

Ldr2Hdr: On-the-fly Reverse Tone Mapping of Legacy Video and Photographs

Allan G. Rempel¹ Matthew Trentacoste¹ Helge Seetzen^{1,2} H. David Young¹
Wolfgang Heidrich¹ Lorne Whitehead¹ Greg Ward²
1) The University of British Columbia, 2) Dolby Canada



Figure 1: Output of our on-the-fly reverse tone mapping algorithm. Left: low dynamic range input image. Center left: a visualization of the brightness enhancement function computed by our method. Center right: two virtual exposures of the resulting HDR image with a contrast of 9300 : 1. Right: the same image shown on an HDR display, using a 10% semi-transparent filter to show details in bright regions.

Abstract

New generations of display devices promise to provide significantly improved dynamic range over conventional display technology. In the long run, evolving camera technology and file formats will provide high fidelity content for these display devices. In the near term, however, the vast majority of images and video will only be available in low dynamic range formats.

In this paper we describe a method for boosting the dynamic range of legacy video and photographs for viewing on high dynamic range displays. Our emphasis is on real-time processing of video streams, such as web streams or the signal from a DVD player. We place particular emphasis on robustness of the method, and its ability to deal with a wide range of content without user adjusted parameters or visible artifacts. The method can be implemented on both graphics hardware and on signal processors that are directly integrated in the HDR displays.

CR Categories: I.4.9 [IMAGE PROCESSING AND COMPUTER VISION]: Applications.

Keywords: Image and Video Processing – High Dynamic Range/Tone Mapping; Methods and Applications – Signal Processing.

1 Introduction

It is widely acknowledged that the dynamic range of the traditional imaging pipeline is severely limited compared to the abilities of the human visual system. This observation has led to active research on high dynamic range (HDR) image sensors (e.g. [Acosta-Serafini et al. 2000]), video standards such as the 10-bit log H.264 (AVC), and HDR file formats (e.g. [Ward and Simmons 2005; Mantiuk et al. 2004a; Mantiuk et al. 2006]). HDR displays were first demonstrated by Seetzen and co-workers [2004; 2003]. Over the past few years, major display manufacturers such as LG Philips, Samsung, and AUO have started to work on this technology, and have presented prototype displays (e.g. [Philips 2006]), indicating that HDR display technology is now getting ready for the consumer market.

Despite the increasing availability of HDR content, legacy low dynamic range (LDR) images and video will represent the majority of content in the near term. Therefore, methods must be found for displaying such content on HDR screens. Recent user studies show that perceived image quality increases with display dynamic range [Seetzen et al. 2006; Yoshida et al. 2006], that is, with *simultaneous* increases of *both* brightness and contrast. This suggests that there should be better solutions for presenting LDR content on HDR screens than by emulating LDR display hardware. Instead, we desire a method that expands the dynamic range of legacy images and video to more closely match the capabilities of the new output devices. As observed by Banterle et al. [2006], this task can be thought of as a ‘reverse tone mapping’ operator, which takes an LDR image and estimates the HDR image it represents.

In most cases, it will be impossible for the reverse tone mapping to be an exact mathematical inverse: the HDR information that may have been present in the original scene was lost during capture or encoding, and thus cannot be recovered in full detail. However, it is possible to use heuristics to generate plausible images with better visual quality than the LDR input image. In this sense, reverse tone mapping is similar to other restorative techniques, including colorization of greyscale movies and images and scratch removal.

In this paper, we describe a reverse tone mapping algorithm that is specifically tailored toward on-the-fly processing of video streams without precomputation. We assume that the input LDR content has been optimized for viewing on regular TV screens, which is the case both for footage shot on video, and for film content that has been re-mastered for DVD distribution. While offline reverse tone mapping is an interesting research topic, our work is focused on methods that can be integrated directly into the display hardware to deal with video streams from conventional video cameras or DVD players. To support this goal, our method was designed to have the following properties:

- The algorithms are well suited for implementation on GPUs, as well as signal processors or field programmable gate arrays (FPGAs), which can be located in the display itself. The algorithms are efficient enough to work in real-time on dynamic HDTV resolution video streams.
- No user input is required. All parameters can be chosen up front based on the hardware characteristics of the display.
- The method is robust in the sense that it does not produce disturbing artifacts. The visual quality of the HDR output image is at least as good as that of the input image for a very large range of content.

