

Globally Optimized Linear Windowed Tone Mapping

Qi Shan, *Student Member, IEEE*, Jiaya Jia, *Member, IEEE*, and Michael S. Brown, *Member, IEEE*

Abstract—This paper introduces a new tone mapping operator that performs local linear adjustments on small overlapping windows over the entire input image. While each window applies a local linear adjustment that preserves the monotonicity of the radiance values, the problem is implicitly cast as one of global optimization that satisfies the local constraints defined on each of the overlapping windows. Local constraints take the form of a guidance map that can be used to effectively suppress local high contrast while preserving details. Using this method, image structures can be preserved even in challenging high dynamic range (HDR) images that contain either abrupt radiance change, or relatively smooth but salient transitions. Another benefit of our formulation is that it can be used to synthesize HDR images from low dynamic range (LDR) images.

Index Terms—High dynamic range image, tone mapping, display algorithms, image enhancement, filtering.

1 INTRODUCTION

HIGH dynamic range compression—or tone mapping—embodies techniques that map HDR images to low dynamic range displays with only moderate contrast. While a global linear scaling can be used to compress the dynamic range, e.g., from 10,000:1 to 200:1, the result will be undesirable since image structures are linearly flattened, causing the loss of visual information mostly in highlights and shadows. Several tone mapping algorithms have been proposed that either adjust the global tone response curve [14] or locally reproduce tonal values [7], [9], [16]. It is also generally noticed ([9], [16]) that local operators, which reproduce the tonal values in a spatially variant manner, perform more satisfactorily than global operators in terms of the detail preservation and compression ratio of the dynamic range.

The objective in this paper is to develop a new tone reproduction operator. We observe that HDR in an image can typically be categorized into two types—1) regions exhibiting significant high dynamic range but with smooth radiance transition (Fig. 1b) and 2) regions exhibiting sharp and significant local radiance change among neighboring pixels (Fig. 1a). Our tone reproduction operator provides a unified framework to effectively address these two types of HDRs. Our approach maintains *local structures*, including sharp edges and smooth color transitions at a perceptual level while not introducing visible artifacts such as halos.

Our approach uses a window-based tone mapping method in which a *global optimization* problem is solved that

satisfies *local* constraints. Specifically, our method operates on windows; in each window, a linear function is used to constrain the tone reproduction in order to naturally suppress strong edges while retaining weak ones. The high dynamic range is compressed by solving an image-level optimization problem that integrates all window-based constraints.

Although our method is regarded as a local operator, it does not involve scale decomposition, layer separation, or image segmentation, and therefore, is resistant to the artifacts known with these procedures. Within our global optimization framework, the overall tone mapping effect is nonlinear and spatially variant, and adapts to rich and diverse structures of images. Our method has a closed-form solution that reaches the global optimum. Due to the optimality, any local tone adjustment in our method gives rise to a global effect, where errors are minimized and distributed across the entire image. Moreover, our method provides flexibility for image quality adjustments with only a few parameters, adaptive to various user requirements pertaining to the range compression ratio and the level of details that should be preserved.

In addition to high dynamic range compression, our method also contributes a unified framework for tone enhancement of ordinary images by solving the same optimization problem. An application of synthesizing an HDR image from a single low dynamic range (LDR) image is presented to simulate the high contrast environment in indoor and outdoor scenes containing shiny light sources or obscure shadows. The effectiveness of our HDR synthesis is evaluated by comparison with ground truth images.

The rest of the paper is organized as follows: Section 2 reviews related work. Section 3 presents our algorithm of HDR compression. Section 4 describes the parameter setting and shows our HDR compression results. Our algorithm is applied to ordinary image enhancement and HDR image synthesis in Section 5. Section 6 discusses the proposed approach and summarizes our work.

2 PREVIOUS WORK

Tone mapping operators have attracted broad interests recently. Surveys of these methods can be found in [5], [25], [21]. We review HDR compression methods in two categories: global operators and local operators. The global tone

- Q. Shan is with the Department of Computer Science and Engineering, University of Washington, Box 352350, Paul G. Allen Center 378, Seattle, WA 98195-2350. E-mail: shanqi@cs.washington.edu.
- J. Jia is with the Department of Computer Science and Engineering, The Chinese University of Hong Kong, Shatin, N.T., Hong Kong. E-mail: leojia@cse.cuhk.edu.hk.
- M.S. Brown is with the School of Computing, National University of Singapore, AS6-#05-15, Computing 1, 13 Computing Drive, Singapore 117417, Republic of Singapore. E-mail: brown@comp.nus.edu.sg.

Manuscript received 25 Feb. 2009; revised 19 June 2009; accepted 7 July 2009; published online 21 July 2009.

Recommended for acceptance by A. Hertzmann.

For information on obtaining reprints of this article, please send e-mail to: tcvg@computer.org, and reference IEEECS Log Number TVCG-2009-02-0043. Digital Object Identifier no. 10.1109/TVCG.2009.92.



Fig. 1. HDR regions exhibiting abrupt and significant radiance change (a) and smoother local transition (b).

mapping operators map the radiance values in a spatially invariant manner. Histogram-based and sigmoid transfer-function-based algorithms, such as gamma correction, are two main categories of global operators. For global operators, the key problem is to establish a one-to-one or onto mapping from high dynamic radiance values to the low dynamic ones. Sigmoidal compression is to redistribute radiance values in an HDR image using a sigmoid curve. This results in low compression at the low and high ends of the radiance range, and more compression in the mid-range. Other similar global functions with similar properties are discussed in detail in [21, Chapters 6 and 7]. In [14], [6], nonlinear global operators were proposed, where global tone reproduction curves with different measures of image qualities are estimated to reduce the dynamic range in a computationally efficient way. These methods do not need manual tuning and are capable of producing visually satisfactory results for a large set of examples. However, these approaches may fail for HDR images, where strong local contrast is presented (i.e., a local region contains both high and low radiance values). Further, finding a perfect tone reproduction curve is difficult [4] for several images.

Inspired by how the human visual system adapts to local regions of scene luminance [21], more sophisticated local tone reproduction methods have been proposed. Many of these approaches involve decomposing, compressing, and recombining layers in the HDR images. Krawczyk et al. [13] first segmented an HDR image and then applied different tone reproduction curves to the segments to achieve spatially variant HDR compression. With appropriate algorithm design and implementation, no seams will be present in the tone mapping results.

In [24], an image is decomposed into a reflectance image R and an illuminance image L , according to the observation that the high dynamic range is generally caused by the illuminance, while reflectance is unlikely to produce high contrast. However, accurately separating the two layers from a single image is nontrivial. To simplify the intrinsic image decomposition, it is assumed that L contains low-frequency information while R may have spatially abrupt change in values [11]. Assuming that the surface reflectance and illumination properties are known, Tumblin [25] compressed high dynamic range only in the illuminance image. In [7], an image is decomposed into a base layer and a detail layer. The base layer is obtained using a bilateral filter and the detail layer is computed by subtracting the base layer from the input image. Linear scaling in the logarithm domain of the base layer is performed to compress the dynamic range. The final result is the combined detail layer and the compressed base layer.

Other scale-decomposition-based techniques decompose the HDR images into several layers [12], [19], [26], imitating

the human vision adaptation to local changes. In [16], a multiscale approach using symmetrical analysis-synthesis filter banks and automatic gain control was proposed to compress the HDR images in subbands. This method, along with those of [20], [2], incorporated schemes to avoid the halting artifacts. This type of artifacts, as explained in [16], can be considered as signal distortions that occurs when combining multiple compressed subband layers. It however can often be reduced with careful selection of parameters.

Recently, in [8], a new edge-preserving multiscale image decomposition method was proposed. Results demonstrate that this image decomposition method outperforms bilateral filtering in producing range compressed images in several cases.

Gradient domain dynamic range compression [9] does not directly process the image radiance. Instead, the tone mapping operator attenuates the image gradients with large magnitude while magnifying the small ones. The final result is computed by solving a Poisson equation. Lischinski et al. [17] presented an interactive method to locally adjust the tonal values and other visual parameters. In their method, the user uses a set of brushes to impose constraints on the image. Influence functions are computed to confine the modification of tonal values.

Different from these methods, our local operator does not require multiscale decomposition or segmentation of the images into binary or fractional maps. This avoids problems associated with layer decomposition. We directly process the radiance map, where a *global optimization framework* is developed to naturally constrain the respective pixel values. In addition, our approach does not require any smoothness constraints to be imposed on the final output.

Inferring an HDR image from an LDR image is known as “inverse tone mapping” [3]. Single image methods can only achieve an approximation due to the information loss. In [3], high-luminance areas are estimated from an LDR image. An inverse tone mapping operator, which reverses the one in [20], is applied to enhance luminance in the detected areas. Meylan et al. [18] proposed detecting the specular highlights in the LDR images and boosting the intensity in these regions. The LDR images generated in the tone mapping stage produce small errors in the dynamic range expansion stage. In [1], psychophysical studies showed that linear contrast scaling works well in most cases for expanding the dynamic range of some ordinary images. Rempel et al. [22] proposed a real-time method to synthesize HDR images by enhancing the brightness gradation in the saturated regions. Li et al. [16] addressed a special HDR compression-expansion problem (a process called “companding,” for example, to turn a 12 bit/channel image into an 8 bit/channel TIFF, and later convert it back to a good approximation of the original 12-bit image) by adopting a feedback-loop scheme. This method reconstructs an HDR image from an LDR one that was tone mapped with their subband technique. This approach, however, cannot be used to expand the dynamic range of an arbitrary LDR image.

