

Visualizing High Dynamic Range Images in a Web Browser

Rafal Mantiuk and Wolfgang Heidrich
The University of British Columbia

Abstract. We present a technique for embedding high dynamic range (HDR) images with interactive exposure control in a web browser. Multiple exposures are encoded as a combination of several basis images, which are stacked together and displayed with different opacity values. The combination of such basis images yields the resulting exposure image. An intuitive interface based on JavaScript callback functions provides control over exposure setting. The technique allows us to encode a series of 24 or more exposure images using only 3 basis images, without noticeable artifacts. A toolkit for generating web pages with an HDR image viewer is provided.

1. Introduction

A web search for images using the keyword “high dynamic range” yields over 6 million hits though none of these images contain high dynamic range data. High dynamic range (HDR) images are represented as floating point numbers or encoded using a semi-logarithmic transfer function [Mantiuk et al. 04][Reinhard et al. 05, Ch. 3] so that their pixel values are related to radiance or luminance units. Their dynamic range, which is the ratio between the brightest and the darkest spot in an image, often exceeds the contrast of a display and therefore cannot be faithfully reproduced on a computer screen. For this reason, most of the images tagged as “HDR” on the Internet are conventional images produced by tone-mapping algorithms [Reinhard et al. 05, Ch. 6–8][Myszkowski et al. 08, Ch. 6] that reduce original image contrast to

that of the display.

A windowing technique is an intuitive method for viewing images of higher dynamic range than that of a display without introducing any contrast distortions. In this approach, a selected luminance range (window) of an image’s entire dynamic range is mapped to the range of the display. The user can adjust the window by moving it either towards brighter or darker tones. Windowing is a standard procedure for viewing CT images in medical applications. In photographic applications it is also known as exposure control and is available in recent HDR-enabled photo-editing software, such as Adobe Photoshop CS2.

In this paper we present a method that utilizes the windowing technique in a web browser to display an HDR image. The method minimizes the amount of data that needs to be transferred over the Internet and embeds the windowing algorithm in a web page. All changes in exposure settings are immediately reflected in the displayed image without any communication with a server. The method employs alpha blending and JavaScript for user control, which are natively supported by all major web browsers. Although there are Java Applets for viewing HDR panoramas [Dersch 03] and it is not difficult to develop Adobe Flash components or browser plug-ins with such functionality, these techniques are more cumbersome, less portable and take longer time to start in a web-browser. We could run our viewer on a range of devices, including mobile phones and PDAs, on all major platforms (Linux, Microsoft Windows and Mac OS X) and under a variety of web-browsers (Firefox, Opera, Internet Explorer, Safari) and in all cases the viewer took very short time to start was very responsive. Moreover, our approach does not require devising a custom HDR image compression for transferring HDR images to a browser, as we can transfer HDR images as a series of JPEG files. Figure 1 shows an example of a web page with the high dynamic range image viewer, in which the user can change the displayed dynamic range by scrolling the semi-transparent window at the top of an image.

The primary purpose of the presented method is visualizing HDR images with non-distorted contrast in a web browser. This is especially desirable for browsing web-galleries of light probes used in image based lighting [Debevec 02] and web-based viewers of medical images. Another attractive application is a general purpose viewer for camera RAW and HDR images.

In a naive implementation of web-based windowing, a separate image could be sent to the web browser for each exposure selection. Assuming that the typical granularity of exposure control is $1/3$ of an f-stop unit¹ and that many HDR images span a range of 16 f-stops, 48 separate “exposure images” would

¹f-stop is a unit commonly used in cameras that represents exposure adjustment and is also known as the exposure value (EV). Adding 1 f-stop to the exposure settings doubles and subtracting 1 f-stop halves the amount of captured light. When photometric units (e.g. luminance) are plotted on the \log_2 axis, each logarithmic unit corresponds to 1 f-stop.