- In particular, the output video stream is temporally coherent. Colors and intensities do not change abruptly unless they do so in the input image.

2 Related Work

Advances in imaging technology result in continuous improvements of image quality for new content. However, once data has been captured and encoded, its quality is frozen and determined by the state-of-the-art of imaging technology at the time of capture. Often it is undesirable, or even impossible, to later recapture scenes with new hardware or techniques, so the existing data must be enhanced in order to take advantage of improved output devices. Prominent examples are the colorization of black-and-white imagery [Welsh et al. 2002; Levin et al. 2004], removal of scratches and repair of damaged image regions, or the removal of unwanted objects [Hirani and Totsuka 1996; Bornard et al. 2002; Bertalmio et al. 2000; Sun et al. 2005].

The problem of estimating HDR images from LDR photographs has also recently received some attention. Li et al. [Li et al. 2005] discuss the problem of HDR *companding*, i.e. the process of first compressing the contrast through tone-mapping, and later expanding the contrast back to the original range. The expansion step assumes that the intermediate LDR image has been created with the specific contrast reduction algorithm they propose. Meylan et al. [2006; 2007] detect specular highlights in LDR images and boost the intensity in these regions. The stability of this segmentation across frames in video sequences is unclear, and presents a potential source of artifacts. It is also worth noting that current HDR displays built on the principles described by Seetzen et al. [2004] are not very good at reproducing very small bright areas such as isolated highlights.

In work parallel to ours, Akyüz et al. [2007] determined through psychophysical studies that in many circumstances a linear contrast scaling works surprisingly well for mapping LDR content onto HDR screens. Unlike Meylan et al., this approach does not attempt to specifically boost the contrast in saturated regions such as highlights. Our method can be thought of as a combination of a simple linear scaling with an additional brightness enhancement in the proximity of saturated regions, while taking the necessary steps to make both operations robust.

Also partially in parallel to our work, Banterle et al. [2006] found a way to approximately invert Reinhard et al.'s photographic tone mapping operator [2002]. While their approach produces beautiful HDR images, it is not suitable for our application for a variety of reasons. First, it is too slow for real-time purposes. Second, it requires the image-specific adjustment of parameters, which cannot be tolerated in our setting. Perhaps most importantly, Banterle et al.'s algorithm relies on segmentation of the images into regions, which can result in temporal artifacts due to image noise, compression artifacts, and other minute differences between frames.

In contrast, our method is simple enough to be implemented on GPUs, or on signal processors or FPGAs directly built into the display, such as the ones integrated into the BrightSide Technologies DR37-P HDR display [BrightSide 2005]. While our algorithm has a few parameters, these depend only on display resolution and output dynamic range, and can therefore be fixed for each display model. Finally, we pay attention to temporal coherence and robustness of the method for large variety of content.

3 Reverse Tone Mapping

The design of our reverse tone mapping algorithm is inspired by both the image encoding process in LDR cameras, as well as the darkroom techniques of dodging and burning [Adams 1983], that

have also inspired Reinhard et al.'s tone mapping algorithm [2002] and its inverse [Banterle et al. 2006]. Dodging refers to the process of blocking light from certain image areas during the exposure of a print, while illuminating other regions. On photographic paper, which has a negative response, the dodged image regions will be reproduced in brighter tones compared to the unprocessed image regions. The sharpness of the dodging process can be controlled by adjusting the distance between the mask and the print material: smaller distances result in sharper transitions [Adams 1983].

Our method consists of two components, as depicted in Figure 2. First, an inverse gamma stage maps pixel values into linear luminance in order to compensate for the non-linear encoding of intensities in LDR images. We also apply an image filter to suppress noise, as well as quantization and image compression artifacts. The second component is the computation of a *brightness enhancement* function, which is used to increase brightness in image regions containing saturated pixels, i.e. pixels with at least one channel near the maximum value. The brightness enhancement function is mostly smooth, but may contain an *edge-stopping* function that ensures that the brightness enhancement ends at strong edges. In the terminology of darkroom techniques, one can think of the smooth brightness enhancement as a dodging operation with a smooth boundary, while the edge-stopping function introduces a sharp boundary at strong edges between dark and bright regions.

Both components are designed to be temporally coherent and robust under noise, as detailed in the following.