Comparing our work with previous inverse tone mapping methods, our algorithm not only makes the results look natural by enhancing pixel intensity more significantly in highlights, but also reaches high radiance similarity with the ground truth HDR images. Synthesizing an HDR image in our system can be regarded as a partial reversion of HDR compression. It is naturally accomplished in our unified framework.

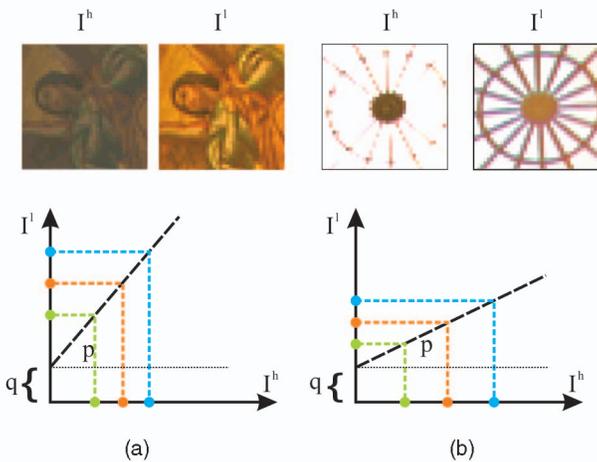


Fig. 2. Illustration of the local window operator. (a) and (b) Two configurations that I^h undergoes linear mappings (represented as the slanted dashed lines) with different coefficients p and q to produce I^l . q controls the base radiance level and p manipulates the local contrast in the window. Setting large p makes the local details in I^l be enhanced, as shown in (a), whereas a small p suppresses local high contrast, as shown in (b).

3 ALGORITHM AND IMPLEMENTATION

Given an input HDR image with radiance map I^h , the high dynamic range compression operator $f(\cdot)$ computes the radiance map in a low dynamic range $I^l = f(I^h)$. An intuitive approach is to decrease the large radiance values and increase small ones. To correctly retain the local structures, however, the monotonic mapping of radiance should be satisfied regarding each pixel's neighborhood in a small local region, similar to the monotone characteristic of the global tone mapping curves. Therefore, if we zoom into the window w_i with pixel i in center, the linear function

$$I^l(j) = p_i I^h(j) + q_i, \quad j \in w_i \quad (1)$$

provides a simple yet complete representation satisfying the local monotonic constraint. The coefficients p_i and q_i , within the local window, control the mapping in two ways. The value of q determines the base radiance level while p represents the slope of the linear function. The terms q and p directly control the local contrast as shown in Fig. 2. If $p > 1$, as shown in Fig. 2a, the local contrast is enhanced, making the result enhance details in dark regions. Fig. 2b shows another example, where small positive p reduces image contrast and improves the visibility of structures in bright regions. Note that we do not add regularization terms, such as the smoothness constraints, to the overall equations. This is because p_i s and q_i s are computed in overlapping windows that contain multiple pixels. In this case, a smoothness-like constraint across pixels is naturally enforced in flat regions, while in regions with strong structures, such as edges, contrast is reasonably maintained.

Combining all local linear equations defined on individual windows, we can reconstruct an image by minimizing

$$\sum_i \sum_{j \in w_i} (I^l(j) - p_i I^h(j) - q_i)^2, \quad (2)$$

where i sums over all pixels in the HDR image. However, directly minimizing (2) cannot compress the high dynamic range because a trivial solution exists. For example, $p_i = 1$

and $q_i = 0$ for all i result in $I^l = I^h$, which is the solution that produces zero error in minimizing (2). As a result, additional constraints are needed. Note that p_i in (1) is the parameter directly controlling the change of local contrast. We propose guiding its value in each window to reduce the global contrast while maintaining the local visual information. We express the final objective function to be minimized as

$$f = \sum_i \left(\sum_{j \in w_i} (I^l(j) - p_i I^h(j) - q_i)^2 + \varepsilon c_i^{-2} (p_i - c_i)^2 \right), \quad (3)$$

where c_i is a preset positive value to guide the modification of local contrast. The collective c_i in the image space form a *guidance map*. By appropriately setting the values in this map, we can suppress strong local contrasts and elevate weak ones. A detailed description of the guidance map construction is given in Section 4. The term $c_i^{-2} (p_i - c_i)^2$ in (3) is the squared relative error from the guidance map and c_i^{-2} is used to normalize the cost introduced by the difference between p_i and c_i . The term ε is a weight to balance the two terms in (3). We set its value to 0.1 in all our experiments.

The benefit of introducing the guidance map c is twofold. First, although the value of c_i influences local contrast, the configuration of c_i does not require high accuracy compared to direct modification of the radiance values in I^h . This is because in our optimization framework, ε is only a small weight, making the guidance c_i act as a soft constraint. The linear mapping term $\sum_{j \in w_i} (I^l(j) - p_i I^h(j) - q_i)^2$ also constrains modification of radiance at each pixel to be similar to the change in the neighboring pixels. This helps preserve image structure. Second, by modifying the guidance map c , our system can readily be applied to other tone management applications, such as ordinary image enhancement and HDR image synthesis. Examples are shown in Section 5.

3.1 Optimization and Implementation

Directly minimizing f in (3) is not trivial given the large set of unknowns. Methods such as gradient decent is not desirable because I^l , p_i , and q_i are highly coupled and the energy function is not convex. However, as each pair of the linear coefficients $[p_i, q_i]$ is only defined in a single window, minimizing (3) can be expressed as

$$\arg \min_{p, q, I^l} f = \arg \min_{I^l} \sum_i \arg \min_{p_i, q_i} f_i, \quad \text{where} \quad (4)$$

$$f_i = \left(\sum_{j \in w_i} (I^l(j) - p_i I^h(j) - q_i)^2 + \varepsilon c_i^{-2} (p_i - c_i)^2 \right).$$

Using the above expression, we first compute (p_i^*, q_i^*) , i.e., the optimal solution of (p_i, q_i) , by setting the partial derivatives of function f_i with respect to p_i and q_i to zeros and then computing the optimal \hat{I}^l by solving a linear system. The computation process is described in the Appendix. We will discuss in Section 4 the configuration of the guidance map c for naturally compressing the high dynamic range.

The local window size can be adjusted in our system. In Fig. 3, we show tone mapping results computed by setting the window sizes from 3×3 to 15×15 . For all examples shown in this paper, we found that window size 3×3 is suitable in terms of maintaining edge sharpness and

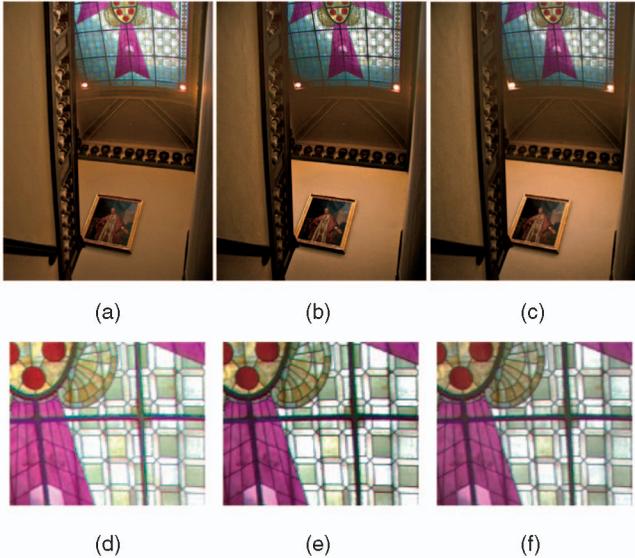


Fig. 3. Tone mapping results with window sizes (a) 3×3 , (b) 7×7 , and (c) 15×15 . (d)-(f) Close-ups of (a)-(c), respectively.

computational efficiency. However, other sizes can also be used.

The above operations process the image's luminance channel. After obtaining the tone mapped radiance map \hat{I}^l , the RGB channels are reconstructed using the method of Schlick [23]:

$$I_k^l(i) = (I_k^h(i)/I^h(i))^s \times \hat{I}^l(i), \quad k \in \{r, g, b\}, \quad (5)$$

where I_k^l is one of the RGB color channels in the image result and I_k^h is one of the color channels in the input HDR image. s represents the saturation factor. Larger s produces a more saturated result. The underlying idea of introducing (5) is to preserve the ratios between the luminance and color channels so that the hue of the range compressed image can be similar to that of the input HDR image. In our experiments, $s \in [0.4, 0.6]$ produces chromatically natural image results. The skeleton of our algorithm is given in Table 1.

