Tone Mapping High Dynamic Range Videos using Wavelets

Pixar Technical Memo 12-01

Qi Shan[1,2] Mark Meyer[1] Tony DeRose[1] John Anderson[1]
[1]Pixar Animation Studios [2]University of Washington

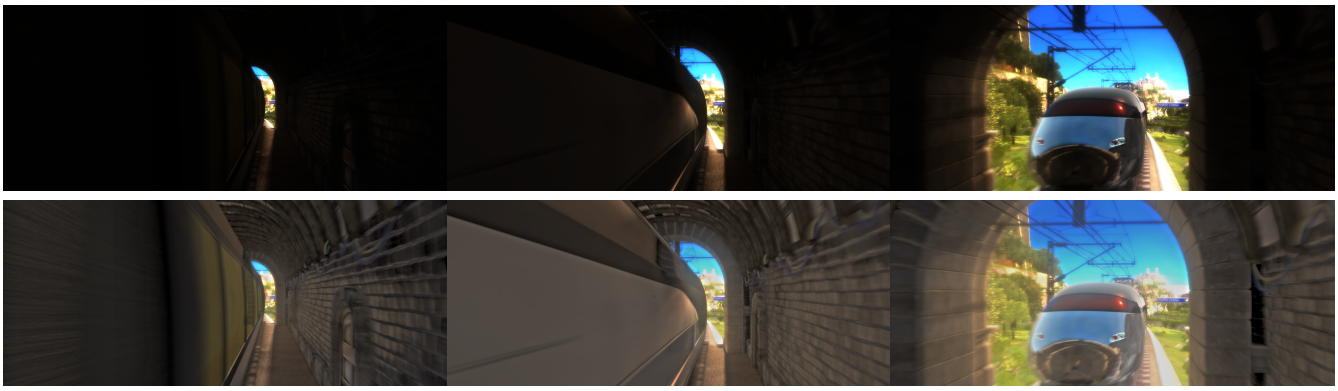

Figure 1: *A short HDR video sequence of a train traveling through a tunnel. The brightness difference between the tunnel and skylight causes dramatic dynamic range variation in time. Top row: three frames from an input high dynamic range video sequence, displayed by linearly scaling all colors to* $[0, 1]$*. Bottom row: corresponding frames from the tone mapped version using our system; notice the textures in dark regions are revealed and the temporal dynamic range variations are well preserved. All video and images* © *Disney/Pixar unless otherwise noted.*

Abstract

We propose a novel 3D wavelet based tone mapping framework for high dynamic range videos. Still image tone mapping methods can be applied to videos in a frame by frame fashion, but they often exhibit haloing artifacts and do not guarantee temporal coherence, resulting in flickering. Directly extending wavelet analysis/synthesis to time solves the flickering problem but leads to ghosting. Our method, based on a new variant of edge avoiding wavelets in time, addresses both flickering and ghosting. We also reduce haloing artifacts through the use of a new gain control method. Our method produces significantly higher quality results compared to existing state-of-the-art techniques when applied to both live action and computer graphics video.

Keywords: High Dynamic Range, Video, Tone Mapping

1 Introduction

High dynamic range (HDR) techniques allow perceptually more accurate depictions of the real world given an image whose range of intensities greatly exceeds the capabilities of a display device. Recent studies show that HDR imagery significantly improves user experience by increasing both brightness and contrast [Yoshida et al. 2006; Rempel et al. 2007; Akyuz et al. 2007; Nakamae et al. 1990].

Displaying HDR still images on low dynamic range (LDR) monitors has been extensively explored in the past decade [Fattal et al. 2002; Durand and Dorsey 2002; Reinhard et al. 2002; Li et al. 2005; Qiu et al. 2006; Mantiuk et al. 2008; Fattal 2009; Kirk and O'Brien 2011]. These techniques are generally referred to as tone mapping operators, and can be further classified into global and local methods. A global operator maps colors using a single spatially constant tone curve or color lookup table [Qiu et al. 2006; Mantiuk et al. 2008], while a local operator has a spatially varying tone curve that depends on local pixel values [Fattal et al. 2002; Reinhard et al. 2002]. It is widely accepted that global operators are more computationally efficient and have fewer artifacts, while local operators in general are more effective at compressing dynamic range while maintaining local contrast.

