CGI2015 manuscript No.
(will be inserted by the editor)

Optimal Exposure Compression for High Dynamic Range Content

Kurt Debattista · Thomas Basfhord-Rogers · Elmedin Selmanovic ·
Ratnajit Mukherjee · Alan Chalmers

Abstract High dynamic range (HDR) imaging has become one of the foremost imaging methods capable of capturing and displaying the full range of lighting perceived by the human visual system in the real world. A number of HDR compression methods for both images and video have been developed to handle HDR data, but none of them has yet been adopted as the method of choice. In particular the backwards-compatible methods that always maintain a stream/image that allow part of the content to be viewed on conventional displays make use of tone mapping operators which were developed to view HDR images on traditional displays. There are a large number of tone mappers, none of which is considered the best as the images produced could be deemed subjective. This work presents an alternative to tone mapping based HDR content compression by identifying a single exposure that can reproduce the most information from the original HDR image. This single exposure can be adapted to fit within the bit depth of any traditional encoder. Any additional information that may be lost is stored as a residual. Results demonstrate quality is maintained as well, and better, than other traditional methods. Furthermore, the presented method is backwards-compatible, straightforward to implement, fast and does not require choosing tone mappers or settings.

1 Introduction

Traditional imaging techniques are incapable of accurately capturing or displaying the wide range of lighting that exists in the real world. The areas of the image outside the limited range, or Low (sometimes also termed Standard) Dynamic Range (LDR), of traditional cameras and displays, are either under or over exposed. High Dynamic Range (HDR) imaging technologies allow for the capture, storage, processing and delivery of a wider range of real-world lighting to provide an enhanced experience.

Unlike LDR data, in order to account for the wider dynamic range HDR delivers, uncompressed data is not stored as a byte per channel. Typical raw HDR images or frames would require the use of 32-bit floating point values per channel. This equates to 96-bits per pixel (bpp) when compared with the 24 bpp required by traditional LDR images. At an HD resolution of 1,920 × 1,080 this is approximately 24MB per frame. These sizes make raw HDR data difficult to manage and handle efficiently. A number of HDR compression methods do exist. Many of these are backwards-compatible, that is the content that is compressed can be played back by either an HDR display or by a traditional viewer for which only a certain aspect of the full range is displayed. The backwards-compatible methods employ a family of methods, known as tone mappers, to compress the dynamic range of HDR content into LDR. The tone mapped LDR content is then encoded using a traditional LDR encoder and, generally, a secondary stream is also compressed to account for aspects of the missing content. While the results of such methods are typically fairly satisfactory the use of tone mappers may complicate the process unnecessarily. There are a vast number of tone mappers in existence and it is unclear

---

Kurt Debattista, Thomas Bashford-Rogers, Ratnajit Mukherjee, Alan Chalmers
WMG
University of Warwick

Elmedin Selmanovic
University of Sarajevo

which of these methods are the best as a large number of evaluation methods appear to produce different results [1]. All the tone mappers have different settings so it is difficult for a non-expert user to set the correct parameters consistently. Related work also provides reason to believe that tone mappers are not necessarily superior to viewing a single exposure of an HDR image [2,3], have an affect on the compressed content [4] and also on visual attention [5].

In this paper a straightforward HDR compression method is proposed which extracts a single optimal exposure that is able to fit within a single stream that a traditional encoder can encode. The philosophy of this process is akin to what modern automatic cameras do and identify the best exposure to be taken without the user needing to struggle with complex settings available on more professional cameras. A secondary stream, if required, is used to store the residual. In order to avoid confusion, the extraction of the optimal exposure is termed Optimal Exposure Extraction (OEE) and the overall compression Optimal Compression (OC). OC is straightforward to implement, it is backwards-compatible, it is computationally fast, it produces similar quality (and sometimes better) results when compared with other backwards-compatible methods and is not burdened by the issues that arise from tone mappers in general.

2 Background and Related Work

A number of image formats have emerged to handle HDR images. These include the Radiance .hdr/.pic format [6] that requires 32 bpp, the OpenEXR format that can store full or half float for 96 bpp or 48 bpp respectively and the LogLUV format that supports 24 bpp and 32 bpp [7]. These formats are frequently compressed using lossless compression methods to achieve modest gains in terms of storage. However, such methods are still insufficient to handle HDR still images and video data efficiently.

Another aspect to consider about HDR imaging is that HDR content cannot be natively displayed on LDR displays. A single exposure of the HDR image can be displayed on an LDR display; or tone mapping operators can be applied to the HDR content to convert it to LDR content that is suitable to be viewed on a traditional LDR display [1]. There are a large number of tone mappers in existence, and they predominantly attempt to convey the same perceptual response to the viewer as the original HDR content. Tone mapping methods attempt to convert the pixels in the image such that the luminance range is compressed into one that is compatible with the viewing device, typically reducing the content to 8-bit depth. Global tone mappers do this by applying a fixed function to all the pixels in the image; local tone mappers adjust each individual pixel according to its neighbourhood. As no objective metrics exist to evaluate tone mappers a large number of subjective methods have been presented but no consensus has yet been reached on the best possible methods and practice. An alternative has been proposed [8]. This method attempts to automatically select an exposure of an HDR image with a goal and algorithm not dissimilar to OEE. However, this work was targeted at computer graphics imagery, worked only on a sample of the pixels, was based on computing histograms from multiple colour channels and was not dedicated or applied to compression.