4 GUIDANCE MAP CONFIGURATION IN HDR COMPRESSION

In the linear optimization framework introduced in Section 3, the definition of the guidance map c principally determines the quality of HDR compression. In this section, we describe how to obtain a suitable guidance map c .

Recall that the value p_i controls the compression ratio of the dynamic range. To preserve the visual information in the tone mapped image, we reduce the value of p_i where the local contrast is large, and increase it otherwise. A measure of the local contrast in each window w_i is the standard deviation σ_i of I^h . However, setting c_i inversely proportional to σ_i cannot produce a satisfying result. An example is illustrated in Fig. 4, where (a) shows the input HDR image and (b) shows the variance map of I^h . The variance map is overly sensitive to the textures, edges, and even the noise in the HDR image. The HDR compression result illustrated in Fig. 5a is produced by setting $c_i = (\sigma_i^{\beta_2} + \kappa)^{-1}$, where $\beta_2 = 0.75$ to attenuate the variance and κ is a small

TABLE 1
Skeleton of Our Tone Mapping Algorithm

1.	Input: An HDR image with radiance map I^h .
2.	Main Steps:
2.1	Generate the guidance map c according to (6) in Section 4.
2.2	Construct matrix S and B according to (18) in the Appendix.
2.3	Compute \hat{I}^l by solving the linear system defined in (17) in the Appendix.
2.3	Restore the RGB channels I_k^l in the tone mapped result by (5).
3.	Output: An ordinary image with color channels I_k^l .

constant value. Although the high dynamic range of the original image in Fig. 4a is reduced in the result shown in Fig. 5a, there are noticeable blemishes in the result especially in the highlighted region. These artifacts are caused by the susceptibility of the variance to the local structures.

To improve the visual quality of the result, we compute σ_i on the Gaussian filtered images to reduce noise. Two other visual factors, the local mean μ_i and the image radiance $I^h(i)$, are incorporated into the guidance map with the following considerations: First, μ_i and $I^h(i)$ represent the image radiance information. Their maps globally reflect image structures. By taking μ_i and $I^h(i)$ into the guidance formation, the overexaggeration of structures using only local contrast can be prevented and the tonal value

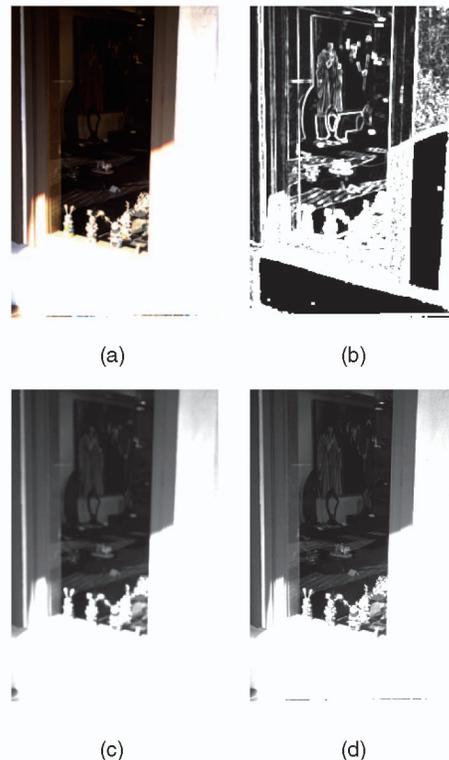


Fig. 4. "Window" example. (a) The input HDR image (image courtesy of Shree Nayar). (b) The truncated variance map of I^h on all local windows w . (c) The truncated map of mean values of I^h . (d) The truncated radiance map I^h .

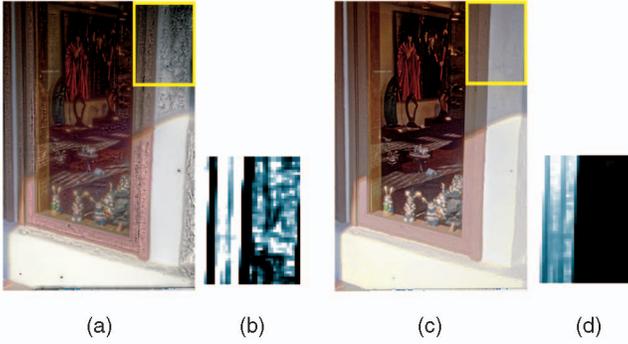


Fig. 5. Illustration of the guidance map configuration. (a) The tone mapping result when c_i is set inversely proportional to the local standard deviation of I^h in each window. There are visible artifacts on the wall. (b) Magnified guidance map c . The noise is amplified. (c) The tone mapping result when c_i is configured according to (6). (d) Magnified guidance map. Note that the noise is suppressed.

adjustment will depend more on the input radiance in a global manner. Second, we observed that the contrast of the radiance within a local window depends largely on the absolute variance values. For instance, in Fig. 4a, the variance of the radiance inside the dark window is as low as 0.1 whereas the radiance variance in the bright window-sill is approximately 10^4 . These observations indicate that incorporating μ_i and $I^h(i)$ into the guidance map construction helps suppress contrast in bright regions.

The maps of μ_i and $I^h(i)$ are illustrated in Figs. 4c and 4d, which emphasize the salient radiance change between the bright and dark regions while not making it highly sensitive to the isolated noise. We define c_i as

$$c_i = (\mu_i^{\beta_1} \sigma_i^{\beta_2} I^h(i)^{\beta_3} + \kappa)^{-1}, \quad (6)$$

where κ is a small weight set to 0.05 in all our experiments to prevent c_i from being divided by zero. Normally $\beta_1 \in [0.4 - 0.9]$, $\beta_2 \in [0.1 - 0.4]$, and β_3 is fixed to 0.1. We show an HDR compression result in Fig. 5c using the guidance map constructed using (6). Comparing the magnified guidance maps shown in Figs. 5b and 5d, it is apparent that Fig. 5d contains much less visual artifacts.

Fig. 6 shows the terms we introduced to construct the guidance map c . Fig. 6a shows an input HDR image. The sky is bright and the hill is relatively dark. Figs. 6b and 6c

show the image brightness and the mean radiance in all 3×3 windows. Fig. 6d shows our measure of local contrast by the standard deviation of radiance σ_i (images are scaled to aid visualization). This emphasizes strong gradients or edges. The term β_2 can be used to control the compression ratio on strong edges, which, in turn, makes it possible to vacate more space for revealing fine local details. We show our tone mapping result in Fig. 6f, computed using the guidance map shown in Fig. 6e.

In order to preserve the details in both bright and dark regions, our algorithm balances local contrast and radiance values in the entire image. Intuitively, the guidance map constructed from a medium to low dynamic range image should be flatter in structure than that from an HDR image. To validate this, we show one example in Fig. 7. The guidance map of an input HDR image (Fig. 7a) is shown in Fig. 7b. For comparison, the guidance map of the HDR compressed image (Fig. 7c) is shown in Fig. 7d constructed using the same parameters. The guidance maps are linearly scaled for illustration. It can be observed that the guidance values in Fig. 7d have much smaller variations.

4.1 Parameter Settings

The parameters β_1 and β_2 can be tuned to produce different tone mapping results. We assign the values based on two principles. First, the sum of the three β s influences global dynamic range compression. Larger value makes the result visually flatter. Second, increasing β_2 enhances compression on strong edges, thus, leaving more space for less salient structures.

It is notable that, in our algorithm, small modifications to β_1 and β_2 do not largely affect the result. The default values $\beta_1 = 0.6$ and $\beta_2 = 0.2$ are already capable of producing visually satisfactory results for most HDR images we have tested. To illustrate the effect of different parameter settings, we show an example in Fig. 8. Fig. 8a shows the original HDR image. Our HDR compression results are shown in Figs. 8b, 8c, 8d, and 8e using different parameter settings. In all these images, the structures on the lamp and the desk corner are augmented. Close-ups show that the text on the book is better enhanced with a larger β_2 .

Fig. 9 shows another example. Image (a) shows an input HDR image. The maps of μ and σ are shown in Figs. 9b and 9c. The HDR compression results using different parameters

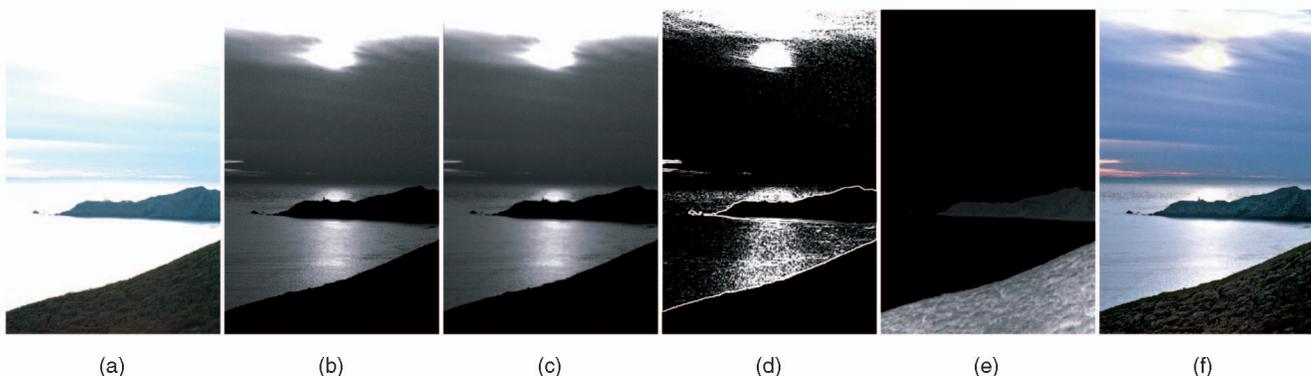


Fig. 6. Illustration of the guidance map configuration. (a) An input HDR image. (b) Original radiance map $I^h(i)$. (c) Map of mean radiance μ_i . (d) Map of standard deviation σ_i . (e) The visualized guidance map c . (f) The tone mapping result using the guidance map shown in (e).

