

Perception-motivated High Dynamic Range Video Encoding

Rafał Mantiuk, Grzegorz Krawczyk, Karol Myszkowski, and Hans-Peter Seidel
MPI Informatik



Figure 1: Gray-scale frames selected from a captured high dynamic range video sequence and perceptually lossless encoded using our technique. Refer to the inset windows and notice the possibility of full visible luminance range exploration in the video.

Abstract

Due to rapid technological progress in high dynamic range (HDR) video capture and display, the efficient storage and transmission of such data is crucial for the completeness of any HDR imaging pipeline. We propose a new approach for inter-frame encoding of HDR video, which is embedded in the well-established MPEG-4 video compression standard. The key component of our technique is luminance quantization that is optimized for the contrast threshold perception in the human visual system. The quantization scheme requires only 10–11 bits to encode 12 orders of magnitude of visible luminance range and does not lead to perceivable contouring artifacts. Besides video encoding, the proposed quantization provides perceptually-optimized luminance sampling for fast implementation of any global tone mapping operator using a lookup table. To improve the quality of synthetic video sequences, we introduce a coding scheme for discrete cosine transform (DCT) blocks with high contrast. We demonstrate the capabilities of HDR video in a player, which enables decoding, tone mapping, and applying post-processing effects in real-time. The tone mapping algorithm as well as its parameters can be changed interactively while the video is playing. We can simulate post-processing effects such as glare, night vision, and motion blur, which appear very realistic due to the usage of HDR data.

CR Categories: I.3.3 [Computer Graphics]: Picture/Image Generation—Display algorithms; I.4.2 [Image Processing and Comp. Vision]: Compression (Coding)—Approximate methods

Keywords: high dynamic range, HDR video, tone mapping, luminance quantization, video compression, MPEG-4, DCT encoding, video processing, visual perception, adaptation

1 Introduction

The range of luminance values in real world scenes often spans many orders of magnitude, which means that capturing those values in a physically meaningful way might require high-dynamic range (HDR) data. Such HDR data is common in surveillance, remote

sensing, space research, and medical applications (e.g., CT scanning). HDR images are also generated in scientific visualization and computer graphics (e.g., as a result of global illumination computation). Many practical applications require handling of HDR data with high efficiency and precision in all stages of the HDR imaging pipeline from acquisition, through storage and transmission, to HDR image display. We briefly discuss all these stages in the context of video, where all frames contain HDR information (see Figure 1). Efficient HDR video encoding and playback are the main focus of this paper.

In recent years significant progress has been made in the development of HDR video sensors such as Lars III (Silicon Vision), Autobrite (SMal Camera Technologies), HDRC (IMS Chips), LM9628 (National), Digital Pixel System (Pixim). Since HDR cameras are still relatively expensive, HDR video is often captured using low dynamic range sensors. A basic principle here is that registered images, which are captured with different exposures, are fused into a single HDR image [Burt and Kolczynski 1993; Debevec and Malik 1997]. This can be done using beam splitters and projecting the resulting image copies to multiple image detectors with preset different exposures [Saito 1995]. In the solutions with a single image detector, sampling in the exposure domain is performed at the expense of either spatial or temporal resolution. For example pixels can be exposed differently by placing a fixed mask [Nayar and Mitsunaga 2000] or an adaptive light modulator [Nayar and Branzoi 2003] with spatially varying transparencies adjacent to the image detector array. In the temporal domain, the exposure can be changed rapidly for subsequent frames, which after their registrations (needed to compensate for camera and object motion) are fused together into HDR frames [Kang et al. 2003].

On the other end of the HDR imaging pipeline the problem of displaying images on devices with limited dynamic range arises. The HDR data compression for accommodating the range limitations is called tone mapping (refer to a recent survey on tone mapping algorithms [Devlin et al. 2002]). Simple tone mapping algorithms, which do not analyze local image content but instead apply the same tone reproduction curve globally for all pixels, can easily be performed in real-time on modern CPUs [Kang et al. 2003; Drago et al. 2003]. Even more advanced algorithms involving different processing, which might depend on local image content, can be executed at video rates using modern graphics hardware [Goodnight et al. 2003]. For sequences with rapid changes in scene intensity, the temporal response of the *human visual system* (HVS) should be modeled. Models of dark and light adaptation [Ferwerda et al. 1996] have already been incorporated into global tone mapping algorithms [Pattanaik et al. 2000; Durand and Dorsey 2000], but similar extensions for local methods remain to be done (at present only

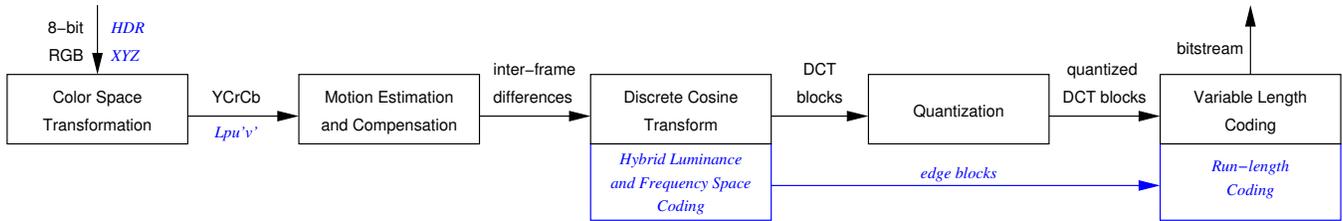


Figure 2: Simplified pipeline for the standard MPEG video encoding (black) and proposed extensions (blue and italic) for encoding high-dynamic range video. Note that *edge blocks* are encoded together with DCT data in the HDR flow.

static images are handled by these methods). It should be noted that the choice of an appropriate tone mapping algorithm and its parameter may depend not only on a particular application, but also on the type of display device (projector, plasma, CRT, LCD) and its characteristics, such as reproduced maximum contrast, luminance range, and gamma correction. Also, the level of surround luminance, which decides upon the HVS adaptation state and effective contrast on the display screen, is important in this context [Ferwerda et al. 1996; CIE 1981]. This means that the visual quality of displayed content can be seriously degraded when already tone mapped images and video are stored/transmitted without any prior knowledge of their actual display conditions. The importance of HDR data has therefore increased significantly, and it will continue to increase as displays covering a luminance range of $0.01\text{--}10,000\text{ cd/m}^2$ become available [Setzen et al. 2004].