3.1 Inverse Gamma and Contrast Scaling

The first step of our method is to map the non-linearly represented pixel values from the LDR image into a linear luminance range with absolute per-pixel luminance values. We then map the pixel values to absolute output intensities on the HDR display. An optional image filter reduces noise and quantization artifacts.

Linearization of pixel values involves compensating for non-linearities in the LDR representation by applying the inverse of the encoding function. We are primarily interested in applying our method to video footage that has been processed for viewing on normal television sets. Standard video and television formats use a gamma curve of 2.2 [ITU 1990; Stokes et al. 1996; Poynton 2003], which itself is meant to compensate for non-linearities in conventional display technologies such as CRTs.¹ By inverting this gamma curve we obtain pixel values that are approximately proportional to the luminance in the original scene.

Although we focus on the processing of dynamic content like video and television material, we also experimented with digital photographs. Unlike most digital video sources, digital photo cameras often apply additional non-linear transformations on the typically linear raw sensor measurements. The goal of these transformations is to simulate response curves similar to those of analog film [Adams 1983], as well as to compress the contrast of digital cameras (often 1 000 : 1 or more) to that of conventional screens (300 : 1 – 400 : 1). To account for the additional non-linearities in digital photographs, one could analyze the image header, which typically encodes camera type and settings. Using this information, one could choose between a set of precomputed response functions. In our experiments, we found that a simple gamma curve produces credible images without visible artifacts even for photographic material.

Contrast scaling. Once we have linearized the pixel values, we map them to absolute luminance values to be shown on the HDR

¹Newer display technologies such as LCD panels usually simulate a gamma curve comparable to that of CRTs for backwards compatibility.

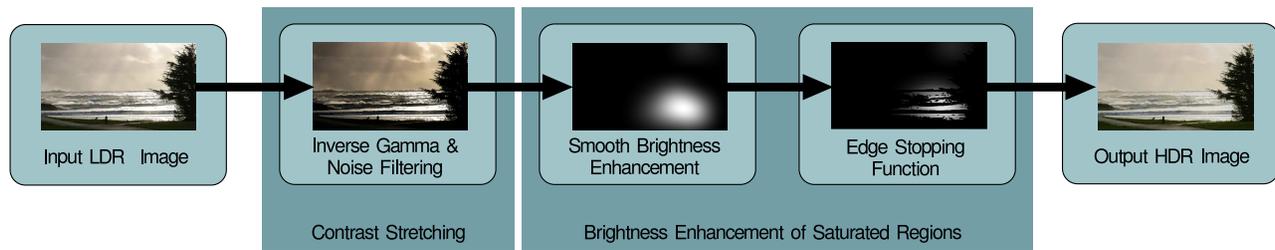


Figure 2: Overview of the method.

display. On conventional LDR screens, the contrast of the imagery is always stretched to the full capabilities of the display device, subject to user preference. As a consequence, the same content can be presented with a contrast of 150 : 1 on one display, and 500 : 1 on another display; the contrast is simply uniformly stretched over the feasible range of the output device.

This observation suggest that we can increase the dynamic range for presentation on HDR screens simply by further linear stretching of the image contrast. A recent perceptual study independently performed by Akyüz et al. [2007] confirms this concept. However, we find that this stretch must not be too large, otherwise images will start looking unnatural. In particular, bright objects may appear to ‘glow’ in certain situations for large increases in contrast.

We ran a large number of tests with different scaling factors, and determined that stretching the contrast up to about 5 000 : 1 yields an improved image quality without introducing artifacts. This threshold is conservative: for many images much larger scale factors produce outstanding results. However, above this threshold, a small percentage of images will start degrading in visual quality. Since we desire a robust and temporally coherent algorithm without the need for user controlled parameters, we select a conservative mapping. On our reference HDR display, we fix the black level of the image at about 0.3 cd/m^2 , which provides a deep black under normal viewing conditions. We then scale the white point to $1\ 200 \text{ cd/m}^2$, significantly brighter than regular displays. These numbers may have to be adjusted for other HDR displays, but they are independent of the images, thereby ensuring temporal coherence without user input.

Noise filtering and quantization reduction. The contrast stretching and non-linear mapping of pixel values amplify quantization artifacts and noise. LDR input images are usually quantized to 256 pixel values, while over 1 000 different values are required to cover the dynamic range of HDR displays at the precision of Just Noticeable Difference (JND) steps [Seetzen et al. 2004]. Lossy video compression can further reduce the number of available intensity levels in a local image region.

