Visual Information Processing Laboratory

School of Computer Science and Information Technology

The University of Nottingham

Technical Report

Report-VIPLAB-02-2005

15 November 2005

Image Contrast Gain Control by Linear Neighbourhood Embedding

Jian Guan and Guoping Qiu

School of Computer Science and Information Technology
The University of Nottingham

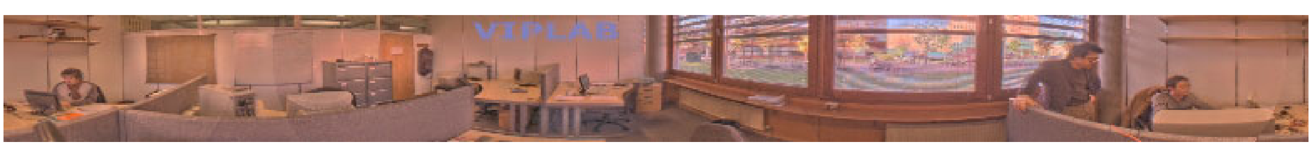

Image Contrast Gain Control by Linear Neighbourhood Embedding

Jian Guan and Guoping Qiu

School of Computer Science and Information Technology
The University of Nottingham

Abstract

In this paper, we present a method that adaptively computes a contrast gain control map for the image through the use of a novel technique termed linear neighborhood embedding (*LNE*) *which first computes a locally linear relation for each pixel and its neighbors and then embeds these relations globally in the gain map image. We borrow the "think globally fit locally" concept and computational techniques from locally linear embedding* (*LLE*) *and compute the gain control image in closed forms by solving constrained optimization problems. We constrain the gain map locally following a gain contrast control mechanism similar to that found in the visual cortex to ensure that weak local contrasts are boosted and strong local contrasts are compressed, and propagate these local constraints globally following the original image pixels' locally linear relations. We have applied our technique to compress high dynamic range images for reproduction in low dynamic range media and to enhance ordinary digital photographs. Results demonstrate that our technique is capable of preserving local details while avoiding artifacts such as halo*

1. 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 and image enhancement has to be applied to improve the image. Although the situations can be remedied by using high dynamic range (HDR) imaging technology [4 – 21] where the so called HDR radiance maps record the actual dynamic range of the scenes, there is still the problem of faithfully reproducing the image in conventional low dynamic range (LDR) reproduction media such as print paper and CRT.

In this paper, we present a technology that adaptively computes a contrast gain control image for a given image through the use of a novel technique termed linear neighborhood embedding (LNE). LNE computes the gain map by solving constrained optimization problems using the concept and computational techniques similar to those of a recent high dimensionality reduction method called local linearly embedding (LLE) [22]. We introduce the idea of "think locally, fit globally" which computes the gain control image at local and global scales simultaneously. We show that our technology can be applied to compress the dynamic range of HDR radiance maps for visualization in LDR reproduction devices and to enhance low quality ordinary LDR digital photographs.

In section 2, we briefly review previous work. In section 3, we introduce our linear neighborhood embedding technology for adaptively computing a contrast gain control image for a given image. In section 4, we present implementation details. We present our results in section 5 and conclude our presentation in section 6.

2. Previous Work

In the past several years, there has been tremendous interest in high dynamic range imaging, from capturing [9 – 13], to storage [20] to display [3-8, 10-12, 14-19, 21]. Although new high dynamic range display devices are being developed, e.g., [19], the dominant reproduction media for HDR images are traditional monitors and print paper having a limited dynamic range between only one to two orders of magnitude. Therefore, how to best compress the dynamic range of the HDR radiance map so that the information in the original scene can be faithfully reproduced in LDR devices is one of the important technical challenges in HDR imaging workflow, and a variety of HDR image display techniques have appeared in the literature. Extensive reviews can also be found in previous works such as [14-18, 21].

There are mainly two categories of techniques: global (spatially invariant) and local (spatially variant). In a global method, a nonlinear function $\boldsymbol{T}$ is used to map each pixel of the input image $I(x, y)$ independently to form an output image $I'(x, y) = \boldsymbol{T}(I(x, y))$. The key to this type of method is to find the appropriate mapping function $\boldsymbol{T}$, which can be some power functions [5, 18] or mapping curves derived from the image's histogram [8]. The advantage of the global method is that it is computational efficient but it can destroy local contrasts making the image look washed out. In a local method, how a pixel is mapped will depend on both the pixel value and its spatial context. Local methods are significantly more complex and can be very difficult to implement in practice. Earlier methods such as [3] introduce "halo" artifacts, improved methods such as [12] can reduce the visibility of halos, and newer methods such as [15, 21] can avoid halo artifacts.

