

A nonlinear quantization scheme for two-layer HDR image encoding

Ishtiaq Rasool Khan^{1,2}

Received: 8 August 2014 / Revised: 2 October 2015 / Accepted: 5 November 2015 / Published online: 18 November 2015
© Springer-Verlag London 2015

Abstract Two-layer schemes provide an effective method of encoding high dynamic range images with backward compatibility. The first layer is the tone-mapped low dynamic range version of the original image, used for visualization. The residual information that cannot be preserved in the first layer is stored in the second layer, which itself is generally encoded as an image of a fixed bit-depth. Any further details that cannot be preserved in the second layer are discarded. In this paper, we present a nonlinear quantization algorithm that can significantly enhance the amount of details that can be preserved in the second layer, and therefore improve the encoding efficiency. The proposed technique can be incorporated in any existing two-layer encoding method and leads to significant improvement in their performance.

Keywords HDR image coding · Nonlinear quantization · Backward compatibility · Two-layer encoding · Image representation

1 Introduction

Real-world scenes have much higher dynamic range compared to what can be shown on a paper or screen. Several tone-mapping operators (TMOs) have been presented to generate low dynamic range (LDR) versions of the high dynamic range (HDR) images for purpose of displaying. Additional information contained in HDR images, which is lost during

tone-mapping, can be quite useful in certain applications, such as high-quality CG rendering for gaming and virtual reality, sensing, scene analysis, and surveillance. Encoding higher dynamic range is not only important for these applications, but also for better visualization experience in future, when HDR displays would possibly become more prevalent. Several HDR image encoding formats have been developed, including RGBE [1], LogLuv [2], and IEEE 16 bit standard float format used by Industrial Light and Magic (ILM) [3]. A detailed discussion of different HDR image formats can be found in [4].

Some lossy compression techniques for HDR image and video have also been proposed [5–7]. An interesting idea proposed first in [8] is to encode HDR images in two layers. The first layer contains a tone-mapped LDR image, and the residual information lost in the process is encoded in the second layer. A major advantage of this approach is its backward compatibility to the existing applications and displays designed for the LDR images. Several variants of this scheme have been proposed. Ward presented his JPEG-HDR [9] extension to the existing JPEG standard. In this format, a tone-mapped LDR image is encoded as a standard JPEG image, and the ratio of the luminance channels of the original HDR and the tone-mapped LDR image is encoded as another JPEG image in the wrappers of the first image.

Some two-layer schemes use inverse tone-mapping functions to reduce the dynamic range of the residual information to improve coding efficiency (details preserved at the same data size). Mantiuk et al. [10] proposed HDR-MPEG format for HDR video and image compression. They form a 256-bin histogram, with each bin corresponding to one of the 0–255 levels of an 8-bit-per-channel LDR image. Corresponding HDR luminance values that map to the same LDR luminance level are collected in these bins. The average bin values are used to reconstruct an approximate HDR image

✉ Ishtiaq Rasool Khan
irkhan@kau.edu.sa; irkhan@uj.edu.sa

¹ Faculty of Computing and Information Technology,
King Abdulaziz University, Jeddah, Saudi Arabia

² Faculty of Computing and Information Technology,
University of Jeddah, Jeddah, Saudi Arabia

from the tone-mapped LDR version. Deviation of the reconstructed HDR image from the original is stored as the residual image, which has reduced dynamic range as compared to the ratio image in JPEG-HDR [9]. This improves the coding efficiency at a small cost of slight overhead of storing a 256-value lookup table.

Several other algorithms have been developed using the concept of the second layer [11–20]. Okuda et al. [11, 12] suggested using smooth functions, such as the Hill function, to approximate the average bin values, to avoid unnecessary energy yielded in the residual data at higher frequencies, due to possible sharp jumps in the lookup tables of [10]. Khan used piecewise linear functions [13, 14] and cubic splines [15] to approximate the inverse tone-mapping operators accurately and to get rid of any unwanted jumps.

Theoretically, it is possible to keep adding layers until the complete information in the original HDR image is preserved. Another way to preserve information is to encode the second layer losslessly, as done in [16] using lossless JPEG2000. These approaches can achieve backward compatibility, but the size of the encoded data becomes very large, and the decoder becomes relatively complex. In general, multiple-layer approaches are restricted only to two layers with fixed bit-depths, to achieve backward compatibility, good compression, and simpler decoding operation. ISO is considering some of these to become part of its JPEG standard, under a new name JEG-XT [20].