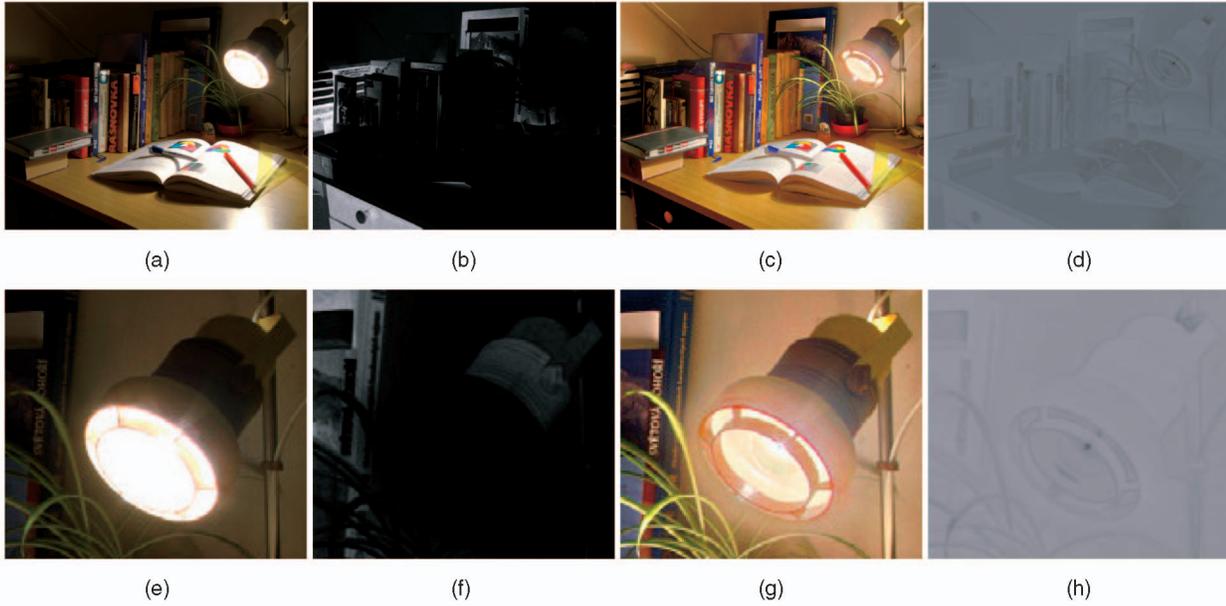


Fig. 7. Lamp desk. (a) An input HDR image. (b) The guidance map c generated from (a) using our method. (c) A tone mapping result by our method. (d) The guidance map c generated from (c) using the same parameters. It is structurally flat comparing to the map shown in (b). (e)-(h) Close-ups of (a)-(d).

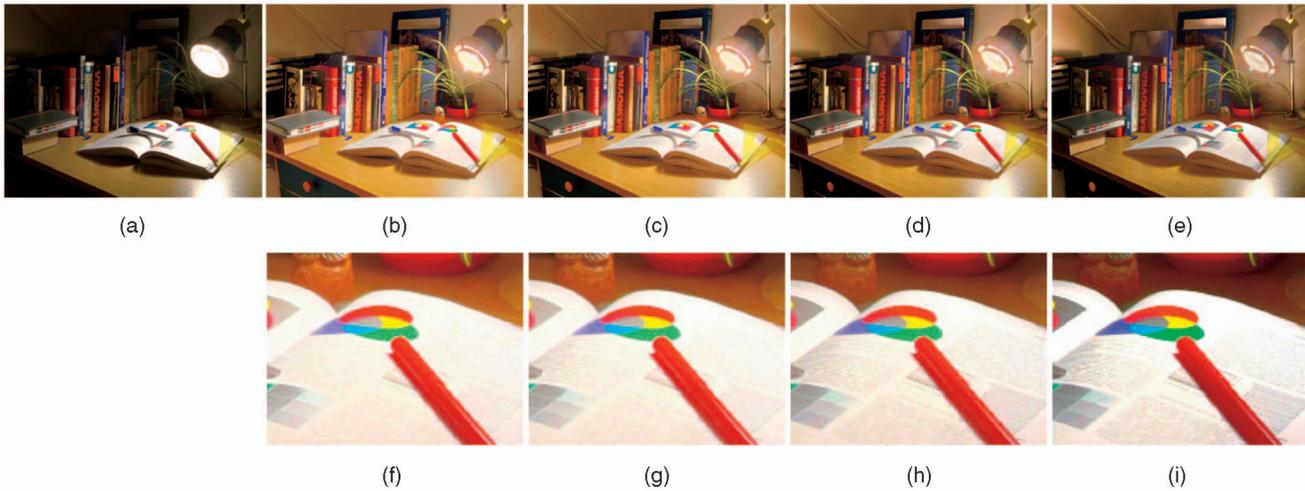


Fig. 8. HDR compression with different parameter settings. (a) The original HDR image, displayed with linear scaling. (b) $\beta_1 = 0.8, \beta_2 = 0.0$. (c) $\beta_1 = 0.7, \beta_2 = 0.1$. (d) $\beta_1 = 0.6, \beta_2 = 0.2$. (e) $\beta_1 = 0.5, \beta_2 = 0.3$. The tone mapping results with different parameter settings. (f)-(i) Close-ups.

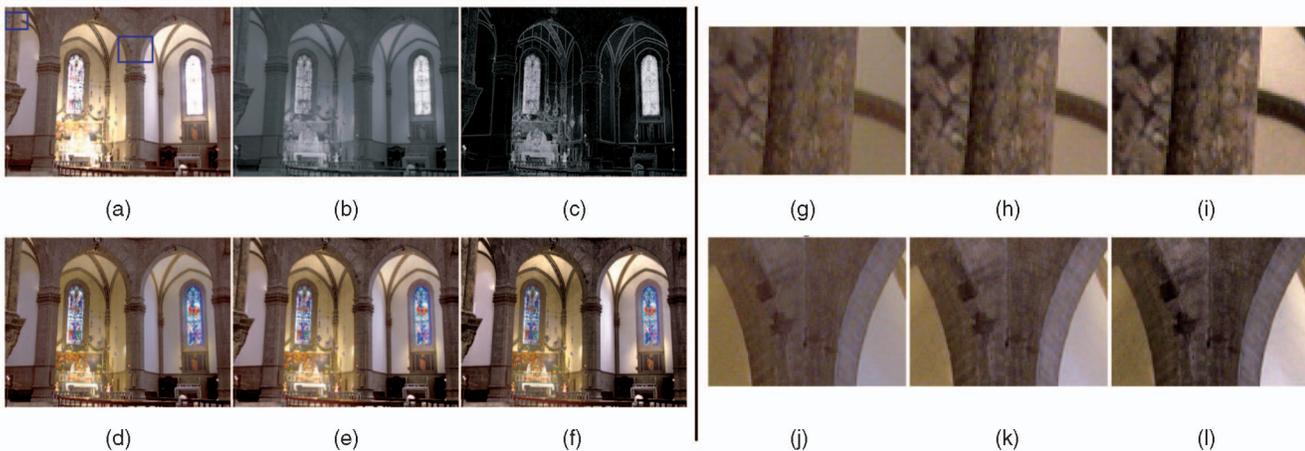


Fig. 9. HDR compression with different settings of β_1 and β_2 . (a) The input HDR image. (b) Visualization of the map $\log_2(\mu + 1)$. (c) Visualization of the map $\log_2(\sigma + 1)$. (d) $\beta_1 = 0.8, \beta_2 = 0.0$. (e) $\beta_1 = 0.7, \beta_2 = 0.1$. (f) $\beta_1 = 0.6, \beta_2 = 0.2$. The tone mapping results with different parameter settings. We magnify two regions and compare them in (g)-(l). Local contrast is enhanced more significantly using larger β_2 .



Fig. 10. HDR compression with different β_1 s. (a) The HDR images under low exposures. (b) The HDR images under high exposures. (c) $\beta_1 = 0.2$. (d) $\beta_1 = 0.4$. (e) $\beta_1 = 0.6$. (f) $\beta_1 = 0.8$. Our HDR compression results. In all these results, we fix $\beta_2 = 0.1$. Since $\beta_1 + \beta_2$ controls the overall compression ratio, larger $\beta_1 + \beta_2$ makes the result exhibit more details in both bright and dark regions.

are shown in Figs. 9d, 9e, and 9f. In all these images, details are preserved. We show the magnified regions in Figs. 9g, 9h, 9i, 9j, 9k, and 9l extracted from Figs. 9d, 9e, and 9f. The local contrast is better preserved with larger β_2 .

