Fattal, Dani Lischinski and Michael Werman, F. Durand and J. Dorsey and P. Debevec.

It is straightforward that our approach can also be applied to the enhancement of ordinary (8 bit/pixel) images. Fig. 7 shows an example of applying our method to image enhancement.

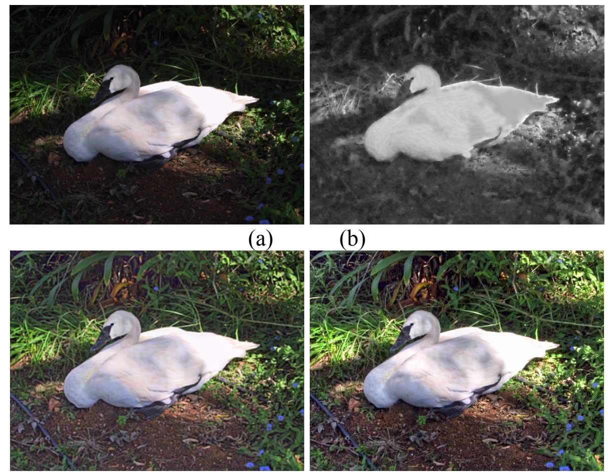

Fig. 7 (a) Original image, (b) The gain map image (c) Our result (d) Result of gradient domain technique [2]. Image data and gradient domain result courtesy of Raanan Fattal, Dani Lischinski and Michael Werman.

V. CONCLUDING REMARKS

In this paper, we have presented a novel technique for the display of high dynamic range images. Our novel approach computes a gain map image using a computational method of machine learning and achieves effective dynamic range compression by suppress strong contrasts and boost weak contrasts. We have presented experimental results which demonstrate that our method is effective both for display HDR images and for image enhancement.

REFERENCES

1 E. Reinhard, G. Ward., S. Pattanaik and P. Debevec, **High Dynamic Range Imaging**, Morgan Kaufmann Publishers, 2005

2 R. Fattal, D. Lischinski and M. Werman, "Gradient domain high dynamic range compression", In Proceedings of SIGGRAPH 2002, pages 249–256, 2002.

3 F. Durand and J. Dorsey, "Fast bilateral filtering for the display of high dynamic range image", In Proceedings of SIGGRAPH 2002, pages 257–265, 2002

4 G Qiu, J. Guan, J. Duan and M. Chen, "Tone mapping for HDR image using optimization – A new closed form solution", ICRP 2006, 18th International Conference on Pattern Recognition, 20 - 24 August 2006, Hong Kong

5 G Qiu and J. Guan, "Color by Linear Neighborhood Embedding", ICIP2005, IEEE International Conference on Image Processing, Genova, Italy, September 11 - 14, 2005

6 L. K. Saul and S. T. Roweis, "Think Globally, Fit Locally: Unsupervised Learning of Low Dimensional Manifolds", The Journal of Machine Learning Research, 4 119-155, 2003