In the past, sophisticated algorithms for addressing such artifacts have been developed, for example the method by Daly and Feng [2004]. Due to our real-time constraints, we employ a simple modified spatial Bilateral filter [Tomasi and Manduchi 1998] to alleviate this situation without blurring sharp features. We use Gaussian functions for both the spatial and the photometric term in the Bilateral filter. However, we stretch the variance of the photometric term linearly with the stretch introduced by the non-linear intensity mapping for the local pixel value, such that the photometric variance is always 2 quantization levels. The effect of this approach is similar to performing the Bilateral filter with uniform variance *before* the intensity mapping. However, our ordering makes better use of the available bit depth so that the operations can be performed in fixed point arithmetic. Despite this optimization, the Bilateral filter remains an expensive operation, so that its spatial extent needs to be kept to around 4 pixels, depending on the computing platform.

3.2 Brightness Enhancement

The second component of our method is a brightness enhancement function, which we introduce with the goal of increasing the luminance of the output image in regions where at least one color channel is saturated. In these regions, information has been lost because the scene intensity was outside the capabilities of the camera or recording medium at the time of capture. While this information cannot be restored exactly, we can attempt to approximate the visceral response associated with the higher contrast and overall brightness in the original scene.

Our approach for achieving this goal is to compute a function that can be multiplied with the linearized intensity image to form a high dynamic range image. This has to be done carefully in order to avoid introducing spatial or temporal artifacts. Our brightness enhancement function is primarily smooth, but may contain sharp edges in areas of strong image gradients in the original image. The result is an increase in brightness not just for pixels with saturated color channels, but for a whole neighborhood surrounding such pixels, much like the traditional darkroom technique of dodging with a mask positioned far from a the photographic paper.

Smooth Brightness Enhancement. The smooth portion of the brightness enhancement function is computed by first determining a binary mask of pixels where at least one color channel exceeds a certain threshold intensity. Video formats typically use a white level of 235, meaning that full white of reflective objects corresponds to that pixel value [ITU 1990]. However, video streams also contain larger, ‘super-saturated’ pixel values corresponding to specular highlights or light sources. We found that using a threshold value of 230 works well for separating the saturated from the super-saturated regions in the presence of lossy video compression. For photographs, we found a threshold of 254 to be adequate in the presence of artifacts introduced by lossy compression.

From this mask, we derive a smooth brightness enhancement function by blurring the mask with a large kernel of approximately Gaussian shape. The exact size of the blur kernel depends on the display dimensions and anticipated range of viewing distance. For example, on a 37" HDR display with a resolution of 1920×1080 pixels, we use a blur with standard deviation of 150 pixels, which corresponds to 1.2° at a viewing distance of 3 m. As a result, the spectrum of the blur filter primarily contains low angular frequencies of 0.5 cycles per degree or less, to which the human visual system is not very sensitive [van Nes and Bouman 1967; Pattanaik et al. 1998]. By designing the brightness enhancement function in this way, we make sure that it can be used to increase the intensity in the neighborhood of saturated image regions without introducing visible artifacts.

We apply the smooth brightness enhancement function by linearly mapping its values to a range of $[1 \dots \alpha]$ and then multiplying it onto the result from the inverse gamma stage. The value of the brightness amplification factor α is chosen based on the capabilities of the target HDR display. In our experiments with the Brightside DR37-P, we used a value of 4, corresponding to a peak intensity of

$4 \cdot 1\,200 = 4\,800 \text{ cd/m}^2$. Due to the large blur radius, this peak intensity is only reached for large saturated regions.

Edge-Stopping Function. The smooth brightness enhancement function by itself stretches the global contrast, and yields images that appear more crisp than the stretched contrast images when viewed on an HDR display. However, this function cannot enhance local contrast around sharp edges. To further improve appearance under such conditions, we introduce an edge-stopping function that limits the influence of the brightness enhancement to image regions that are not separated by strong edges from the saturated pixels.

We compute this binary edge stopping function using a flood fill algorithm that uses the initial binary saturation mask as a seed. The flood fill proceeds outwards from these saturated pixels until it reaches pixels with a large gradient magnitude, or the boundary of the area of influence for the smooth brightness enhancement function. We estimate image gradients using divided differences, but for robustness we use a wide baseline of 5 pixels, such that we obtain thick edges that reliably prevent the flood fill algorithm from leaking across the edges.