Tone mapping of video sequences, especially local methods, has to date received less attention [Pattanaik et al. 2000; Ramsey et al. 2004; Lee and Kim 2007; Mantiuk et al. 2008]. As sources of HDR video become more widespread, video tone mapping will become increasingly important. These new sources include HDR video cameras for capturing live action, as well as computer animation and video games that are moving to HDR rendering pipelines to create more realistic worlds. Physical based lighting in particular leads to HDR imagery, often stored in OpenEXR format, where each color channel is represented as 32-bit floating point numbers. To display HDR video on an LDR monitor, manually tuned global tone response curves are often applied. Examples of feature films using global tone mapping include "Up", "Toy Story 3", "Cars 2", "Monster House", and "Alice in Wonderland".

Temporal coherence for global operators is guaranteed if the tone curve is constant in time as well as space, whereas local methods present significant challenges. Applying a local operator independently to each frame will, in general, create flickering.

In this paper we present a novel local video tone mapping operator that was inspired by the still image methods of [Li et al. 2005] and [Fattal 2009]. The method is temporally coherent, preserves local detail better than global methods, and in practice runs in linear time. We demonstrate the effectiveness of the method on both live action and CGI video sequences.

Contributions of our work include:

- Reduced haloing compared to [Li et al. 2005] due to a new sub-band gain control method.
- Temporal coherence due to a new edge avoiding wavelet inspired band decomposition method.

2 Related Work

In this section, we briefly review related work on still image and video tone mapping.

2.1 Tone mapping for still images

Tone mapping of still HDR images has been an active area of research for well over a decade [Tumblin and Rushmeier 1993; Ferwerda et al. 1996]. The purpose of tone mapping depends on the particular application. Early techniques were mainly designed to simulate human visual perception by compressing the dynamic range of pixel intensities while maintaining local contrast, in order to reveal details in both dark and bright regions [Reinhard et al. 2002; Durand and Dorsey 2002; Fattal et al. 2002; Li et al. 2005]. More recent work focuses on changing the tonal values of an image to achieve different artistic looks [Lischinski et al. 2006; Kirk and O'Brien 2011]. Interested readers may refer to [Reinhard et al. 2005] for a comprehensive survey.

Tone mapping operators are generally categorized into global and local operators. Global operators apply a spatially uniform tone reproduction curve to the entire image. For example, gamma curves ($f(x) = x^{\alpha}$, where usually $\alpha = 1/2.2$) and sigmoid curves ($f(x) = \frac{x}{1+x}$). Histogram equalization or matching a reference histogram leads to more content adaptive curves. Advanced global techniques estimate optimal tone curves according to image statistic properties [Qiu et al. 2006; Cvetkovic et al. 2008; Mantiuk et al. 2008]. Once a particular tone curve is computed, global operators map each pixel using this curve without considering its neighboring information.

Local operators map colors by taking its surrounding pixels into consideration. It is generally believed that local operators are more effective in perceptually matching a real scene observed by human perception. Scale separation is usually performed by using bilateral filtering [Durand and Dorsey 2002], wavelets [Li et al. 2005], or edge avoiding wavelets [Fattal 2009][Paris et al. 2011]. Optimization based algorithms work in a way that each pixel proposes an intended adjustment, usually pixels in dark areas propose to be brighter and pixels in saturated areas propose to be dimmer. All the proposed adjustments are then optimized as a whole to obtain a natural looking result [Fattal et al. 2002][Shan et al. 2010]. Local operators have a reputation of introducing visual artifacts, such as haloing and over-sharpening.

2.2 Video tone mapping

Tone mapping techniques for video has attracted much less attention. Early work by Pattanaik et al. [2000] observes the fact that visual adaptation occurs within the retina, which is a chemical reaction and has a delayed effect over time, and proposes to simulate this effect with a frame by frame key value. Several techniques extend the still image tone mapping of [Reinhard et al. 2002] to video by smoothing the log mean luminance in temporal domain [Kang et al. 2003; Ramsey et al. 2004; Krawczyk et al. 2005]. These methods usually assume an input sequence has smooth dynamic range change over time - rapid range change can cause flickering. Lee and Kim [2007] extends the gradient domain tone mapping work of [Fattal et al. 2002] by adding temporal smoothness constraints from optical flow computations. Mantiuk et al. [2008] propose to filter per-frame tone curves to prevent flickering artifacts. Wang et al. [2010] propose an interactive video tone mapping technique which combines Gaussian Mixture model based classification and smoothing filtering of exposure map. The result quality is sensitive to the classification. While not strictly speaking a tone mapping paper, Tocci et al. [2011] describe a hardware system for capturing high dynamic range video sequences.

