Prediction-Guided Quantization for Video Tone Mapping

Agnès Le Dauphin^a, Ronan Boitard^{a,b}, Dominique Thoreau^a, Yannick Olivier^a, Edouard Francois^a and Fabrice LeLéannec^a

^aTechnicolor, 975 Av. des Champs Blancs, 35576 Cesson-Sevigne, France ^bIRISA, 263 Avenue du Général Leclerc, 35000 Rennes, France

ABSTRACT

Tone Mapping Operators (TMOs) compress High Dynamic Range (HDR) content to address Low Dynamic Range (LDR) displays. However, before reaching the end-user, this tone mapped content is usually compressed for broadcasting or storage purposes. Any TMO includes a quantization step to convert floating point values to integer ones. In this work, we propose to adapt this quantization, in the loop of an encoder, to reduce the entropy of the tone mapped video content. Our technique provides an appropriate quantization for each mode of both the Intra and Inter-prediction that is performed in the loop of a block-based encoder. The mode that minimizes a rate-distortion criterion uses its associated quantization to provide integer values for the rest of the encoding process. The method has been implemented in HEVC and was tested over two different scenarios: the compression of tone mapped LDR video content (using the HM10.0) and the compression of perceptually encoded HDR content (HM14.0). Results show an average bit-rate reduction under the same PSNR for all the sequences and TMO considered of 20.3% and 27.3% for tone mapped content and 2.4% and 2.7% for HDR content.

Keywords: HDR, Video Tone Mapping, HEVC, Quantization, Perceptual Encoding, Video Compression

1. INTRODUCTION

Tone Mapping Operators (TMOs) compress High Dynamic Range (HDR) images or video sequences to a lower dynamic range. They aim at preserving details of HDR content despite its retargeting. This operation is mandatory to display any content on a traditional Low Dynamic Range (LDR) device. However, before reaching the end-user, a tone mapped video sequence is usually compressed using a codec (coder-decoder) for distribution or storage purposes. Consequently, by optimizing the pair TMO/codec, one can reduce the entropy of a tone mapped video sequence and/or increase the compression efficiency of the codec.

Two different approaches have been designed to minimize the entropy of a tone mapped video sequence^{1,2}. The first technique optimizes the mapping performed by a TMO to achieve high compression ratio¹. However, this technique modifies the rendering of the tone mapped content, hence this technique will not be able to preserve any artistic intent. The second method reduces the entropy of a video sequence while preserving a given rendering. This technique, called Motion-Guided Quantization (MGQ), relies on a pre-processing step to adapt the quantization operation of any TMO². However, in its current implementation, the MGQ has two main limitations. First, this solution does not take into account the spatial and temporal coding structure of a codec³. Second, the Intra-prediction process does not benefit from the adaptive quantization. To summarize, the MGQ adapts the quantization only for one prediction mode while in a codec several prediction modes are tested, the best one being selected using a rate-distortion function⁴.

In this work, we propose to implement the Adaptive Quantization (AQ) in the loop of an encoder. Our method, called Prediction-Guided Quantization (PGQ), aims at increasing the compression efficiency of HDR and tone mapped LDR content. To better understand the proposed technique, we provide the necessary background in Section 2. Then, we present our method that implements the adaptive quantization at the encoding stage. In Section 4, we compare, in terms of compression efficiency, our method with a rounding-based quantization when compressing LDR and HDR video content. Before concluding, we discuss the different choices made to encode HDR content.

Further author information: (Send correspondence to Ronan Boitard)

Ronan Boitard: E-mail: ronan.boitard@technicolor.com, Telephone: +332 99 27 95 93



Figure 1. Block-based encoder workflow.

2. BACKGROUND

TMOs convert HDR content into LDR content. However, video content needs to be compressed to reduce the amount of data transmitted or stored. In this section, we first briefly introduce some key aspects of the ITU-T H.265 / MPEG-H Part 2 High Efficiency Video Codec (HEVC)³. Then, we describe the different operations required to tone map an HDR content. Finally, we present the MGQ technique as implemented by Boitard et al.².