HDR compression methods for both still images and video can be broadly divided into two categories, those that are backwards-compatible and those that are not. The backwards-compatible methods produce a format which can be, partially, directly viewed by a traditional LDR viewer without any modifications to the legacy software. The content that an LDR player displays for the backwards-compatible method is an LDR stream (or image) which is sub-part of the full stream (or image). Alternately, if a specialised player is available, the HDR content can be extracted; typically, by inverting the tone mapping process and using information imbedded in the format in addition to the video stream. The non-backwards compatible methods cannot be displayed with existing LDR viewers and instead use proprietary viewers to display the HDR content on either an LDR or HDR display. The method proposed here is a backwards-compatible method.

2.1 Still Images

There are a number of backwards-compatible lossy methods for compressing still HDR images. Ward and Simmons [9] method for JPEG compression of HDR images, tone maps the HDR image, and creates a ratio image by dividing the original HDR by the tone mapped image. The ratio image is subsequently divided with the original HDR image to produce a modified tone mapped image. The modified tone mapped image is compressed using a traditional JPEG encoder and the ratio image is compressed to 64KB to fit within the JPEG subband.

Okuda and Adami [10] presented a backwards-compatible HDR encoding method in which the original HDR image is encoded with a tone mapper that uses a sigmoid with parameters identified via an optimisation function. Residuals are computed from the reconstructed tone mapped image and stored using wavelets. The decoding process reverses the sigmoid computation based on the identified parameters and re-combines this with the residuals.

Xu, Pattanaik and Hughes [11] presented a non-backwards compatible HDR method that used JPEG2000. The method involves using the native 16 bit per channel storage facility of JPEG2000. In order to make use of this, the method converts the floating point values of the original HDR image to 16 bit integers. The results are encoded using the JPEG2000 except for the wavelet domain sub-band quantisation where the perception-related factor is omitted because HDR images were scene referred as opposed to being display referred.

2.2 HDR Video Compression

The backwards-compatible methods for video include Mantiuk et al. [12] and Lee and Kim [13]. These two methods follow the same overall method proposed by Ward and Simmons [9] for still images. Mantiuk et al. (2006) tone map the image, restore the tone mapped frame backwards to a compatible colour space to compare with the original, generate a mapping from LDR to HDR and compute a residual representing differences between the reconstructed HDR and original HDR in terms of luminance. The tone mapped frame and the residual frame are temporally compressed using LDR video compression methods and a reconstruction function is also compressed using lossless encoding. The decoding involves reconstructing the HDR from the tone mapped version using the reconstruction function and re-combining with the residual.

Lee and Kim's method [13] follows a similar method to Ward and Simmons [9]. The HDR frames are tone mapped using a temporally coherent tone mapper that extends the gradient domain tone mapping operator. A residual is constructed by computing the logarithm of the division of the original HDR luminance by the decoded LDR luminance. This stream is cross-bilaterally filtered [14] with the original to reduce noise. The tone mapped and residual stream are encoded separately. To reduce distortions of both the TM and reconstructed HDR sequences the quantisation parameters of the LDR and ratio sequences was controlled. They are then reconstructed by decoding the two streams and re-combining them.

A number of non-backwards compatible methods for HDR video compression have also been proposed. Mantiuk, Krawczyk, Myszkowski and Seidel [15] suggested an early method for compressing HDR videos. They mofidied the capabilities of the MPEG-4 video codec and extended it to work with HDR video data. The main characteristic of the proposed algorithm is quantisation of luminance where errors were kept below the just noticeable threshold values of the human visual system. To facilitate HDR data, MPEG-4 data structures were expanded from 8 to 11 bits and an efficient coding scheme for DCT blocks was introduced.

Adaptive bit-depth transformation of HDR data was explored by Motra and Thoma [16] and Zhang et al. [17]. Motra and Thoma [16] transformed HDR images to LogLuv format, which they have optimised for 16 bit floating point numbers. Then quantisation errors were minimised by adaptively utilising levels which were left unused after transformation. Zhang et al. [17] extended the method by optimising bit-depth quantisation via the Lloyd-Max algorithm. In addition invisible high frequency noise was reduced by transforming frames into the wavelet domain where a contrast sensitivity function weighted wavelet subbands.

3 Optimal Exposure Compression

The backwards-compatible methods use some form of tone mapping to compress the luminance range of an HDR stream or still image to LDR before encoding it. This enables the encoded still-image/stream to be backwards compatible and it makes it possible to use legacy viewers. However, tone mapping can result in different types of artefacts, requires a choice of tone mapper and an understanding of the settings. The presented method proposes an alternative technique towards backwards-compatible HDR compression. Instead of relying on tone mapping to produce an LDR image/frame, the best continuous luminance range is extracted, and this is augmented with additional information to store and eventually reconstruct the original HDR image/frame.