3 System Overview

Our system uses a wavelet sub-band framework inspired by [Li et al. 2005]. This framework can be divided into three stages:

Analysis: The input image is scale-separated into different bands $B_i(x, y)$, where each band contains different frequency information about the image (level 1 is the coarsest level, level N is the finest). We use a redundant endpoint interpolating linear B-spline wavelet formulation, often referred to as stationary wavelets, where the images are not downsampled when going to coarser levels. Our low-pass basis function is $[0.5, 1, 0.5]$ and the high-pass basis function is $\frac{1}{3}[1, -6, 10, -6, 1]$ as described on page 211-212 of [Stollnitz et al. 1996].

Gain Control: Adapts each level of the wavelet pyramid to achieve range compression:

$$B_i'(x, y) = G_i(x, y)\ B_i(x, y)$$

where $G_i(x, y)$ is the *Gain Map*.

Synthesis: The output image is reconstructed from the modified bands.

The advantage of a wavelet based technique is two-fold. First, it does not involve solving a linear system and is computationally efficient. Formally the complexity is $O(nN)$, where n is the total number of pixels and N is the number of bands. In theory $N = log(n)$, but in practice we only compute a constant number (9 or fewer) of bands. Thus in practice the complexity is linear in both computational time and space. This is an important feature as we are dealing with hundreds of frames of 1920×1080 film resolution. Second, there is a large amount of research work on the mathematics of wavelets, which provides a rich set of useful tools and analysis.

The following subsections describe our formulation for **Gain Control** and **Temporal Coherence**.

3.1 Gain Control

Gain Control is used in the second stage to compress different scales in the pyramid to achieve range reduction. [Li et al. 2005] defines the *Gain Map*, $G_i(x, y)$, as follows:

$$G_i(x, y) = m_i \left(\frac{A_{ag}(x, y) + \epsilon}{\delta} \right)^{(\gamma - 1)} \quad (1)$$

$$A_i(x, y) = GaussianBlur(|B_i(x, y)|) \quad (2)$$

$$A_{ag}(x, y) = \sum A_i(x, y) \quad (3)$$

where m_i is a level dependent multiplier (we use $m_i = [1.0, 0.7, 0.4, 0.4, 0.4, \cdots]$ from the finest level to the coarsest level), and ϵ is a small constant (such as 10^{-6}) to prevent numerical precision issues. γ and δ are two user specified control parameters - constant for all levels.

Unfortunately, using this gain function can result in halo artifacts as shown in Figure 2 (a) around the army men. We observe the

(a) [Li et al. 2005] (b) Our Method

Figure 2: *Still image tone mapping with the proposed gain control. (Left) The gain control of [Li et al. 2005] contains halo artifacts around the army men while (Right) our method has significantly reduced the halos.*

fact that if we change the gain function and make it more constant-like for coarser signals, the haloing can be largely suppressed. The linearity of the gain function in Eq. (1) is mainly controlled by γ. Specifically, as $\gamma \to 1$, $(\gamma - 1) \to 0$, which makes the gain curve approach a constant function. We propose to use the following equation for our *Gain Map*:

$$G'_i(x,y) = \left(\frac{A(x,y) + \epsilon}{\delta'_i} \right)^{(\gamma'_i - 1)}, \tag{4}$$

where

$$\begin{aligned} \gamma'_i &= \min(\gamma + 0.05(N-i), 0.9), \\ \delta'_i &= \left(\frac{\sum_{x,y} A(x,y) \left(\frac{A(x,y)+\epsilon}{\delta} \right)^{\gamma - 1}}{\sum_{x,y} A(x,y) \left(A(x,y) + \epsilon \right)^{\gamma'_i - 1}} \right)^{\frac{1}{1-r'_i}}, \end{aligned} \tag{5}$$

The idea is to increase γ' by 0.05 in each coarser level, and adjust δ' accordingly to make the energy match by satisfying $\sum_{x,y} A(x,y) G_{ag}(x,y) = \sum_{x,y} A(x,y) G'_i(x,y)$.

