

Gradient Domain High Dynamic Range Compression

Raanan Fattal

Dani Lischinski

Michael Werman

School of Computer Science and Engineering*
The Hebrew University of Jerusalem

Abstract

We present a new method for rendering high dynamic range images on conventional displays. Our method is conceptually simple, computationally efficient, robust, and easy to use. We manipulate the gradient field of the luminance image by attenuating the magnitudes of large gradients. A new, low dynamic range image is then obtained by solving a Poisson equation on the modified gradient field. Our results demonstrate that the method is capable of drastic dynamic range compression, while preserving fine details and avoiding common artifacts, such as halos, gradient reversals, or loss of local contrast. The method is also able to significantly enhance ordinary images by bringing out detail in dark regions.

CR Categories: I.3.3 [Computer Graphics]: Picture/image generation—display algorithms, viewing algorithms; I.4.3 [Image Processing and Computer Vision]: Enhancement—filtering, grayscale manipulation, sharpening and deblurring

Keywords: digital photography, high dynamic range compression, image-based rendering, image processing, signal processing, tone mapping

1 Introduction

High dynamic range (HDR) radiance maps are becoming increasingly common and important in computer graphics. Initially, such maps originated almost exclusively from physically-based lighting simulations. Today, however, HDR maps of real scenes are very easy to construct: all you need is a few differently exposed photographs of the scene [Debevec and Malik 1997], or a panoramic video scan of it [Aggarwal and Ahuja 2001a; Schechner and Nayar 2001]. Furthermore, based on recent developments in digital imaging technology [Aggarwal and Ahuja 2001b; Nayar and Mitsunaga 2000], it is reasonable to assume that tomorrow's digital still and video cameras will capture HDR images and video directly.

HDR images have many advantages over standard low dynamic range images [Debevec and Malik 1997], and several applications have been demonstrated where such images are extremely useful [Debevec 1998; Cohen et al. 2001]. However, HDR images also pose a difficult challenge: given that the dynamic range of various common display devices (monitors, printers, etc.) is much smaller than the dynamic range commonly found in real-world scenes, how can we display HDR images on low dynamic range (LDR) display

devices, while preserving as much of their visual content as possible? This is precisely the problem addressed in this paper.

The problem that we are faced with is vividly illustrated by the series of images in Figure 1. These photographs were taken using a digital camera with exposure times ranging from 1/1000 to 1/4 of a second (at f/8) from inside a lobby of a building facing glass doors leading into a sunlit inner courtyard. Note that each exposure reveals some features that are not visible in the other photographs¹. For example, the true color of the areas directly illuminated by the sun can be reliably assessed only in the least exposed image, since these areas become over-exposed in the remainder of the sequence. The color and texture of the stone tiles just outside the door are best captured in the middle image, while the green color and the texture of the ficus plant leaves becomes visible only in the very last image in the sequence. All of these features, however, are simultaneously clearly visible to a human observer standing in the same location, because of adaptation that takes place as our eyes scan the scene [Pattanaik et al. 1998]. Using Debevec and Malik's method [1997], we can compile these 8-bit images into a single HDR radiance map with dynamic range of about 25,000:1. However, it is not at all clear how to display such an image on a CRT monitor whose dynamic range is typically below 100:1!

In this paper, we present a new technique for *high dynamic range compression* that enables HDR images, such as the one described in the previous paragraph, to be displayed on LDR devices. The proposed technique is quite effective, as demonstrated by the top image in Figure 2, yet it is conceptually simple, computationally efficient, robust, and easy to use. Observing that drastic changes in luminance across an HDR image must give rise to luminance gradients of large magnitudes, our approach is to manipulate the gradient field of the luminance image by attenuating the magnitudes of large gradients. A new, low dynamic range image is then obtained by solving a Poisson equation on the modified gradient field.

The current work is not the first attempt to tackle this important problem. Indeed, quite a few different methods have appeared in the literature over the last decade. A more detailed review of previous work is provided in Section 2. In this paper we hope to convince the reader that our approach has some definite advantages over previous solutions: it does a better job at preserving local contrasts than some previous methods, has fewer visible artifacts than others, and yet it is fast and easy to use. These advantages make our technique a highly practical tool for high dynamic range compression.

We do not attempt to faithfully reproduce the response of the human visual system to the original high dynamic range scenes. Nevertheless, we have achieved very good results on HDR radiance maps of real scenes, HDR panoramic video mosaics, ordinary high-contrast photographs, and medical images.