2.1 HEVC

HEVC is the successor of the ITU-T H.264 / MPEG-4 Part 10 'Advanced Video Coding' (AVC) codec⁵. Developed by the Joint Collaborative Team on Video Coding (JCT-VC), it was released in January 2013 and is reported to double AVC compression ratio. The HEVC test Model (HM) is currently in its version 14.

HEVC is a block-based codec that exploits both spatial and temporal correlations between the code values of pixels to achieve high compression ratios. Figure 1 gives an overview of the different processes and their relationship inside a block-based encoder. To exploit these correlations, blocks are predicted using two types of prediction: Intra or Inter. Intra-prediction relies on spatial correlation to predict the current block using blocks already decoded in the current frame. Inter-prediction exploits the temporal correlation by predicting the current block using blocks from a set of previous/subsequent decoded frames. Using a rate-distortion function, the best predictor is selected among the Inter and Intra-prediction modes. The predicted block is then subtracted from the current block to obtain the residual block.

To encode the residuals, the current block is first converted to the frequency domain using a frequency transform (T). The resulting harmonic coefficients are then quantized (Q) before being fed to the entropy coder. The bitstream is composed of the encoded residuals and the prediction side-data to reproduce the same prediction at the decoder side. Note that in order to use the same pixels value at the decoder and encoder side, the blocks are reconstructed at the encoder side to be used for the prediction process of future blocks.

An extension for HEVC is currently being issued to support high bit-depth and full chroma sampling. This extension is called Range Extensions (RExt) and is in its version 7.

2.2 Tone Mapping

The conversion from HDR content to LDR ones is performed by Tone Mapping Operators (TMOs). In HDR imaging, the pixels represent the physical scene luminance (expressed in cd/m^2) stored as floating point values. In the case of LDR imaging, the pixels are assigned code values corresponding to a display-dependent color space



Figure 2. Workflow of the three steps needed to perform a tone mapping operation. \mathbf{F}^{t} is the \mathbf{t}^{th} frame of the video sequence.

such as sRGB⁶. Tone mapping an HDR image amounts to retargeting physical values with unlimited bit-depth to a constrained space. Figure 2 illustrates the HDR to LDR conversion. First, the mapping operation, which is the core of a TMO, compresses HDR values to fit in the range [0-1]. Secondly, the gamma encoding redistributes the tonal level closer to how our eyes perceive them (usually $\gamma = 1/2.2$). Finally, the quantization step converts floating point values to integer code values corresponding to the used bit-depth (i.e. $[0; 2^n - 1]$ for n bits). This operation consists in scaling the gamma encoded values of the current frame ($\mathbf{F}_{\mathbf{g}}^{\mathbf{t}}$) to the maximum value desired (i.e. $2^n - 1$) and then rounding them to the nearest integer:

$$\mathbf{F}_{\mathbf{d}}^{\mathbf{t}} = Q(\mathbf{F}_{\mathbf{g}}^{\mathbf{t}}) = \lfloor (2^{n} - 1)\mathbf{F}_{\mathbf{g}}^{\mathbf{t}} + 0.5 \rfloor = \lfloor \mathbf{F}_{\mathbf{s}}^{\mathbf{t}} + 0.5 \rfloor$$
(1)

where $Q(\cdot)$ represents the quantization operation, n the used bit-depth and $\lfloor \cdot \rfloor$ the rounding to the nearest lower integer. $(2^n - 1)\mathbf{F}_{\mathbf{g}}^{\mathbf{t}}$ (respectively $\mathbf{F}_{\mathbf{d}}^{\mathbf{t}}$) represents the unquantized (respectively quantized) gamma encoded frame. For the sake of clarity, the scaled gamma encoded frame $\mathbf{F}_{\mathbf{s}}^{\mathbf{t}}$ will be preferred to the notation $(2^n - 1)\mathbf{F}_{\mathbf{g}}^{\mathbf{t}}$. The quantization is performed for each channel of the frame $\mathbf{F}_{\mathbf{g}}^{\mathbf{t}}$ separately, regardless of its representation, i.e. RGB, YCbCr etc.