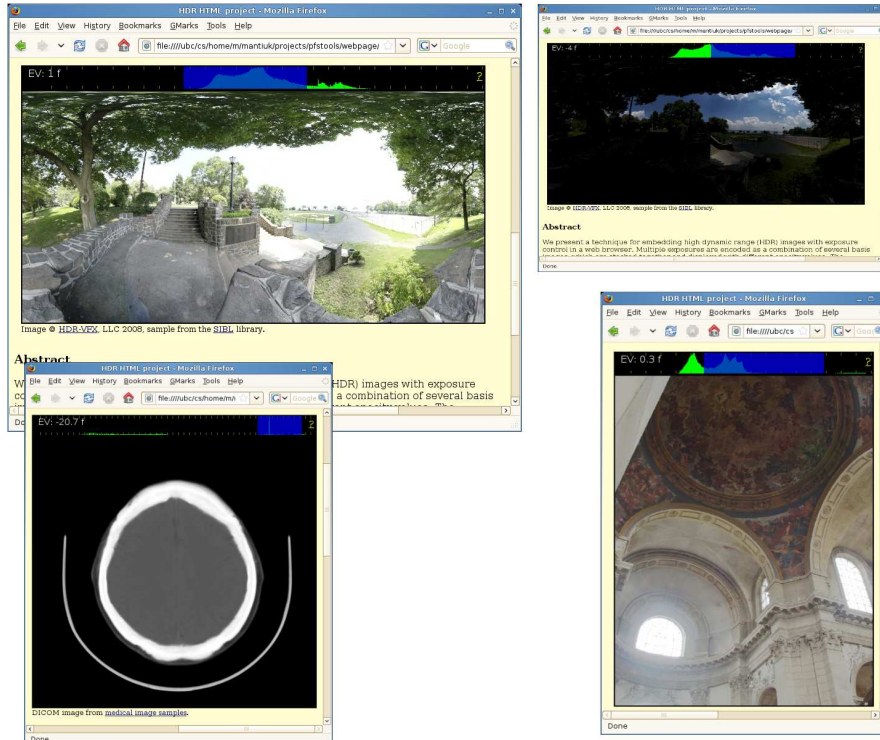


Figure 1. Multi-exposure HDR viewer shown in HTML web pages. The JavaScript code replaces basis images and updates opacity settings for the image stack to provide fine control over exposure settings. Only a few JPEG basis images need to be transferred, and the entire exposure control code is executed in the web browser. Luminance histogram is shown on top of each image. The semi-transparent blue region represents the currently displayed dynamic range window, which can be scrolled with a mouse.

need to be stored on a server and then transferred. Such an approach is clearly impractical. There are also techniques for encoding HDR images as a pair of 8-bit images [Cohen et al. 01, Ward and Simmons 05], but reconstructing an arbitrary exposure from such a pair of images requires image operations that cannot be realized using HTML components. Therefore instead, we employ the *opacity* property of the Cascading Style Sheets 3 (CSS3). Multiple exposures are encoded as a pseudo-linear combination of several basis images, which are stacked together with different opacity factors. This lets us encode 24 exposure images using as few as 3 low dynamic range basis images, thus greatly reducing the amount of data that needs to be sent to a web browser.

2. Windowing technique

Even though the entire dynamic range for most HDR images cannot be reproduced on a computer screen, we can select a window of that range and map it to the available luminance range of a display. To do so, we need to apply an inverse display model that converts linear luminance (or radiance) values L of an HDR image to 3×8-bit pixel values V , which can be shown on a display:

$$V = IDM(L, e) = \left(\left[\frac{L \cdot e - b}{1 - b} \right]_0^1 \right)^{\frac{1}{\gamma}}, \quad (1)$$

where $[\cdot]_0^1$ denotes clamping values in the range from 0 to 1. The variable e represents the exposure or the position the dynamic range window that should be shown. Since all displays have limited contrast due to light leakage and ambient light reflected from a screen, the inverse display model above contains a black level factor b , which elevates the minimum luminance that is produced by the display. Although modern LCD displays can achieve a contrast up to 3000:1, the peak luminance in a typical office setting does not exceed 100 cd/m^2 , and the black level is strongly elevated by reflected ambient light and is usually close to 1 cd/m^2 . This reduces the contrast ratio to 100:1. For convenience, we assume an effective contrast of 128 : 1, which corresponds to 7 f-stops, resulting in the black level factor $b = 1/128$. The majority of computer displays mimic the sRGB standard and have a *gamma correction* value close to $\gamma = 2.2$. Equation 1 is independently applied to three color channels (red, green and blue).