An important problem is HDR image encoding, which for storage and transmission efficiency usually relies on the luminance and color gamut quantization. While a number of successful encodings for still HDR images have been developed [Ward Larson 1998; Bogart et al. 2003], no efficient inter-frame encoding of HDR video has been proposed so far. This work is an attempt to fill this gap. We chose the MPEG-4 standard as our framework for the HDR video encoding. This allowed us to exploit all the strengths of this well-established standard, as well as making our implementation more simple. In the future this may also lead to backward compatibility between low- and high-dynamic range contents. A number of MPEG-4 extensions are needed to accommodate HDR data. To obtain visually lossless encodings we introduce a novel HDR luminance quantization scheme in which the quantization errors are kept below the just noticeable threshold values imposed by the HVS. This also requires extending MPEG-4 data structures from 8 to 11 bits. Additionally we introduce an efficient coding scheme for discrete cosine transform (DCT) blocks with high contrasts. Also, we investigate the applicability of standard MPEG-4 weighting matrices (used for the DCT quantization and tuned for typical display luminance ranges) in the context of HDR data compression. We use graphics hardware to support HDR video decoding, tone mapping, and on-the-fly effects, which rely on HDR pixel information such as glare, light and dark adaptation, and motion blur.

The remainder of the paper is organized as follows. In Section 2 we briefly discuss the MPEG-4 standard and we propose extensions needed to accommodate the HDR data. Section 3 provides some implementation details concerning our HDR video encoder and player. In Section 4 we discuss the compression performance results and we describe a client side post-processing of HDR video including real-time rendering of special effects and tone mapping. Since HDR video is not well established, in Section 5 we discuss some of its possible applications in the context of techniques presented in this paper. In Section 6 we conclude this paper and propose some future work directions.

2 Encoding of High-Dynamic Range Video

In this section we introduce a novel algorithm for encoding HDR video. Although the choice of a video compression method often depends on the application, we focus on a general encoding algorithm that is effective for storage/transmission (utilizes inter-frame compression) and at the same time does not introduce perceivable artifacts (is visually lossless). Moreover, we do not consider multi-pass approaches, where the encoding is adaptively tuned for video contents, since those would limit possible applications (e.g. a real-time video capture and encoding is possible only in a single-pass). As a framework for HDR video encoding we selected the MPEG-4 standard, which is state-of-the-art in general video encoding for low dynamic range (LDR) video. Recent studies demonstrate that wavelet transforms extended into the temporal domain and coupled with motion prediction can also be successfully applied for LDR video compression (e.g. [Shen and Delp 1999]), but no wavelet-based standard utilizing inter-frame compression has been established so far.

The scope of required changes to MPEG-4 encoding is surprisingly modest. Figure 2 shows a simplified pipeline of MPEG-4 encoding, together with proposed extensions. A standard MPEG-4 encoder takes as an input three 8-bit RGB color channels whereas our encoder uses HDR XYZ color space since it can represent the full color gamut and the complete range of luminance the eye can adapt to. To improve the compression ratio, the MPEG-4 encoder transforms RGB data to the $Y C_B C_R$ color space. Our encoder stores color information using a perceptually linear $u'v'$ color space, in a similar way as it is realized in the LogLuv image encoding [Ward Larson 1998]. The color space $u'v'$ offers similar compression performance as $Y C_B C_R$ and can represent the complete color gamut [Nadanau 2000]. Furthermore, an 8-bit encoding of $u'v'$ values does not introduce perceivable artifacts as the quantization error is below a Just Noticeable Difference for a skilled observer [Hunt 1995, Section 8.6]. Real-world luminance values are mapped to 11-bit integers using a perceptually conservative function, which we derive in Section 2.1. We choose an 11-bit representation of luminance as it turns out to be both conservative and easy to introduce to the existing MPEG-4 architecture.

The next stage of MPEG-4 encoding involves motion estimation and compensation (refer to Figure 2). Such inter-frame compression results in significant savings in bit-stream size and can be easily adapted to HDR data. After the motion compensation stage, inter-frame differences are transformed to a frequency space by the Discrete Cosine Transform (DCT). The frequency space offers a more compact representation of video and allows perceptual processing.

A perceptually motivated quantization of DCT frequency coefficients is the lossy part of the MPEG-4 encoding and the source of the most significant bit-stream size saving. Although the MPEG-4 standard assumes only the quantization of LDR data of a display

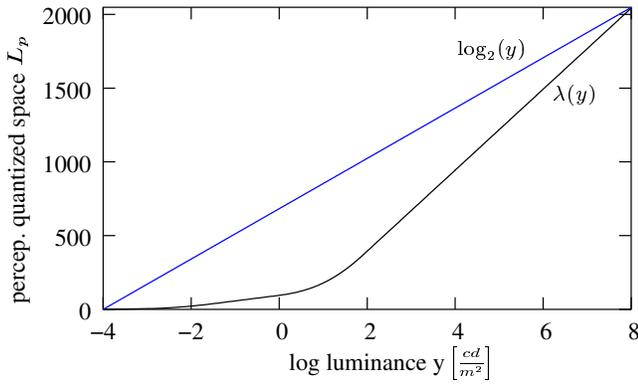


Figure 3: Shape of the luminance-to-integers mapping function λ compared with a logarithmic compression (e.g., LogLuv format). The function λ allocates more values to the mesopic and photopic ranges, where the human eye is the most sensitive to contrast.

device, in Section 2.2 we generalize the quantization method to the full range of visible luminance in HDR video.

Due to quantization of DCT coefficients, noisy artifacts may appear near edges of high-contrast objects. While this problem can be neglected for LDR data, it poses a significant problem for HDR video, especially for synthetic sequences. To alleviate this, in Section 2.3 we propose a hybrid frequency and luminance space encoding, where sharp edges are encoded separately from smoothed DCT data.

In the following sections we describe our extensions to the MPEG-4 format, which are required for efficient HDR video encoding. For detailed information on the MPEG-4 encoding refer to the standard specification [ISO-IEC-14496-2 1999].

2.1 Perceptual Quantization of Luminance Space

The key factor that affects the encoding efficiency of HDR images is the format in which luminance is stored. As visible luminance can span 12–14 orders of magnitude [Hood and Finkelstein 1986], the obvious choice would be encoding luminance using floating point values. Such an approach was taken in the OpenEXR [Bogart et al. 2003] format for still images. Unfortunately floating point values do not compress well, so an improved encoding efficiency can be expected if integer values are used. In this section we propose such a luminance-to-integers mapping, which takes into account the limitations of human perception.