The final edge stopping function is cleaned up with a morphological *open* operator and, blurred slightly to suppress aliasing, before it is multiplied with the smooth brightness enhancement function. As before, the result is linearly mapped to a range of $[1 \dots \alpha]$, and then multiplied onto the color image. Figure 3 shows an LDR input image, the smooth brightness enhancement function, and the one including the edge stopping term.

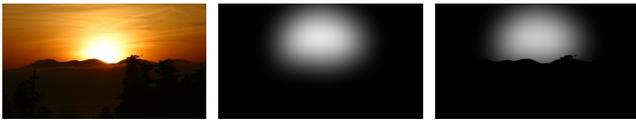


Figure 3: Smooth brightness enhancement function (center) and brightness enhancement with edge stopping function (right) for an input image (left).

Implementation Using Image Pyramids. Although the brightness enhancement algorithm calls for an image blur with a large radius, it can be implemented in a highly efficient manner using image pyramids [Burt and Adelson 1983]. Figure 4 illustrates this hierarchical version of the algorithm. The large Gaussian blur is implemented by successively downsampling the mask representing the saturated pixel regions (1), and then upsampling it again with nearest-neighbor interpolation, while applying a small (3×3 pixel) Gaussian blur at each level (2). To compute the edge stopping function, we first generate the edge image at the highest resolution as described above, and then downsample into an image pyramid (3). The actual edge stopping function is then created through a sequence of upsampling and dilation operations from the lowest resolution binary saturation mask (4). As with the smooth brightness enhancement function, the upsampling for the edge stopping function uses nearest neighbor interpolation. The dilation operators use a 3×3 binary mask (4), but stop at pixels that are marked as an edge in the edge image of the corresponding resolution (3).

Note that the radius of the dilation at each level (4) is the same as that of the blur on the corresponding level of the upsampling pyramid (2), so that the blur and the edge stopping function propagate outwards at the same speed. Note also that the edge stopping function can have hard edges in smooth image regions. However, these edges are outside the area of influence of the blur function, and thus do not create discontinuities in the final image.

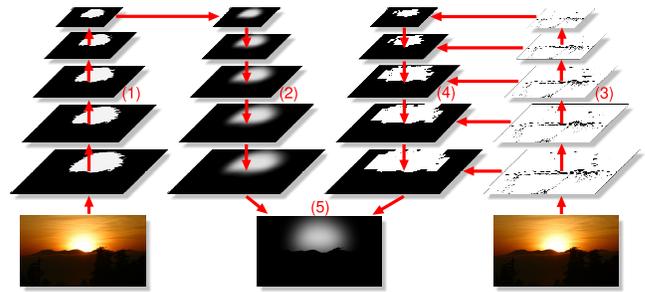


Figure 4: Diagram of the brightness enhancement algorithm using image pyramids (see text for details).

4 Evaluation and Results

The algorithms presented in this paper are based on experiments with large set of sample images, verified by viewing the resulting images on a commercial HDR display, the Brightside DR37-P. Results from early versions of the algorithm have been shown to several thousand people on various occasions, including trade shows.

For a quantitative analysis we applied Mantiuk et al.’s visible difference predictor (VDP) [2004b], an HDR variant of the original VDP work by Daly et al. [Daly 1993]. The VDP by Mantiuk et al. computes the likelihood of a human observer detecting the difference between two images, based on a side-by-side comparison of individual image regions. In contrast to other image metrics, the HDR VDP accounts for HDR-specific effects such as veiling glare, and its performance has been demonstrated with a perceptual study [Mantiuk et al. 2004b].

Since the VDP assumes adaptation to the individual image regions, it is possible to directly compare the linearized input LDR image to the HDR image created by reverse tone mapping. In this fashion, we can validate that the saturation extension does not introduce artifacts. Figure 5 shows an example of the comparison, where image regions have been colored according to their detection probability.



Figure 5: Colored results of the VDP comparison between LDR and HDR image.

Note that the VDP is not suitable for identifying the exact shape of the differences; it simply indicates regions in which problems may occur. As expected, the probability of detection for the smooth saturation extension is low, even in direct side-by-side comparison as measured by the VDP. On the other hand, the edge stopping function does enhance the local contrast, which can be detectable *in side-by-side comparisons*, with detection probabilities of 25% or more for 0.27% of the pixels. Visual inspection shows that these

local contrast enhancements do not negatively affect image quality, but in many cases actually improve it.