3.1 Motivation

A single exposure of the HDR image, without tone mapping, presents the user with a more readily understandable image when this is viewed on an LDR display; tone mapped images not processed properly may be considered unrealistic by individuals used to seeing traditional images consisting of single exposures. Strgar Kurečić et al. [3] surveyed 100 amateur and professional photographers on tone mapped images and only 4% of them considered such images realistic compared to a single exposure image and, the majority (*ca.*70%), found them to be artificial, unreal or exaggerated. Although there are many different types of tone mappers there is no consensus on which the best one is; a number of evaluation studies have been conducted and they differ on the results as can be seen in Banterle et al.'s overview of evaluation methods [1]. While there is little doubt that with the right settings and right tone mapper an expert can create realistic tone mapped images, the vast plethora of tone mappers and settings is definitely an issue for the general user; furthermore, there is no guarantee that even the best tone mappers consistently produce better results than a single exposure. Akyuz et al. [2] found no significant difference between single exposed images and tone mapped images in an experiment that compared tone mapped images, single exposure images and HDR images. Similarly, a recent study [18] of 38 participants found no significant difference among tone mapped images and a single exposure, chosen as the zero exposure, when compared with the reference HDR; this is not necessarily the best exposure to show, yet the participants still did not prefer the tone mapped images over the single exposure. There is also evidence that tone mapped images can change the visual attention of an image [19] and of the compressed content [4]. Furthermore, different tone mappers can perform better on different images/frames or even on different parts of the same image / frame as demonstrated by Banterle et al. [20] whereby hybrid TMOs outperformed dedicated TMOs in participant experiments. The choice of tone mappers and the setting of the individual parameters for any given tone mapper is thus quite a difficult task for non-experts. A correctly chosen single exposure corresponds to the type of images users expect to see from an imaging system and avoids the artefacts common to tone mapping algorithms.

The proposed method avoids the problems with tone mapping by extracting a single exposure designed to fit the size of the encoder. The size of the extracted range equates to the bit-depth supported by a given encoder. Typically this will be 8-bit for most encoders but support for other 10-,12-,14-bit profiles do exist and the method natively adapts to be able to support these profiles. If the HDR content is not a very high dynamic range, the residuals would be very small, so the size of the final compressed image/video would be relatively small too.

3.2 Method

Figure 1 illustrates the encoding method and Figure 2 the decoding method that define OC. The encoding process commences by identifying a single exposure of the HDR image. The single exposure is computed by selecting the contiguous area of luminance to fit within a required bit depth; this is typically the bit depth permissable by the LDR encoder. In the method presented here and in the subsequent section, the logarithm of the luminance rather than the luminance itself, is used to conform with the human visual system's response to luminance. Thus, it is the log encoded largest contiguous area of luminance which fits in a single exposure or to occupy the bit-depth of the encoder that is identified. Once the single exposure is chosen it is compared to the original HDR and residuals are computed.

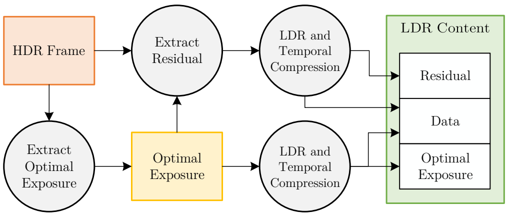

Fig. 1 OC encoding process.

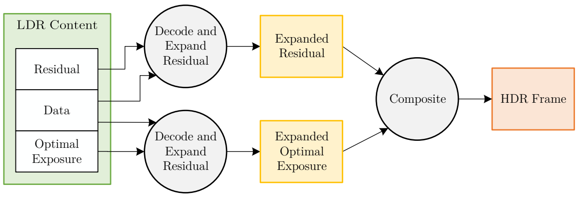

Fig. 2 OC decoding process.

The contiguous area of luminance is computed by maximising the luminance for a number of pixels such that the luminance fits within the encoders bit depth:

$$max_E(f(I(E))) \tag{1}$$

where the function $f()$ counts the number of well exposed pixels in an HDR image I at exposure E. The function $f()$ is defined as follows:

$$f(I(E)) = \sum_{p\ \epsilon\ pixels} \begin{cases} 1 & \text{if,}(2^{BD}-1) \times I_p(E)\ \epsilon\ [a \dots b] \\ 0 & \text{otherwise} \end{cases} .$$

This calculates for each pixel in the image, p, if the pixel value at the current exposure $I_p(E)$ scaled by the bit depth BD of the encoder is within a predetermined acceptable range $[a \dots b]$. Section 3.3 presents an algorithm for computing Equation 1.

For still images, the chosen single exposure is compressed via a traditional LDR encoder (for example, but not limited to, JPEG) and will constitute the body of the file. For video streams, the chosen single exposure is encoded via a traditional LDR encoder (for example, but not limited to, h.264).

After computing the optimal exposure frame, an HDR frame is reconstructed from the optimal exposure frame and the difference between the original and the reconstructed HDR constitute the residuals. The residuals are stored in another channel, or in a subband in the case of images, after quantisation and compression. The residuals are stored in a single context for images and a single stream for video. Alternatively, the residuals may also be stored in two separate sets, representing the higher dynamic range and the lower dynamic range. Values in the higher dynamic range can be quantised more aggressively due to the human visual systems ability to notice changes in luminance at lower values more than at higher values. For the case of this paper we use one residual which is log encoded to account for such characteristics of the human visual system. The data for the chosen exposure, as well as any other practical arguments required for reconstruction are stored as part of the header or a separate stream. The choice of the single exposure takes temporal data into account via temporal filtering of the chosen exposure locations, to ensure the encoded LDR stream does not contain sudden jumps in luminance or flickering.

The decoding procedure on a traditional LDR viewer will show only the single exposure image that has been stored in the encoded still image/stream (see Figure 2). When viewed on a specialised HDR viewer, the single exposure is scaled back up to the original values and the residuals are composited back onto the image.