2 Previous work

In the past decade there has been considerable work concerned with displaying high dynamic range images on low dynamic range dis-

*e-mail: {raananf | danix | werman}@cs.huji.ac.il

¹All of the images in this paper are provided at full resolution on the proceedings CD-ROM (and also at <http://www.cs.huji.ac.il/~danix/hdrc>), as some of the fine details may be difficult to see in the printed proceedings.



Figure 1: A series of five photographs. The exposure is increasing from left (1/1000 of a second) to right (1/4 of a second).

play devices. In this section we provide a brief review of previous work. More detailed and in-depth surveys are presented by DiCarlo and Wandell [2001] and Tumblin *et al.* [1999].

Most HDR compression methods operate on the luminance channel or perform essentially the same processing independently in each of the RGB channels, so throughout most of this paper we will treat HDR maps as (scalar) luminance functions.

Previous approaches can be classified into two broad groups: (1) global (spatially invariant) mappings, and (2) spatially variant operators. DiCarlo and Wandell [2001] refer to the former as TRCs (tone reproduction curves) and to the latter as TROs (tone reproduction operators); we adopt these acronyms for the remainder of this paper.

The most naive TRC linearly scales the HDR values such that they fit into a canonic range, such as $[0, 1]$. Such scaling preserves relative contrasts perfectly, but the displayed image may suffer severe loss of visibility whenever the dynamic range of the display is smaller than the original dynamic range of the image, and due to quantization. Other common TRCs are gamma correction and histogram equalization.

In a pioneering work, Tumblin and Rushmeier [1993] describe a more sophisticated non-linear TRC designed to preserve the apparent brightness of an image based on the actual luminances present in the image and the target display characteristics. Ward [1994] suggested a simpler linear scale factor automatically determined from image luminances so as to preserve apparent contrast and visibility around a particular adaptation level. The most recent and most sophisticated, to our knowledge, TRC is described by Ward Larson *et al.* [1997]. They first describe a clever improvement to histogram equalization, and then show how to extend this idea to incorporate models of human contrast sensitivity, glare, spatial acuity, and color sensitivity effects. This technique works very well on a wide variety of images.

The main advantage of TRCs lies in their simplicity and computational efficiency: once a mapping has been determined, the image may be mapped very quickly, e.g., using lookup tables. However, such global mappings must be one-to-one and monotonic in order to avoid reversals of local edge contrasts. As such, they have a fundamental difficulty preserving local contrasts in images where the intensities of the regions of interest populate the entire dynamic range in a more or less uniform fashion. This shortcoming is illustrated in the middle image of Figure 2. In this example, the distribution of luminances is almost uniform, and Ward Larson’s technique results in a mapping, which is rather similar to a simple gamma correction. As a result, local contrast is drastically reduced.

Spatially variant tone reproduction operators are more flexible than TRCs, since they take local spatial context into account when deciding how to map a particular pixel. In particular, such operators can transform two pixels with the same luminance value to different display luminances, or two different luminances to the same display intensity. This added flexibility in the mapping should make it possible to achieve improved local contrast.

The problem of high-dynamic range compression is intimately related to the problem of recovering reflectances from an image [Horn 1974]. An image $I(x, y)$ is regarded as a product

$$I(x, y) = R(x, y) L(x, y),$$

where $R(x, y)$ is the reflectance and $L(x, y)$ is the illuminance at each point (x, y) . The function $R(x, y)$ is commonly referred to as the *intrinsic image* of a scene. The largest luminance variations in an HDR image come from the illuminance function L , since real-world reflectances are unlikely to create contrasts greater than 100:1². Thus, dynamic range compression can, in principle, be achieved by separating an image I to its R and L components, scaling down the L component to obtain a new illuminance function $\tilde{L}$, and re-multiplying:

$$\tilde{I}(x, y) = R(x, y) \tilde{L}(x, y).$$

Intuitively, this reduces the contrast between brightly illuminated areas and those in deep shadow, while leaving the contrasts due to texture and reflectance undistorted. Tumblin *et al.* [1999] use this approach for displaying high-contrast synthetic images, where the material properties of the surfaces and the illuminance are known at each point in the image, making it possible to compute a perfect separation of an image to various layers of lighting and surface properties.