Figure 2. Example of the windowing technique: images on the left and right visualize two exposures of an HDR image and the two curves on each plot represent tone-curves used to generate these images. The plots show how the input luminance factor (HDR pixel values) is mapped to display luminance (top) or pixel values (bottom). Note that the windowing technique does not distort image contrast: the tone curves in the upper plot are straight lines with the slope equal 1.

Figure 2 shows two exposure images rendered from the same HDR image. The exposure correction introduces a horizontal shift on the log-luminance axis. Although linear interpolation between luminance values (top plot) can produce intermediate exposure values, this is not the case for the gamma-corrected pixel values (bottom plot), which are available in a web-browser. Producing intermediate exposure images using a stack of semi-transparent images is not a trivial problem, since the images in such a stack are gamma corrected and limited to the range from 0 to 1. The overexposed pixels (values > 1) in images with the opacity factor less than one will contribute erroneously to the rendered results in the sense of reducing contrast in highlights. Minimizing these distortion is the main problem that we solve below.

3. Alpha blending in the web browser

The CSS3 *opacity* property for HTML elements [W3-CSS3] acts as a multiplier of pixel values, which are then combined with the pixel values of underlying elements or with the background. If we stack N images $B_{1..N}$ on top of each other and set their opacities to $c_{1..N}$, the resulting pixel values can be found with the recursive formula [W3-AC]:

$$\begin{aligned} I_0 &= b_0 \\ I_i &= c_i B_i + (1 - c_i) I_{i-1} \quad \text{for } i = 1..N, \end{aligned} \tag{2}$$

where b_0 is the background color and B_1 is the bottommost image. All c and B values must be within the range 0 to 1 (negative opacity is not allowed).

4. Optimal encoding of multiple exposures

The previous sections explain how to generate a single “exposure image” for a given dynamic range window and how the CSS3 opacity property modifies pixel values in an image stack. In this section we consider the best way to encode a set of exposures with the least number of semi-transparent images. Our task is to find the basis images $B_{i=1..N}$ and opacity coefficients $c_{i,k}$ (one for each basis i and for each exposure k), so that the blended image stack (Equation 2) is as close as possible to the actual exposure image (Equation 1).

Since image blending is spatially independent, we can operate on tone-curves instead of pixel values of B_i , thus greatly reducing the number of optimized variables. For each basis image B_i , we assign a tone-curve that maps input luminance (radiance) L to pixel values of the basis image:

$$T_i(L) = B_i. \tag{3}$$

If we assume that the background color is black (we found no improvement for different backgrounds), a tone-curve for exposure k resulting from alpha-blending becomes:

$$\begin{aligned} H_0(L, k) &= 0 \\ H_i(L, k) &= c_{i,k} T_i(L) + (1 - c_{i,k}) H_{i-1}(L, k) \quad \text{for } i = 1..N. \end{aligned} \quad (4)$$

If we summarize all the assumptions, our problem can be formulated as minimizing the differences between the ideal tone-curves for exposures e_k (Equation 1) and the tone-curves resulting from alpha blending (Equation 4):

$$\begin{aligned} \arg \min_{c_{i,k}, T_i} \sum_{k=1}^E \int_L (IDM(L, e_k) - H_N(L, k))^2 dL \\ \text{subject to } 0 \geq B_i \geq 1 \\ 0 \geq c_{i,k} \geq 1 \quad \text{for } i = 1..N, \quad k = 1..E. \end{aligned} \quad (5)$$

The above objective function integrates squared differences over the range of luminance values, L , and then sums the errors from each exposure $e_{k=1..E}$. The above problem is non-linear due to clipping in the inverse display model (IDM), however the optimum solution can be found efficiently using sequential quadratic programming. To solve for T_i , $c_{i,k}$, we used the algorithm available in the Matlab Optimization Toolbox (Version 4.0, R2008a). We did not observe local minima for any number of basis images that we considered (up to 5).