3.3 An algorithm for computing Optimal Exposure Extraction

Algorithm 1 Optimal $(hdr, bitDepth)$

```
 1: $hdrLum \leftarrow \mathrm{lum}(hdr)$
 2: $DR \leftarrow \log_2(\frac{\max(hdrLum)}{\min(hdrLum(\text{ for } hdrLum>0))})$
 3: $FD \leftarrow \frac{2\times \mathrm{IQR}(\log_2(hdrLum))}{\sqrt[3]{n}}$
 4: $binNo \leftarrow \frac{\max(\log_2(hdrLum))-\min(\log_2(hdrLum))}{FD}$
 5: $hst \leftarrow \mathrm{hist}(\log_2(hdrLum), binNo)$
 6: $best \leftarrow 0$
 7: $bstBnd \leftarrow 1$
 8: $step \leftarrow round(\frac{binNo}{DR\times bitDepth})$
 9: for $i = 1$ to $binNo - step$ do
10:     $count = \mathrm{sum}(hst(i : i + step))$
11:     if $count > best$ then
12:         $best \leftarrow count$
13:         $bstBnd \leftarrow i$
14:     end if
15: end for
16: $lb \leftarrow 2^{\frac{bstBnd\times DR}{binNo}} + \log_2(\min(hdrLum))$
17: return $lb$
```

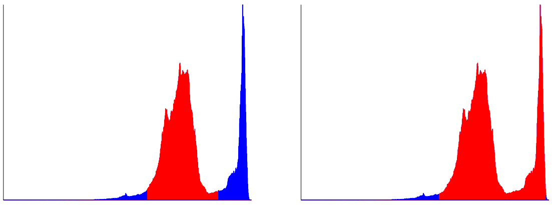

(a) 8 bit (b) 12 bit

Fig. 3 Histogram showing luminance for basement scene shown for different bit depths. Red bars mark the area chosen by the proposed method.

Algorithm 1 provides an overview of the method used to solve Equation 1, although other solutions are possible. The HDR image (*hdr*) and the bit depth (*bitDepth*) are input as parameters. The luminance for the HDR image is first computed (*hdrLum*) as is the dynamic range (*DR*). The number of bins to construct the histogram (*binNo*) is calculated by means of the FreedmanDiaconis rule [21] (see Line 3) which is a relatively robust method suitable for finding the size of a number of bins in the histogram, where IQR() computes the interquartile range. Starting at the first bin, the value of all the bins within a given range is checked. This value then represents the current maximum and is stored (*best*). The process then cycles through all the bins doing the same thing (calculating total overall luminance in that range) and checking if the new value is greater than the stored maximum. If it is it becomes the new maximum. The point in the bin representing the minimum luminance *lb* is stored. This value is sufficient to be able to identify the range of luminance that the optimal compression method will make use of. The end of the range representing the maximum luminance of the chosen range could also be stored, or could be calculated later by taking into account *lb* and the bit depth. The presented algorithm is linear in the number of bins chosen. Once the luminance range is chosen, the backwards-compatible optimal channel is reconstructed per channel. Figure 3 shows how the luminance of the Basement image (Figure 5.d) for four different bit depths. The red bars demonstrate the selected luminance range for the given bit depth. Note that Algorithm 1 is presented with readability in mind and optimisations should be applied in a given implementation.

4 Results

Results are presented for OC when compared to a number of other backwards compatible methods introduced in Section 2. All the presented methods consist of dual streams, one which encodes the backwards compatible

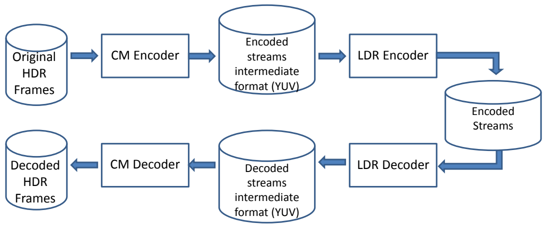

Fig. 4 Pipeline used for comparing results across compression methods (CMs).

stream and another secondary stream. In addition they usually maintain some extra per scene or per frame parameters in a third stream; this amount of data is of the order of a few bytes per frame and can be considered negligible. The backwards compatible methods chosen for comparison were: the rate-distortion method [13] (Rate Distortion), HDR MPEG [22] (HDR MPEG), a video version of the JPEG HDR method [9] (Ward) and a straightforward inverse tone mapping method (Inverse). The choice of methods of HDR MPEG, Ward and Rate Distortion represent state of the art backwards compatible methods; Ward was originally proposed for static images, however it produced good results when modified for video. This version performs better than the other backwards compatible methods and thus an alternative version of Ward was added to the results. The alternative version (Ward (OEE)) was selected to use OEE instead of its traditional tone mapping technique to have another method which is not just using tone mapping. The motivation behind this is to show that OC gives good comparable results to more complex methods that may also be using OEE; effectively this serves highlight that the contribution of this work is in the combination of OC via OEE and not just OEE. Inverse, like OC, is a relatively straightforward method, and was chosen to broadly represent the methods that tone map and then inverse tone map the HDR content and to act as an alternative to optimal's choice. Implementation details of the methods are presented below.