Unfortunately, computing such a separation for real images is an *ill posed* problem [Ramamoorthi and Hanrahan 2001]. Consequently, any attempt to solve it must make some simplifying assumptions regarding R , L , or both. For example, *homomorphic filtering* [Stockham 1972], an early image enhancement technique, makes the assumption that L varies slowly across the image, in contrast to R that varies abruptly. This means that R can be extracted by applying a high-pass filter to the logarithm of the image. Exponentiating the result achieves simultaneous dynamic range compression and local contrast enhancement. Similarly, Horn [1974] assumes that L is smooth, while R is piecewise-constant, introducing infinite impulse edges in the Laplacian of the image’s logarithm. Thus, L may be recovered by thresholding the Laplacian. Of course, in most natural images the assumptions above are violated: for example, in sunlit scenes illuminance varies abruptly across shadow boundaries. This means that L also has high frequencies and introduces strong impulses into the Laplacian. As a result, attenuating only the low frequencies in homomorphic filtering may give rise to strong “halo” artifacts around strong abrupt changes in illuminance, while Horn’s method incorrectly interprets sharp shadows as changes in reflectance.

More recently, Jobson *et al.* [1997] presented a dynamic range compression method based on a multiscale version of Land’s “retinex” theory of color vision [Land and McCann 1971]. Retinex estimates the reflectances $R(x, y)$ as the ratio of $I(x, y)$ to its low-pass filtered version. A similar operator was explored by Chiu *et al.* [1993], and was also found to suffer from halo artifacts and dark bands around small bright visible light sources. Jobson *et al.* compute the logarithm of the retinex responses for several low-pass filters of different sizes, and linearly combine the results. The linear combination helps reduce halos, but does not eliminate them entirely. Schlick [1994] and Tanaka and Ohnishi [1997] also experimented with spatially variant operators and found them to produce halo artifacts.

Pattanaik and co-workers [1998] describe an impressively comprehensive computational model of human visual system adaptation

²For example, the reflectance of black velvet is about 0.01, while that of snow is roughly 0.93.



Figure 2: Belgium House: An HDR radiance map of a lobby compressed for display by our method (top), the method of Ward Larson *et al.* (middle) and the LCIS method (bottom).

and spatial vision for realistic tone reproduction. Their model enables display of HDR scenes on conventional display devices, but the dynamic range compression is performed by applying different gain-control factors to each bandpass, which also results in halos around strong edges. In fact, DiCarlo and Wandell [2001], as well as Tumblin and Turk [1999] demonstrate that this is a fundamental problem with any multi-resolution operator that compresses each resolution band differently.

In order to eradicate the notorious halo artifacts Tumblin and Turk [1999] introduce the *low curvature image simplifier* (LCIS) hierarchical decomposition of an image. Each level in this hierarchy is generated by solving a partial differential equation inspired by anisotropic diffusion [Perona and Malik 1990] with a different diffusion coefficient. The hierarchy levels are progressively smoother versions of the original image, but the smooth (low-curvature) regions are separated from each other by sharp boundaries. Dynamic range compression is achieved by scaling down the smoothest version, and then adding back the differences between successive levels in the hierarchy, which contain details removed by the simplification process. This technique is able to drastically compress the dynamic range, while preserving the fine details in the image. However, the results are not entirely free of artifacts. Tumblin and Turk note that weak halo artifacts may still remain around certain edges in strongly compressed images. In our experience, this technique sometimes tends to overemphasize fine details. For example, in the bottom image of Figure 2, generated using this technique, certain features (door, plant leaves) are surrounded by thin bright outlines. In addition, the method is controlled by no less than 8 parameters, so achieving an optimal result occasionally requires quite a bit of trial-and-error. Finally, the LCIS hierarchy construction is computationally intensive, so compressing a high-resolution image takes a substantial amount of time.

3 Gradient domain HDR compression

Informally, our approach relies on the widely accepted assumptions [DiCarlo and Wandell 2001] that the human visual system is not very sensitive to absolute luminances reaching the retina, but rather responds to local intensity ratio changes and reduces the effect of large global differences, which may be associated with illumination differences.

Our algorithm is based on the rather simple observation that any drastic change in the luminance across a high dynamic range image *must* give rise to large magnitude luminance gradients at some scale. Fine details, such as texture, on the other hand, correspond to gradients of much smaller magnitude. Our idea is then to identify large gradients at various scales, and attenuate their magnitudes while keeping their direction unaltered. The attenuation must be progressive, penalizing larger gradients more heavily than smaller ones, thus compressing drastic luminance changes, while preserving fine details. A reduced high dynamic range image is then reconstructed from the attenuated gradient field.