The individual stages of the algorithm are designed for robustness and temporal coherence. For example, while we use thresholding of the image to determine a saturation mask, this mask is aggressively low-pass filtered before use as a brightness enhancement function. This approach not only ensures that the spatial frequencies are not disturbing to humans, but it also makes the method extremely stable under temporal fluctuations caused by noise or compression artifacts.

For the edge stopping function, the wide baseline gradient estimation proves robust under noise and thus ensures that the edge information is effective at stopping the flood fill algorithm. The specific edge shape varies with the chosen edge threshold, but both our visual evaluation and the experiments with the VDP indicate that the precise shape is perceptually masked by significant veiling glare. The key characteristic of the edges is temporal coherence, which we ensure with the robust edge thresholding.

Since the brightness enhancement function is monochromatic, we always reproduce the colors from the LDR image. In the future, it may be interesting to develop methods that adapt the color in saturated regions when different color channels saturate in different regions. However, this is a potentially unstable operation, which is why we opted to maintain the colors from the LDR image instead.

In Figure 6 we show a subset of the test images we used in our comparison. Due to limitations of the print medium, we cannot present the generated HDR imagery here directly. Instead, we show split-image representations with two virtual exposures, as well as images tone-mapped with Reinhard et al.'s photographic tone mapping operator [Reinhard et al. 2002]. The split-image representation shows the contrast in the output images, while the tone-mapped representation is another way of verifying that our method does not introduce artifacts, since high spatial frequencies introduced by our algorithm would be preserved by this operator. The HDR versions of these images are included on the conference DVD. The video demonstrates the performance of the method on dynamic content.

5 Conclusions

In this paper we have presented a method for on-the-fly expansion of the dynamic range of legacy, low dynamic range, video content for viewing on HDR displays. The method is robust and temporally coherent, and does not require image-specific parameter adjustment. As such the method is well-suited for integration directly in HDR display hardware, where it can be used to process video streams from legacy sources such as television or DVDs.

In this paper we have focused on the space of real-time solutions to the problem of expanding the dynamic range of LDR imagery. More sophisticated, albeit slower methods would be interesting to explore in the future.

6 Acknowledgments

We would like to thank the anonymous reviewers for their valuable comments. This work was supported by an NSERC Special Opportunities Grant and an NSERC/MITACS Industrial Postgraduate Scholarship held by the first author.

References

ACOSTA-SERAFINI, P., MASAKI, I., AND SODINI, C., 2000. Single-chip imager system with programmable dynamic range. U.S Patent 6,977,685, Feb.

ADAMS, A. 1983. *The Print*. Little, Brown, and Co.

AKYÜZ, A., REINHARD, E., FLEMING, R., RIECKE, B., AND BUELTHOFF, H. 2007. Do HDR displays support LDR content? A psychophysical evaluation. *ACM Trans. Graph. (Proc. ACM Siggraph)* 26, 3, (preceding paper in these proceedings).

BANTERLE, F., LEDDA, P., DEBATTISTA, K., AND CHALMERS, A. 2006. Inverse tone mapping. In *Proc. of GRAPHITE '06*, 349–356.

BERTALMIO, M., SAPIRO, G., CASELLES, V., AND BALLESTER, C. 2000. Image inpainting. In *Proc. of ACM Siggraph 2000*, 417–424.

BORNARD, R., LECAN, E., LABORELLI, L., AND CHENOT, J.-H. 2002. Missing data correction in still images and image sequences. In *Proc of MULTIMEDIA '02*, 355–361.

BRIGHTSIDE, 2005. DR37-P product description. <http://brightsidetech.com/products/dr37p.php>.

BURT, P., AND ADELSON, E. 1983. The laplacian pyramid as a compact image code. *IEEE Transactions on Communication* 31, 4, 532–540.

DALY, S., AND FENG, X. 2004. Decontouring: Prevention and removal of false contour artifacts. In *Proc. of Human Vision and Electronic Imaging IX*, SPIE, vol. 5292, 130–149.

DALY, S. 1993. The visible differences predictor: An algorithm for the assessment of image fidelity. *Digital Images and Human Vision*, 179–206.

HIRANI, A., AND TOTSUKA, T. 1996. Combining frequency and spatial domain information for fast interactive image noise removal. In *Proc. of ACM Siggraph '96*, 269–276.