4.1 Results' Method

The method of testing results for video is outlined in Figure 4. The chosen sequences of 120 frames each compressed at 24 frames per second are shown in Figure 5, Each method compresses a series of individual HDR frames using the proposed method resulting into two YUV blocks of frames which are then compressed by a traditional encoder at various bit rates. They are then decoded by the traditional encoder's decoding counterpart to produce decoded YUV blocks. The resultant YUV blocks are then decoded using the appropriate method to produce a new set of HDR frames. The original and output HDR frames are compared for each frame across all bit rates and for all scenes using PSNR across all channels.

Results are presented across a wide range of bit rates using two different LDR encoders: h.264 representing current technology and the upcoming standard h.265/HEVC. For all methods the same bit rate is used for both the backwards compatible and residual streams. Bit rates were computed by changing quantisation parameters for settings of $1, 2, 5, 10, 15, 20, 25, 30, 35$. x264 and x265 were used for compression with parameters set to the default maximum quality compression preset settings (very slow).

4.2 Compression Methods

All compression methods were implemented from the source materials by ourselves in Matlab and used the exact same framework for undertaking results. OC is present in two guises which have the optimal backwards-compatible stream compressed at the traditional 8-bit and 10-bit referred to OC and OC (10-bit) respectively. For Rate Distortion the original video tone mapper employed in their article was used [23]. The optimisation for bit rate originating from this method is not employed to maintain equality in the encoding of all methods since the actual encoder and decoder was being tested. A fixed bit rate for both streams was used; in addition the benefits of the bit rate optimisation used in rate distortion could potentially be employed by all methods. For HDR MPEG and Ward, Reinhard's TMO was used [24] due to it frequently doing well in comparison tests [1]. Inverse also made use of the Reinhard TMO and at the decoding stage made use of the inverse Reinhard tone mapper [25,26]. The Ward method is represented with two sets of results. In the more traditional version (Ward) the Reinhard tone mapper was used. For Ward (OEE) the tone mapping part is replaced with OEE. Secondary streams for Ward, Ward(OEE) and Inverse are encoded using log encoding to maintain similarity with OC as no specific solutions for storing secondary streams were given in the original work. The use of log encoding was found to produce better results than alternative methods of storing the residuals.

4.3 Quantitative Results

Results for all the methods across all eight scenes are shown in Figure 6 for h.264 and Figure 7 for the

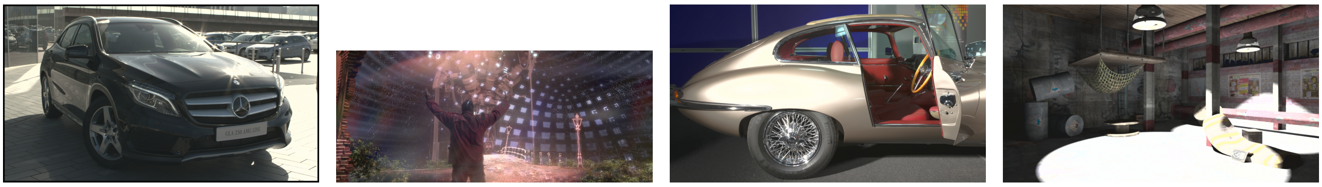

(a) Mercedes; DR = 4.28 (b) Tears; DR = 5.54 (c) Jag; DR = 5.35 (d) Base; DR = 8.37

(e) Machine; DR = 5.62 (f) Seine; DR = 6.30 (g) Disco; DR = 5.52 (h) Weld; DR = 6.48

Fig. 5 Scenes used for video results. Dynamic range (DR) calculated as $log_{10}(\frac{max-min}{min})$.

h.265/HEVC results. Results demonstrate that OC and OC (10-bit) outperform the other methods. The Ward method also does well, and in particular Ward (OEE) produces results that further demonstrate how useful the use of computing the optimal can be; they also demonstrate that there is no need of complex encoding and decoding methods to take advantage of OEE, as the straightforward method OC competes favourably with Ward (OEE). Rate Distortion also does well for most scenes. We expect HDR MPEG to perform better if the data was calibrated to real world luminance.

4.4 Still Image Results

The method can also be applied directly to compress still images using JPEG compression. Following Ward and Simmons' JPEG HDR method [9] the backwards compatible component constitutes the image part of the JPEG format and the residual is stored in the subband. Results, see Figure 9, are presented for OC, versus the JPEG HDR method using Reinhard's TMO (JPEG HDR) and the JPEG HDR method using OEE (JPEG HDR (OEE)). Results represent PSNR at various compression qualities controlled by the traditional JPEG quality parameter for values of $1, 4, 16, 32, 48, 64, 80, 96$ for encoded then decoded HDR images compared against the original uncompressed HDR image. 20 HDR images, shown in Figure 8 were used to obtain the results; OEE was used to choose the exposure shown for each image. The PSNR results are averaged across all images. As with the results for videos OC performs best, closely followed by JPEG HDR (OEE) and then JPEG HDR.

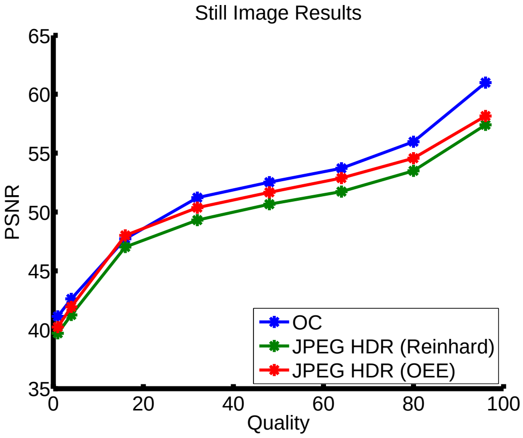

Fig. 9 JPEG Results. Quality represents different JPEG quality encodings values. Results averaged over 20 HDR images shown in Figure 8.