It should be noted that all of our computations are done on the logarithm of the luminances, rather than on the luminances themselves. This is also the case with most of the previous methods reviewed in the previous section. The reason for working in the log domain is twofold: (a) the logarithm of the luminance is a (crude) approximation to the perceived brightness, and (b) gradients in the log domain correspond to ratios (local contrasts) in the luminance domain.

We begin by explaining the idea in 1D. Consider a high dynamic range 1D function. We denote the logarithm of this function by $H(x)$. As explained above, our goal is to compress large magnitude changes in H , while preserving local changes of small magnitude, as much as possible. This goal is achieved by applying an appropriate spatially variant attenuating mapping Φ to the magnitudes of the derivatives $H'(x)$. More specifically, we compute:

$$G(x) = H'(x) \Phi(x).$$

Note that G has the same sign as the original derivative H' everywhere, but the magnitude of the original derivatives has been altered by a factor determined by Φ , which is designed to attenuate large derivatives more than smaller ones. Actually, as explained in Section 4, Φ accounts for the magnitudes of derivatives at different scales.




Figure 3: (a) An HDR scanline with dynamic range of 2415:1. (b) $H(x) = \log(\text{scanline})$. (c) The derivatives $H'(x)$. (d) Attenuated derivatives $G(x)$; (e) Reconstructed signal $I(x)$ (as defined in eq. 1); (f) An LDR scanline $\exp(I(x))$: the new dynamic range is 7.5:1. Note that each plot uses a different scale for its vertical axis in order to show details, except (c) and (d) that use the same vertical axis scaling in order to show the amount of attenuation applied on the derivatives.

We can now reconstruct a reduced dynamic range function I (up to an additive constant C) by integrating the compressed derivatives:

$$I(x) = C + \int_0^x G(t) dt. \quad (1)$$

Finally, we exponentiate in order to return to luminances. The entire process is illustrated in Figure 3.

In order to extend the above approach to 2D HDR functions $H(x, y)$ we manipulate the gradients ∇H , instead of the derivatives. Again, in order to avoid introducing spatial distortions into the image, we change only the magnitudes of the gradients, while keeping their directions unchanged. Thus, similarly to the 1D case, we compute

$$G(x, y) = \nabla H(x, y) \Phi(x, y).$$

Unlike the 1D case we cannot simply obtain a compressed dynamic range image by integrating G , since it is not necessarily integrable. In other words, there might not exist an image I such that $G = \nabla I$! In fact, the gradient of a potential function (such as a 2D image) must be a *conservative field* [Harris and Stocker 1998]. In other words, the gradient $\nabla I = (\partial I / \partial x, \partial I / \partial y)$ must satisfy

$$\frac{\partial^2 I}{\partial x \partial y} = \frac{\partial^2 I}{\partial y \partial x},$$

which is rarely the case for our G .

One possible solution to this problem is to orthogonally project G onto a finite set of orthonormal basis functions spanning the set of integrable vector fields, such as the Fourier basis functions [Frankot and Chellappa 1988]. In our method we employ a more direct and more efficient approach: search the space of all 2D potential functions for a function I whose gradient is the closest to G in the least-squares sense. In other words, I should minimize the integral

$$\iint F(\nabla I, G) dx dy, \quad (2)$$

where $F(\nabla I, G) = \|\nabla I - G\|^2 = \left(\frac{\partial I}{\partial x} - G_x\right)^2 + \left(\frac{\partial I}{\partial y} - G_y\right)^2$.

According to the Variational Principle, a function I that minimizes the integral in (2) must satisfy the Euler-Lagrange equation

$$\frac{\partial F}{\partial I} - \frac{d}{dx} \frac{\partial F}{\partial I_x} - \frac{d}{dy} \frac{\partial F}{\partial I_y} = 0,$$

which is a partial differential equation in I . Substituting F we obtain the following equation:

$$2 \left(\frac{\partial^2 I}{\partial x^2} - \frac{\partial G_x}{\partial x} \right) + 2 \left(\frac{\partial^2 I}{\partial y^2} - \frac{\partial G_y}{\partial y} \right) = 0.$$

Dividing by 2 and rearranging terms, we obtain the well-known Poisson equation

$$\nabla^2 I = \text{div } G \quad (3)$$




Figure 4: Gradient attenuation factors used to compress the Belgium House HDR radiance map (Figure 2). Darker shades indicate smaller scale factors (stronger attenuation).

where ∇^2 is the Laplacian operator $\nabla^2 I = \frac{\partial^2 I}{\partial x^2} + \frac{\partial^2 I}{\partial y^2}$ and $\text{div } G$ is the divergence of the vector field G , defined as $\text{div } G = \frac{\partial G_x}{\partial x} + \frac{\partial G_y}{\partial y}$. This is a linear partial differential equation, whose numerical solution is described in Section 5.