It is well known that the HVS sensitivity depends strongly on the adaptation (background) luminance level. The sensitivity is often measured in psychophysical experiments as the just noticeable luminance threshold Δy that can be distinguished on the uniform background luminance y [Hood and Finkelstein 1986]. The results of such experiments are reported as *contrast versus intensity* $cvi(y) = \Delta y/y$ or *threshold versus intensity* $tvi(y) = \Delta y$ curves [Ferwerda et al. 1996; CIE 1981], which are different representations of the same data. Our goal is to model a luminance-to-integers mapping function, so that the quantization error is always lower than the threshold of perceivable luminance Δy .

To find the best mapping function λ from the luminance space Y to our perceptually quantized luminance space L_p , we start with an inverse mapping $y = \psi(l)$

$$\psi : L_p \rightarrow Y \left[\frac{cd}{m^2} \right] \text{ where } L_p = [0, 2^{nbits} - 1] \quad (1)$$

The maximum quantization error, due to rounding to integer values, for each integer luminance l is given by

$$e(l) = \max\{|\psi(l+0.5) - \psi(l)|, |\psi(l) - \psi(l-0.5)|\} \quad (2)$$

We can approximate the maximum quantization error from the Taylor series expansion by

$$e(l) \approx 0.5 \frac{d\psi(l)}{dl} \quad (3)$$

where $\frac{d\psi}{dl}$ is a derivative of the function ψ . We want to find such a function ψ that the quantization error $e(l)$ is lower than the maximum perceivable luminance threshold $e(l) \leq tvi(y_{adapt})$, where $tvi()$ is a *threshold versus intensity* function. To change this inequality to an equality, we introduce a variable $f \geq 1$

$$e(l) = f^{-1} \cdot tvi(y_{adapt}) \quad (4)$$

For simplicity we assume visual adaptation to a single pixel¹, thus

$$y_{adapt} = y = \psi(l) \quad (5)$$

From equations 3, 4 and 5, we can write a differential equation

$$\frac{d\psi(l)}{dl} = 2 \cdot f^{-1} \cdot tvi(\psi(l)) \quad (6)$$

Boundary conditions for the above equation are given by the minimum and maximum visible luminance²: $\psi(0) = 10^{-4} cd/m^2$ and $\psi(l_{max}) = 10^8 cd/m^2$, where $l_{max} = 2^{nbits} - 1$ is the maximum integer value we encode. Now we can numerically solve this two point boundary problem using for example the *shooting method* [Press et al. 1993]. This way we find both the integer-to-luminance mapping function $y = \psi(l)$, as well as the variable f . The variable f indicates how much lower than the $tvi()$ function the maximum quantization errors are, or how conservative our mapping is. This gives us a trade-off between the number of bits and the quality of the luminance mapping.

We experimented with two t.v.i. curves: The Visibility Reference Function defined in [CIE 1981] and a t.v.i. function introduced to computer graphics by Ferwerda et al. [1996]. Although the function proposed by the CIE standard does not show a discontinuity at the transition between rod and cone vision, to solve equation 6 we used Ferwerda's more conservative t.v.i. function.

As the function ψ is strictly monotonic, a reverse function $l = \lambda(y)$ can be found as well. The reverse function λ for 11-bit luminance coding is plotted in Figure 3. A similar function to λ but with no constraints on the quantization errors, called the capacity function, was derived in the context of tone mapping by Ashikhmin [2002]. A natural representation of $\psi(l)$ is a lookup table as the number

¹There are significant arguments in psychophysical literature for localized adaptation of the human eye [Shapley and Enroth-Cugell 1984]. Although the eye does not adapt to the area of a single pixel, such an assumption is often taken [Daly 1993; Sezan et al. 1987]. In terms of quantization this is a conservative assumption as well because the threshold luminance for the pixel should be the lowest when the eye is adapted to that pixel [Nadenu 2000, Section 2.6.2].

²Some sources suggest the minimum level of adaptation at $10^{-6} cd/m^2$. However it is difficult to achieve such an adaptation even under laboratory conditions [Hood and Finkelstein 1986]. Therefore experimental data is usually missing for such small luminance levels. The *contrast versus intensity* curve in Figure 4 shows a perceivable threshold over 1,000% for the luminance of $10^{-4} cd/m^2$. This means that luminance below that value does not contain any meaningful information that should be encoded in HDR video.

# of bits to encode luminance	f	VDP $P > 0.75$	VDP $P > 0.95$
8 bits	1.67	6%–47.9%	0%–8%
9 bits	3.38	0%–6.3%	0%
10 bits	6.67	0%	0%
11 bits	13.26	0%	0%

Table 1: Precision of the luminance-to-integers mapping for increasing number of bits used to encode luminance. Value $f = 1.67$ means that the maximum quantization error of any mapped luminance y equals $\frac{tvi(y)}{1.67} \left[\frac{cd}{m^2} \right]$ for 8 bits and is twice reduced with every additional bit. Two rightmost columns contain the responses of the Visual Difference Predictor for several quantized frames taken from different animations. The percent values denote relative area of the image for which artifacts will be visible with the probability P greater than 0.75 and 0.95.

of discrete values is usually below several thousand. Binary search can be used to find values of the inverse function $\lambda(y)$. Alternately, an analytic function can be fitted to the data.

Table 1 shows the values of the f variable for a different number of bits used for luminance encoding. Surprisingly, 8-bit encoding seems to guarantee visually lossless quantization of high dynamic range luminance! Such a low number of required bits can be explained by the shape of the c.v.i. curve, which gives the lowest contrast detection of 6% (see Figure 4), while the threshold of 1% is usually assumed in the image processing literature [Sezan et al. 1987]. This comes from the fact that c.v.i. functions have been measured for a particular stimulus (usually a round disk on a uniform background) and thus the thresholds may not be directly applicable to complex stimuli like video. We suggest that the thresholds of the c.v.i. functions should be lowered by a constant value, so that they are below 1% for the photopic conditions. This way the c.v.i. curves still predict loss of sensitivity for low luminance levels and the thresholds are consistent with image processing standards. Such conservative assumptions on the visibility thresholds are met if 10 or more bits are used to encode luminance. The results of the Visible Difference Predictor [Daly 1993](Table 1) further confirm that a 10-bit quantization does not result in visible artifacts.