A classification of two-layer schemes can be done into two broad categories: (1) which contrive new TMOs to obtain the first layer that can better preserve the details, and hence reduce the leak of information to the residual, and (2) which formulate better inverse tone-mapping functions for reconstruction. JPEG is generally used to encode the first layer, because it is good at preserving the visual features. However, the residual information in the second layer, whose spectrum may be very different from that of the original image, can be treated as random data with a goal to preserve as much information as possible in a 2D structure (image) of a given bit-depth. There is not much work available in

the literature in this direction. In this paper, we present an additional step of nonlinear quantization in encoding of the second layer, which results in preserving a significant amount of additional details in an image of a given bit-depth.

As in all existing methods, we encode the first layer as a JPEG image for backward compatibility with non-HDR displays and applications. The residual data are encoded as a lossless PNG or a lossy JPEG image of 8 bits per channel. In the former case, the only loss occurs due to quantization, which our proposed nonlinear quantization scheme attempts to minimize. The proposed technique can be seen as pre-processing of the second layer before encoding and can be used with any existing two-layer algorithm to improve its performance. There is a slight overhead of encoding 256 values of the quantization levels, but it is negligible as compared to the overall size of the image. We present several experimental results to show effectiveness of the proposed technique.

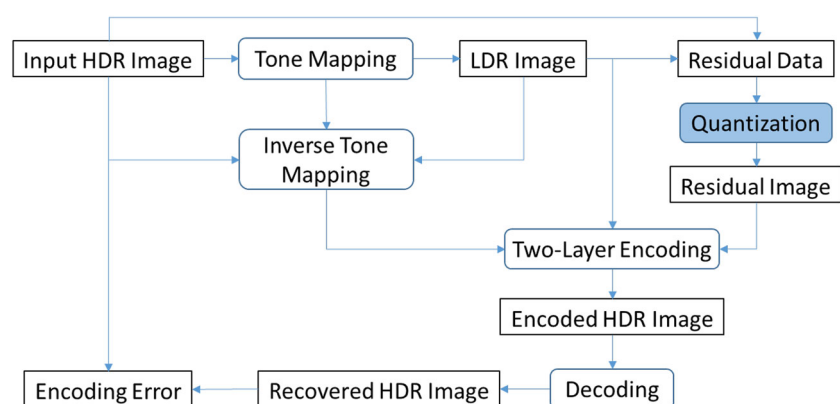
2 Two-layer encoding

Structure of a typical two-layer encoding scheme is shown in Fig. 1. An input image is tone-mapped to an LDR image. The residual data are encoded as another LDR image. Optionally, an inverse model that can approximately map the first LDR image back to HDR can also be constructed. In this case, the residual data would be different of the original from the reconstructed image and have reduced dynamic range. A nonlinear quantization scheme, such as one proposed in this paper, can be used to contain higher amount of details in the same bit-depth image. In this section, we will explain this model and particularly highlight the quantization step, which is the major contribution to this work.

2.1 Tone-mapping

Several TMOs have been presented to convert an HDR image to an LDR version for displaying. This LDR image forms the

Fig. 1 Structure of a typical two-layer HDR encoding scheme. *Rounded edges rectangles* are used to distinguish operations from the data



first layer in a two-layer scheme. It is recommended to use only global TMOs to get the first LDR layer in two-layer HDR encoding schemes, because their curves can be represented mathematically more accurately and in simpler forms. A study on impact of TMOs used to generate the first layer on coding efficiency of a two-layer scheme was presented in [10], in which local operators showed significantly poor performance compared to the global operators.

2.2 The residual image

In general, TMOs are constructed only for the luminance channel of the HDR images, whereas the chroma channels are kept unchanged. Therefore, the residual image is in general a single-channel image, although some works form three-channel residual images as well ([10] for example). In this paper, we will use single-channel residual image; however, the proposed technique can be applied to three-channel images as well. Denoting luminance of the HDR image as Y_{HDR} , and the luminance of its tone-mapped LDR version as Y_{LDR} , the residual data R can be written as:

$$R = \log(Y_{HDR}) - f(\log(Y_{LDR})), \tag{1}$$

where $f(\cdot)$ is the inverse TMO.