4 Gradient attenuation function

As explained in the previous section, our method achieves HDR compression by attenuating the magnitudes of the HDR image gradients by a factor of $\Phi(x, y)$ at each pixel. We would like the attenuation to be progressive, shrinking gradients of large magnitude more than small ones.

Real-world images contain edges at multiple scales. Consequently, in order to reliably detect all of the significant intensity transitions we must employ a multi-resolution edge detection scheme. However, we cannot simply attenuate each gradient at the resolution where it was detected. This could result in halo artifacts around strong edges, as mentioned in Section 2. Our solution is to propagate the desired attenuation from the level it was detected at to the full resolution image. Thus, all gradient manipulations occur at a single resolution level, and no halo artifacts arise.

We begin by constructing a Gaussian pyramid $H_0, H_1, \dots, H_d$, where H_0 is the full resolution HDR image and H_d is the coarsest level in the pyramid. d is chosen such that the width and the height of H_d are at least 32. At each level k we compute the gradients using central differences:

$$\nabla H_k = \left(\frac{H_k(x+1, y) - H_k(x-1, y)}{2^{k+1}}, \frac{H_k(x, y+1) - H_k(x, y-1)}{2^{k+1}} \right).$$

At each level k a scaling factor $\phi_k(x, y)$ is determined for each pixel

based on the magnitude of the gradient there:

$$\varphi_k(x, y) = \frac{\alpha}{\|\nabla H_k(x, y)\|} \left(\frac{\|\nabla H_k(x, y)\|}{\alpha} \right)^\beta.$$

This is a two-parameter family of functions. The first parameter α determines which gradient magnitudes remain unchanged (multiplied by a scale factor of 1). Gradients of larger magnitude are attenuated (assuming $\beta < 1$), while gradients of magnitude smaller than α are slightly magnified. In all the results shown in this paper we set α to 0.1 times the average gradient magnitude, and β between 0.8 and 0.9.

The full resolution gradient attenuation function $\Phi(x, y)$ is computed in a top-down fashion, by propagating the scaling factors $\varphi_k(x, y)$ from each level to the next using linear interpolation and accumulating them using pointwise multiplication. More formally, the process is given by the equations:

$$\begin{aligned}\Phi_d(x, y) &= \varphi_d(x, y) \\ \Phi_k(x, y) &= L(\Phi_{k+1})(x, y) \varphi_k(x, y) \\ \Phi(x, y) &= \Phi_0(x, y)\end{aligned}$$

where d is the coarsest level, Φ_k denotes the accumulated attenuation function at level k , and L is an upsampling operator with linear interpolation. As a result, the gradient attenuation at each pixel of the finest level is determined by the strengths of all the edges (from different scales) passing through that location in the image. Figure 4 shows attenuation coefficients computed for the Belgium House HDR radiance map.

It is important to note that although the computation of the gradient attenuation function is done in a multi-resolution fashion, ultimately only the gradients at the finest resolution are manipulated, thus avoiding halo artifacts that typically arise when different resolution levels are manipulated separately.

5 Implementation

In order to solve a differential equation such as (3) one must first specify the boundary conditions. In our case, the most natural choice appears to be the Neumann boundary conditions $\nabla I \cdot \mathbf{n} = 0$ (the derivative in the direction normal to the boundary is zero). With these boundary conditions the solution is now defined up to a single additive term, which has no real meaning since we shift and scale the solution in order to fit it into the display device limits.

Since both the Laplacian ∇^2 and div are linear operators, approximating them using standard finite differences yields a linear system of equations. More specifically, we approximate:

$$\nabla^2 I(x, y) \approx I(x+1, y) + I(x-1, y) + I(x, y+1) + I(x, y-1) - 4I(x, y)$$

taking the pixel grid spacing to be 1 at the full resolution of the image. The gradient ∇H is approximated using the forward difference

$$\nabla H(x, y) \approx (H(x+1, y) - H(x, y), H(x, y+1) - H(x, y)),$$

while for $\text{div } G$ we use backward difference approximations

$$\text{div } G \approx G_x(x, y) - G_x(x-1, y) + G_y(x, y) - G_y(x, y-1).$$

This combination of forward and backward differences ensures that the approximation of $\text{div } G$ is consistent with the central difference scheme used for the Laplacian.