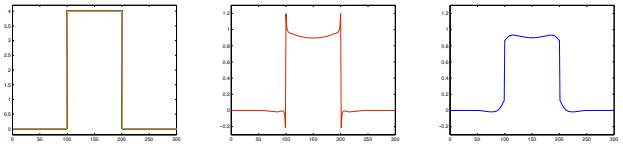

(a) Input (b) [Li et al. 2005] (c) Our Method

Figure 3: *Gain control. (a) Input HDR signal. (b) Tone mapped signal using [Li et al. 2005]. (c) Tone mapped signal using our gain control.*

Fig. 3 shows how a one dimensional signal in range $[0, 4]$ gets tone mapped using the gain control of [Li et al. 2005] and our method. In both cases, we use linear B-Spline wavelets as described in Sec. 3. Note that our result (Figure 3(c)) has less severe halos than the result of using [Li et al. 2005] (Figure 3(b)).

Fig. 2 shows still image tone mapping results using [Li et al. 2005] and our gain control. [Li et al. 2005]'s result has noticible halos around the army man. Although our gain control is not guaranteed to prevent halo artifacts, in all of our tests, it produces less halos than the gain control approach in [Li et al. 2005].

3.2 Temporal coherence

Applying still image tone mapping operators to video in a frame by frame fashion can cause flickering. We show the flickering resulting from applying [Li et al. 2005] to each frame of the plane chase sequence (from Figure 5) in the supplementary video. A natural

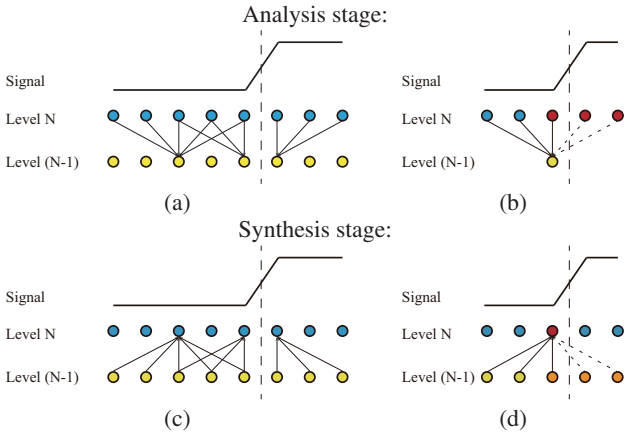

Figure 4: *Illustration of temporal edge avoidance. (a) Redundant wavelet analysis process - data from blue points are used to create the next level data (yelow points). (b) Zooming in on one point we see that it uses ghost particles (red) instead of the particles across the discontinuity. (c) The synthesis process works similarly to the analysis. (d) The synthesis process of one point again shows the use of ghost particles (orange). Notice how both the analysis and synthesis processes do not use information from across the discontinuity.*

idea to prevent flickering is to extend our spatial wavelets to the temporal domain, creating a 3D volume (x,y,t). However, this leads to ghosting, where the pixel values from previous or subsequent frames can influence the current frame. Examples of ghosting are shown in Figure 5(d) around the wing tips and in Figure 6(b) where characters are visible through the door.

To overcome the ghosting problem we attempted to use the edge avoiding wavelet scheme described in [Fattal 2009]. The results are shown in Figure 5 (e) and Figure 6 (c). Note that ghosting artifacts are reduced, but still visible. We therefore needed a temporal edge avoiding scheme that penalizes data across edges even more than the technique of [Fattal 2009].

We designed a new edge avoiding scheme tailor-made for the task of video tone mapping. It decomposes a signal into resolution bands like wavelets, and reconstructs a tone mapped signal from these resolution bands in a wavelet style. Note that unlike wavelets it isn't just a change of basis - as our synthesis process is not an exact inverse of the analysis process. However, our scheme is effective for achieving temporal coherence during tone mapping.