One type of successful local methods that avoid halo estimate edge preserving gain control image. The key observation is that the gain map should have sharp edges at the same points that the original image does, thereby preventing halos [15, 17, 21]. In [17], the authors compute a gain map using the bilateral filter and have achieved very good results. In [15], the authors compute the gain map in the gradient domain that compresses strong edges and boosts weak edges. The reduced low dynamic range image is then retrieved by solving Poisson's equation using an approximated solution which achieves satisfactory results with reasonable computational cost. In [21], the authors introduce a method that computes a gain map for multiscale subband images and they have successfully avoided halos that are often associated with multiscale methods. We have developed a method that computes a gain map image in the intensity image domain by solving constrained optimization problems and we have applied the method to compress high dynamic range image for display and to the enhancement of ordinary images and have achieved good results.

3. Contrast Gain Control by LNE

For a given image $I(x, y)$, we seek a gain map image $G(x, y)$ to produce a new image $I'(x, y) = I(x, y) + G(x, y)$ for high dynamic range compression and/or image enhancement. Our basic idea is illustrated in Figure 1, 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. We sparsely constrain the gain map image locally by using a "contrast gain control" mechanism similar to that found in the visual cortex [25] and propagate these sparse local gain controls globally throughout the entire image following the locally linear pixel neighborhood relations of the original image. We compute the locally linear relations and solve the global embedding problem by borrowing the computational techniques of the "think globally, fit locally" framework [22]. In fact, as will become clear later in the paper, our work can be more accurately described as "think locally, fit globally" strategy.

3.1. Gain Map by Neighborhood Embedding

An image $I(x, y)$ can be regarded as a product of the surface reflectance image $R(x, y)$ and illumination image $L(x, y)$ [1], we can write (in logarithmic domain):

$$I(x,y) = R(x,y) + L(x,y) \quad (1)$$

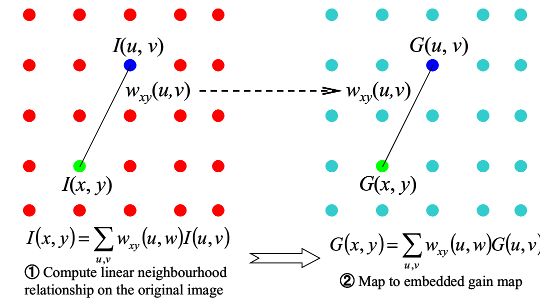

Figure 1: Schematic of linear neighborhood embedding (LNE) and its application to computing the gain map image. The locally linear neighborhood relations of the original image are used to compute the gain map image subject to some appropriate constraints.

The largest variation of the image comes from the illumination image since real world reflectance images are unlikely to have contrast greater than 100:1. The illumination image tends to vary slowly across the image [2] and large intensity variations tend to mostly occur between regions rather than locally within a small area. If the image can be separated into the product of R and L, then high dynamic range compression can be achieved by compressing the illumination image L only. This way, the large contrast between brightly illuminated regions and dark shadow areas can be reduced while the contrast caused by reflectance of local texture features is preserved. However, separating the reflectance from the illumination is an ill-posed problem and it is almost impossible to completely separate L from R.

Assume that we can compute a gain map $G(x, y)$ for the image $I(x, y)$ to produce a new image $I'(x, y)$ for applications such as high dynamic range compression and/or image enhancement, we can write

$$I'(x,y) = R(x,y) + L(x,y) + G(x,y) \quad (2)$$

The gain map should serve at least two objectives, at a coarse scale it should reduce the high contrast between very brightly illuminated regions and dark shadow areas; and at a local scale, it should enhance or preserve the details within local regions. As well as achieving these two objectives simultaneously, it must avoid introducing artifacts such as halo. Therefore, $G(x, y)$ should compress strong edges caused by illumination variations (high dynamic range compression), preserve or boost local edges (image enhancement) and protect original edge signs (avoiding halo).

We have developed a "think locally, fit globally" technology to compute the gain map image to achieve these goals simultaneously based on following rationale.

Because human visual system is most sensitive to local intensity variations, then it is most important to preserve maybe even enhance local intensity contrasts. However, processing local contrasts will have to be done at a global scale in the sense that when putting locally processed regions together, the whole image should preserve the original scene's visual integrity. Therefore, we cannot process local regions independently, but rather, local regions contrasts should be processed such that they will fit together globally. Equally important, the gain map should protect the directions of the intensity changes across the whole image to avoid halo which also requires that local processing fit into global relations. Therefore, the gain map image must be computed at local and global scales simultaneously.