At the boundaries we use the same definitions, but assume that the derivatives around the original image grid are 0. For example, for each pixel on the left image boundary we have the equation $I(-1, y) - I(0, y) = 0$.

The finite difference scheme yields a large system of linear equations — one for each pixel in the image, but the corresponding matrix has only five nonzero elements in each row, since each pixel is coupled only with its four neighbors. We solve this system using the *Full Multigrid Algorithm* [Press et al. 1992], with Gauss-Seidel smoothing iterations. This leads to $O(n)$ operations to reach an approximate solution, where n is the number of pixels in the image. Another alternative is to use a “rapid Poisson solver”, which uses the fast Fourier transform to invert the Laplacian operator. However, the complexity with this approach would be $O(n \log n)$.

As mentioned earlier, our method operates on the luminances of an HDR radiance map. In order to assign colors to the pixels of the compressed dynamic range image we use an approach similar to those of Tumblin and Turk [1999] and Schlick [1994]. More specifically, the color channels of a pixel in the compressed dynamic range image are computed as follows:

$$C_{\text{out}} = \left(\frac{C_{\text{in}}}{L_{\text{in}}} \right)^s L_{\text{out}}$$

for $C = R, G, B$. L_{in} and L_{out} denote the luminance before and after HDR compression, respectively, and the exponent s controls the color saturation of the resulting image. We found values between 0.4 and 0.6 to produce satisfactory results.

6 Results

Multiple exposure HDRs. We have experimented with our method on a variety of HDR radiance maps of real scenes. In all cases, our method produced satisfactory results without much parameter tweaking. In certain cases we found that the subjective quality of the resulting image is slightly enhanced by running a standard sharpening operation. The computation times range from 1.1 seconds for an 512 by 384 image to 4.5 seconds for an 1024 by 768 image on a 1800 MHz Pentium 4.

The top row in Figure 5 shows three different renderings of a “streetlight on a foggy night” radiance map³. The dynamic range in this scene exceeds 100,000:1. The left image was produced using the method of Ward Larson *et al.* [1997], and the right image⁴ was produced by Tumblin and Turk’s [1999] LCIS method. The middle image was generated by our method. The left image loses visibility in a wide area around the bright light, details are lost in the shadowed regions, and the texture on the ground is washed out. The LCIS image (right) exhibits a grainy texture in smooth areas, and appears to slightly overemphasize edges, resulting in an “embossed”, non-photorealistic appearance. In our image (middle) smoothness is preserved in the foggy sky, yet at the same time fine details are well preserved (tree leaves, ground texture, car outlines). Our method took 5 seconds to compute this 751 by 1130 image, while the LCIS method took around 8.5 minutes.

The second row of images in Figure 5 shows a similar comparison using an HDR radiance map of the Stanford Memorial church⁵. The dynamic range in this map exceeds 250,000:1. Overall, the same observations as before hold for this example as well. In the left image the details in the dark regions are difficult to see, while the skylight and the stained glass windows appear over-exposed. In the LCIS image (right) the floor appears slightly bumpy, while our image (middle) shows more details and conveys a more realistic impression.

The last row of images⁶ in Figure 5 and the first row in Figure 6 show several additional examples of HDR compression by

³Radiance map courtesy of Jack Tumblin, Northwestern University.

⁴Image reprinted by permission, © 1999 Jack Tumblin [1999].

⁵Radiance map courtesy of Paul Debevec.

⁶Source exposures courtesy of Shree Nayar.




Figure 5: The top two rows compare results produced by our method (middle column) to those of Ward Larson *et al.* (left column) and those of Tumblin and Turk (right column). The differences are discussed in Section 6. The bottom row shows three more examples of results produced by our method (the thumbnails next to each image show some of the LDR images from which the HDR radiance map was constructed).