5 Discussion and Limitations

The presented results demonstrate the usefulness of OC and OEE. In the presented results, it is OC that leads in terms of quality and followed closely by Ward (OEE) (this also applies to JPEG HDR (OEE) which is the still image version). This serves to highlight the effectiveness of OEE but also demonstrate that by itself it is insufficient and the straightforward OC implementation combines well to provide better results than when using OEE with other more established methods.

In terms of computation OC is one of the fastest. Rate distortion requires the use of a bilateral filter and a complex video tone mapping operator. HDR MPEG requires relatively computationally expensive wavelet computations. Ward may not be as complicated but as with the others relies on the use of tone mappers which may be quite computationally expensive too, although it does provide flexibility in the choice of tone mappers

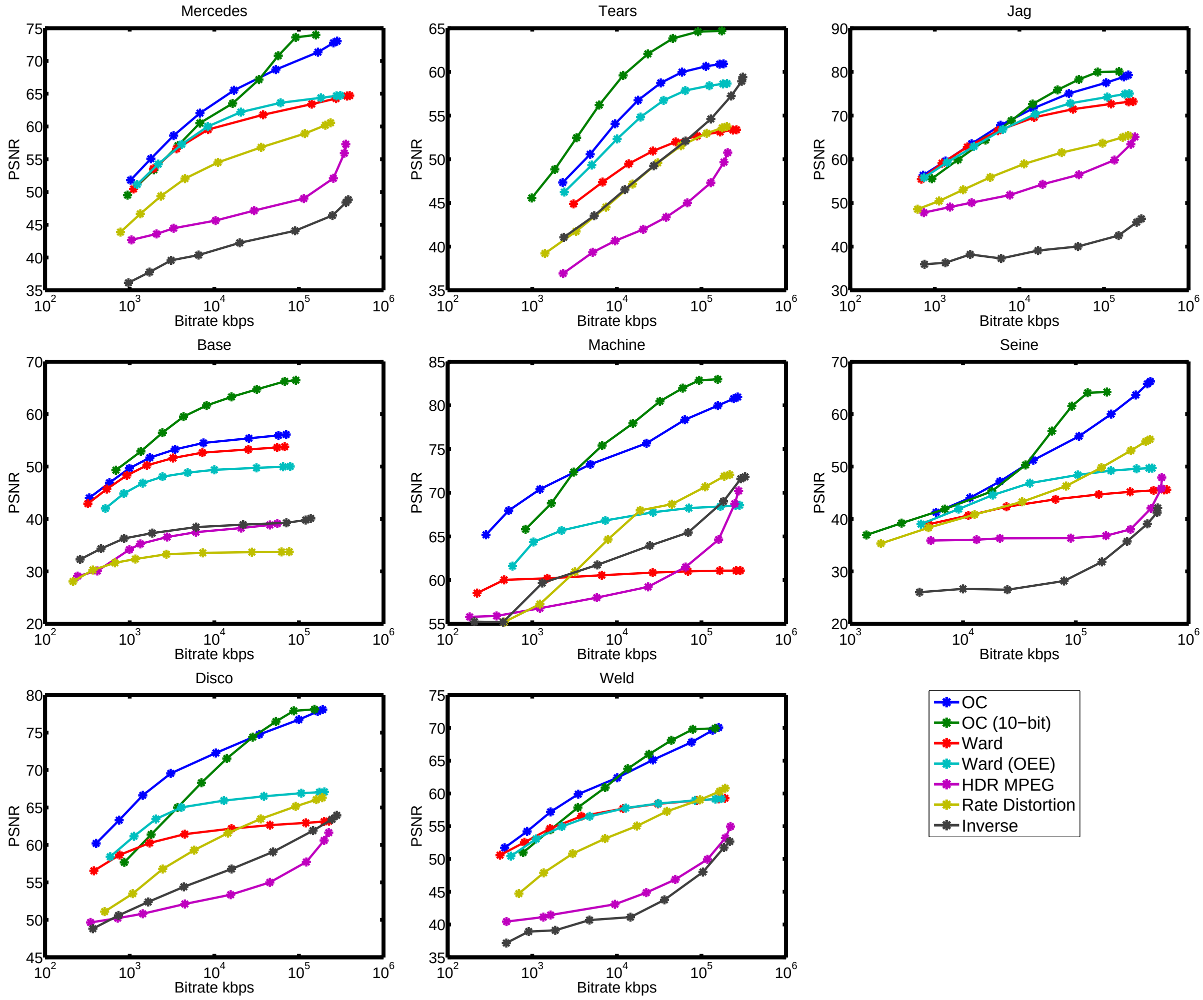

Fig. 6 h.264 Results.

which can be seen as a positive compared to OC and most other methods. Inverse can be quite straightforward but cannot make use of all tone mappers, only those that have an inverse; moreover, it did not perform too well. It must be remembered that computation speed plays a fundamental role in activities such as real life transmission of HDR content which is fundamental if HDR is to play a role in broadcast media.