Our method first "think locally", it strategically picks a few local patches of the image, finds the largest and the smallest intensities within each patch and processes the min/max pixel pairs according to a contrast gain control mechanism similar to that found in the visual cortex. The processed min/max pixel pairs from these local patches are then "fit globally" to whole image. To achieve such "think locally, fit globally" strategy, we borrow the concept and computational techniques similar to those of the "think globally, fit locally" framework – locally linear embedding (LLE) [22].

At a coarse scale, the regional intensity variations within the image may be characterized by locally linear pixel relations. One possible approach to characterizing the locally linear spatial pixel relation is to construct a linear model that reconstructs the pixels from a linear combination of their neighbors. To build such model, we can perform following constrained optimization:

Minimizing

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

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 from (2) it is clear that the gain map G should follow the variations of I. Therefore, we can construct G by embedding W_{xy}'s in the gain map by solving following constrained optimization problem:

Minimizing

$$E(G) = \sum_{x,y} \left(G(x,y) - \sum_{u,v} w_{xy}(u,v) G(u,v) \right)^2 \qquad (4)$$

Subject to

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

where g_i's are pre-determined values of the gain map at location (x_i, y_i).

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, G should follow I locally and globally as well. However, because the local weights only capture the linear relations, they can only capture the slow changing components of I that should reflect regional variations of the image. Local scale processing can be realized by setting the constraints locally. Informally, we can view (4) as globally fits the local constraints $G(x_i, y_i)$ to the whole image globally. Clearly, G will be greatly dependent on the numbers, the locations and the values of the constraints $G(x_i, y_i)$ and setting the constraints is a crucial step in our algorithm.

3.2. Setting Constraints

To get some insight into how the constraints should be set, it may be helpful to consider how the human visual system performs gain control. The visual world is hugely complex and the visual system is constantly exposed to scenes of huge dynamic ranges. Apparently, the visual system performs various types of automatic control to cope with the high dynamic ranges and noisy environments to keep it within the optimal operating range. At a global level, there seems to be a gain control mechanism at the retina, where the photoreceptors rapidly adapt to the ambient light level. In computational vision, this type of gain control is normally modeled by taking the log of the input intensity. It is widely accepted that the logarithm of the luminance is a (crude) approximation to the perceived brightness. Therefore, displaying the log image directly should give reasonably correct brightness but such image will be lack of details[1]. It is also widely accepted that human visual system is not very sensitive to absolute luminance reach the retina, but rather responds to local

[1] Logarithm itself incurs no information loss, the details are destroyed by numerical quantization. When the log image is scaled down to within the dynamic range of the reproduction devices, pixels with similar values will be quantized to the same level thus destroying details.

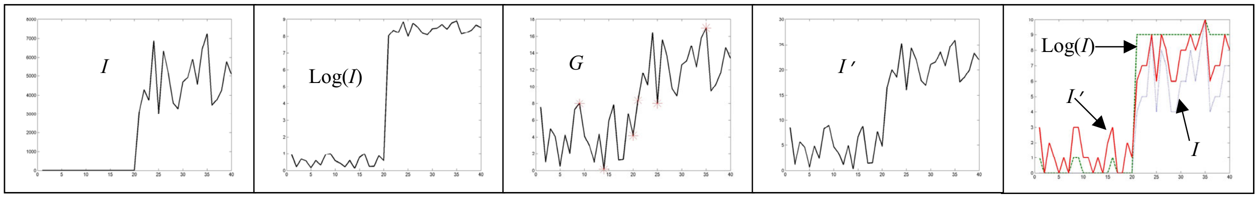

(a) (b) (c) (d) (e)

Figure 2. (a) The original high dynamic range signal with a dynamic range of 6218 : 1. (b) Logarithm of (a). (c) The gain map, where the * are the constrained points. (d) The result signal = (b)+(c). (e) Linearly scale (a), (b) and (d) to the same dynamic range of 10 : 1. It is seen that in the result signal, local details are well preserved (slightly enhanced) in the low dynamic range image.

intensity variations. In the visual cortex, there are neurons that operate a gain control mechanism known as “contrast gain control” in which moderately low contrasts are boosted while high contrasts are compressed [25].