Figure 6: Top row: more examples of HDR radiance map compression. Our method successfully combines features that can only be captured using very different exposures into a single image. Second row: an HDR panoramic video mosaic. The remaining images demonstrate that our method can also be used for ordinary image enhancement. See Section 6 for more detailed explanations.

our method. Next to each image there are thumbnail images showing some of the exposures used to construct the HDR map. Note that our method manages to combine in a realistic manner details that can only be captured with very different exposures. The top right image in Figure 6 (a stained glass window in the National Cathedral in Washington, DC) was made using only the two exposures shown on its left⁷. These two exposures are roughly four stops apart; they were taken by a professional photographer, who manually blended these two images together in order to display both the bright window and the dark stone surfaces simultaneously (<http://users.erols.com/maxlyons>). Our method achieves a similar effect automatically, while revealing more detail in the dark regions.

HDR panoramic video mosaics. A popular way to acquire a panoramic image is to scan a scene using a video camera and then construct a mosaic from the video frames. If we let the camera's auto-exposure control set the correct exposure for each frame, each scene element is imaged at multiple aperture settings and we can construct an HDR as in [Debevec and Malik 1997]. Using specialized hardware [Schechner and Nayar 2001; Aggarwal and Ahuja 2001a] also produced HDR panoramic video mosaics.

The second row in Figure 6 shows an HDR panorama compressed by our method. The top left image simulates what the panorama would have looked like with an exposure suitable for the left part of the panorama, while the one below it simulates an exposure suitable for the right part of the panorama. Clearly, none of these two exposure settings yield satisfactory results. With our HDR compression method we were able to obtain the panorama on the right, in which detail is visible across the entire field of view.

LDR image enhancement. Our method can also be used to enhance ordinary (LDR) images. By attenuating strong gradients and rescaling the reconstructed image back to the original 0..255 range, small contrasts in dark regions become easier to see.

The five images of the Notre Dame de Paris (Figure 6) demonstrate image enhancement using our method. The top left image is the original; the top right image is the result produced by our method. The bottom row shows the best results we could obtain with gamma correction (left), histogram equalization (middle), and contrast limited adaptive histogram equalization [Pizer et al. 1987] (right). Notice that our result brings out more details from the shadowed areas, while maintaining good contrasts elsewhere (brick wall in the foreground, fine details on the building). Adaptive histogram equalization is almost as good, but it introduces halo artifacts in the sky along the roofs and the treetop on the right.

The bottom row of Figure 6 shows two more examples. The example on the left is a typical example of an image containing sunlight and shadows. Again, our method (on its right) succeeds in bringing out the details from the shadowed areas.

The pair on the right shows a dark, low contrast fluoroscopic femur image. After enhancement using our method the bone structure is visible much more clearly (note that the femur canal becomes clearly visible).

7 Conclusions and Future Work

We have described a new, simple, computationally efficient, and robust method for high dynamic range compression, which makes it possible to display HDR images on conventional displays. Our method attenuates large gradients and then constructs a low dynamic range image by solving a Poisson equation on the modified gradient field.

Future work will concentrate on the many different exciting possible applications of the construction of an image from modified gradient fields. Preliminary results show promise in denoising, edge manipulation and non-photorealistic rendering from real images. In addition, we would like to extend our work so as to in-

corporate various psychophysical properties of human visual perception in order to make our technique more useful for applications such as lighting design or visibility analysis.

Acknowledgments

We would like to thank Paul Debevec, Max Lyons, Shree Nayar, Jack Tumblin, and Greg Ward for making their code and images available. Thanks also go to Siggraph's anonymous reviewers for their comments. This work was supported in part by the Israel Science Foundation founded by the Israel Academy of Sciences and Humanities.