The residual R in Eq. (1) is floating point data, which needs to be rounded to an integer of a given bit-depth and saved in the second-layer image. For the sake of discussion, we will assume a bit-depth of eight, which is a popular choice among the existing methods. An extension to larger bit-depths is straightforward. If R' is the residual data rounded to 8-bit integers, the recovered luminance Y'_{HDR} can be written as:

$$\log(Y'_{HDR}) = R' + f(\log(Y_{LDR})). \tag{2}$$

The difference between the actual and the recovered luminance is

$$\log(Y'_{HDR}) - \log(Y_{HDR}) = R' - R, \tag{3}$$

which implies that the information lost in two-layer encoding schemes is lost only in encoding the residual image, and if the residual image is encoded losslessly, the HDR image can be perfectly recovered. The loss is due to either quantization or the lossy encoding (such as JPEG) and can be reduced by the following three methods:

- *Choice of a suitable TMO:* Some operators can better preserve the HDR image features and thus reduce the information content in the residual data.
- *Construction of inverse TMO:* Theoretically, if an inverse function can recover the HDR image from the first layer LDR image, the residual data will not be produced.

However, tone-mapping is many-to-one mapping, and therefore, a perfect inverse cannot be constructed. Several techniques can however be employed to recover an approximation (see [10–20] for example).

- *Effective encoding of the second layer:* The basic idea is to preserve as much details of the residual data as possible in an image of a given bit-depth [21]. We present a nonlinear quantization scheme in the next subsection that can achieve this goal.

2.3 Nonlinear quantization

Residual data are in floating point format and have quite high dynamic range. Simple scaling and rounding operation to bring it in 0–255 integer range is equivalent to using 256 equidistant quantization levels, which is not an effective approach and causes large quantization errors. We have shown a typical case in Fig. 2, in which accumulated quantization errors of all pixels falling in each of 256 bins are plotted. Note that the error at the mid-range of intensities is much higher than in the low and high intensity regions. This is in fact scene dependent. For this particular image, majority of the pixels have intensities in the middle range, and this can be different for a different image.

To reduce the overall error, our approach is to spread the error somewhat uniformly among the bins. Since total error of a bin is approximately proportional to the number of pixels inside it, one simple and intuitive approach to do so would be to use nonuniform bin widths, such that the number of pixels in each bin is nearly equal. However, this can make some bins unreasonably wider and hence lead to larger individual pixel errors and accumulated bin errors. A typical case has been shown in Fig. 2 (dotted curve). A more sophisticated clustering model (such as k -means) can significantly reduce the overall error, but can also lead to higher errors of the individual bins and pixels. Moreover, the clustering algorithms are generally very slow. Even on reasonably advanced machines,

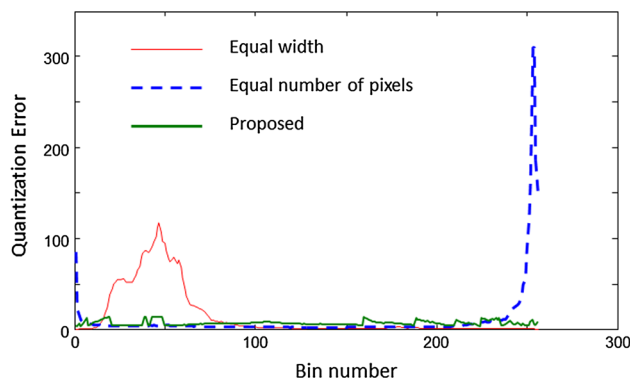


Fig. 2 Accumulated quantization errors in each bin, when they are equally spaced, have same number of pixels, and are placed using proposed algorithm

clustering of an image having a few millions of pixels into at least 256 bins (for bit-depth of 8) can take time in hours. We propose the following simple algorithm to reduce and spread the error over the bins almost uniformly:

1. Sort all pixel values in ascending order and place them in a single bin with edges at the minimum and the maximum values.
2. Divide the bin into two by placing an additional edge at the bin center, which we define as mean of the pixel values inside the bin. Some advantage in terms of speed can be obtained by placing the new edge at the middle of two existing edges, but a slight degradation in performance (increase in error) is observed.
3. Calculate the errors for each new bin. For a typical image, total error of two new bins is almost 50.5% of the error in the parent bin. This figure is based on the average reduction in all 255 split operations for the *Memorial* HDR image. Here we define the bin error as sum of the Euclidean distance of each pixel in the bin from the bin center.
4. Finally, we perform an additional refinement step. Pixels lying near the boundary of each bin are moved to the right or left neighboring bin if doing so reduces their contributions in the total error. Bin centers are then updated, and the procedure is repeated a few times until a local minimum is achieved and no significant further reduction in the error is obtained. Note that this local minimization does not distort the state of nearly uniform spread of the error obtained by the split operations given above. This cannot be guaranteed for a global optimization method.