The problem of perception-based image data quantization that minimizes contouring artifacts has been extensively studied in the literature [Sezan et al. 1987; Lubin and Pica 1991] but mostly for LDR imaging. A simpler mapping function for HDR images than the one derived above is used in the LogLuv format [Ward Larson 1998]. LogLuv uses a logarithmic function to map from luminance values to 15-bit integers. The quantization error of such mapping against a range of visible luminance is shown in Figure 4. For comparison, a function of maximum perceivable luminance contrast at a particular adaptation luminance (c.v.i.) was plotted in the same figure. Ward’s mapping function is well aligned to the c.v.i. curve at high luminance values. However, the logarithmic mapping is too conservative for scotopic and mesopic conditions. As a result, a significant amount of bits is wasted to encode small contrast changes at low luminance, which are not visible to the human observer. We propose a more effective mapping from luminance to discrete values, which is in a better agreement with human perception.

The best quantization accuracy can be expected for those HDR data, which are calibrated in terms of luminance values. Even a rough calibration is sufficient for the 11-bit encoding, which as can be seen in Figure 4 leads to conservative values for the perceivable contrast threshold. For non-calibrated data, it is expected that prior to the encoding step a multiplier value is set by the user for each video sequence to adjust its pixel values to a reasonable luminance range.

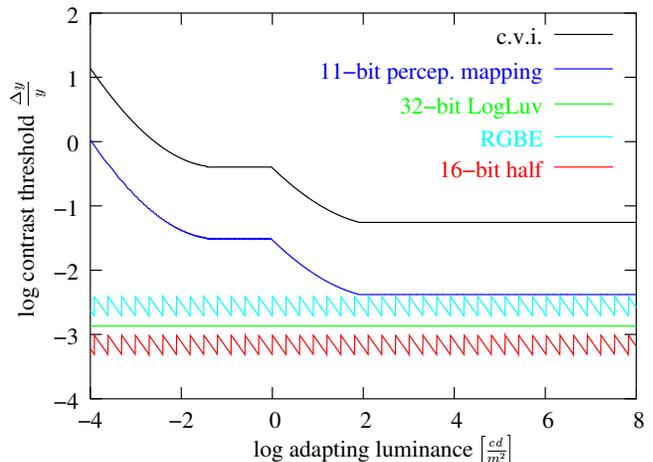


Figure 4: Quantization error of popular luminance encoding formats compared with Ferwerda’s contrast versus intensity function (c.v.i.). *11-bit perceptual mapping* refers to the mapping function λ derived in Section 2.1. *32-bit LogLuv* denotes 15-bit encoding of luminance used in the LogLuv TIFF format [Ward Larson 1998]. *RGBE* refers to the Radiance format [Ward 1991] and *16-bit half* denotes 16-bit float encoding used in OpenEXR [Bogart et al. 2003]. The edgy shape of both *RGBE* and *16-bit half* is caused by rounding the mantissa. Unlike other functions, a curve of the proposed perceptual mapping is aligned to the thresholds of visible contrast (t.v.i.) for the full range of visible luminance. The y-axis can be interpreted both as the lowest perceivable contrast (about 6% for the c.v.i. at 100 cd/m^2) and maximum quantization error.

2.2 Quantization of Frequency Components

In the previous section we derived a perceptual quantization strategy for luminance values. Such a quantization depends on the HVS response to contrast at different illumination levels. However, the loss of information in the human eye is limited not only by the thresholds of luminance contrast but also by the spatial configuration of image patterns (visual masking) [Daly 1993]. To take full advantage of those HVS characteristics, MPEG encoders apply the DCT to each 8×8 pixel block of an image. Then each DCT frequency coefficient is quantized separately with the precision that depends on the spatial frequency it represents. As we are less sensitive to high frequencies [Van Nes and Bouman 1967], larger loss of information for high frequency coefficients is allowed. In this section we show that the MPEG-4 quantization strategy for frequency coefficients can be applied to HDR data.

In MPEG encoders, the quantization of frequency coefficients is determined by a quantization scale q_{scale} and a weighting matrix W . Frequency coefficients F are changed into quantized coefficients \hat{F} using the formula

$$\hat{F}_{ij} = \left\lceil \frac{F_{i,j}}{W_{i,j} \cdot q_{scale}} \right\rceil \text{ where } i, j = 1..8 \quad (7)$$

The brackets denote rounding to the nearest integer and i, j are indices of the DCT frequency band coefficients. The weighting matrix W usually remains unchanged for whole video or a group of frames, and only the coefficient q_{scale} is used to control quality and bit-rate. Note that the above quantization can introduce noise in the signal that is less than half of the denominator $W_{i,j} \cdot q_{scale}$.

Both our HDR perceptually quantized space L_p (Section 2.1) and

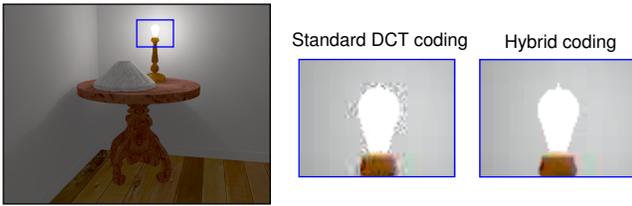


Figure 5: Quality comparison of the standard DCT coding of the block and our hybrid frequency and luminance space coding. Quantized DCT blocks show artifacts at sharp edges, which are not visible for the hybrid encoding. The hybrid encoding increased size of the bit-stream by 7%.

the gamma corrected $Y_C B_C R_C$ space of LDR pixel values are approximately perceptually uniform [Nadenau 2000, Section 7.2.2]. In other words, the same amount of noise results in the same visible artifacts regardless of the background luminance. If quantization adds noise to the signal that is less than half of the denominator of equation 7, quantizing frequency coefficients using the same weighting matrix W in both spaces introduces artifacts, which differ between those spaces by a roughly constant factor. Therefore to achieve the same visibility of noise in the HDR space as in LDR space, the weighting matrix W should be multiplied by a constant value. This can be achieved by setting a proper value of the coefficient q_{scale} .

The default weighting matrices currently used in MPEG-4 for quantization [ISO-IEC-14496-2 1999, Section 6.3.3] are tuned for typical CRT/LCD display conditions and luminance adaptation levels around $30\text{--}100\text{ cd/m}^2$. Contrast sensitivity studies [Van Nes and Bouman 1967] demonstrate that the HVS is the most sensitive in such conditions and the corresponding threshold values essentially remain unchanged across all higher luminance adaptation values. On the other hand, the threshold values significantly increase for the lower luminance adaptation levels. This means that MPEG-4 weighting matrices are conservative for HDR data. More effective and still conservative quantization can be expected if separate weighting matrices are used for lower luminance levels. However, this requires additional storage overhead, as updated matrices have to be encoded within the stream. Moreover, such adaptive quantization requires multi-pass encoding, which restricts possible applications. Another solution is prefiltering of input images to remove imperceptible spatio-temporal frequencies [Border and Guillot 2000]. In this work we do not investigate those approaches and we always use a single weighting matrix.