References

- AGGARWAL, M., AND AHUJA, N. 2001. High dynamic range panoramic imaging. In *Proc. IEEE ICCV*, vol. I, 2–9.
- AGGARWAL, M., AND AHUJA, N. 2001. Split aperture imaging for high dynamic range. In *Proc. IEEE ICCV*, vol. II, 10–17.
- CHIU, K., HERF, M., SHIRLEY, P., SWAMY, S., WANG, C., AND ZIMMERMAN, K. 1993. Spatially nonuniform scaling functions for high contrast images. In *Proc. Graphics Interface '93*, Morgan Kaufmann, 245–253.
- COHEN, J., TCHOU, C., HAWKINS, T., AND DEBEVEC, P. 2001. Real-time high-dynamic range texture mapping. In *Rendering Techniques 2001*, S. J. Gortler and K. Myszkowski, Eds. Springer-Verlag, 313–320.
- DEBEVEC, P. E., AND MALIK, J. 1997. Recovering high dynamic range radiance maps from photographs. In *Proc. ACM SIGGRAPH 97*, T. Whitted, Ed., 369–378.
- DEBEVEC, P. 1998. Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *Proc. ACM SIGGRAPH 98*, M. Cohen, Ed., 189–198.
- DICARLO, J. M., AND WANDELL, B. A. 2001. Rendering high dynamic range images. In *Proceedings of the SPIE: Image Sensors*, vol. 3965, 392–401.
- FRANKOT, R. T., AND CHELLAPPA, R. 1988. A method for enforcing integrability in shape from shading algorithms. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 10, 4 (July), 439–451.
- HARRIS, J. W., AND STOCKER, H. 1998. *Handbook of Mathematics and Computational Science*. Springer-Verlag.
- HORN, B. K. P. 1974. Determining lightness from an image. *Computer Graphics and Image Processing* 3, 1 (Dec.), 277–299.
- JOBSON, D. J., RAHMAN, Z., AND WOODILL, G. A. 1997. A multi-scale Retinex for bridging the gap between color images and the human observation of scenes. *IEEE Transactions on Image Processing* 6, 7 (July), 965–976.
- LAND, E. H., AND MCCANN, J. J. 1971. Lightness and Retinex theory. *Journal of the Optical Society of America* 61, 1 (Jan.), 1–11.
- NAYAR, S. K., AND MITSUNAGA, T. 2000. High dynamic range imaging: Spatially varying pixel exposures. In *Proc. IEEE CVPR*.
- PATTANAIK, S. N., FERWERDA, J. A., FAIRCHILD, M. D., AND GREENBERG, D. P. 1998. A multiscale model of adaptation and spatial vision for realistic image display. In *Proc. ACM SIGGRAPH 98*, M. Cohen, Ed., 287–298.
- PERONA, P., AND MALIK, J. 1990. Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 12, 7 (July), 629–639.
- PIZER, S. M., AMBURN, E. P., AUSTIN, J. D., CROMARTIE, R., GESELOWITZ, A., GREER, T., TER HAAR ROMENY, B., ZIMMERMAN, J. B., AND ZUIDERVELD, K. 1987. Adaptive histogram equalization and its variations. *Computer Vision, Graphics, and Image Processing* 39, 3 (Sept.), 355–368.
- PRESS, W. H., TEUKOLSKY, S. A., VETTERLING, W. T., AND FLANNERY, B. P. 1992. *Numerical Recipes in C: The Art of Scientific Computing*, 2nd ed. Cambridge University Press.
- RAMAMOORTHY, R., AND HANRAHAN, P. 2001. A signal-processing framework for inverse rendering. In *Proc. ACM SIGGRAPH 2001*, E. Fiume, Ed., 117–128.
- SCHECHNER, Y. Y., AND NAYAR, S. K. 2001. Generalized mosaicing. In *Proc. IEEE ICCV*, vol. I, 17–24.
- SCHLICK, C. 1994. Quantization techniques for visualization of high dynamic range pictures. In *Photorealistic Rendering Techniques*, Springer-Verlag, P. Shirley, G. Skaas, and S. Müller, Eds., 7–20.
- STOCKHAM, J. T. G. 1972. Image processing in the context of a visual model. In *Proceedings of the IEEE*, vol. 60, 828–842.
- TANAKA, T., AND OHNISHI, N. 1997. Painting-like image emphasis based on human vision systems. *Computer Graphics Forum* 16, 3, 253–260.
- TUMBLIN, J., AND RUSHMEIER, H. E. 1993. Tone reproduction for realistic images. *IEEE Computer Graphics and Applications* 13, 6 (Nov.), 42–48.
- TUMBLIN, J., AND TURK, G. 1999. LCIS: A boundary hierarchy for detail-preserving contrast reduction. In *Proc. ACM SIGGRAPH 99*, A. Rockwood, Ed., 83–90.
- TUMBLIN, J., HODGINS, J. K., AND GUENTER, B. K. 1999. Two methods for display of high contrast images. *ACM Transactions on Graphics* 18, 1 (Jan.), 56–94.
- WARD LARSON, G., RUSHMEIER, H., AND PIATKO, C. 1997. A visibility matching tone reproduction operator for high dynamic range scenes. *IEEE Transactions on Visualization and Computer Graphics* 3, 4, 291–306.
- WARD, G. J. 1994. A contrast-based scalefactor for luminance display. In *Graphics Gems IV*, P. S. Heckbert, Ed. Academic Press Professional, 415–421.

⁷Exposures courtesy of Max Lyons, © 2001 Max Lyons.