ITU. 1990. ITU-R BT.709, basic parameter values for the HDTV standard for the studio and for international programme exchange. Standard Recommendation 709, International Telecommunication Union.

LEVIN, A., LISCHINSKI, D., AND WEISS, Y. 2004. Colorization using optimization. *ACM Trans. Graph. (Proc. ACM Siggraph)* 23, 3, 689–694.

LI, Y., SHARAN, L., AND ADELSON, E. 2005. Compressing and companding high dynamic range images with subband architectures. *ACM Trans. Graph. (Proc. ACM Siggraph)* 24, 3, 836–844.

MANTIUK, R., KRAWCZYK, G., MYSZKOWSKI, K., AND SEIDEL, H.-P. 2004. Perception-motivated high dynamic range video encoding. *ACM Trans. Graph. (Proc. of ACM Siggraph)* 23, 3, 733–741.

MANTIUK, R., MYSZKOWSKI, K., AND SEIDEL, H. 2004. Visible difference predictor for high dynamic range images. In *Proc. IEEE International Conference on Systems, Man and Cybernetics*, 2763–2769.

MANTIUK, R., EFREMOV, A., KRAWCZYK, G., MYSZKOWSKI, K., AND SEIDEL, H.-P. 2006. Backward compatible high dynamic range MPEG video compression. *ACM Trans. Graph. (Proc. of ACM Siggraph)* 25, 3, 713–723.

MEYLAN, L., DALY, S., AND SÜSSTRUNK, S. 2006. The reproduction of specular highlights on high dynamic range displays. In *Proc. of the 14th Color Imaging Conference*.