2.3 Hybrid Frequency / Luminance Space Encoding

In the previous section we showed that the quantization of DCT coefficients can be safely applied to the perceptually quantized HDR space thus greatly reducing the size of the video stream. Unfortunately the DCT is not always an optimal representation for HDR data. HDR images can contain sharp transitions from low to extremely high luminance values, for example at the edges of light sources. Information about sharp edges is encoded into high frequency DCT coefficients, which are coarsely quantized. This results in visible noisy artifacts around edges, as can be seen in Figure 5. This is especially pronounced in the case of synthetic images, which often contain sharp luminance transitions between neighboring pixels. To solve this problem we propose a hybrid encoding, which stores separately low-frequency data in DCT blocks and elevation of sharp edges in “edge blocks”.

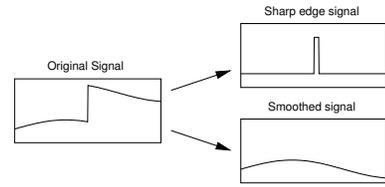


Figure 6: Decomposition of a signal into sharp edge and smoothed signals.

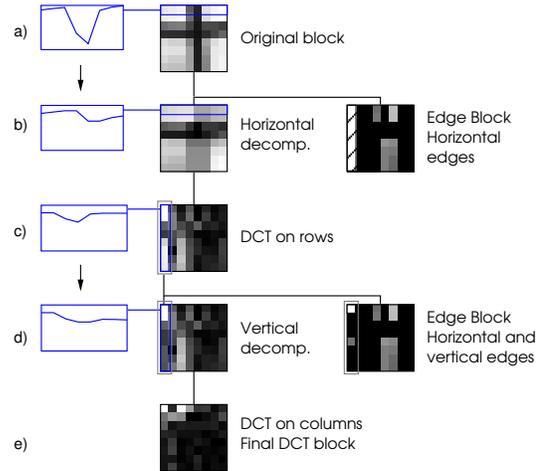


Figure 7: Steps of a hybrid frequency and luminance space coding of a single 8×8 block. Blue insets on the left show a cross-section of the first row (a and b) and the first column (c and d) of the block values. Note how the curves are smoothed as edges are removed from the block, resulting in lower values for the high frequency DCT coefficients. See text for detailed description.

Figure 6 illustrates how, in case of 1D data, input luminance that contains a sharp edge can be split into two signals: One piece-wise constant that contains the sharp edge alone and another that holds slowly changing values. The original signal can be reconstructed from those two signals. Due to the fact that sharp edges occur in sequences relatively infrequently, the signal that stores them can be effectively encoded. The second signal no longer contains large values of high frequency coefficients and can be transformed into a compact DCT representation.

A process of hybrid encoding of a single 8×8 block is shown in Figure 7. The original block (7a) contains a part of a stained glass from the “Memorial Church” HDR image [Debevec and Malik 1997]. To isolate sharp edges from the rows of this block, we use a simple local criterion: If two consecutive pixels in a row differ by more than a certain threshold (discussed in the next paragraph), they are considered to form an edge. In such case the difference between those pixels is subtracted from all pixels in the row, starting from the second pixel of that pair up to the right border of the block. The difference itself is stored in the edge block at the position of the second pixel of that pair. The algorithm is repeated for all 8 rows of the block. This step is shown in Figure 7b. After the rows have been smoothed, they can be transformed to DCT space (Figure 7c). Due to the fact that the smoothed and transformed rows contain large values only for the DC frequency coefficients, only the first column containing those coefficients has to be smoothed in order to eliminate sharp edges along the vertical direction. We process that column in the same way as the rows and place resulting “edges” in the first column of the edge block (Figure 7d). Finally, we can

q_{scale}	1-5	6	7	8	9-31
Threshold inter	n/a	936	794	531	186
Threshold intra	n/a	n/a	919	531	186

Table 2: Threshold contrast values of a sharp edge above which artifacts caused by DCT quantization can be seen. The values can be used to decide whether a sharp edge should be coded in a separate edge block. The thresholds are given for different compression quality factors q_{scale} and for both intra- and inter-encoded frames (since MPEG-4 uses different weighting matrices to quantize intra- and inter-encoded frames). Note that for $q_{scale} \leq 5$ noisy artifacts are not visible and no hybrid encoding is necessary.

apply a vertical DCT (Figure 7e).

Most of the values of the resulting edge blocks are equal to zero and can be compressed using a run-length encoding. However, because this is still more expensive in terms of bit-rate than encoding DCT blocks alone, only the edges that are the source of visible artifacts should be coded separately in edge blocks. The threshold contrast value that an edge must exceed to cause visible artifacts depends on the maximum error of the quantization (refer to Section 2.2) and can be estimated. Table 2 shows such thresholds for MPEG-4 standard quantization matrices and 11-bit encoded luminance in the L_p space (refer to Section 2.1). The thresholds were found for an estimated quantization error greater than 1 Just Noticeable Difference (JND), where 1 JND equals 13.26 units of the L_p space (see Table 1). Note that the lowest threshold equals 186, which corresponds to the local luminance contrast 1:30 for mesopic and 1:5 for photopic range (see Figure 3). Because such high contrast between neighbouring pixels rarely occurs in low dynamic range images, hybrid coding shows visible improvement of quality for high contrast HDR video.

The proposed hybrid block coding improved quality of encoded sequences at the cost of a larger bit-stream (see Figure 5). The artifacts that the hybrid coding can eliminate are mostly visible in synthetic and non-photorealistic images, since those often contain smooth surfaces that do not mask noise. Such artifacts can not be eliminated in post-processing, like blocky artifacts of the DCT. The hybrid coding gives additionally more localized control over the quality than q_{scale} factor. This way, it is possible to remove salient high frequency artifacts while the overall quality is kept the same. Although the hybrid encoding is not strictly necessary to encode HDR video, it solves the problem of encoding high values of frequency coefficients, which would otherwise require extended variable-length coding tables. We noticed that using the standard MPEG-4 variable-length coding of AC coefficients is sufficient for HDR video when the hybrid block coding is used.

3 Implementation

In this section we outline technical details of our implementation of HDR compression and playback.