Bin errors for our example image obtained by using the above algorithm are shown in Fig. 2 (thick solid curve), which are quite low in magnitude and spread almost uniformly as compared to other two cases. Some further refinements in the algorithm can be made to obtain more evenly spread error curves, but that increases unnecessary complexity without much reduction further in the overall error.

2.4 Encoding and decoding

As mentioned earlier, we encode the first layer as JPEG at the highest quality for better visualization. The second layer can be encoded using some lossy or lossless format. We recommend lossless formats for two reasons: (1) Information is packed more tightly as result of the quantization step, and there is not much room for compression; (2) JPEG tries to preserve visual features, but spectrum of the original image and the residual data can be very different, and therefore, unexpected artefacts can be observed in certain regions. The parameters of the inverse TMO and the quantization levels of the proposed scheme are encoded losslessly. For

decoding, the front and the residual layers are transformed using the inverse TMO and the nonlinear quantization levels, respectively, and added together to recover the HDR image. Denoting these transformations by τ and κ , respectively, the decoding process, based on Eq. (2), can be written mathematically as:

$$\begin{aligned} Y'_{\text{HDR}} &= \exp(\kappa(R') + \tau(\log(Y_{\text{LDR}}))) \\ &= \exp(\kappa(R')) \cdot \exp(\tau(f(\log(Y_{\text{LDR}})))) \end{aligned} \quad (4)$$

Due to the additional overhead of storing quantization levels and more tightly packed (and hence less compressible) data, the size of the encoded image can increase. However, the amount of additional details preserved as result of it is quite significant, as will be shown in the experimental results reported in the next section.