These important properties of the human visual system may provide us with some important clues to set the local constraints for equation (4) to achieve our goals. Firstly, we should set the constraints $G(x_i, y_i)$ sparely scattered across the image to enable the embedding process to enforce G to follow the variations of I thus protecting the visual integrity of the original image. Secondly, we should set the values of these constraints such that they will boost weak local contrasts and compress local high contrasts. In this paper, we use following strategy to set the constraints:

Constraints Setting Procedure

i) Segment the image into coarse regions

ii) Within each segmented region, randomly pick non-overlapping patches of certain size (in this paper 17x17 patch size is used, other patch sizes are possible).

iii) For each patch P_i, find the smallest and the largest pixel within the patch

$$\left(x_{\max(P_i)}, y_{\max(P_i)}\right) = \arg\left(\max\left\{I(x,y); \forall(x,y) \in P_i\right\}\right)$$

$$\left(x_{\min(P_i)}, y_{\min(P_i)}\right) = \arg\left(\min\left\{I(x,y); \forall(x,y) \in P_i\right\}\right)$$

and compute the patch's contrast as

$$C(P_i) = I\left(x_{\max(P_i)}, y_{\max(P_i)}\right) - I\left(x_{\min(P_i)}, y_{\min(P_i)}\right) \tag{5}$$

iv) Within each patch, set two constraints for G, one at $(x_{\max(Pi)}, y_{\max(Pi)})$ and the other at $(x_{\min(Pi)}, y_{\min(Pi)})$ according to

$$G\left(x_{\min(P_i)}, y_{\min(P_i)}\right) = 0 \qquad G\left(x_{\max(P_i)}, y_{\max(P_i)}\right) = \alpha\left(\frac{C(P_i)}{\alpha}\right)^{\beta} - C(P_i) \tag{6}$$

where α determine which local contrasts are boosted or compressed. If the local contrasts are smaller than α then they are boosted and if they are greater than α then they are compressed (assuming $0 < \beta < 1$). A similar function was used in [15] and [21] for gradient domain gain contrast control and subband domain gain control. In all results presented in this paper, α is set to 0.3 times the average local patch contrast of the image and β between 0.6 – 0.8.

From the above contrast setting procedure, we can make following observations. The purpose of coarsely segmenting the image is to identify the strongest changes across image regions. Because the edges across these regions are so important in protecting the visual integrity of the image, we do not want to set constraints across these strong edges, rather we want to set constraints away from these strong edges and prefer to find the gain map pixels for these strong edges through embedding. For each local patch, we only fix two pixels for the gain map, one at the smallest pixel's position and the other at the largest pixel position. Basically, we fix the smallest pixel and increase or decrease the largest pixel. Although we in effect assume that logarithm will make the smallest pixel in each patch have the right brightness (gain map value equals zero) and set the gain map value at the largest pixel position, these constraints are rather mild because, firstly, only two pixels within a 17x17 window are fixed (in practice <0.1% pixels are used as constraints), secondly, these pixels are only “borrowed” to constrain the optimization and over 99.9 % of the pixels are found through global embedding. The gain map pixels within each local patch are not only affected by the constraints of its own patch but are also influenced by ALL constraints from different patches, therefore, these locally set constraints will propagate globally. In fact, in the final gain map, the pixels at the constraint positions can be replaced by averaging or majority voting thus removing the hard constraints initially put on for computational purpose.

3.3. Algorithm Summary

The LNE gain map construction method and its application can be summarized in following steps:

1 Compute the logarithm of the input I

2 Compute linear pixel neighborhood relations according to (3)

3 Set gain map constraints according to the ***constrain setting procedure***

4 Compute the gain map G according to (4)

5 Compute the output image $I' = I + G$ and linearly scale it to within the dynamic range of the reproduction media.

The procedure of our method is illustrated in Figure 2.

4. Implementation

Implementation of the LNE gain map algorithm is relatively straightforward. Because it is only necessary to segment the images into coarse regions and accurate image segmentation is not necessary, we first down scale the original image by a factor of 4 in both dimensions, and then use the mean shift robust computational segmentation method [23] to automatically segment the image into regions. Since our main purpose is to identify where the strong edges are likely situated, we then rescale the segmented image into full resolution and produce "thick" boundaries where constraint pixels should not be placed across. Inside each segmented region, we place a 17 x 17 window at non-overlapping positions and at each position identify the smallest and the largest pixels within the window and set the constraints for G at these min and max locations according to (6).

To solve the constrained least squares fit problem of (3), we follow the computational method of LLE [22] 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 [22] 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 (4), 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 [24, 27] which will lead to a complexity scales as $O(m)$.