A limitation of the results presented here may be attributed to the lack of comparison with other tone mapping methods. The results were intended to be a comprehensive comparison of backwards-compatible methods and in this regard this has been achieved and demonstrated; using different tone mapping methods would lead to a significant increase in the number of tests which would have made achieving such comprehensive results impractical. Secondly, the choice of the vast number of tone mappers and possible settings would have been a very complex issue. Importantly, this is the exact point why OC was developed. Furthermore, the chosen tone mapper (Reinhard) is considered one of the best and is one of the most popular. If other tone mappers were chosen and produced slightly superior results the same issues with such tone mapping based methods remain. Also, certain other tone mappers can be significantly more complex even in terms of computation. OC provides an alternative which is straightforward, fast and maintains good quality.

6 Conclusions

The work presented here differs from the current standard use of tone mappers as the basis of backwards-compatible HDR image and video encoding methods. The proposed method does not require the choice of tone mappers, and works automatically for any scene without any complex settings. The method is straightforward to implement and is relatively fast at both the

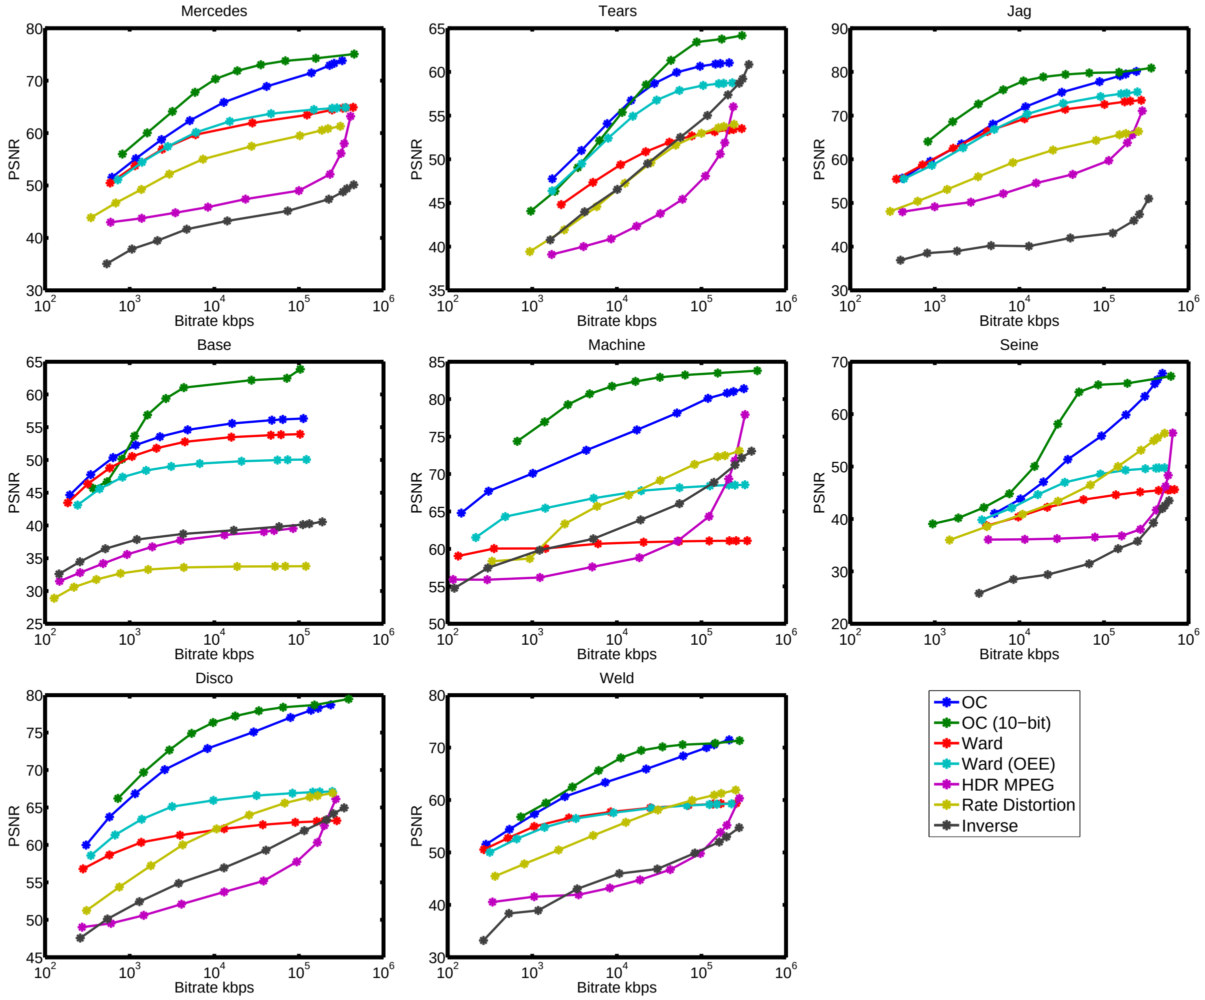

Fig. 7 h.265 Results.

encoding and decoding stage. While tone mappers do provide a very important role for HDR imaging their use in compression is useful but, on occasions, it can be cumbersome and problematic. We hope OC can be considered as an alternative to the status quo of current HDR compression methods.

Acknowledgements

Debattista is partially supported by a Royal Society Industrial Fellowship.

