Compression Efficiency of HDR/LDR Content

Maryam Azimi, Ronan Boitard, Basak Oztas, Stelios Ploumis, Hamid Reza Tohidypour
Electrical and Computer Engineering Department,
University of British Columbia
Vancouver, Canada
maryama@ece.ubc.ca

Mahsa T. Pourazad
TELUS Communications Inc.,
The Institute for Computing, Information and Cognitive Systems (ICICS),
University of British Columbia
Vancouver, Canada
pourazad@icics.ubc.ca

Panos Nasiopoulos
Electrical and Computer Engineering Department, the Institute for Computing, Information and Cognitive Systems (ICICS),
University of British Columbia
Vancouver, Canada
panos@ece.ubc.ca

***Abstract*— High Dynamic Range (HDR) imaging is capable of delivering a wider range of luminance and color gamut compared to Standard Dynamic Range (SDR), offering to viewers a visual quality of experience close to that of real-life. In this study, we evaluate the quality of coded original HDR streams and HDR streams reconstructed from SDR videos and metadata, both compressed by the HEVC standard. Our evaluations have shown that the single HDR approach is largely preferred over the SDR counterpart.**

Keywords*— *HDR, SDR, dynamic range, video quality assessment, subjective evaluation, video compression, HEVC, video quality metrics

I. Introduction

High Dynamic Range (HDR) imaging techniques can capture and reproduce a wide range of luminance values, far beyond the capabilities of Standard Dynamic Range (SDR) imagery [1]. In fact, by combining different exposures of a recorded scene [2], HDR acquisition techniques can capture more than what the Human Visual System (HVS) can perceive. Furthermore, as HDR content stores the physical light intensities (cd/m^2) of a scene (in floating point values), HDR displays can achieve a one-to-one relationship between the displayed luminance and the recorded scene. Consequently, HDR content is usually considered more natural and of better quality than its SDR counterpart.

Following the increasing interest of broadcasting companies in HDR technology, the Motion Picture Experts Group (MPEG) has recently issued a Call for Evidence (CfE) [3] to assess the feasibility of broadcasting HDR content to the end-user. The main goal of this CfE is to assess the impact of encoding HDR content in terms of bit-rate overhead and compression pipeline changes. Indeed, existing video compression standards such as the ITU-T H.265/MPEG-H Part 2 High Efficiency Video Codec (HEVC) [4][5] cannot directly compress HDR videos since they rely on pixels represented by integer code values. Consequently, the traditional distribution pipeline will need to be updated. Over the last decade, several solutions have been proposed to encode HDR content [6]. However, these solutions have a relatively large impact on the distribution architecture.

In this paper, we study and compare the performance of two HDR encoding schemes that have no impact on this architecture: Scheme I: Perceptual encoding: conversion of HDR video content from floating point values to integer pixel values through perceptual encoding [7] [8] and compression of the content based on HEVC standard; Scheme II: tone mapping the HDR content and generating metadata, compression of the tone mapped version using HEVC, and then inverse tone mapping the decoded video using the metadata to reconstruct the HDR stream. These two schemes only require pre and post distribution processing, and are defined as possible HDR encoding schemes in the recent MPEG CfE for HDR [3]. The ideal way of comparing the performance of these two schemes, is through subjective tests. On the other hand, subjective evaluations are not always available or possible to obtain, depending on the application. For this reason, video objective quality metrics are employed to simulate the observers' opinion. However for HDR video content, there is no standardized objective quality metric and the performance of the existing quality metrics are still under investigation [9][10]. Very limited attention has been given to assessing the relationship between HDR objective metrics (HDR-VDP-2 [11], PU-PSNR [12], etc.) and subjective studies. In our study we compare the performance of two HDR encoding schemes both subjectively and objectively. We perform a comparative study to evaluate the correlation between objective metrics and a subjective evaluation of compressed HDR content.

The rest of this paper is organized as follows: Section II provides an overview of schemes for compressing HDR video streams, Section III describes the test procedure, while Section IV presents the reults and discusses the findings. Section V concludes the paper.

II. HDR Video Compression Schemes

In our study we compare the performance of two different schemes for HDR video compression. The following subsections elaborate on these two schemes.