Our edge avoiding band decomposition introduces ghost particles to disconnect neighboring pixels with large intensity discontinuity in time. Suppose we have an 1D image of 8 pixels in level N denoted as $I_N(1), \cdots, I_N(8)$, as shown in Figure 4(a) with blue dots. Let $I^l_{N-1}(j), j \in 1, 2, \ldots, 8$ be the low-pass signal at level $N-1$, as shown in yellow dots. Classic redundant wavelet analysis uses a convolution process to compute a lower resolution band [Fowler 2005], which can be written as:

$$I^l_{N-1}(j) = I_N(j) \oplus f^l_N = \sum_{p=-k}^{k} I_N(j-p) f^l_N(p), \tag{6}$$

where $\oplus$ is a convolution operator and f^l_N is the low-pass filter at level N, and k is half the width of f^l_N. Similarly for the high-pass signal I^h_{N-1}, the process is

$$I^h_{N-1}(j) = I_N(j) \oplus f^h_N, \tag{7}$$

where f_N^h is the high-pass filter.

Our edge avoiding decomposition is defined with an edge avoiding convolution process, denoted with the operator $*$. We modify the analysis process into:

$$I_{N-1}^l(j) \quad = \quad I_N(j) * f_N^l = \sum_{p=-k}^{k} I'_N(j-p) f_N^l(p), \quad (8)$$

where

$$I'_N(j-p) = \begin{cases} I_N(j-p), & \text{if } edge(j, j-p, N) = 0 \\ I_N(j), & \text{if } edge(j, j-p, N) = 1 \end{cases} \quad (9)$$

$edge(x, y, q)$ is a binary function that indicates if there is an edge between point x and y in level q. The idea is to use the value of $I_N(j)$ to shadow $I_N(j-p)$ if there is an edge between point $(j-p)$ and j in level N. There are multiple choices for an edge detector. In our implementation, we found a simple thresholding approach is enough and works surprisingly well. We define the edge function as:

$$edge(x, y, q) = \begin{cases} 0 & \text{if } |I_q^l(x) - I_q^l(x)| < \alpha \\ 1 & \text{otherwise} \end{cases} \quad (10)$$

where α is a fixed constant threshold (set to 0.5 in our results). The reconstruction process is also an analog to the synthesis process of classic redundant wavelets, defined as:

$$I_N^r = \frac{1}{2} \left(I_{N-1}^l * f_p + I_{N-1}^h * f_q \right), \quad (11)$$

where I_N^r denotes the reconstructed signal, f_p is a reconstruction filter corresponding to the low-pass filter, f_q is the corresponding high-pass reconstruction filter. Recall that we use B-Spline wavelets in both space and time. So in our case $f_p = [0.5, 1, 0.5]$, $f_q = \frac{1}{3}[1, -6, 10, -6, 1]$. In short, our edge avoiding decomposition scheme shares a same framework with classic redundant wavelet analysis/synthesis process, but replace the convolution process with our edge avoiding convolution.

To make the math less confusing, we show the analysis process of one point $I_{N-1}^l(5)$ in Fig. 4 (b). It is still a convolution process, as in Eq. (8). But it does not take the value from $I_N(6)$ or $I_N(7)$, due to the discontinuity. Instead, the edge avoiding convolution process uses the value from $I_N(5)$ to shadow both $I_N(6)$ or $I_N(7)$, as shown in red dots in (b). We call the red dots for $I_N(6)$ and $I_N(7)$ ghost particles, borrowing the idea from fluid simulation. So in this case $I_{N-1}^l(5) = f(2)I_N(3) + f(1)I_N(4) + f(0)I_N(5) + f(-1)I_N(5) + f(-2)I_N(5)$.

The idea behind the design of our edge avoiding scheme is intuitive. If there is a large discontinuity between two points, we don't want any color information to propagate, in order to protect against ghosting artifacts. Note that is a major difference between our scheme and the one in [Fattal 2009]. We show the results of applying our temporal edge avoiding decomposition in Figure 5(f) and Figure 6(d).

4 Results

We present our results in images as well as in the accompanying video. Please view the images digitally in full resolution for best quality.

We compare the proposed method with two state-of-the-art video tone mapping techinques described in [Ramsey et al. 2004] and [Mantiuk et al. 2008], by reproducing [Ramsey et al. 2004] and using the implementation of [Mantiuk et al. 2008] in the online source code of pfstmo-1.4 [Mantiuk et al. 2007]. We tried a large set of parameters and show the best results that we found.