2.3 Motion-Guided Quantization

As video sequences are bound to be distributed to the end user, they will undergo compression. Compressing video content for mass distribution inevitably results in a loss of information which leads to a decreased subjective quality. As storage disc capacity and broadcaster's bandwidth are limited, compressed content bit-rates need to at least be below this limitation. Consequently, the level of entropy that any content can achieve is as important as the compression efficiency of a codec. Moreover, even if the tone mapping achieves the best subjective quality envisioned, it is more than likely that the compression stage will impair this quality. By optimizing the pair TMO/codec, one can reduce the entropy of tone mapped video content. Two major approaches have been considered so far. The first consists in modifying the mapping performed by a TMO to reduce the entropy of a sequence⁷. However, this modification alters the visual perception of this content and can impair any artistic intent or desired rendering.

The second technique proposes to adapt the quantization to increase temporal correlation². As the quantization in LDR imagery has been designed to be visually lossless (the number of quantization bins is higher than the number of Just Noticeable Difference in an 8 bits video), modifying it will not alter the rendering. Consequently, the Motion-Guided Quantization (MGQ) technique performs an Adaptive Quantization (AQ) to minimize the distortion between successive frames through a motion compensation. It uses the predicted frame $\mathbf{F}_{\mathbf{p}}$ to adapt the quantization of the current scaled gamma encoded frame $\mathbf{F}_{\mathbf{s}}$:

$$\mathbf{F}_{\mathbf{d}}^{\mathbf{A}\mathbf{Q}} = AQ(\mathbf{F}_{\mathbf{s}}) = \begin{cases} [\mathbf{F}_{\mathbf{s}}] & \text{if } 0 \leq \mathbf{F}_{\mathbf{s}} - \mathbf{F}_{\mathbf{p}} < \delta \\ [\mathbf{F}_{\mathbf{s}}] & \text{if } -\delta < \mathbf{F}_{\mathbf{s}} - \mathbf{F}_{\mathbf{p}} < 0 \\ [\mathbf{F}_{\mathbf{s}} + 0.5] & otherwise \end{cases}$$
(2)

where $AQ(\mathbf{F_s})$ represents the Adaptive Quantization step while $\lfloor . \rfloor$ (respectively $\lceil . \rceil$) represents the rounding to the nearest lower (respectively higher) integer. Recall that $\mathbf{F_s}$ is expressed with floating point values while $\mathbf{F_p}$ with integer ones. Both frames values range from 0 to $2^n - 1$. The proposed method allows a trade-off between the compression efficiency and the distortion between the unquantized and quantized tone mapped video through the δ parameter (see Boitard et al.²).

3. PREDICTION-GUIDED QUANTIZATION

Although the MGQ technique increases the compression efficiency, it remains sub-optimum with respect to the prediction process performed in HEVC. Indeed, the MGQ does not follow the temporal and spatial coding structure³ used in a codec as it only uses motion compensation between successive frames. Furthermore, the Intra-prediction process does not benefit from the Adaptive Quantization (AQ). In addition, the MGQ adapts the quantization only for one prediction mode while in a codec several prediction modes are tested, the best one being selected using a rate-distortion function⁴. Finally, performing a motion estimation can be quite complex for a simple tone mapping operation.

To address these limitations, we propose to implement the AQ in the loop of the encoder. Our Prediction-Guided Quantization (PGQ) technique aims at reducing more efficiently the entropy of any tone mapped video content. As a codec greatly relies on the prediction to remove redundant data, increasing the quality of the prediction provides a higher compression efficiency⁷. The AQ is tuned separately for each of the available prediction modes. Henceforth, the selected mode and its associated quantization will depend on the codec rate-distortion optimization function.

To implement our approach inside the coding loop of a block-based encoder (say HEVC), the computations within this codec need to be performed with floating point values rather than integers. That is why, the unquantized tone mapped frame (\mathbf{F}_{s} , stored as floating point values) has to be quantized inside the codec. To this end, we need to perform the AQ when the residuals are computed. Figure 3 illustrates where the AQ needs to be added to the traditional workflow of a block-based encoder. As HEVC is a block-based codec, the input frame that requires to be quantized inside the codec \mathbf{F}_{s} is divided into blocks \mathbf{B}_{s} . Each value of \mathbf{B}_{s} is adaptively quantized thanks to the AQ, according to the prediction block \mathbf{B}_{p} (based on previously quantized blocks). The computed values are attributed to the quantized tone mapped block \mathbf{B}_{d}^{AQ} . The best predictor is then selected among the Inter and Intra-prediction using a rate-distortion function. Finally, the predicted block \mathbf{B}_{p} is subtracted from the current quantized block \mathbf{B}_{d}^{AQ} to obtain the residual block \mathbf{B}_{res} that will be encoded.