Fig. 10 shows two examples, where β_1 is set differently. The input HDR images under certain exposures are shown in Figs. 10a and 10b. We fix $\beta_2 = 0.1$ and only modify β_1 . Figs. 10c, 10d, 10e, and 10f, respectively, show our HDR compression results with the increased sum of all β s. It can be noticed that the original high dynamic range is compressed with gradually increased ratio, and more details are maintained in both bright and dark regions using larger β_1 .

4.2 Running Time

Our algorithm is implemented using Matlab on a desktop PC with an Intel Core2Duo 2.4 GHz CPU. The running time to process different images is closely related to the image resolution and is plotted in Fig. 11. Taking Fig. 12 as an example, the results can be obtained within 2 seconds given the image containing 512×768 pixels and with window size 3×3 . We have found that using larger window size for w needs more computation due to solving larger linear systems. To produce the results shown in Figs. 3a, 3b, and 3c, our method spends 3, 25, and 150 seconds, respectively, in computation.

4.3 More Results and Comparison

In this section, we show more HDR compression results and compare them with other state-of-the-art techniques.

In Fig. 12, the tone mapping results of the HDR image “memorial church” are shown. The images in Figs. 12a, 12b, and 12c are published in [7], [9], and [16], respectively. Our result is illustrated in Fig. 12d with parameters $\beta_1 = 0.7$, $\beta_2 = 0.2$, and $\beta_3 = 0.1$. The radiance of the dark regions, for instance, the upper left corner of image, is elevated. The brightness of the stained glass

window and the clerestory is reduced, revealing clear structural details. The textures and paintings on the wall and ceiling are also preserved. The magnified local regions are shown in Figs. 12e, 12f, 12g, and 12h.

HDR compression results of the “chairs” example by our method and the gradient-based method [9] are shown in Figs. 13a and 13b, respectively. More results using our optimization method are shown in Figs. 14 and 15. The dynamic range is compressed from about $10^5 : 1$ to $255 : 1$. All structures, especially those originally buried in shadow and highlights, are naturally enhanced. In Fig. 16, we show a comparison with the result of Farbman et al. [8].

5 APPLICATIONS

5.1 Ordinary Image Enhancement

Our method can be extended to perform automatic enhancement of ordinary LDR images in order to improve the visibility of dark and heavily saturated regions containing structural information. The configuration of c_i , in this case, is similar to that introduced in Section 4. Since

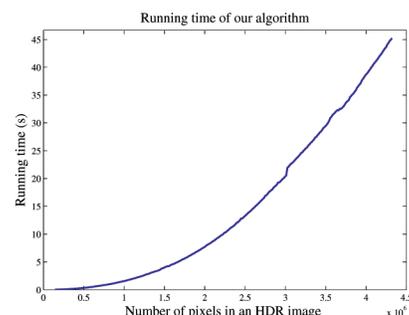


Fig. 11. Plot of running time of our tone mapping algorithm with respect to different image resolutions.

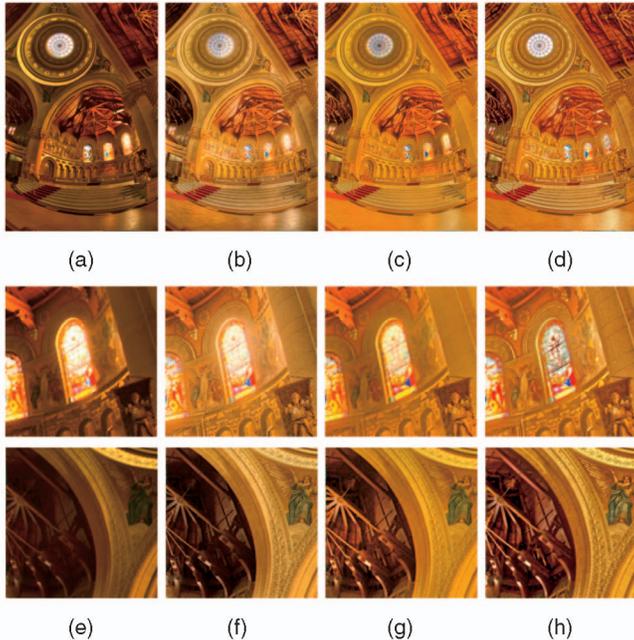


Fig. 12. Memorial church. (a) Tone mapping result of Durand and Dorsey [7]. (b) Tone mapping result of Fattal et al. [9]. (c) Result by the multiscale method using one aggregated gain map [16]. (d) Our result with parameters $\beta_1 = 0.7, \beta_2 = 0.2$, and $\beta_3 = 0.1$. (e)-(h) Magnified regions from (a)-(d), respectively. (HDR image courtesy of Paul Debevec.)

the ordinary image has much lower contrast, the values of β_1, β_2 , and β_3 are set smaller. Fig. 18 shows results with different parameter settings. In our experiments, the default values $\beta_1 = 0.4, \beta_2 = 0.2$, and $\beta_3 = 0.05$ work well for most examples.

Two examples are shown in Fig. 17 using the default parameter setting. In Fig. 17a, the mountain and its reflection in water are hardly seen. After our image enhancement, the trees and the green patterns on the mountain together with the distorted reflection are visible in Fig. 17b. Figs. 17c and 17d show another example. The lichen on the tree in the original



Fig. 13. Chairs. (a) Our tone mapping result with parameters $\beta_1 = 0.6, \beta_2 = 0.2$, and $\beta_3 = 0.1$. (b) The tone mapping result published in the project Web page of [9]. (HDR image courtesy of Shree Nayar.)

image is almost unidentifiable while our result clearly enhances its details.

5.2 HDR Image Synthesis

In this section, we present a method for synthesizing an HDR image from a single LDR image. In our optimization framework, this process can be regarded as spatially variant dynamic range expansion. Note that in our HDR compression model, we minimize the following least square energy according to the local constraint in each window i :

$$f_i = \sum_{j \in w_i} (I^l(j) - p_i I^h(j) - q_i)^2 + \varepsilon c_i^{-2} (p_i - c_i)^2. \quad (7)$$

Similarly, in the LDR expansion, we take I^l as input and compute I^h by simulating the global contrast expansion using the same model. We rewrite (7) as

$$p_i^2 \left(\sum_{j \in w_i} \left(I^h(j) - \frac{1}{p_i} I^l(j) + \frac{q_i}{p_i} \right)^2 + \varepsilon c_i^{-2} \left(1 - \frac{c_i}{p_i} \right)^2 \right), \quad (8)$$

by taking p_i^2 out of the outermost parenthesis. Note that there does not exist a closed-form solution in minimizing (8) because p_i and $I^h(j)$ are both unknowns and cannot be separated in the proposed two-stage optimization. For computation simplicity, we approximate the weight p_i^2 in (8) by c_i^2 . As a result, the HDR synthesis can be achieved by minimizing the energy:



Fig. 14. Street. HDR images ©Industrial Light & Magic. The top row shows the input HDR image under different exposures. The bottom row shows our HDR compressed image result and its close-ups, with parameters $\beta_1 = 0.6, \beta_2 = 0.2$, and $\beta_3 = 0.1$. The radiance of the sun is largely reduced.



Fig. 15. More high dynamic range compression results by our method. The parameters that produce these results are identical: $\beta_1 = 0.5$, $\beta_2 = 0.2$, and $\beta_3 = 0.1$.

$$\begin{aligned}
 E' &= \sum_i c_i^2 \left(\sum_{j \in w_i} \left(I^h(j) - \frac{1}{p_i} I^l(j) + \frac{q_i}{p_i} \right)^2 + \varepsilon c_i^{-2} \left(1 - \frac{c_i}{p_i} \right)^2 \right) \\
 &= \sum_i c_i^2 \left(\sum_{j \in w_i} \left(I^h(j) - \frac{1}{p_i} I^l(j) + \frac{q_i}{p_i} \right)^2 + \varepsilon (c_i^{-1} - p_i^{-1})^2 \right).
 \end{aligned} \tag{9}$$

Denoting $p'_i = 1/p_i$, $q'_i = -q_i/p_i$, and $c'_i = 1/c_i$, (9) can be written as

$$f' = \min_{p, q, I^h} \sum_i c_i'^{-2} \left(\sum_{j \in w_i} (I^h(j) - p'_i I^l(j) - q'_i)^2 + \varepsilon (p'_i - c'_i)^2 \right), \tag{10}$$

in a form similar to (3) defined for HDR compression. Thus, the HDR synthesis can be dealt with within our framework by similarly employing the linear optimization.