Figure 5 and the video show a "plane chase" sequence. There is a large brightness difference between the characters and the sky. This sequence is interesting and challenging for two reasons: fast object motion and rapid local dynamic range variation. The gun firing at the end of sequence introduces local brightness variation in time, while the sky remains relatively constant. The results of [Ramsey et al. 2004] are shown in Figure 5(b). In addition to the flickering artifacts, the results of [Ramsey et al. 2004] appear dark on the characters. The details on the tow truck (fourth image from left) are crushed. Similar detail loss can be seen in Figure fig:gun-fire-chase(c), where [Mantiuk et al. 2008] is applied. Notice the tail and the wheels of the plane, as well as the hood of the tow truck. Note that both [Ramsey et al. 2004] and [Mantiuk et al. 2008] are global tone mapping operators. It is widely accepted that local operators are more effective at preserving local details. The results in Figure 5(e) are generated with our technique but replacing our temporal edge avoiding scheme with the edge avoiding wavelets of [Fattal 2009]. They preserve the details in the dark regions, but suffer from ghosting artifacts. Our results (Figure 5(f)) maintain the details without flickering or ghosting.

Figure 6 and the video show a "sliding door opening" sequence. It is an indoor sequence where the objects in front of the sliding door get much less illumination. We show the results of both [Ramsey et al. 2004] and [Mantiuk et al. 2008] in the video. The details in the results of [Ramsey et al. 2004] get crushed and lost in dark regions. The results of [Mantiuk et al. 2008] suffer from contrast loss and look foggy. Our result has a natural look and reveals details on the door.

The "train tunnel" sequence in Figure 1 and the video is challenging due to the motion of train and the dramatic brightness variation over time. Our results reveal interesting details inside the tunnel, do not suffer from either flickering or ghosting, and preserve the temporal dynamic range variations of the sequence.

The accompanying video also shows our results on a live action video sequence (HDR video data courtesy of Grzegorz Krawczykw). Notice how the text of the road sign is revealed using our method.

Performance: There are 7 detail coefficient layers in each band and one low-pass layer at the coarsest level, therefore the memory cost of processing a video sequence is $n(7N + 1)$, where n is the total number of luminance values and N is the number of bands (in our experiments N is typically 4 or 5). It takes a few seconds to process a frame of 1920×1080 resolution on average with our current single thread C++ implementation on a Linux box with a 2.83 GHz Intel Xeon CPU. Since the tone mapping process is highly parallelizable, a GPU implementation should greatly reduce the running time.

5 Discussion and Future Work

We have presented a 3D wavelet based local tone mapping operator for videos. The method is temporally coherent, preserves local detail better than global methods, and is relatively efficient - running in linear time in practice. Our technique addresses the problem of haloing through an improved gain control method for both still images and video. Temporal coherence is achieved using our novel edge avoiding wavelet inspired band decomposition in time. Results on video sequences with fast motion and rapid dynamic range

Figure 5: *Tone mapping results on the "plane chase" sequence. (a) Five input HDR frames displayed in linear color space. (b) Results of our implemention of [Ramsey et al. 2004]. (c) Results of our implementation of [Mantiuk et al. 2008]. (d) Results of applying standard wavelets in the temporal domain. (e) Results of applying edge avoiding wavelets in [Fattal 2009] in the temporal domain. (f) Results from our method. Full video sequences can be found in the supplementary video.*

change show significant quality improvement over state-of-the-art video tone mapping techniques.

There are several limitations and areas for improvement with the proposed approach. First, although our gain control method has empirically proven to reduce haloing, there are no guarantees that the results will be halo free. Second, although the memory requirements are linear in the size of the input sequence, we would like to find ways of reducing the total in-core memory. One approach might be to develop non-redundant wavelets that are translation invariant in both time and space. This would make the total memory requirements equal to the size of the input. Another possibility would be to develop a streaming version to allow arbitrarily long video sequences to be processed. Third, it should be possible to derive a variant of our edge avoiding temporal decomposition so that the analysis and synthesis steps are inverses of one another, thereby putting the method on firmer theoretical ground. Finally, our technique is largely automatic (with only a few parameters) and does not provide much local artistic control. In the future we intend to develop an interactive tool for artists.

References

AKYUZ, A. O., FLEMING, R., RIECKE, B. E., REINHARD, E., AND BULTHOFF, H. H. 2007. Do hdr displays support ldr content? a psychophysical evaluation. In *SIGGRAPH*.

CVETKOVIC, S. D., KLIJN, J., AND H.N., P. 2008. Optimization of tone-mapping functions in video cameras for high dynamic range images. In *Real-time Image Processing*.