4. RESULTS

In this section, we propose to evaluate the efficiency of our PGQ technique in two different scenarios: the compression of tone mapped video content and of perceptually encoded HDR video content. We assess the efficiency of our technique by comparing the average bit-rate reduction under the same PSNR with a simple rounding. In addition, the overhead in complexity is also assessed.

4.1 Compression of Tone Mapped Video

To test the compression of tone mapped video content, we used the HM10.0 in its 8-bits main profile configuration (4:2:0 chroma format). Our input content is an HDR video that we tone mapped in floating point values in the range [0;1] and transposed in the YC_bC_r representation⁶. To assess the efficiency of our technique, we quantize the current frame to tone map \mathbf{F}_s^t in three different ways: rounding outside the codec, inside the codec and the PGQ. We perform our tests on eight HDR videos, four TMOs and two different configurations. Several TMOs are used to prove that the PGQ improves the compression efficiency for any TMO. One operator⁸ is optimized for the compression of HDR images while the other three focus on the subjective quality⁹⁻¹¹. The two profiles correspond to the Intra Only and Random Access configurations¹².

In a first test, we evaluated the difference in complexity between our solution and a simple rounding. Results show that our method increases the encoding time by 48% (Intra Only) and 38% (Random Access). This overhead is mostly due to disk access of input files that are four times bigger (32 bits versus 8 bits) as well as the



Figure 3. Propagation of the floating point values in a block-based encoder workflow. AQ stands for Adaptive Quantization, \mathbf{B}_{res} corresponds to the residual block, \mathbf{B}_{p} represents the prediction block, \mathbf{B}_{s} is the unquantized tone mapped block and \mathbf{B}_{d}^{AQ} is the adaptively quantized block. \mathbb{R} corresponds to floating point values and \mathbb{N} to integers.



Figure 4. Architecture of the tests performed to assess the complexity and compression efficiency of the PGQ method when compressing tone mapped video content.

use of floating point values in the codec. Note that the encodings were computed on a grid-cluster, consequently the complexity assessment is only informative.

In a second test, we compared the PGQ with a simple rounding inside the codec. This test assess the efficiency of our technique as well as the complexity of the Adaptive Quantization. Table 1 provides the average bit-rate reduction under the same PSNR¹³. Results show that the PGQ requires smaller bit-rates to achieve the same quality (average reduction of 20.3% for Intra Only and 27.3% for Random Access). We observe that the complexity is similar for both rounding. We also tested with $\delta = 1$ and up to 3.2% reduction of the total bit-rate is achieved.

Finally, Table 2 compares the results of the PGQ to those of the MGQ². For this test, only three HDR

	Average BD-rate reduction $(\%)$			Complexity (%)
	Y PSNR	U PSNR	V PSNR	Complexity (70)
Intra Only $\delta = \infty$	-20.3	-33.6	-26.1	100
Random Access $\delta = \infty$	-27.3	-51.3	-41.8	77
Intra Only $\delta = 1$	-3.2	-9.4	-4.7	

Table 1. Average BD-rate reduction¹³ under the same PSNR and assessment of the complexity overhead.

		MGQ BD-rate (%)	PGQ BD-rate (%)
Sun		-12.8	-44.3
Tunne	1	-10.4	-36.0
Studen	nts	-5.6	-26.1
Mean		-9.6	-35.5

Table 2. Average BD-rate reduction under the same PSNR between the MGQ and the PGQ for three HDR videos, for the LDR compression for the Random Access configuration. The sequences Sun and Tunnel have been tone mapped with Ramsey et al.¹¹ operator, while the sequence Students has been tone mapped with Farbman et al.¹⁰ operator.

sequences and two TMOs were considered. Results show that the PGQ can save in average 35.5% of the total bit-rate where the MGQ only achieved 9.6%.

4.2 Compression of HDR content

We also evaluated the PGQ when compressing HDR video content. Three main techniques have been designed to compress native HDR content represented by floating point values¹⁴.