Similar to that of HDR compression, we construct a guidance map c' incorporating the image color and structure information to appropriately constrain the LDR expansion in (10). Specifically, we use the local contrast measure σ'_i in each window w_i of I^l , the local mean μ'_i of I^l , and the image radiance $I^l(i)$ to construct the guidance map, encouraging the

exaggeration of local structures in appropriate degrees. The new guidance map for (10) is expressed as

$$c'_i = (\mu_i'^{\beta_1} \sigma_i'^{\beta_2} I^l(i)^{\beta_3} + \kappa),$$

where $\kappa = 1$ to maintain or expand the image contrast. The values of β_1, β_2 , and β_3 are defined similar to those in



(a)

(b)



(c)

(d)



(a)

(b)

Fig. 16. Comparison with the result of Farbman et al. [8]. (a) The tone mapped image shown in [8]. (b) Our tone mapping result.

Fig. 17. Ordinary image enhancement. (a) and (c) The original images. (b) and (d) The automatically enhanced images showing more balanced local details and global luminance.



Fig. 18. Image enhancement with different parameter settings. We fix $\beta_3 = 0.05$ in this example. Larger parameter values enhance more structural details. (a) Input. (b) $\beta_1 = 0.2, \beta_2 = 0.1$. (c) $\beta_1 = 0.3, \beta_2 = 0.1$. (d) $\beta_1 = 0.4, \beta_2 = 0.2$. (e) $\beta_1 = 0.5, \beta_2 = 0.2$.

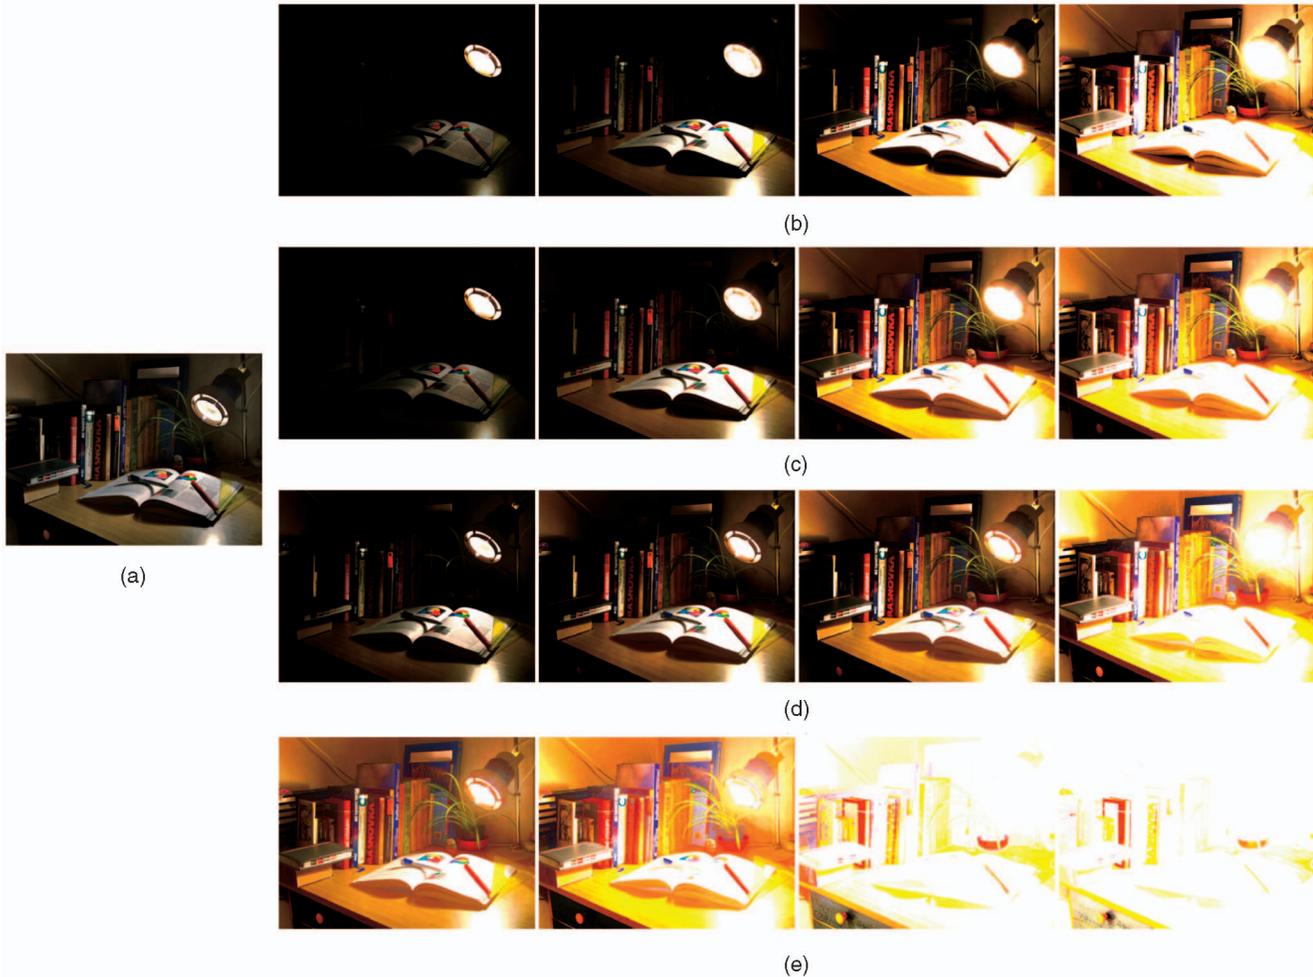


Fig. 19. HDR image synthesis. (a) The input LDR image I^l tone mapped from an HDR image I^h displayed in (b). (b) The original HDR image I^h displayed under different exposures. (c) The resynthesized HDR image $I^{h'}$ by our algorithm (displayed under different exposures). (d) The synthesized HDR image using the LDR2HDR method [22]. (e) The synthesized HDR image using linear scaling.

Section 4.1. In our experiments, β_3 is fixed to 0.1, $\beta_1 \in [0.4 - 0.9]$, and $\beta_2 \in [0.1 - 0.4]$.

Equation (10) can be optimized by the algorithm described in the Appendix. Since (10) is slightly different from (3), there exist variable differences on the expansion of Δ_i and S introduced in (14) and (18). Specifically, in the HDR synthesis, we have new Δ_i and s_{kj} defined as

$$\Delta_i = \sigma_i^2 + \left(\frac{\varepsilon}{m_i} \right),$$

$$s_{kj} = \sum_{i \in \{k, j\} \subset w_i} c_i^{-2} \left(\delta_{kj} - \frac{1}{m_i \Delta_i} ((I^h(k) - \mu_i)(I^h(j) - \mu_i) + \Delta_i) \right).$$

All other variables are unchanged. The radiance expansion can be computed by solving the linear system of (17) similarly.

We denote the HDR synthesis result by $I^{h'}$ from the LDR image I^l . To evaluate the efficacy of our algorithm, we assume that I^l is originally tone mapped from an HDR image I^h . By comparing I^h and $I^{h'}$, we are able to compute the ground truth expansion errors. In our experiments, we use Peak Signal-to-Noise Ratio (PSNR) to measure the ground truth error between the original HDR image I^h and resynthesized image $I^{h'}$.

Two examples are shown in Figs. 19 and 20. In Fig. 19, the input LDR image I^l shown in Fig. 19a is compressed in dynamic range from an HDR image I^h illustrated in Fig. 19b

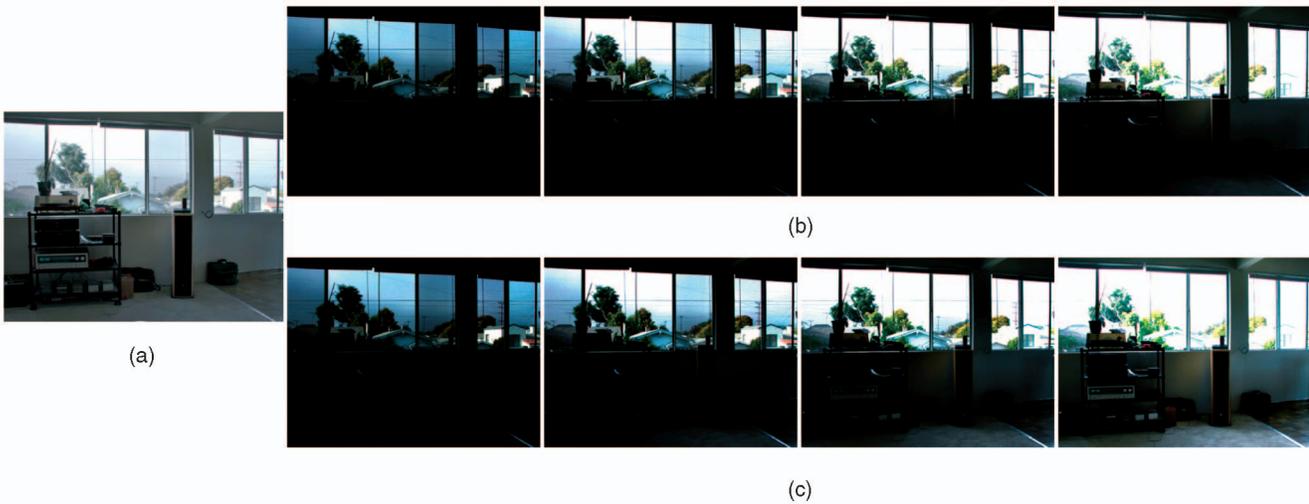


Fig. 20. Indoor example. (a) The input LDR image I^l compressed from an HDR image. (b) The original HDR image I^h shown under different exposures. (c) The resynthesized HDR image $I^{h'}$ from I^l shown under different exposures.