Although we could find optimal basis images and opacity values for each individual image, this is impractical since generating a single piece of HDR browser code would require solving an optimization problem. Instead, we pre-compute an optimal solution that can be reused for any number of images. This can be achieved under two assumptions: Firstly, all luminance values have equal importance (Equation 5 is not weighted by a luminance histogram to find an image-specific solution). Secondly, the solution is found for a fixed segment of the dynamic range and is reused for the next segment. As a consequence of this approach, a certain number of tone-curves (basis images) can be shared by two neighboring segments, thereby reducing the number of required basis images. For example, if an image spans 16 f-stops of the dynamic range, and we use 2 unique and 1 shared basis for each 8 f-stop segment, we need a total of $2 \times 2 + 1 = 5$ basis images.

In our practical solution we use 8 f-stop segments and assume that one tone curve is shared, so that $T_N(L) = T_1(L \cdot 2^{-8})$. We approximate a tone-curve with a piece-wise linear function consisting of 40 segments. The exposures are selected with the same 1/3 f-stop granularity found in most photographic cameras, so that:

$$e_k = 2^{-(k-1)/3} \quad \text{for } k = 1..E. \quad (6)$$

Overall, we need to optimize for 39 tone-curve nodes $\times (N - 1)$ non-shared basis images $+ (N - 1)$ opacity coefficients $\times (3 \times 8 = 24)$ exposures. The integral over the luminance range is approximated with discrete summation in the logarithmic luminance domain. Although we find an optimal solution for luminance/luma values, we apply the same tone-curves to the red, green and blue color channels.

5. Examples

The tone-curves reconstructed from a combination of optimal basis tone-curves (images) and optimal opacity settings are shown in Figure 3. Each plot represents a solution for different number of basis images (tone-curves). A smaller number of basis images results in fewer JPEG images that need to be transferred to and combined in a web browser, but it also produces a larger reconstruction error. The shape of the basis tone-curves (blue lines in Figure 3) is similar but not identical to the ideal exposure tone curves (green dashed lines in Figure 3). The difference is especially significant for a lower number of bases. Two basis functions are in general not sufficient to correctly reproduce darker tones (see Figure 3 top-left), but three or more basis tone-curves can already produce a set of exposures that is difficult to discern from the exact solution. Although the PSNR values shown in Figure 3 may seem low, the approximation errors are not objectionable since all errors come from brightening, darkening and contrast change operations, for which the human eye is not very sensitive. This is demonstrated in Figure 4, in which exposure images reconstructed using 3 and more basis images are very close to those produced by the inverse display model. The most noticeable difference is the color clipping, which is present in the reference exposure images, but is reduced in the reconstructed images. The reason for that is the smooth round-off of the reconstructed tone-curves for the maximum pixel values (refer to Figure 3).

6. Toolkit

We developed a toolkit for generating HTML pages with a multi-exposure HDR viewer from HDR images. The toolkit uses precomputed optimal basis tone-curves to generate a set of JPEG images used as basis images for reconstructing multiple exposures. Then, it generates HTML and JavaScript code that offers intuitive control over exposure settings, as demonstrated in Figure 1. The JavaScript code updates opacity coefficients and replaces basis images in the image stack embedded in the web page. The resulting HTML can be used as a stand-alone web page or pasted in another web page. The




Figure 3. Tone-curves resulting from an exact exposure rendering (green dashed lines) compared with the reconstructed tone-curves (red solid lines). The reconstructed tone-curves are a combination of the basis tone-curves (blue lines) with different opacity settings. For clarity, only 9 exposures (from 25) are shown on the plots. The dashed-blue line represents the basis tone-curve that is shared with the next 8-fstop segment of the dynamic range (is a shifted copy of the leftmost basis tone-curve).

URL link to the toolkit can be found in the Web Information section below.