The first main technique^{15, 16} uses a ratio-based method and requires two codec instances: one to encode an LDR tone mapped version while the second encodes a quantized ratio between the original HDR sequence and its associated tone mapped LDR sequence (from 8 to 14 bits). At the decoder side, the decoded LDR video is multiplied by the decoded ratio video to reconstruct the HDR content.

The second main technique^{8, 17, 18} relies on an inverse tone mapping operator (operation that transforms LDR 8-bits data into floating-point HDR data). This approach is scalable and requires also two codec instances: one to compress a tone mapped LDR version of the original HDR video and a second that encodes only the residuals between the original HDR video and an inverse tone mapped version of the decoded LDR tone mapped video.

The third main technique^{19,20} is based on a perceptual curve, also called inverse Electro-Optic Transfer Function $(EOTF^{-1})$, that transforms HDR sequence (floating-point data) into a high bit-depth representation (integer data with quantization step not visible by the human eye). The generated LDR video is directly compressed using a single high-bit depth codec. At the decoder side, the decoded sequence is computed by the inverse of the perceptual curve to reconstruct the HDR sequence.

We argue that the PGQ method can be used for any of those three methods but we tested it only with the third one. Figure 5 illustrates the workflow to encode HDR video using a single layer. First the $EOTF^{-1}$ encodes perceptually an HDR sequence before quantizing it to a targeted bit-depth. The sequence is then encoded, decoded, inverse quantized and finally converted back in physical floating point values using the EOTF to obtain the HDR reconstructed video. The codec needs a high bit-depth, that is to say superior to 10 bits¹⁸.

For our tests, we first convert the HDR video to the XYZ color space²¹. We encoded the HDR video with the Perceptual Quantizer^{22,23} (PQ) $EOTF^{-1}$ using a 12 bits quantization. The PQ curve encodes luminance values ranging from 0.005 to 10,000 cd/m^2 . Then, since the YC_bC_r representation is tied to a restricted color space^{6,24} and cannot represent all the colors present in HDR content, we transposed the HDR video in the YD_zD_x representation (currently under standardization by SMPTE²⁵) with a 4:4:4 chroma format. YD_zD_x is a color difference encoding where D_x corresponds to the weighted difference between X' and Y'(respectively between Z' and Y' for the D_z):

$$D_x = \frac{c_1 Z' - Y'}{2.0}, \quad with \quad c_1 = \frac{2763}{2800},$$
(3)



Figure 6. Architecture of the tests performed to assess the complexity and compression efficiency of the PGQ method when compressing perceptually encoded HDR video content.

	Average BD rate $(\%)$			Complexity (%)
	Y PSNR	U PSNR	V PSNR	Complexity (70)
Intra Only	-2.4	-5.2	-10.5	114
Random Access	-2.7	-7.7	-13.1	101

Table 3. Average BD-rate reduction under the same PSNR between the PGQ and a simple rounding inside the codec when compressing perceptually encoded HDR content.

$$D_z = \frac{c_2 X' - Y'}{2.0}, \quad with \quad c_2 = \frac{2741}{2763},$$
(4)

where X', Y' and Z' are the perceptually encoded HDR values in the range [0;1] while D_z and D_x are the color difference encodings of the chroma in the range [-0.5;0.5]. The whole workflow is described in Figure 6.

We tested the PGQ on five HDR sequences²⁶ using the HM14.0-RExt7.0. We first assessed the difference in complexity between a codec using floating point values and 16 bits integer, that is to say comparing the rounding quantization inside and outside the codec. The complexity ratio is 119% and 103% for the Intra Only and the Random Access configurations. This increase of complexity is less important than during the LDR encoding as the input file is on 16 bits instead of 8 bits. Then we compared the PGQ with performing the rounding inside the codec. Results are provided in Table 3. They show an average bit-rate reduction under the same PSNR of 2.4% and 2.7% for the luma and ranging from 5.2% to 13.1% for the chroma. The bit-rate reduction is hence less important than during the LDR encoding. This comes from the use of 12 bits integer values instead of 8 bits integer values. Indeed, since the quantization is performed on more bits, a difference of 1 in the quantized value gives a less important change compared to a rounding, with respect to the signal dynamics.