Our HDR encoder / decoder is based on the XviD library³, which is an open source implementation of the Simple Profile ISO MPEG-4 standard [ISO-IEC-14496-2 1999]. We extended this implementation to support an encoding of DCT coefficients using more than 8-bits per color channel (NOT_8_BIT). This let us encode perceptually quantized luminance (L_p , refer to Section 2.1) represented as 11-bit integers. The two color channels $u'v'$ are sub-sampled to half of the resolution of the original image and encoded with 8-bit precision. The hybrid encoding (refer to Section 2.3) is applied only to

³XviD project home page: <http://www.xvid.org>

the luminance channel. The edge blocks are encoded in the video stream together with DCT blocks. To reduce impact on the stream size, only those edge blocks are encoded that are not empty (less than 7% for our test sequences).

To playback an HDR video we created a player capable of decoding, tone mapping, and applying post-processing effects in real-time. To achieve such performance we had to overcome the bottleneck problem of CPU-to-GPU memory transfer. A naive approach would be transferring HDR frames to the GPU as 16- or 32-bit floating point RGB textures. Instead, we send data in the $L_p u'v'$ format (11-,8-,8-bit, refer to the previous paragraph). The $L_p u'v'$ format gives a gain of 20-40% of a texture size without any visible degradation of quality. Color conversion from the $L_p u'v'$ to RGB format is implemented effectively using fragment shaders and thus lowering CPU load on MPEG decoding.

To tone map frames we employed a simple lookup table approach. Because the number of possible values of the quantized luminance L_p is small (2048 for 11 bits), we tone map only corresponding real-world luminance values and send them to the graphics hardware as a 1D texture. We later use dependent texture lookups to find the values of tone mapped pixels. Tone mapping parameters and computationally expensive variables, like logarithmic mean luminance, are provided within the bit-stream as an annotation script. In this way any global tone mapping operator can be implemented with a marginal effect on performance. On a Pentium IV 2.4GHz processor and an ATI Fire GL X1 graphic card we were able to decode and display about 30 frames per second for a sequence of the resolution 640×480 .

4 Applications and Results

In our experiments with HDR video encoding and playback we used computer graphics animations, panoramic images, and video captured using specialized HDR cameras. The OFFICE sequence is an example of indoor architectural walkthrough rendered using global illumination software with significant changes of illumination levels between rooms (Figure 10). The camera panning was simulated for the CAFETERIA panorama obtained using a Spheron PanoCam camera. The scene contains both a dim cafeteria interior and a window view in a sunny day (Figure 8). To capture natural grayscale sequences we used a Silicon Vision Lars III HDR video camera, which returned linear radiance values. The LIGHT sequence shows a direct view of halogen lamp which illuminates objects with different reflectance characteristics (Figure 9).

As we discussed in Section 2.1, our perceptual quantization strategy for luminance values performs the best for HDR video calibrated in terms of luminance values. Such calibrated data are immediately available for our computer animations resulting from the global illumination computation. We also performed a calibration procedure for the Lars III HDR video camera, using a Kodak GrayCard with 80% reflectance. For the remaining video material we assigned a common sense luminance level for selected scene regions and we then rescaled all pixel intensities accordingly.

4.1 HDR Encoding Performance

To give an overview of the capabilities of the proposed HDR video encoding, we compared its compression ratio with state-of-the-art LDR video compression and existing intra-frame HDR encoding.

Although LDR and HDR video compression store a different amount of information and their performance cannot be matched,

Video Clip	MPEG-4		HDR Enc.		OpenEXR	
	ratio	bpp	ratio	bpp	ratio	bpp
OFFICE hq	0.54	0.27	1.00	0.51	32.17	16.27
OFFICE lq	0.51	0.05	1.00	0.10		
LIGHT hq	0.56	0.71	1.00	1.25	22.56	28.25
LIGHT lq	0.57	0.10	1.00	0.18		
CAFETERIA hq	0.63	0.12	1.00	0.19	142.58	27.40
CAFETERIA lq	0.54	0.05	1.00	0.09		

Table 3: Comparison of compression performance of LDR MPEG-4, the proposed HDR encoding, and the OpenEXR format. "ratio" is a relative bit-stream size increase or decrease compared to our encoding. "bpp" denotes bits per pixel. "hq" and "lq" next to the video clip name means high quality and low quality respectively. There are empty entries for low quality OpenEXR because this format does not support lossy compression. The proposed HDR encoding gives about half of the compression ratio of MPEG-4. High compression gain of MPEG-4 and HDR encoding for the CAFETERIA video clip can be explained by efficient motion compensation in camera panning.

such comparison can give a general notion of the additional overhead required to store HDR data. To compare the performance of LDR and HDR encoding, each test sequence was compressed using our HDR encoder, decompressed, and tone mapped to LDR format. Then the same sequence was tone mapped, encoded to MPEG-4 using the FFmpeg⁴ encoder, and decoded. The quality of the resulting frames from both LDR and HDR encoding was measured using the Universal Quality Index [Wang and Bovik 2002], which gives more reliable quality measure than PSNR. Next, we matched pairs of LDR and HDR streams that had a similar quality index, and compared their sizes. The results are shown in Table 3.

The OpenEXR format offers nearly lossless encoding (up to quantization precision of 16-bit floating point numbers) and intra-frame compression, i.e., each frame is compressed separately. The performance of such compression can be expected to be below that of inter-frame DCT based encoding used in our encoder. However, the OpenEXR format is commonly used for storing animation frames and we decided to include it in the performance summary in Table 3.

4.2 HDR Video Player

In order to benefit from HDR information encoded in our video stream we have developed a video player with extended capabilities. The new functionality includes a convenient dynamic range exploration tool, tone mapping with adaptation to different viewing conditions and display devices, and a client side post processing. We briefly describe each of these extensions.

A dynamic range data exploration tool allows the user to view a selected range of luminance in a rectangular window displayed on top of the video (see Figures 8-10). The user can move the window interactively and choose what part of a dynamic range should be mapped linearly to the display for closer inspection. As the smaller range of luminance is chosen, the tool can reveal quantization artifacts, especially in darker regions of the scene. This is a correct side-effect of our compression because luminance is mapped to a finite number of integer values. Note however that the quantization artifacts are always below the threshold that can be seen in the real world by the human eye.

⁴FFmpeg project home page: <http://ffmpeg.sourceforge.net/>



Figure 8: CAFETERIA sequence, dynamic range $-1.9 \div 3.6[\log cd/m^2]$. The background frame is clamped to a displayable range. Our dynamic range exploration tool, visible as two windows, shows a luminance range $-1.0 \div 1.0[\log cd/m^2]$ (lower right) and a high luminance range $1.0 \div 3.0[\log cd/m^2]$ (upper left). Details in these windows are not visible in LDR video. The source panorama courtesy of Spheron, Inc.