DURAND, F., AND DORSEY, J. 2002. Fast bilateral filtering for the display of high-dynamic-range images. In *SIGGRAPH*.

FATTAL, R., LISCHINSKI, D., AND WERMAN, M. 2002. Gradient domain high dynamic range compression. In *SIGGRAPH*.

FATTAL, R. 2009. Edge-avoiding wavelets and their applications. In *SIGGRAPH*.

FERWERDA, J. A., PATTANAIK, S. N., SHIRLEY, P., AND GREENBERG, D. P. 1996. A model of visual adaptation for realistic image synthesis. In *SIGGRAPH*.

FOWLER, J. E. 2005. The redundant discrete wavelet transform and additive noise. *IEEE Signal Processing Letters 12*, 9, 629–632.

KANG, S. B., UYTTENDAELE, M., WINDER, S., AND SZELISKI, R. 2003. High dynamic range video. *TOG 22*, 3, 542–547.

KIRK, A. G., AND O'BRIEN, J. F. 2011. Perceptually based tone mapping for low-light conditions. In *SIGGRAPH*.

KRAWCZYK, G., MYSZKOWSKI, K., AND SEIDEL, H.-P. 2005. Perceptual effects in real-time tone mapping. In *SCCG*.

LEE, C., AND KIM, C.-S. 2007. Gradient domain tone mapping of high dynamic range videos. In *ICIP*.