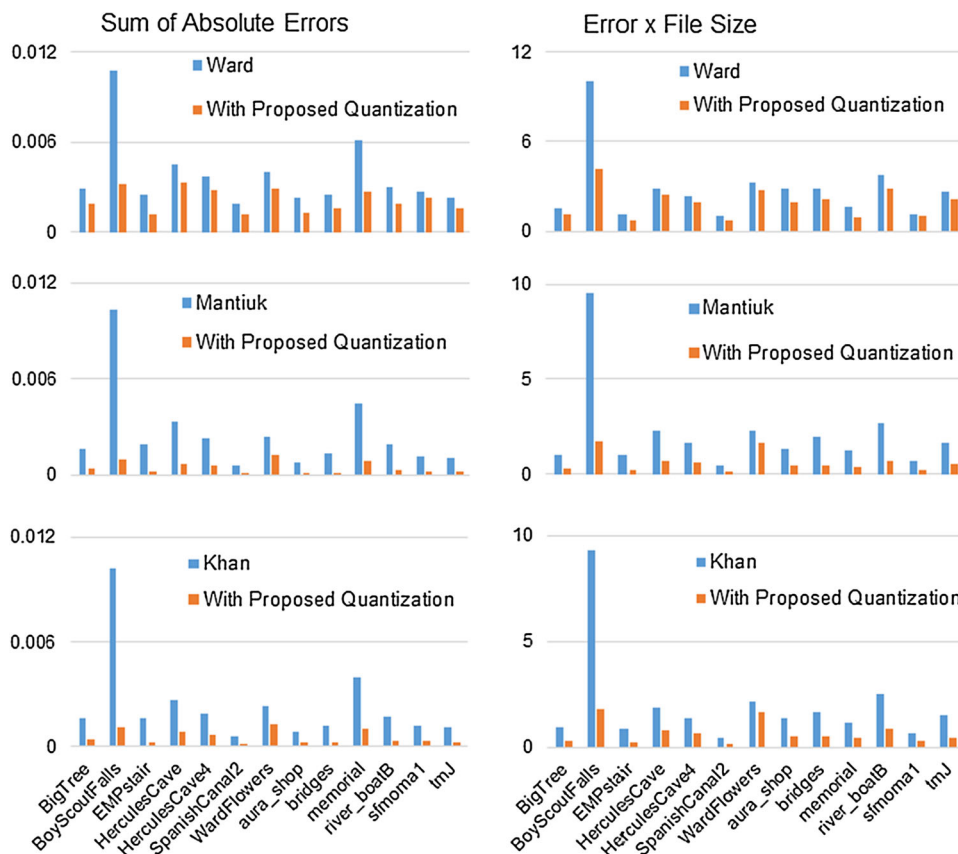
3 Experimental results

The presented nonlinear quantization scheme can work with any of the existing two-layer schemes and improve its coding performance by a better ‘packing’ of the residual data in the second layer. Here, we present some results to demonstrate the impact of our method on the performance of JPEG-HDR by Ward et al. [9], HDR-MPEG by Mantiuk et al. [10], and the method of Khan [14]. The photographic TMO by Reinhard et al. [22] is used to create the first layer, which is encoded as a JPEG image. The second layer is encoded in lossless PNG format for the reasons mentioned earlier, and also to avoid any artefacts of lossy encodings and to observe only the impact of quantization. For the same reason, we carried out our implementations of all algorithms without any post-processing steps. The bin centers (which are quantization levels for the second layer) and the inverse tone-mapping parameters (for [10] and [14]) are encoded losslessly using zlib deflate algorithm, which combines Huffman coding and LZ77 compression.

We carried out experiments using a large number of existing HDR images. The results shown here describe the average results obtained for 13 randomly selected images taken from the accompanying DVD of [23]. We compare the performance of all algorithms with and without using our quantization step using different metrics.

One of the simplest yet very effective metric that can be used for comparing the original image with the reconstructed is sum of absolute values of the individual pixel deviations. Figure 3 (left) shows the improvement in terms of this metric when the second layers of [9, 10] and [14] are quantized with the proposed method before encoding. It can be clearly seen that reduction in the error is quite significant for all three methods. On average, the reduction in error of [9, 10], and [14] was 45, 82, and 77%, respectively, when these

Fig. 3 Improvement in the performance of existing methods when they are implemented with the proposed quantization scheme. Two metrics, the sum of the absolute differences (*left*) and “error times size of the encoded data” (*right*), are used for comparison



methods were implemented with the proposed quantization step.

We calculated the above errors in the CIE LAB space as well, which is considered to have a closer match to the HVS compared to the RGB space. We converted original and the recovered images to the LAB color space and calculated the sum of individual pixel errors in the luminance channel. Reduction in errors were observed to be 39.6, 83, and 79% in this case for the three algorithms, respectively.

The quantization step proposed here is meant for packing the data more tightly in the given bit-depth, which leaves lesser room for further compression. For this reason, and due to additional overhead of saving quantization levels, the overall size of the encoded data increases. It can be an interesting way to look at the product of two parameters (size and error), to observe the performance of the encoding methods. We calculated this metric (sum of absolute differences times the total size of the encoded data) for the three methods considered in the study above and observed very significant improvement in performance when they were implemented with the proposed quantization step. Improvement was more significant for the methods of Mantiuk (71.3%) and Khan (65.9%) as compared to the improvement in the method of Ward (32.6%). Results for individual images are shown in Fig. 3 on right.

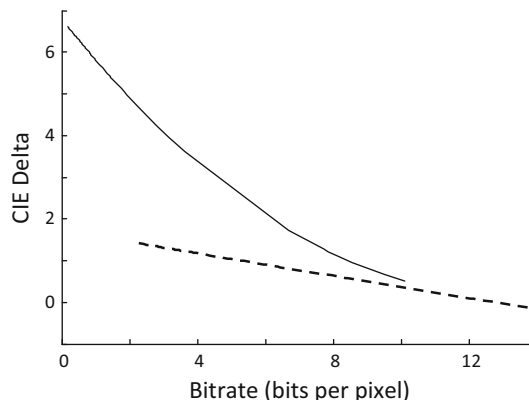


Fig. 4 Performance comparison of the JPEG-XT profile C with (*dotted line*) and without (*solid line*) the proposed quantization

JPEG-XT [20] is a new proposed standard for encoding images, which adds certain new capabilities to the legacy JPEG standard. JPEG-XT provides an option to encode HDR images, and the proposed profiles A, B, and C have publically available implementations. We added our nonlinear quantization step to the profile C and observed the effect on performance. In JPEG-XT, the second layer is encoded in lossy format. We used different quality settings for the front and back layers in our experiments to do encoding at dif-

ferent bitrates. Mean square errors in LAB space are shown at different bitrates in Fig. 4 for the profile C of JPEG-XT implemented with and without our proposed extension. Significant reduction in error can be observed at all bitrates.