Most of the tone mapping operators have one or more parameters that can be tuned by the user to match her or his taste. As we have source HDR data available, we can give the user freedom to control those parameters on the fly. The user can switch to different tone mapping operators as well. Alternatively, a video stream can be accompanied with an annotation script, which contains the best tone mapping and its parameters for each scene shot. In our video player we implemented the logarithmic mapping [Drago et al. 2003], the global version of photographic tone reproduction [Reinhard et al. 2002], and the perception inspired tone mapping introduced by Pattanaik et al. [2000]. These algorithms are extended with the simulation of the temporal adaptation mechanisms as described in Ferwerda et al. [1996] and Pattanaik et al. [2000]. The result of the latter algorithm with additional post-processing [Thomson et al. 2002] is visible in Figure 10.

LDR movie files are usually tuned for a single kind of display device and viewing condition. Since we have real world luminance values in our HDR video stream, our video player can adjust presentation parameters to any existing or future display device. A tone mapping operator with an inverse display model [Pattanaik et al. 2000] was used in our player for such device dependent tuning.

An HDR video stream with real world luminance values makes it possible to add client-side post-processing, which accurately simulates the limitations of human visual perception or video cameras. The filtering of the LDR stream may not give a convincing result due to the lack of crucial luminance information. In our video player we implemented a night vision post-processing [Thomson et al. 2002] and veiling luminance effect [Spencer et al. 1995; Ward Larson et al. 1997]. The result of the first one is visible in Figure 10. Due to very low level of luminance in the office, the scene is displayed with desaturated colors and a bluish cast, i.e., the way it would be perceived by the human eye. Notice however, that the information on the correct color and luminance range is still available, and can be revealed using the dynamic range exploration tool.



Figure 9: LIGHT sequence captured with the HDR video camera, dynamic range $0.3 \div 4.9 [\log cd/m^2]$. Details of the halogen bulb are well preserved despite high luminances. The visible range in exploration tool window is $2.9 \div 4.9 [\log cd/m^2]$.

5 Discussion

Although the HDR video has not been well established so far, many practical applications would benefit greatly by providing more precise, possibly calibrated streams of temporally coherent data. Our HDR video encoding relies on insensitivities of the HVS in terms of luminance and contrast perception, and therefore it is appropriate for all those applications whose goal is to reproduce the appearance of images as perceived by the human observer in the real world. This assumption matches well to such applications as realistic image synthesis in computer graphics, digital cinematography, documenting reality, tele-medicine, and some aspects of surveillance.

For many applications linear HDR data encoding is possibly required (e.g., dynamic HDR environment maps used for scene re-lighting in computer graphics). Linear or logarithmic HDR video encoding might be desirable in remote sensing, space research, and typical computer vision applications such as monitoring, tracking, recognition, and navigation. For such applications our perception-based luminance quantization algorithm (Section 2.1) is less useful while high-contrast DCT encoding (Section 2.3) might be still applicable.

For other applications, custom quantization algorithms can be required, for example to match sensor characteristics used to acquire HDR data in medical applications. In such a case our approach to quantization (Section 2.1) can be easily adapted.

Note that though the original purpose of our luminance quantization is encoding HDR video, the proposed luminance-to-integer mapping function can be used for static images as well. Also, in the context of global tone mapping algorithms our quantization scheme leads to a small lookup table, which can be applied in any application that requires the perceptual match between the appearance of real world scenes and displayed images. There is no need to perform tone mapping in the continuous luminance space, since luminance values differing less than the quantization error in our luminance encoding cannot be perceived anyway. Then such luminance differences should not be visible in the tone mapped image as well [Ferwerda et al. 1996; Ward Larson et al. 1997].



Figure 10: OFFICE sequence with simulated low-level lightning, dynamic range $-4.0 \div 0.2 [\log cd/m^2]$. The main frame is tone mapped using the Pattanaik et al. [2000] algorithm. Lack of colors and the bluish cast are due to the night vision post-processing as proposed by Thompson et al. [2002]. The exploration window reveals color and details in the $-2.2 \div -1.2 [\log cd/m^2]$ range. The scene model courtesy of VRA, GmbH.

6 Conclusions

In this paper, we have presented a technique for encoding high-dynamic range (HDR) video, which requires only modest extensions of the MPEG-4 compression standard. The central component of our technique is a perception-based HDR luminance-to-integer encoding which requires only 10–11 bits to encode the full perceivable luminance range (12 orders of magnitude) and ensures that the quantization error is always below visibility thresholds. Also, we have proposed an efficient scheme for handling the DCT blocks with high contrast information by decomposing them into two layers of LDR details and HDR edges, which are separately encoded. The size of a HDR video stream encoded by our technique increases less than two times with respect to its LDR version.

The strengths of our video encoding method can be fully exploited for HDR displays, but our method can be beneficial for LDR displays as well. HDR information makes it possible to adjust tone mapping parameters for any display device and surround lighting conditions, which improves the video reproduction quality. Also, using our luminance quantization, the overhead for arbitrary global tone mapping is negligible and amounts to the cost of a small lookup table computation. We demonstrated that by playing back HDR video on graphics hardware the bandwidth of uploaded frames can be significantly reduced and many realistic effects relying on HDR pixels such as glare and motion blur can be properly simulated on the fly.

As future work we plan to investigate the use of less conservative weighting matrices (refer to Section 2.2), which should lead to a better compression for videos with scotopic levels of lighting. Also, it would be interesting to apply our luminance quantization for local tone mapping in order to faithfully reproduce the local contrast perception of human observers.

7 Acknowledgments

We would like to thank Paul Debevec and Spheron, Inc. for making their HDR images available and Jozef Zajac for modeling and rendering test sequences. Special thanks go to Volker Blanz, Scott Daly, Michael Goesele, and Jeffrey Schonier for their helpful comments concerning this work. We are grateful to Christian Fuchs for his help with the HDR camera.