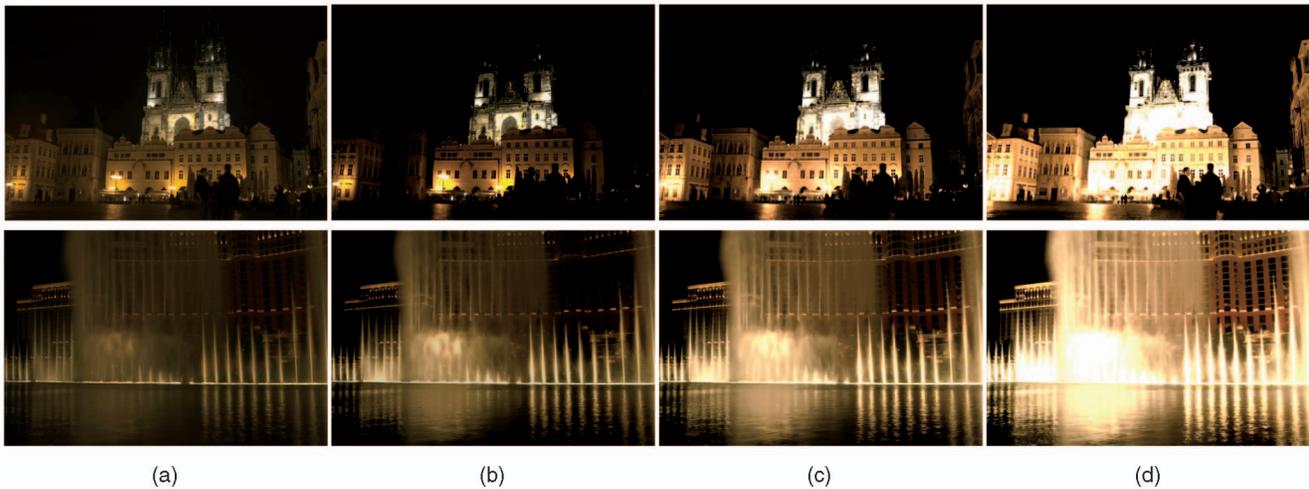


Fig. 21. Night view example. (a) The input LDR image. (b)-(d) Our synthesized HDR scene, shown under different exposures.

under different exposures. We synthesize HDR image $I^{h'}$ from I^l . The parameters are set as $\beta_1 = 0.5$, $\beta_2 = 0.1$, and $\beta_3 = 0.1$ for the example shown in Fig. 19, and $\beta_1 = 0.6$, $\beta_2 = 0.2$, and $\beta_3 = 0.1$ for the examples shown in Figs. 20 and 21. For comparison, in Figs. 19d and 19e, we show the HDR image results produced by the LDR2HDR method [22] and simple linear scaling. The PSNRs of the results in Figs. 19 and 20 are 58.2873 and 57.1476, respectively, with respect to the radiance of the original HDR images.

Fig. 21 shows examples of our HDR image synthesis from ordinary 8-bit LDR images. Fig. 21a shows the input images. Figs. 21b, 21c, and 21d are the synthesized HDR images shown under different exposures. The high-contrast illumination on the buildings and the fountains is naturally synthesized. The dynamic range is increased from 200:1 to $10^6:1$, which is not possible with a single shot of a standard commodity camera. These examples describe dynamic scenes with moving persons or fountains, which are difficult to be generated using traditional methods involving multiple camera shots under different exposures.

6 DISCUSSION AND CONCLUSION

This paper has introduced a novel high dynamic range compression method, which effectively suppresses the global contrast while preserving local image structure details. We proposed a globally nonlinear method that uses overlapping window-based linear functions to reconstruct the image radiance. From a global perspective, the overlapping of the windows makes the modification of a pixel value be confined within a certain range largely depending on its neighborhood. Range compression is achieved in our method using guided linear models. Our method can also be applied to ordinary image enhancement and LDR image expanding in the same framework. The parameter adjustment is easy, where a range of values are suitable for tonal adjustment.

Our method is different from multiscales approaches that require the layer separation and combination steps. In our approach, each p and q locally influence a group of pixels, which, in turn, globally guide the optimization. We note that although our method focuses on the automatic dynamic range compression, it is not difficult to add user interactions for further interactive refinement. To modify the local radiance, the user can use stroke brushes to specify the desired values. The corresponding elements are then

fixed in solving (3). The user can also modify c_i locally to adjust the level of details in the results.

APPENDIX

All the terms in the energy definition in (4) are quadratic making f_i a continuous convex function. We, thus, compute (p_i^*, q_i^*) , the optimal solution of (p_i, q_i) , by taking the partial derivatives of function f_i with respect to p_i and q_i and setting them to zeros, similar to the derivations in [15]:

$$\begin{aligned} \frac{\partial f_i}{\partial p_i} \Big|_{p_i=p_i^*, q_i=q_i^*} &= 2\epsilon c_i^{-2} (p_i^* - c_i) \\ &+ \sum_{j \in w_i} 2(I^l(j) - p_i^* I^h(j) - q_i^*) \cdot (-I^h(j)) = 0, \end{aligned} \quad (11)$$

$$\frac{\partial f_i}{\partial q_i} \Big|_{p_i=p_i^*, q_i=q_i^*} = \sum_{j \in w_i} -2(I^l(j) - p_i^* I^h(j) - q_i^*) = 0. \quad (12)$$

Equations (11) and (12) are linear equations with respect to p_i and q_i , thus, can be written into a matrix form

$$H_i \cdot [p_i^* \ q_i^*]^T = \eta_i, \quad (13)$$

where

$$H_i = \begin{bmatrix} \epsilon c_i^{-2} + \sum_{j \in w_i} I^h(j)^2 & \sum_{j \in w_i} I^h(j) \\ \sum_{j \in w_i} I^h(j) & \sum_{j \in w_i} 1 \end{bmatrix}$$

and

$$\eta_i = \begin{bmatrix} \epsilon c_i^{-1} + \sum_{j \in w_i} I^h(j) \cdot I^l(j) \\ \sum_{j \in w_i} I^l(j) \end{bmatrix}.$$

By solving (13), we obtain

$$[p_i^* \ q_i^*]^T = H_i^{-1} \cdot \eta_i = \frac{1}{m_i \cdot \Delta_i} \begin{bmatrix} 1 & -\mu_i \\ -\mu_i & \Delta_i + \mu_i^2 \end{bmatrix} \cdot \eta_i, \quad (14)$$

where

$$\begin{aligned} m_i &= \sum_{j \in w_i} 1, \quad \mu_i = \frac{1}{m_i} \sum_{j \in w_i} I^h(j), \\ \Delta_i &= \sigma_i^2 + \left(\frac{\epsilon c_i^{-2}}{m_i} \right). \end{aligned}$$

After estimating p_i and q_i , we similarly compute \hat{I}^l , the optimal solution of (3), by taking the partial derivatives of function f with respect to $I^l(k)$ and setting them to zeros:

$$\frac{\partial f}{\partial I^l(k)} \Big|_{I^l=\hat{I}^l} = \sum_{i|k \in w_i} 2(\hat{I}^l(k) - p_i^* I^h(k) - q_i^*) = 0. \quad (15)$$

Since f is also a continuous convex function in differentiability class C^∞ with respect to I^l , the optimal solution of I^l

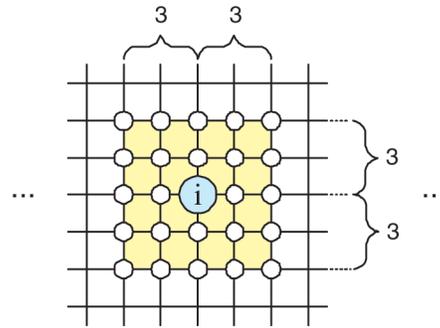


Fig. 22. Illustration of the region containing pixel i . All windows w with size 3×3 and containing i must be within the yellow region. There are at most 24 pixels in relation to i by local windows w . Since there are N pixels in an image, the total number of nonzero elements in matrix S will be less than $25 \times N$.

by solving (15) can be obtained. Combining (14) and (15), we get

$$\begin{aligned} \frac{\partial f}{\partial I^l(k)} \Big|_{I^l=\hat{I}^l} &= \sum_{i|k \in w_i} \left(\hat{I}^l(k) - \frac{1}{m_i \Delta_i} \left(\epsilon c_i^{-1} I^h(k) \right. \right. \\ &+ I^h(k) \sum_{j \in w_i} I^h(j) \hat{I}^l(j) - \mu_i I^h(k) \sum_{j \in w_i} \hat{I}^l(j) \\ &- \mu_i \epsilon c_i^{-1} - \mu_i \sum_{j \in w_i} I^h(j) \hat{I}^l(j) \\ &\left. \left. + (\Delta_i + \mu_i^2) \sum_{j \in w_i} \hat{I}^l(j) \right) \right), \end{aligned} \quad (16)$$

$$= 0.$$

It can be observed that (16) is a linear combination of \hat{I}^l 's. We can rewrite it in a form of a large linear system:

$$S \cdot \hat{I}^l = B, \quad (17)$$

where

$$\begin{aligned} s_{kj} &= \frac{\partial^2 f}{\partial I^l(k) \partial \hat{I}^l(j)} \\ &= \sum_{i|\{k,j\} \subset w_i} \left(\delta_{kj} - \frac{1}{m_i \Delta_i} ((I^h(k) - \mu_i)(I^h(j) - \mu_i) + \Delta_i) \right), \\ b_k &= \sum_{i|k \in w_i} \frac{\epsilon}{m_i \Delta_i c_i} (I^h(k) - \mu_i), \end{aligned} \quad (18)$$

where δ_{kj} is the Kronecker delta. The optimal solution of (3) can be obtained by solving a linear system (17).