References

1. F. Banterle, A. Artusi, K. Debattista, A. Chalmers, Advanced High Dynamic Range Imaging: Theory and Practice, A K Peters/CRC Press, 2011.
2. A. O. Akyüz, E. Reinhard, Noise reduction in high dynamic range imaging, Journal of Visual Communication and Image Representation 18 (5) (2007) 366–376.
3. M. Strgar Kurečić, A. Poljičak, L. Mandić, A survey on the acceptance and the use of hdr photography among croatian photographers, Acta Graphica znanstveni časopis za tiskarstvo i grafičke komunikacije 24 (1-2) (2013) 13–18.
4. M. Narwaria, M. P. Da Silva, P. Le Callet, R. Pepion, et al., Impact of tone mapping in high dynamic range image compression, Procceding of VPQM 2014.
5. M. Narwaria, M. Perreira Da Silva, P. Le Callet, R. Pepion, Tone mapping based hdr compression: Does it affect visual experience?, Signal Processing: Image Communication.
6. G. Ward, A contrast-based scalefactor for luminance display, in: Graphics gems IV, Academic Press Professional, Inc., 1994, pp. 415–421.
7. G. Ward, LogLuv Encoding for Full-Gamut, High-Dynamic Range Images, Journal of Graphics Tools 3 (1) (1998) 15–31.
8. L. Neumann, K. Matkovic, W. Purgathofer, Automatic exposure in computer graphics based on the minimum information loss principle, in: Computer Graphics International, 1998. Proceedings, IEEE, 1998, pp. 666–677.
9. G. Ward, M. Simmons, Subband encoding of high dynamic range imagery, in: Proceedings of the 1st Symposium on Applied perception in graphics and visualization

Fig. 8 Images used for still image results.

- APGV '04, ACM Press, New York, New York, USA, 2004, p. 83.

10. M. Okuda, N. Adami, Two-layer coding algorithm for high dynamic range images based on luminance compensation, Journal of Visual Communication and Image Representation 18 (5) (2007) 377–386.
11. R. Xu, S. N. Pattanaik, C. E. Hughes, High-dynamic-range still-image encoding in jpeg 2000, Computer Graphics and Applications, IEEE 25 (6) (2005) 57–64.
12. R. Mantiuk, K. Myszkowski, H.-P. Seidel, Lossy compression of high dynamic range images and video, in: Electronic Imaging 2006, 2006, pp. 60570V–60570V–10.
13. C. Lee, C. Kim, Rate-distortion optimized compression of high dynamic range videos, in: Proceedings of the 16th European Signal Processing Conference (EUSIPCO 2008), 2008.
14. E. Eisemann, F. Durand, Flash photography enhancement via intrinsic relighting, in: ACM transactions on graphics (TOG), Vol. 23, ACM, 2004, pp. 673–678.
15. R. Mantiuk, G. Krawczyk, K. Myszkowski, H.-P. Seidel, Perception-motivated high dynamic range video encoding, ACM Transactions on Graphics 23 (3) (2004) 733.
16. A. Motra, H. Thoma, An adaptive Logluv transform for High Dynamic Range video compression, in: 2010 IEEE International Conference on Image Processing, 2010, pp. 2061–2064.
17. Y. Zhang, E. Reinhard, D. Bull, Perception-based high dynamic range video compression with optimal bit-depth transformation, in: 8th IEEE International Conference on Image Processing, 2011, pp. 1321–1324.
18. M. Narwaria, M. P. D. Silva, P. Le Callet, R. Pepion, Single exposure vs tone mapped high dynamic range images: A study based on quality of experience, in: Signal Processing Conference (EUSIPCO), 2013 Proceedings of the 22nd European, IEEE, 2014, pp. 2140–2144.
19. M. Narwaria, M. P. Da Silva, P. Le Callet, R. Pepion, Effect of tone mapping operators on visual attention deployment, in: SPIE Optical Engineering+ Applications, International Society for Optics and Photonics, 2012, pp. 84990G–84990G.
20. F. Banterle, A. Artusi, E. Sikudova, T. Bashford-Rogers, P. Ledda, M. Bloj, A. Chalmers, Dynamic range compression by differential zone mapping based on psychophysical experiments, in: Proceedings of the ACM Symposium on Applied Perception, ACM, 2012, pp. 39–46.
21. D. Freedman, P. Diaconis, On the histogram as a density estimator: L 2 theory, Probability theory and related fields 57 (4) (1981) 453–476.
22. R. Mantiuk, A. Efremov, K. Myszkowski, Design and Evaluation of Backward Compatible High Dynamic Range Video Compression, Tech. Rep. April (2006).
23. C. Lee, C.-S. Kim, Gradient Domain Tone Mapping of High Dynamic Range Videos, in: 2007 IEEE International Conference on Image Processing, IEEE, 2007, pp. III – 461–III – 464.
24. E. Reinhard, M. Stark, P. Shirley, J. Ferwerda, Photographic tone reproduction for digital images, ACM Transactions on Graphics 21 (3) (2002) 267–276.
25. F. Banterle, P. Ledda, K. Debattista, A. Chalmers, Inverse tone mapping, Proceedings of the 4th international conference on Computer graphics and interactive techniques in Australasia and Southeast Asia - GRAPHITE '06 (2006) 349.
26. F. Banterle, P. Ledda, K. Debattista, A. Chalmers, M. Bloj, A framework for inverse tone mapping, The Visual Computer 23 (7) (2007) 467–478.