7. Discussion

For most images the proposed technique lets us reduce the number of separately coded exposures from 48 ($16 \text{ f-stops} \times 3 \text{ exposures per f-stop}$) to 5 (2




Figure 4. The result of an exact exposure rendering from an HDR image (top row) compared to the reconstruction from basis images (three lower rows). Three or more basis images per 8 f-stops are sufficient to produce multiple exposures without noticeable artifacts.

× 3 basis per each 8 f-stops - 1 shared basis). This allows for fine control over exposure settings in a web browser without the need to transfer prohibitively large amounts of image data.

Although the CSS3 specification has not been finalized yet (as of Novem-

ber 2008), the *opacity* property is available in most web browsers. The exception is Internet Explorer 7, which still requires proprietary attributes to control opacity. We tested our code for several popular web browsers, including Firefox 3.0, Safari 3.2 and Internet Explorer 7.

The method has been demonstrated to work with the CSS3 opacity property but is also applicable to any 2D rendering engine that supports alpha compositing such as the Scalable Vector Graphics (SVG) renderer. The exposure control, which is the focus of this paper, is just one example of multiple operations that could be encoded using a set of basis images. The same technique could be used to encode other tone and color operations on images such as contrast, white-balance, highlight, and flood-light adjustments, found in typical image editing software.

Acknowledgments.

We would like to thank Cheryl Lau for her helpful comments. This work was supported under the Dolby Research Chair in Computer Science.

References

- [Cohen et al. 01] J. Cohen, C. Tchou, T. Hawkins, and P. Debevec. “Real-time high-dynamic range texture mapping.” In *Proc. of 12th Eurographics Workshop on Rendering*, pp. 313–321, 2001.
- [Debevec 02] P. Debevec. “Image-Based Lighting.” *IEEE Computer Graphics and Applications*, pp. 26–34.
- [Dersch 03] H. Dersch. “Interactive Java Viewer for HDR-panoramas.”, 2003. Available online (<http://webuser.fh-furtwangen.de/~dersch/hdr/HDRViewer.pdf>).
- [Mantiuk et al. 04] R. Mantiuk, G. Krawczyk, K. Myszkowski, and H.P. Seidel. “Perception-motivated high dynamic range video encoding.” *ACM Transactions on Graphics (TOG)* 23:3 (2004), 733–741.
- [Myszkowski et al. 08] Karol Myszkowski, Rafał Mantiuk, and Grzegorz Krawczyk. *High Dynamic Range Video*. Synthesis Lectures on Computer Graphics and Animation, 5., Morgan & Claypool, 2008.
- [Reinhard et al. 05] Erik Reinhard, Greg Ward, Sumanta Pattanaik, and Paul Debevec. *High Dynamic Range Imaging. Data Acquisition, Manipulation, and Display*. Morgan Kaufmann, 2005.

Rafał Mantiuk and Wolfgang Heidrich: Visualizing HDR Images in a Web Browser 11

[W3-AC] W3-AC. “Simple alpha compositing.” <http://www.w3.org/TR/SVG11/masking.html#SimpleAlphaBlending>. The same alpha compositing formulas are used in SVG and CSS3 specifications.

[W3-CSS3] W3-CSS3. “Specification of the CSS3 ‘opacity’ property.” <http://www.w3.org/TR/2008/WD-css3-color-20080721/#opacity>.

[Ward and Simmons 05] Greg Ward and Maryann Simmons. “JPEG-HDR: A Backwards-Compatible, High Dynamic Range Extension to JPEG.” In *Proceedings of the 13th Color Imaging Conference*, pp. 283–290, 2005.

Web Information:

The precomputed basis tone-curves and opacity values and the toolkit for generating HTML pages with a multi-exposure HDR viewer can be found at:

<http://pfstools.sourceforge.net/hdrhtml/>

For more information, please contact the first author:

Rafał Mantiuk, University of British Columbia, 201-2366 Main Mall, Vancouver, BC, Canada, V6T 1Z4
(mantiuk@cs.ubc.ca)

Received [DATE]; accepted [DATE].