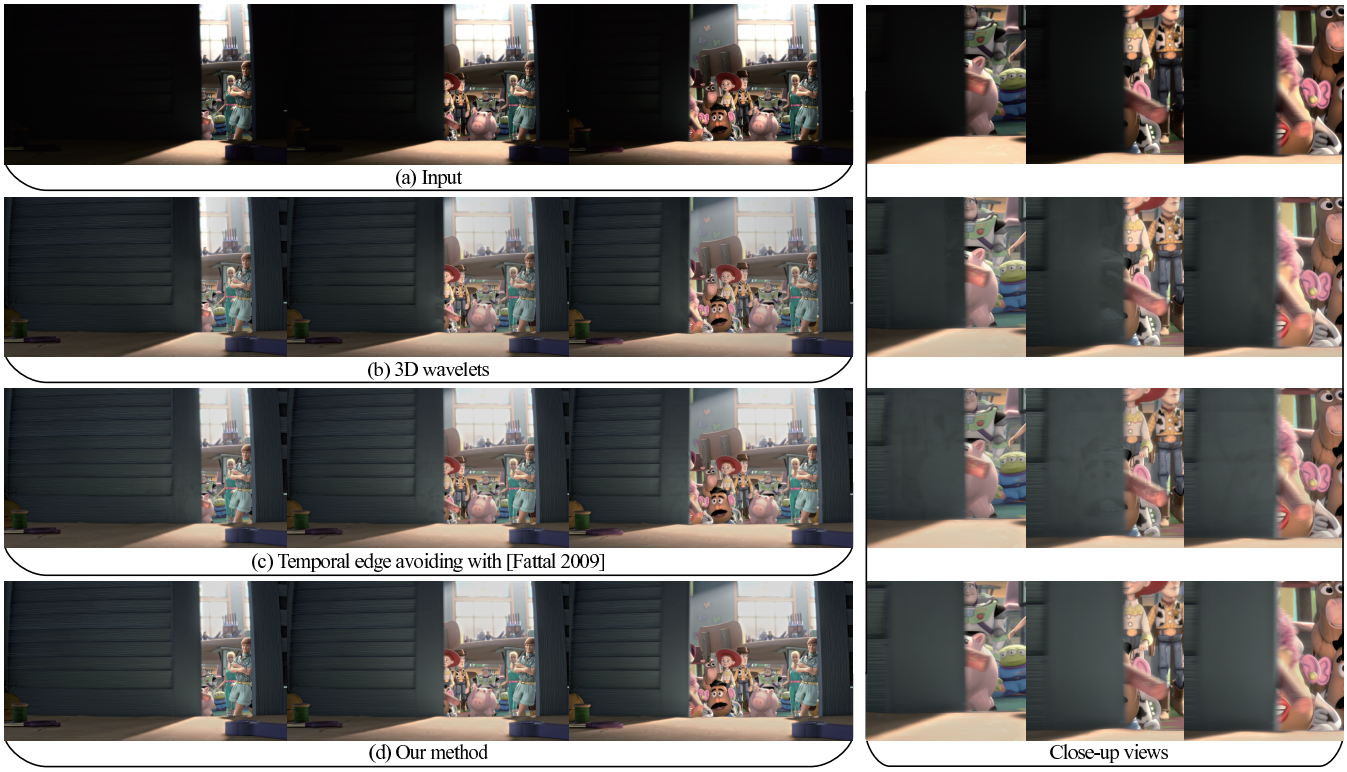

Figure 6: *Ghosting artifacts. (a) Three of the input HDR frames displayed in linear color space. (b) Tone mapped frames using 3D wavelet based algorithm (without temporal edge avoidance). (c) Tone mapped frames of applying edge avoiding wavelets from [Fattal 2009] in time. (d) Our results with temporal edge avoiding wavelets. Close-up views are shown on the right. Ghosting artifacts can be observed through the sliding door in (b) and (c).*

LI, Y., SHARAN, L., AND ADELSON, E. H. 2005. Compressing and companding high dynamic range images with subband architectures. In *SIGGRAPH*.

LISCHINSKI, D., FARBMAN, Z., UYTTENDAELE, M., AND SZELISKI, R. 2006. Interactive local adjustment of tonal values. In *SIGGRAPH*.

MANTIUK, R., KRAWCZYK, G., MANTIUK, R., AND SEIDEL, H.-P. 2007. High dynamic range imaging pipeline: Perception-motivated representation of visual content. In *Human Vision and Electronic Imaging XII*.

MANTIUK, R., DALY, S., AND KEROFSKY, L. 2008. Display adaptive tone mapping. In *SIGGRAPH*.

NAKAMAE, E., KANEDA, K., OKAMOTO, T., AND NISHITA, T. 1990. A lighting model aiming at drive simulators. In *SIGGRAPH*.

PARIS, S., W., S., AND KAUTZ, J. 2011. Local laplacian filters: Edge-aware image processing with a laplacian pyramid. In *SIGGRAPH*.

PATTANAIK, S. N., TUMBLIN, J., YEE, H., AND GREENBERG, D. P. 2000. Time-dependent visual adaptation for fast realistic image display. In *SIGGRAPH*.

QIU, G., GUAN, J., DUAN, J., AND CHEN, M. 2006. Tone mapping for hdr image using optimization a new closed form solution. In *ICPR*.

RAMSEY, S. D., JOHNSON, T., AND HANSEN, C. 2004. Adaptive temporal tone mapping. In *Computer Graphics and Imaging (CGIM)*.

REINHARD, E., STARK, M., SHIRLEY, P., AND FERWERDA, J. 2002. Photographic tone reproduction for digital images. In *SIGGRAPH*.

REINHARD, E., WARD, G., PATTANAIK, S., AND DEBEVEC, P. 2005. *High dynamic range imaging: acquisition, display, and image-based lighting*.

REMPEL, A. G., TRENTACOSTE, M., SEETZEN, H., YOUNG, H. D., HEIDRICH, W., WHITEHEAD, L., AND WARD, G. 2007. Ldr2hdr: On-the-fly reverse tone mapping of legacy video and photographs. In *SIGGRAPH*.

SHAN, Q., JIA, J., AND BROWN, M. 2010. Global optimized linear windowed tone-mapping. *TVCG*.

STOLLNITZ, E. J., DEROSE, T. D., AND SALESIN, D. H. 1996. *Wavelets for Computer Graphics: Theory and Applications*. Morgan Kaufmann.

TOCCI, M., KISER, C., TOCCI, N., AND SEN, P. 2011. A versatile hdr video production system. In *SIGGRAPH*.

TUMBLIN, J., AND RUSHMEIER, H. 1993. Tone reproduction for realistic images. *IEEE Computer Graphics and Applications 13*, 6.

WANG, Z. W., ZHAI, J., ZHANG, T., AND LLACH, J. 2010. Interactive tone mapping for high dynamic range video. In *Acoustics Speech and Signal Processing (ICASSP)*.

YOSHIDA, A., MANTIUK, R., MYSZKOWSKI, K., AND SEIDELY, H.-P. 2006. Analysis of reproducing real-world appearance on displays of varying dynamic range. In *Eurographics*.