5. DISCUSSION

For all the tests presented in this paper, we assessed the compression efficiency by computing the average BDrate reduction under the same PSNR¹³. The PSNR was computed between quantized videos, that is to say between integer values. For the LDR scenario, this is not an issue since LDR displays are addressed by integer values and cannot handle floating point values. Furthermore, as stated in Section 2.3, the LDR quantization is supposed to be visually lossless for traditional LDR luminance range (0.1 to 100 cd/m^2). Regarding the HDR compression, our computation of the PSNR is less valid and only represents the distortion between the original and reconstructed signal. This measurement cannot accurately represent the visual distortion of the video. Another solution could have been to compute the PSNR between the original and the reconstructed HDR video. However, the PSNR assumes the perceptual uniformity of pixel's value which is not the case for HDR pixels. As no consensus on the use of an objective metrics for HDR content has yet been reached, the only way to assess the efficiency of our technique would have been to perform a subjective evaluation. This evaluation requires the use of at least one high bit-depth HDR display. As the only commercially available HDR display²⁷ is limited with respect to the gamut and the sampling of its dynamic range, such an evaluation would most likely be impaired by the limitation of the display. However, it will most probably be possible to evaluate our solution on near future HDR displays. Another aspect of the HDR tests needs to be discussed: the choice of the color difference encoding. Indeed, all the HDR compression tests were performed using YD_zD_x representation. This representation is currently under investigation for HDR compression by the ad hoc MPEG group on Wide Color Gamut and High Dynamic Range. However, it is not the only format under investigation. We argue the PGQ is independent from the chosen format and that test should be conducted for the color difference encoding recommended by the MPEG group. This is also true for the choice of the EOTF.

6. CONCLUSION

In this paper, we pointed out that the implementation of the Adaptive Quantization (AQ) presented in Boitard et al.² was sub-optimum with respect to the prediction process performed in HEVC. We proposed to implement the AQ in the loop of an encoder (HEVC). Our Prediction-Guided Quantization technique adapts the quantization for each Inter and Intra-prediction mode. Our method was tested over two different scenarios: the compression of tone mapped LDR video content and the compression of perceptually encoded HDR content. Results showed that an average bit-rate reduction of 20.3% and 27.3% under the same PSNR was achieved for the LDR encoding (2.4% and 2.7% with the HDR encoding).

REFERENCES

- Koz, A. and Dufaux, F., "A comparative survey on high dynamic range video compression," in [SPIE 8499, Applications of Digital Image Processing XXXV], Tescher, A. G., ed., 84990E (Oct. 2012).
- [2] Boitard, R., Cozot, R., Thoreau, D., and Bouatouch, K., "Motion-Guided Quantization for Video Tone Mapping," in [IEEE International Conference on Multimedia and Expo (ICME)], (2014).
- [3] Sullivan, G. J., Ohm, J.-R., Han, W.-J., and Wiegand, T., "Overview of the High Efficiency Video Coding (HEVC) Standard," *IEEE Transactions on Circuits and Systems for Video Technology* 22, 1649–1668 (Dec. 2012).
- [4] Kim, Y.-H., Yoo, J.-W., Lee, S.-W., Shin, J., Paik, J., and Jung, H.-K., "Adaptive mode decision for h.264 encoder," *Electronics Letters* 40, 1172–1173(1) (September 2004).
- [5] Wiegand, T., Sullivan, G., Bjontegaard, G., and Luthra, A., "Overview of the H.264/AVC video coding standard," *IEEE Transactions on Circuits and Systems for Video Technology* 13, 560–576 (July 2003).
- [6] ITU, "RECOMMENDATION ITU-R BT.709-3," (1998).
- [7] Boitard, R., Thoreau, D., Cozot, R., and Bouatouch, K., "Impact of Temporal Coherence-Based Tone Mapping on Video Compression," in [Proceedings of the 21st European Signal Processing Conference (EU-SIPCO)], (2013).
- [8] Mai, Z., Mansour, H., Mantiuk, R., Nasiopoulos, P., Ward, R., and Heidrich, W., "Optimizing a tone curve for backward-compatible high dynamic range image and video compression.," *IEEE Transactions on Image Processing (TIP)* 20(6), 1558–1571 (2011).