For high dynamic range compression, like all other methods in the literature we only work on the luminance channel, and color of the image is untouched. After compressing the luminance, color is put back to the image in a way similar to those in [15]. For ordinary image enhancement, we also only process the achromatic channel.

5. 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. Figure 3 shows the segmented regions, the locations of the constraint pixels, the gain map image, the log image and gain map modified image of the Memorial Church HDR radiance map. Figure 4 shows a comparison of our result with those of three other recent high dynamic range compression techniques for the Memorial Church image.

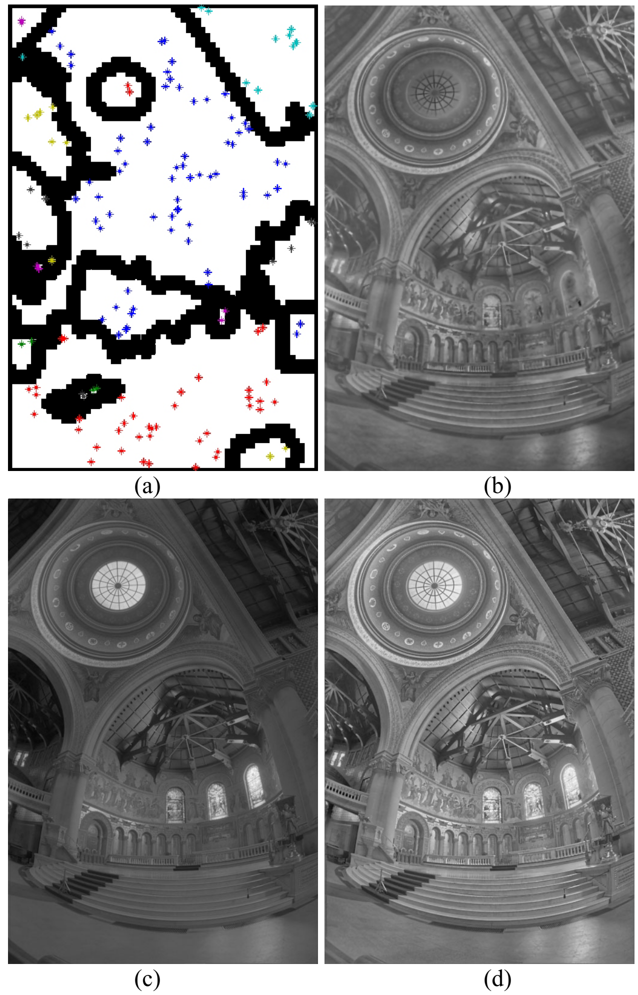

Figure 3. (a) Segmented regions and constraints. The thick black curves shows the region boundaries found by the segmentation algorithm where constraints should not be placed across. Constraints positions are shown in different colors within the regions. (b) The gain map image computed by LNE. (c) The achromatic channel of the log image in which local details is lost due to scaling. (d) Result image = (b) + (c). Radiance map courtesy of Paul Debevec.

Figure 5 shows the segmented coarse regions and the locations of the constraints used for constructing the gain map, the gain map image, our result and the result of the gradient domain method [15] for the Belgium House image. Figure 6 shows more examples of our result as compared with other methods in the literature.

These results demonstrate that our technique works effective in compressing high dynamic range image for display. It is seen that in all our results, local details are well preserved and in some case slightly enhanced. Compared our method with those state of the art shows that our method works equally well. Visually, there are some differences amongst the results of different techniques, such as the overall brightness, contrast and color of the images, however, as has been very well stated in [21], these differences should not be over-interpreted since they may change depending on the detail implementation.