4 Conclusion

We have presented a nonlinear quantization algorithm, which can be applied to the residual layer of any existing two-layer encoding schemes for better preservation of information. Experimental results have been presented to show the improvement achieved by the proposed add-on in the performance of several existing two-layer encoding methods. The additional preserved information is important for both visualization (on HDR displays) and other applications such as surveillance, security, and scene analysis.

Acknowledgments This work was funded by the Deanship of Scientific Research (DSR), King Abdulaziz University, Jeddah, under grant No. (611-398-D1435). The authors, therefore, acknowledge with thanks DSR technical and financial support.

References

- Larson, G.W.: Real Pixels in Graphics Gems II. In: Arvo, J. (ed.) Academic Press (1991)
- Larson, G.W.: LogLuv encoding for full-gamut, high-dynamic range images. *J. Graph. Tools* **3**(1), 15–31 (1998)
- Bogart, R., Kainz, F., Hess, D.: The OpenEXR File Format. SIGGRAPH (2003)
- <http://www.anyhere.com/gward/hdrenc/Encodings.pdf>. Accessed 5 Mar 2014
- Taubman, D.S., Marcellin, M.W.: JPEG 2000: Image Compression Fundamentals, Standards and Practice, Kluwer International Series in Engineering and Computer Science (2001)
- Xu, R., Pattanaik, S.N., Hughes, C.E.: High-dynamic-range still-image encoding in JPEG 2000. *IEEE Comput. Graph. Appl.* **25**(6), 57–64 (2005)
- Mantiuk, R., Krawczyk, G., Myszkowski, K., Siedel, H.P.: Perception-motivated high-dynamic range video encoding. *ACM Trans. Graphics* **23**(3), 733–741 (2004)
- Spaulding, K.E., Joshi, R.L., Woolfe, G.J.: Using a residual image formed from a clipped limited color gamut digital image to represent an extended color gamut digital image. United States Patent 6301393
- Ward, G., Simmons, M.: JPEG-HDR: a backwards-compatible, high dynamic range extension to JPEG. In: Proceedings of the Thirteenth Color Imaging Conference (Nov 2005)
- Mantiuk, R., Efremov, A., Myszkowski, K., Seidel, H.P.: Backward compatible high dynamic range MPEG video compression. In: SIGGRAPH (2006)
- Okuda, M., Adamai, N.: Two-layer coding algorithm for high dynamic range images based on luminance compensation. *J. Vis. Commun. Image Represent.* **18**(5), 377–386 (2007)
- Okuda, M., Adamai, N.: JPEG compatible raw image coding based on polynomial tone mapping model. *IEICE Trans. Fundam. Electron. Commun. Comput. Sci.* **E91-A**(10), 2928–2933 (2008)
- Khan, I.R.: Two layer scheme for encoding of high dynamic range images. In: IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) (2008)
- Khan, I.R.: HDR image encoding using reconstruction functions based on piecewise linear approximations. *Multimedia Systems*, published online (Nov 2014)
- Khan, I.R.: Effect of smooth inverse tone-mapping functions on performance of two-layer high dynamic range encoding schemes. *J. Electron. Imaging* **24**(1), 013024 (2015)
- Iwahashi, M., Kiya, H.: Two layer lossless coding of HDR images. In: IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 1340–1344, Canada (2013)
- Fujiki, T., Adami, N., Jino, T., Okuda, M.: High dynamic range image compression using base map coding. In: Asia-Pacific Signal & Information Processing Association Annual Summit and Conference (APSIPA ASC), (2012)
- Liu, J., Hassan, F., Carletta, J.: Embedding high dynamic range tone mapping in JPEG compression. In: Proc. SPIE 8655, Image Processing: Algorithms and Systems XI, 86550B (2013)
- Korshunov, P., Ebrahimi, T.: Context-dependent JPEG backward-compatible high-dynamic range image compression. *Opt. Eng.* **52**(10), 102006 (2013)
- Richter, T.: On the standardization of the JPEG XT image compression. In: Picture Coding Symposium (PCS), (2013)
- Khan, I.R.: A new quantization scheme for HDR two-layer encoding schemes. In: Poster SIGGRAPH (2014)
- Reinhard, E., Stark, M., Shirley, P., Ferwerda, J.: Photographic tone reproduction for digital images. In: Proceedings of the ACM SIGGRAPH 2002. *ACM Transactions on Graphics* **21**(3), 267–276 (2002)
- Reinhard, E., Ward, G., Pattanaik, S., Debevec, P.: High Dynamic Range Imaging' Acquisition, Display, and Image-based Lighting. Morgan Kaufmann, Los Altos (2005)