- [9] Kiser, C., Reinhard, E., Tocci, M., and Tocci, N., "Real-time Automated Tone Mapping System for HDR Video," Proceedings of the IEEE International Conference on Image Processing, 2749–2752 (2012).
- [10] Farbman, Z., Fattal, R., Lischinski, D., and Szeliski, R., "Edge-preserving decompositions for multi-scale tone and detail manipulation," in [ACM SIGGRAPH 2008 papers on - SIGGRAPH '08], 1, ACM Press, New York, New York, USA (2008).
- [11] Ramsey, S., III, J. J., and Hansen, C., "Adaptive temporal tone mapping," Computer Graphics and Imaging - 2004, 3–7 (2004).
- [12] Bossen, F., Bross, B., Suhring, K., and Flynn, D., "HEVC Complexity and Implementation Analysis," IEEE Transactions on Circuits and Systems for Video Technology 22, 1685–1696 (Dec. 2012).
- [13] Bjontegaard, G., "Improvement of the bd-psnr model." ITU-T SG16, VCEG-AI11 (Jul 2008).
- [14] Touzé, D., Lasserre, S., Olivier, Y., Boitard, R., and François, E., "HDR Video Coding based on Local LDR Quantization," in [HDRi2014 - Second International Conference and SME Workshop on HDR imaging (2014),], (2014).
- [15] Lee, C. and Kim, C.-s., "Rate-Distorsion Optimized Compression of High Dynamic Range Videos," 16th European Signal Processing Conference (EUSIPCO) (Eusipco) (2008).
- [16] Ward, G. and Simmons, M., "Subband encoding of high dynamic range imagery," in [Proceedings of the 1st Symposium on Applied perception in graphics and visualization - APGV '04], 83, ACM Press, New York, New York, USA (2004).
- [17] Koz, A. and Dufaux, F., "Optimized Tone Mapping with Perceptually Uniform Luminance Values for Backward-Compatible High Dynamic Range Video Compression," *IEEE Proc. Visual Communications and Image Processing (VCIP'2012)* (2012).
- [18] Mantiuk, R., Myszkowski, K., and Seidel, H.-p., "Lossy Compression of High Dynamic Range Images and Video," in [SPIE Proceedings Vol. 6057: Human Vision and Electronic Imaging XI], Rogowitz, B. E., Pappas, T. N., and Daly, S. J., eds., 60570V-60570V-10 (Feb. 2006).
- [19] Zhang, Y., Reinhard, E., and Bull, D., "Perception-based high dynamic range video compression with optimal bit-depth transformation," in [2011 18th IEEE International Conference on Image Processing], 1321–1324, IEEE (Sept. 2011).
- [20] Motra, A. and Thoma, H., "An adaptive Logluv transform for High Dynamic Range video compression," in [2010 IEEE International Conference on Image Processing], 2061–2064, IEEE (Sept. 2010).
- [21] Smith, T. and Guild, J., "The c.i.e. colorimetric standards and their use," Transactions of the Optical Society 33(3), 73 (1931).
- [22] Miller, S., Nezamabadi, M., and Daly, S., "Perceptual Signal Coding for More Efficient Usage of Bit Codes," SMPTE Motion Imaging Journal 122, 52–59 (May 2013).
- [23] Kunkel, T., Ward, G., Lee, B., and Daly, S., "HDR and wide gamut appearance-based color encoding and its quantification," in [2013 Picture Coding Symposium (PCS)], 357–360, IEEE, San Jose (Dec. 2013).
- [24] ITU, "Recommendation ITU-R BT.2020," (2012).
- [25] SMPTE ST 2085, "YDzDx color-difference encoding for XYZ signals," in [Society of Motion Picture and Television Engineers], (2014).
- [26] Lasserre, S., Le Léannec, F., and Francois, E., "Description of HDR sequences proposed by Technicolor," in [ISO/IEC JTC1/SC29/WG11 JCTVC-P0228], IEEE, San Jose, USA (2013).
- [27] "Sim2: Hdr47e s 4k technical specifications." http://www.sim2.com/HDR/hdrdisplay/hdr47e_s_4kon01. 05.2012.