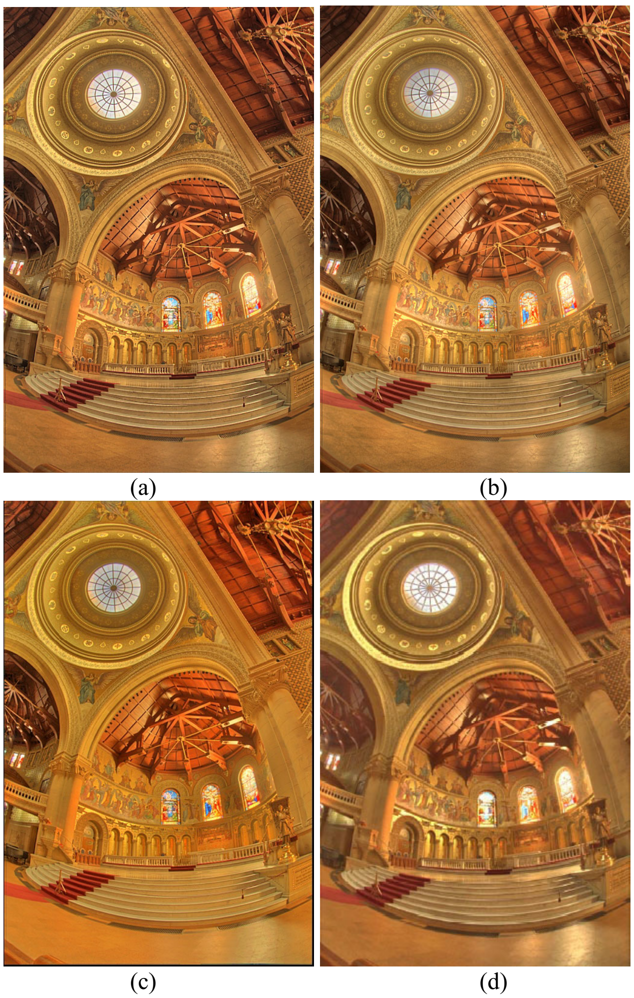

Figure 4. (a) Our result. (b) Result of gradient domain technique [15]. (c) Result of subband technique [21] (d) Results of bilateral filtering method [17]. Radiance map courtesy Paul Debevec.

A very important property of our technique is that it is completely free from halo artifacts since it will not reverse the edges. In fact, the way in which our technique works can be understood in this way: it enhances the contrasts everywhere in the image, but enhances small contrasts more and large contrast less. Therefore, the directions of intensity changes will be protected while the magnitude of changes are either boosted or attenuated. Because strong changes are enhanced relatively less and weak changes are enhanced more, when the image is scaled down to fit the reproduction device, the end effect is that local small details are relatively enhanced while large changes between regions are relatively compressed. Figure 7 shows a scanline form the Belgium House image in Figure 5, which perfectly illustrates the behavior of our technique.

Figure 5 (a) Coarsely segmented regions and the constraints used for constructing the gain map. (b) The gain map image (c) our result (d) Result of gradient domain technique. Radiance map courtesy of Raanan Fattal, Dani Lischinski and Michael Werman.

Image enhancement. Our technique can be equally applied to the enhancement of ordinary images. Figure 8 shows an example of applying our technique to the enhancement of an ordinary image. As a comparison, result of the gradient domain technique is also shown. We see that our technique works well. The two results are slightly different, again, the differences should not be over-interpreted because they may change depending on the details of implementation.

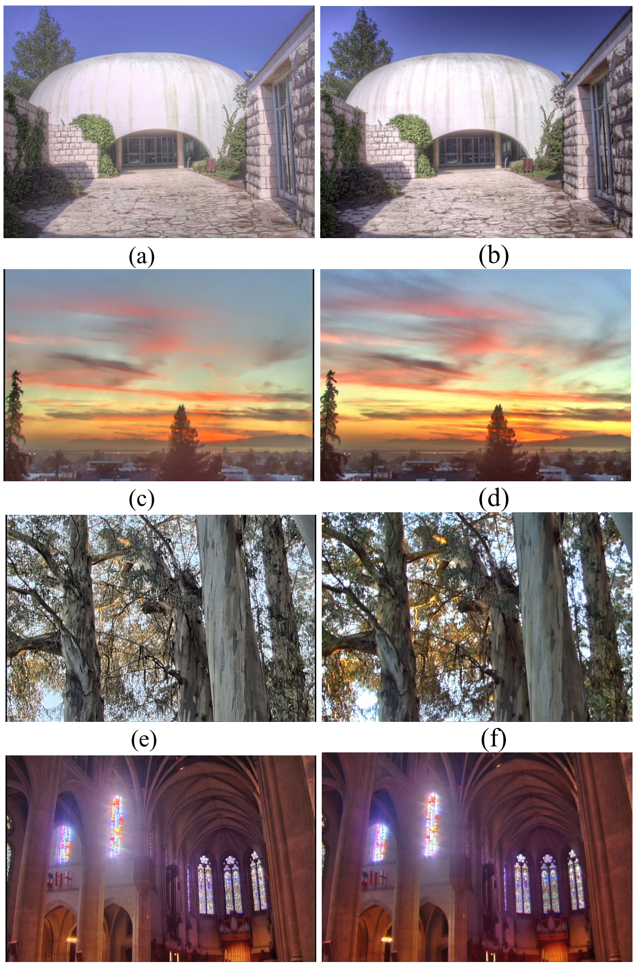
(a) (b) (c) (d) (e) (f) (g) (h)

Figure 6 more examples. Our results (a), (c), (e) and (g). Results of gradient domain technique [15] (b) and bilateral filtering technique [17] (d), (f) and (h). Radiance map courtesy of Raanan Fattal, Dani Lischinski and Michael Werman, and Paul Debevec.