References

- ASHIKHMIN, M. 2002. A tone mapping algorithm for high contrast images. In *Proc. of the 13th Eurographics workshop on Rendering*, 145–156.
- BOGART, R., KAINZ, F., AND HESS, D. 2003. OpenEXR image file format. In *ACM SIGGRAPH 2003, Sketches & Applications*.
- BORDER, P., AND GUILLOTTEL, P. 2000. Perceptually adapted MPEG video encoding. In *IS&T/SPIE Conf. on Hum. Vis. and Electronic Imaging V*, Proc. of SPIE, volume 3959, 168–175.
- BURT, P., AND KOLCZYNSKI, R. 1993. Enhanced image capture through fusion. In *Proc. of International Conference on Computer Vision (ICCV)*, 173–182.
- CIE. 1981. *An Analytical Model for Describing the Influence of Lighting Parameters Upon Visual Performance*, vol. 1. Technical Foundations, CIE 19/2.1. International Organization for Standardization.
- DALY, S. 1993. The Visible Differences Predictor: An algorithm for the assessment of image fidelity. In *Digital Image and Human Vision*, Cambridge, MA: MIT Press, A. Watson, Ed., 179–206.
- DEBEVEC, P., AND MALIK, J. 1997. Recovering high dynamic range radiance maps from photographs. In *Proceedings of SIGGRAPH 97*, Computer Graphics Proceedings, Annual Conference Series, 369–378.
- DEVLIN, K., CHALMERS, A., WILKIE, A., AND PURGATHOFER, W. 2002. Tone Reproduction and Physically Based Spectral Rendering. In *Eurographics 2002: State of the Art Reports*, Eurographics, 101–123.
- DRAGO, F., MYSZKOWSKI, K., ANNEN, T., AND CHIBA, N. 2003. Adaptive logarithmic mapping for displaying high contrast scenes. *Computer Graphics Forum, proceedings of Eurographics 2003* 22, 3, 419–426.
- DURAND, F., AND DORSEY, J. 2000. Interactive tone mapping. In *Rendering Techniques 2000: 11th Eurographics Workshop on Rendering*, 219–230.
- FERWERDA, J., PATTANAIK, S., SHIRLEY, P., AND GREENBERG, D. 1996. A model of visual adaptation for realistic image synthesis. In *Proceedings of SIGGRAPH 96*, Computer Graphics Proceedings, Annual Conference Series, 249–258.
- GOODNIGHT, N., WANG, R., WOOLLEY, C., AND HUMPHREYS, G. 2003. Interactive time-dependent tone mapping using programmable graphics hardware. In *Rendering Techniques 2003: 14th Eurographics Symposium on Rendering*, 26–37.
- HOOD, D., AND FINKELSTEIN, M. 1986. Sensitivity to light. In *Handbook of Perception and Human Performance: 1. Sensory Processes and Perception*, Wiley, New York, K. Boff, L. Kaufman, and J. Thomas, Eds., vol. 1.
- HUNT, R. 1995. *The Reproduction of Colour in Photography, Printing and Television: 5th Edition*. Fountain Press.
- ISO-IEC-14496-2. 1999. *Information technology: Coding of audio-visual objects, Part 2: Visual*. International Organization for Standardization, Geneva, Switzerland.
- KANG, S., UYTENDAELE, M., WINDER, S., AND SZELISKI, R. 2003. High dynamic range video. *ACM Transactions on Graphics* 22, 3, 319–325.
- LUBIN, J., AND PICA, A. 1991. A non-uniform quantizer matched to the human visual performance. *Society of Information Display Int. Symposium Technical Digest of Papers*, 22, 619–622.
- NADENAU, M. 2000. *Integration of Human color vision Models into High Quality Image Compression*. PhD thesis, École Polytechnique Fédérale Lausanne.
- NAYAR, S., AND BRANZOI, V. 2003. Adaptive dynamic range imaging: Optical control of pixel exposures over space and time. In *Proc. of IEEE International Conference on Computer Vision (ICCV 2003)*, 1168–1175.
- NAYAR, S., AND MITSUNAGA, T. 2000. High dynamic range imaging: Spatially varying pixel exposures. In *Proc. of IEEE Conf. on Computer Vision and Pattern Recognition*, 472–479.
- PATTANAIK, S., TUMBLIN, J., YEE, H., AND GREENBERG, D. 2000. Time-dependent visual adaptation for realistic image display. In *Proceedings of ACM SIGGRAPH 2000*, Computer Graphics Proceedings, Annual Conference Series, 47–54.
- PRESS, W., TEUKOLSKY, S., VETTERLING, W., AND FLANNERY, B. 1993. *Numerical Recipes in C*. Cambridge Univ. Press.
- REINHARD, E., STARK, M., SHIRLEY, P., AND FERWERDA, J. 2002. Photographic tone reproduction for digital images. *ACM Transactions on Graphics* 21, 3, 267–276.
- SAITO, K. 1995. Electronic image pickup device. *Japanese Patent 07-254965*.
- SEETZEN, H., HEIDRICH, W., STUERZLINGER, W., WARD, G., WHITEHEAD, L., TRENTACOSTE, M., GHOSH, A., AND VOROZCOVS, A. 2004. High dynamic range display systems. *ACM Transactions on Graphics* 23, 3.
- SEZAN, M., YIP, K., AND DALY, S. 1987. Uniform perceptual quantization: Applications to digital radiography. *IEEE Transactions on Systems, Man, and Cybernetics* 17, 4, 622–634.
- SHAPLEY, R., AND ENROTH-CUGELL, C. 1984. Visual adaptation and retinal gain controls. In *Progress in Retinal Research*, Oxford: Pergamon Press, vol. 3, 263–346.
- SHEN, K., AND DELP, E. 1999. Wavelet based rate scalable video compression. *IEEE Transactions on Circuits and Systems for Video Technology* 9, 1, 109–122.
- SPENCER, G., SHIRLEY, P., ZIMMERMAN, K., AND GREENBERG, D. 1995. Physically-based glare effects for digital images. In *Proceedings of ACM SIGGRAPH 95*, 325–334.
- THOMPSON, W. B., SHIRLEY, P., AND FERWERDA, J. A. 2002. A spatial post-processing algorithm for images of night scenes. *Journal of Graphics Tools* 7, 1, 1–12.
- VAN NES, F., AND BOUMAN, M. 1967. Spatial modulation transfer in the human eye. *Journal of the Optical Society of America* 57, 401–406.
- WANG, Z., AND BOVIK, A. 2002. A universal image quality index. *IEEE Signal Processing Letters* 9, 3, 81–84.
- 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 LARSON, G. 1998. Logluv encoding for full-gamut, high-dynamic range images. *Journal of Graphics Tools* 3, 1, 815–830.
- WARD, G. 1991. Real pixels. In *Graphics Gems II*, J. Arvo, Ed. Academic Press, 80–83.