The matrix S in (17) is symmetric and sparse. In each row i of S , the number of nonzero elements is up to a defined region size according to (17). Given an input image containing N pixels and the size of the window w_i set to $k \times k$, S has less than $(2k-1)^2 \times N$ nonzero elements (a $k=3$ example is shown in Fig. 22). Mathematically, the minimum possible window size is 2×2 , as there are three unknowns in each local window and at least three linear constraints are needed to make the problem well-posed.

If the image has moderate resolution, it is still possible to directly apply the symmetric LQ or generalized minimum residual methods to solve the linear system (17) in Matlab. However, to process an image with resolution higher than $1,000 \times 1,000$, the above methods will take several minutes to converge. To accelerate, a multigrid method [10], same as the one used in [15], is employed that produces a satisfactory \hat{I}^l in only seconds in all our experiments.

ACKNOWLEDGMENTS

The authors would like to thank the associate editor and all reviewers for their constructive comments to improve the manuscript. This work was supported by a grant from the Research Grants Council of the Hong Kong Special Administrative Region, China (Project No. 412708), a grant from SHIAE (No. 8115016), and an AcRF Tier 1 Singapore grant (Project No. R-252-000-333-133).

REFERENCES

- [1] A.O. Akyüz, R. Fleming, B.E. Riecke, E. Reinhard, and H.H. Bühlhoff, "Do HDR Displays Support LDR Content? A Psychological Evaluation," *Proc. ACM SIGGRAPH*, p. 38, 2007.
- [2] M. Ashikhmin, "A Tone Mapping Algorithm for High Contrast Images," *Proc. 13th Eurographics Workshop Rendering*, pp. 145-156, 2002.
- [3] F. Banterle, P. Ledda, K. Debattista, and A. Chalmers, "Inverse Tone Mapping," *Proc. Int'l Conf. Computer Graphics and Interactive Techniques (GRAPHITE)*, pp. 349-356, 2006.
- [4] K. Chiu, M. Herf, P. Shirley, S. Swamy, C. Wang, and K. Zimmerman, "Spatially Nonuniform Scaling Functions for High Contrast Images," *Proc. Conf. Graphics Interface*, pp. 245-253, 1993.
- [5] K. Devlin, "A Review of Tone Reproduction Techniques," Technical Report CSTR-02-005, Dept. of Computer Science, Univ. of Bristol, 2002.
- [6] J. Duan, G. Qiu, and M. Chen, "Comprehensive Fast Tone Mapping for High Dynamic Range Image Visualization," *Proc. Conf. Pacific Graphics*, 2005.
- [7] F. Durand and J. Dorsey, "Fast Bilateral Filtering for the Display of High-Dynamic-Range Images," *Proc. ACM SIGGRAPH*, pp. 257-266, 2002.
- [8] Z. Farbman, R. Fattal, D. Lischinski, and R. Szeliski, "Edge-Preserving Decompositions for Multi-Scale Tone and Detail Manipulation," *Proc. ACM SIGGRAPH*, 2008.
- [9] R. Fattal, D. Lischinski, and M. Werman, "Gradient Domain High Dynamic Range Compression," *Proc. ACM SIGGRAPH*, pp. 249-256, 2002.
- [10] W.W. Hager, S.-J. Huang, P.M. Pardalos, and O.A. Prokopyev, *Multiscale Optimization Methods and Applications (Nonconvex Optimization and Its Applications)*. Springer, 2005.
- [11] B.K.P. Horn, "Determining Lightness from an Image," *Computer Graphics and Image Processing*, vol. 3, pp. 277-299, 1974.
- [12] 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 Trans. Image Processing*, vol. 6, no. 7, pp. 965-976, 1997.
- [13] G. Krawczyk, K. Myszkowski, and H.-P. Seidel, "Computational Model of Lightness Perception in High Dynamic Range Imaging," *Proc. IS&T/SPIE's Ann. Symp. Human Vision and Electronic Imaging*, 2006.
- [14] G.W. Larson, H. Rushmeier, and C. Piatko, "A Visibility Matching Tone Reproduction Operator for High Dynamic Range Scenes," *IEEE Trans. Visualization and Computer Graphics*, vol. 3, no. 4, pp. 291-306, Oct.-Dec. 1997.
- [15] A. Levin, D. Lischinski, and Y. Weiss, "A Closed Form Solution to Natural Image Matting," *Proc. Int'l Conf. Computer Vision and Pattern Recognition (CVPR)*, 2006.
- [16] Y. Li, L. Sharan, and E.H. Adelson, "Compressing and Combanding High Dynamic Range Images with Subband Architectures," *Proc. ACM SIGGRAPH*, pp. 836-844, 2005.
- [17] D. Lischinski, Z. Farbman, M. Uyttendaele, and R. Szeliski, "Interactive Local Adjustment of Tonal Values," *Proc. ACM SIGGRAPH*, pp. 646-653, 2006.
- [18] L. Meylan, S. Daly, and S. Süsstrunk, "The Reproduction of Specular Highlights on High Dynamic Range Displays," *Proc. Color Imaging Conf.*, 2006.
- [19] 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," *Proc. ACM SIGGRAPH*, pp. 287-298, 1998.
- [20] E. Reinhard, M. Stark, P. Shirley, and J. Ferwerda, "Photographic Tone Reproduction for Digital Images," *Proc. ACM SIGGRAPH*, vol. 21, no. 3, pp. 267-276, 2002.
- [21] E. Reinhard, G. Ward, S. Pattanaik, and P. Debevec, *High Dynamic Range Imaging*. Morgan Kaufman, 2005.
- [22] A.G. Rempel, M. Trentacoste, H. Seetzen, H.D. Young, W. Heidrich, L. Whitehead, and G. Ward, "Ldr2hdr: On-the-Fly Reverse Tone Mapping of Legacy Video and Photographs," *Proc. ACM SIGGRAPH*, p. 39, 2007.
- [23] C. Schlick, "A Customizable Reflectance Model for Everyday Rendering," *Proc. Fourth Eurographics Workshop Rendering*, pp. 73-83, 1993.
- [24] T.G. Stockham, "Image Processing in the Context of a Visual Model," *Proc. IEEE*, vol. 60, no. 7, pp. 828-842, 1972.
- [25] J. Tumblin, "Three Methods of Detail-Preserving Contrast Reduction for Displayed Images," PhD thesis, Georgia Inst. of Technology, 1999.
- [26] J. Tumblin and G. Turk, "LCIS: A Boundary Hierarchy for Detail-Preserving Contrast Reduction," *Proc. ACM SIGGRAPH*, pp. 83-90, 1999.



Qi Shan received the joint bachelor's degree from Fudan University in China and the University College Dublin in Ireland in 2006 and the MPhil degree in computer science from the Chinese University of Hong Kong in 2008. Currently, he is working toward the PhD degree in the Department of Computer Science and Engineering at the University of Washington, Seattle. His research interests include computer graphics and vision. He is a student member of the IEEE.



Jiaya Jia received the PhD degree in computer science from the Hong Kong University of Science and Technology in 2004. He joined the Department of Computer Science and Engineering at the Chinese University of Hong Kong in September 2004, where he is currently an assistant professor. His research interests include vision geometry, image/video editing and enhancement, and motion deblurring. He has served on the program committees of ICCV, CVPR, ECCV, and ACCV. He served as the cochair of the Interactive Computer Vision Workshop 2007 (in conjunction with ICCV '07). He is a member of the IEEE.



Michael S. Brown received the BS and PhD degrees in computer science from the University of Kentucky in 1995 and 2001, respectively. He is currently the Sung Kah Kay assistant professor in the School of Computing at the National University of Singapore. His research interests include computer vision, image processing and computer graphics. He served as the general cochair of the 5th ACM/IEEE Projector-Camera-Systems (PROCAMS '08) Workshop collocated with SIGGRAPH 2008 and is an organizer for the eHeritage 2009 Workshop collocated with ICCV 2009. He was an area chair of the CVPR 2009. He is a member of the IEEE.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.