It is interesting to note that the colorization using optimization work of [26] solves a similar optimization as eq. (4). However, they use the W's directly computed using the squared difference or normalized correlation of neighboring pixels, rather than computed using eq. (3). These directly computed weights are not suitable for computing the gain map image because they fail to capture the local changes of the image. Figure 9 shows a scanline and its processed results by using gain maps constructed weights from eq. (3) and those from squared difference and normalized correlation. It is seen that while the result using weights computed using LNE clearly follows the original signal and preserves and enhances local details, those using weights from squared difference and normalized correlation clearly loss track of the signal and smear local details.

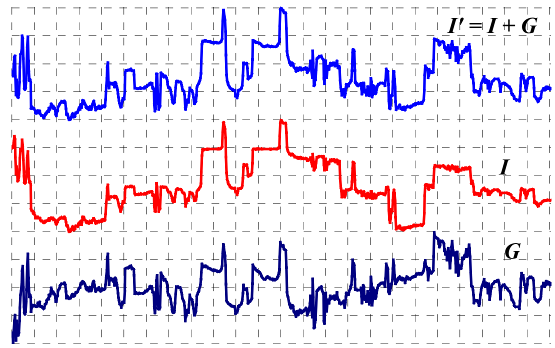

Figure 7 A scanline from the Belgium House image, all three signals are scaled to 0 ~ 255. It is seen that the gain map image (bottom line) strictly follows the changes of the original image (middle line). The result image (top line) and the original image (middle line) have exactly the same edge directions. It is also seen that I' (top line) has more local details than I (middle line)

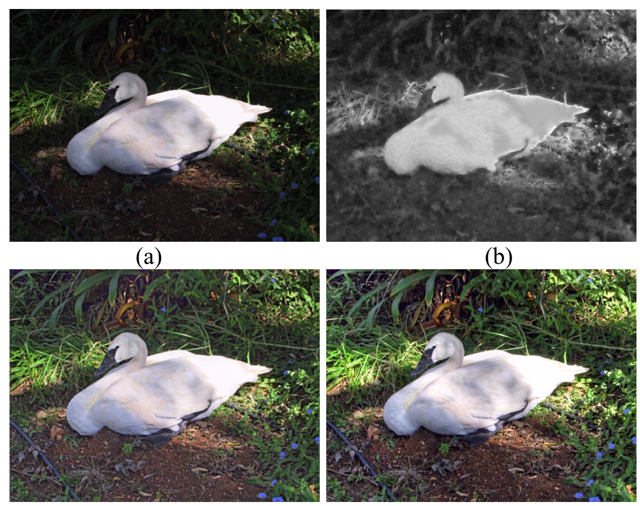
(a) (b) (c) (d)

Figure 8 (a) Original image, (b) The gain map image (c) Our result (d) Result of gradient domain technique [15]. Image data and gradient domain result courtesy of Raanan Fattal, Dani Lischinski and Michael Werman.

6. Concluding remarks

In this paper, we have presented a novel technique for computing a contrast gain control image for a given image. We introduce the idea of "think locally, fit globally" to construct the gain map image to achieve local detail enhancement/preservation and global visual integrity protection simultaneously. We have developed techniques to compute the gain map image in closed form by solving constrained optimization problems. We have demonstrated that image gain map constructed by our method can be used to compress high dynamic range images and enhance ordinary images without introducing visual artifacts.

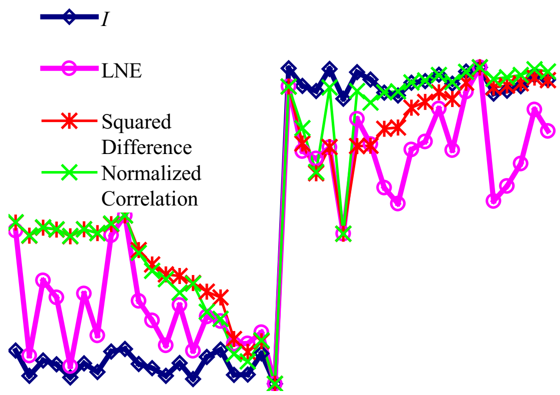

Figure 9 An original image scanline (*I*) and the processed results using gain maps constructed from weights computed in different ways. It clearly shows that using squared difference or normalized correlation between neighborhood pixels to compute the local pixel relations are unsuitable for computing the gain map.