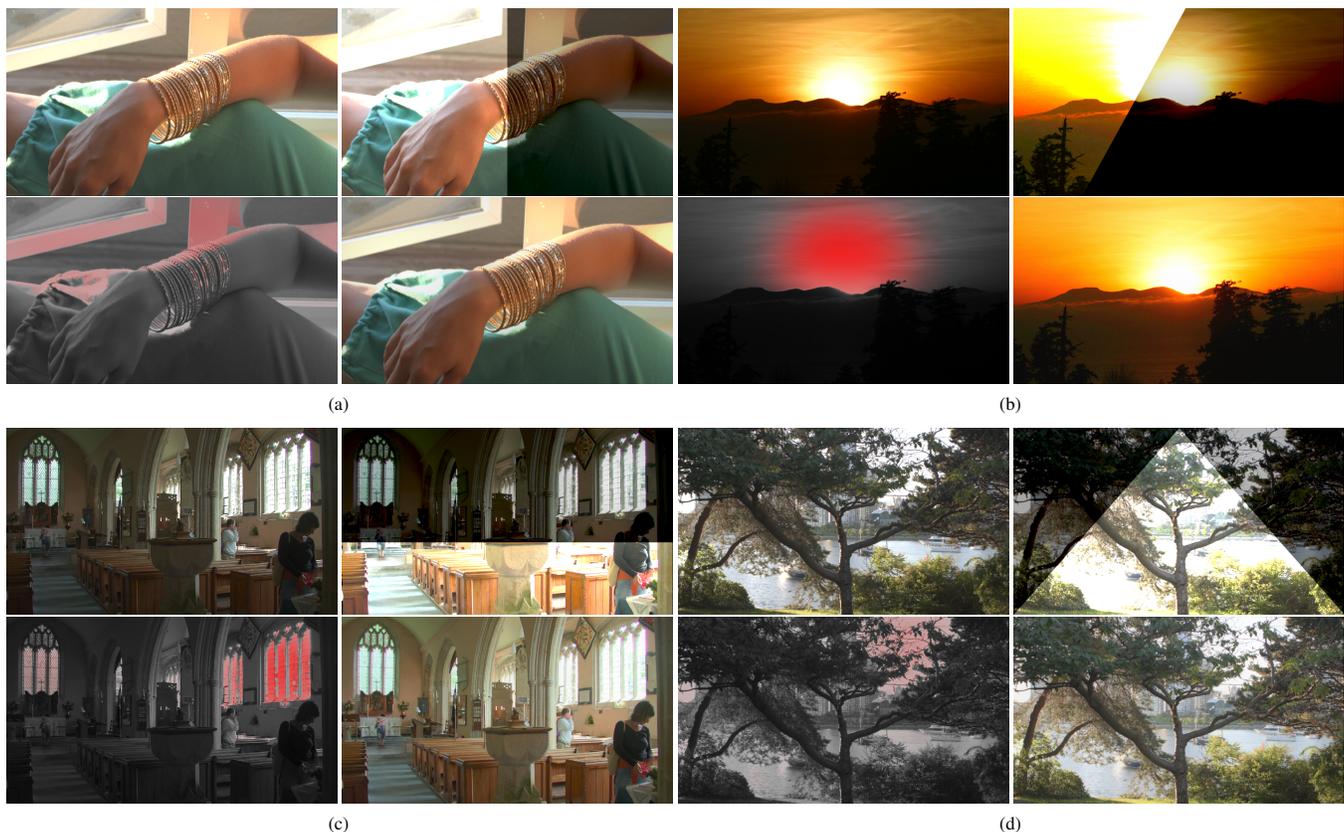


Figure 6: Examples of our algorithm. In each group, we have in clockwise order, from the top left: the original LDR image; a split-screen representation of the HDR output image; a tone-mapped version of the HDR output (using Reinhard et al.'s operator [2002]); and a visualization of the brightness enhancement function overlaid with a grayscale version of the image.

- MEYLAN, L., DALY, S., AND SÜSTRUNK, S. 2007. Tone mapping for high dynamic range displays. In *Proc. of Human Vision and Electronic Imaging XII*, vol. 6492.
- PATTANAİK, S. N., FERWERDA, J. A., FAIRCHILD, M. D., AND GREENBERG, D. 1998. A multiscale model of adaptation and spatial vision for realistic image display. In *Proc. of ACM SIGGRAPH '98*, 287–298.
- PHILIPS, 2006. Philips showcased latest LCD backlighting improvements at FPD 2006. <http://preview.tinyurl.com/yhc4nb>.
- POYNTON, C. 2003. *Digital Video and HDTV: Algorithms and Interfaces*. Morgan Kaufmann.
- REINHARD, E., STARK, M., SHIRLEY, P., AND FERWERDA, J. 2002. Photographic tone reproduction for digital images. *ACM Trans. Graph. (Proc. of ACM Siggraph)* 21, 3, 267–276.
- SEETZEN, H., WHITEHEAD, L., AND WARD, G. 2003. A high dynamic range display using low and high resolution modulators. In *Society for Information Display International Symposium Digest of Technical Papers*, 1450–1453.
- SEETZEN, H., HEIDRICH, W., STUERZLINGER, W., WARD, G., WHITEHEAD, L., TRENTACOSTE, M., GHOSH, A., AND VOROZCOVS, A. 2004. High dynamic range display systems. *ACM Trans. Graph. (Proc. ACM Siggraph)* 23, 3, 760–768.
- SEETZEN, H., LI, H., YE, L., WARD, G., WHITEHEAD, L., AND HEIDRICH, W. 2006. Guidelines for contrast, brightness, and amplitude resolution of displays. In *Society for Information Display (SID) Digest*, 1229–1233.
- STOKES, M., ANDERSON, M., AND MOTTA, R. 1996. Standard default color space for the Internet. Tech. rep., World Wide Web Consortium.
- SUN, J., YUAN, L., JIA, J., AND SHUM, H.-Y. 2005. Video enhancement using per-pixel virtual exposures. *ACM Trans. Graph. (Proc. ACM Siggraph)* 24, 3, 861–868.
- TOMASI, C., AND MANDUCHI, R. 1998. Bilateral filtering for gray and color images. In *Proc. of ICCV '98*, 839.
- VAN NES, F., AND BOUMAN, M. 1967. Spatial modulation transfer in the human eye. *Journal of the Optical Society of America* 57, 401–406.
- WARD, G., AND SIMMONS, M. 2005. JPEG-HDR: A backwards-compatible, high dynamic range extension to JPEG. In *Proc. of Color Imaging Conference (CIC) '05*.
- WELSH, T., ASHIKHMIN, M., AND MUELLER, K. 2002. Transferring color to greyscale images. In *Proc. of ACM Siggraph '02*, 277–280.
- YOSHIDA, A., MANTIUK, R., MYSZKOWSKI, K., AND SEIDEL, H.-P. 2006. Analysis of reproducing real-world appearance on displays of varying dynamic range. In *Proc. of Eurographics 2006*, 415–426.