References

[1] T. Stockham, "Image processing in the context of a visual model", Proceedings of the IEEE, vol. 60, pp. 828-842, 1972

[2] B. K. P. Horn, "Determining lightness from an image", Computer Graphics and Image Processing, vol. 3, pp. 277-299, 1974

[3] K. Chiu, et al, "Spatially nonuniform scaling functions for high contrast images", in Proceedings of Graphics Interface'93, pp. 245 –253

[4] J. Tumblin and H. E. Rushmeier, "Tone reproduction for realistic images", IEEE Computer Graphics and Applications, vol. 13, pp. 42 - 48, 1993

[5] C. Schlick, "Quantization techniques for the visualization of high dynamic range pictures", in Proceedings of 5th Eurographics Workshop on Rendering, June 1994

[6] Greg Ward. "A contrast-based scale factor for luminance display". In Graphics Gems IV, pages 415–421. Academic Press, 1994

[7] J. A. Ferwerda, S. N. Pattanaik, P. Shirley and D. P. Greenberg. " A model of visual adaptation for realistic image synthesis". Proceedings of SIGGRAPH'96, pp. 249-258

[8] G. Ward, 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

[9] P. E. Debevec and J. Malik, "Recovering high dynamic range radiance maps from photographs", In Proceedings of SIGGRAPH'97, pp. 369 – 378

[10] D. J. Jobson, Z. Rahman, and G. A. Woodell, "A Multi-scale 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

[11] S. N. Pattanaik, James A. Ferwerda, Mark D. Fairchild, and Donald P. Greenberg, "A Multiscale Model of Adaptation and Spatial Vision for Realistic Image Display", Proceedings of SIGGRAPH'98, pp. 287-298, Orlando, July 1998

[12] J. Tumblin, J. K. Hodgins and B. K Guenter, "Two methods for display of high contrast images", ACM Trans on Graphics, vol. 18, pp. 56 - 94, 1999

[13] T. Mitsunaga and S. K. Nayar, "High dynamic range imaging: Spatially varying pixel exposures", Proc. CVPR'2000, vol. 1, pp. 472-479, 2000

[14] J. DiCarlo and B. Wandell, "Rendering high dynamic range images", Proc. SPIE, vol.3965, pp. 392 – 401, 2001

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

[16] E. Reinhard, M. Stark. P. Shirley and J. Ferwerda, "Photographic tone reproduction for digital images", In Proceedings of SIGGRAPH 2002, pages 267–276, 2002.

[17] 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

[18] Michael Ashikhmin. "A tone mapping algorithm for high contrast images", in Rendering Techniques '02 (Proceedings of the 13th Eurographics Workshop on Rendering), pages 145–156, June 26–28, 2002.

[19] H. Seetzen, W. Heidrich, W. Stuerzlinger, G. Ward, L. Whitehead, M. Trentacoste, A. Ghosh, A. Vorozcovs, "High dynamic range display systems", ACM Transactions on Graphics (Siggraph 2004), 23(3): 760-768

[20] G. Ward and M. Simmons, "Subband encoding of high dynamic range imagery", Proceedings of the 1st Symposium on Applied perception in graphics and visualization, Pages: 83 – 90, 2004

[21] Y. Li, L. Sharan, E. H. Adelson, "Compressing and companding high dynamic range images with subband architectures", ACM Transactions on Graphics (TOG), 24(3), Proceedings of SIGGRAPH 2005

[22] 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

[23] D. Comaniciu, P. Meer: Mean Shift: A Robust Approach toward Feature Space Analysis, IEEE Trans. Pattern Analysis Machine Intell., Vol. 24, No. 5, 603-619, 2002

[24] W. H. Press, S. A. Teukolsky, et al., "Numerical Recipes in C++ -- The Art of Scientific Computing", Cambridge, MA, Cambridge University Press, 2002

[25] D. J. Heeger, Half-squaring in responses of cat striate cells, Visual Neuroscience, 9:427-443, 1992.

[26] A. Levin, D. Lischinski and Y. Weiss, "Colorization using optimization", SIGGRAPH 2004

[27] W. Hackbusch, Multi-grid Methods and Applications, Springer, Berlin, 1985