A. Scheme I: broadcast perceptually encoded HDR content

The existing video compression standards are not directly applicable to HDR videos as pixels of HDR content represent the physical scene luminance (expressed in cd/m^2) and are stored as floating point values. This representation lacks efficiency in storing, broadcasting and processing. Conversion of floating point physical values to integer values is done through perceptual encoding. Perceptual encoding removes information that would be invisible after decoding (visual noise) and optimizes bit-depth to minimize the visual loss due to the quantization introduced by the codec. Perceptual encoding is particularly useful for storage and distribution applications, as it effectively reduces the required bit-depth without loss of visual quality. One of the first perceptually uniform coding methods of HDR luminance was derived from the threshold versus intensity (t.v.i.) models [7]. The derivation involves rescaling luminance values so that the difference in code values corresponds to the detection threshold throughout the entire encoded range of luminance. Once the HDR content is perceptually coded, it is directly fed to a standard video codec. In our study we follow the workflow suggested by MPEG's CfE (see Fig. 1). First the HDR video data are perceptually encoded using the Perceptual Quantizer (PQ) [8] The video is then converted to the Y'C_bC_r color difference space with a 10-bit quantization before downsampling the chroma channels to 4:2:0 and compression. This data flow is inverted at the decoding stage. Note that by adding metadata to the bitstream, one can tone map the decoded HDR content to address SDR displays [13].

The codec used in this workflow is HEVC (Main 10 profile). HEVC is the most recent video compression standard, which can achieve 50% compression efficiency over its predecessor H.264/AVC [14]. Currently, MPEG is in the process of evaluating the potential use cases and suitability of the HEVC standard for HDR content [3]. Such studies are necessary to reveal the capacity of HEVC for effectively compressing various HDR content formats, and assess whether there is a need for extensions specifically dedicated to HDR content.

B. *Scheme II: broadcast SDR content with metadata*

The second HDR compression scheme relies on a tone mapping operation to transform HDR floating-point data into integer SDR data. Tone mapping is the operation that ensures backward compatibility between HDR content and SDR displays. Note that due to bit-depth limitations, many similar HDR values will be tone mapped to the same SDR value. Consequently, contrast between neighboring pixels as well as between spatially distant areas will be reduced. In general Tone Mapping Operators (TMOs) have various goals, which can range from simulating the human vision to achieving the best subjective quality [1]. Once the HDR content is tone mapped to SDR it is then encoded using a standard video codec while additional information (metadata) is generated to help reconstruct the HDR content at the decoder side [15]. This HDR encoding scheme is usually referred to as backward compatible as the SDR stream can directly address a SDR display. Note that in this scheme, alternatively the SDR content can be generated through grading by an artist, instead of being generated by a TMO. The workflow of Scheme II is similar to the one shown in Fig. 1, however in this case PQ is replaced by a TMO.

III. Test Procedure

A. *Video Data Preparation for Compression*

For our tests, we employ the HDR video dataset provided by Technicolor to the MPEG community [16]. These sequences consist of three HDR videos available in .exr format, with BT.709 color gamut, enclosed in a BT.2020 container. Two of these sequences (FireEater2 and Market3) are natural outdoor videos, the former a night shot and the latter a daylight scene. The last sequence (Tibul2) is a computer-generated HDR video.

In Scheme I, the conversion of HDR videos from floating point values to a format accepted by the HEVC encoder, followes the pipeline described in Fig. 1. as the HDR original sequences are perceptually encoded to 10 bits.

For Scheme II, we need to tone map the HDR sequence. Over the years, many TMOs have been proposed [1], hence choosing one is always a debatable choice. For our objective, the chosen TMO should have two main attributes: being temporally coherent and easily invertible. High temporal coherency is required since it affects the efficiency of inter-prediction and as a result the compression of tone mapped video content as reported in [15]. Being easily invertible is required because we need to inverse tone map each video to reconstruct the HDR sequence. To fulfill these requirements, we selected the camera TMO [17]. An interesting feature of this TMO is that it generates a tone mapped sequence that closely approximates what a traditional camera would have recorded. One major difference from the latter case is that more noise in low light areas would have been captured using a real camera. In other words, using this TMO, we simulate a workflow where a scene captured/recorded by a camera, whose response function is known, is subsequently broadcast to the end-user. At the display end, the camera response function is used to reconstruct physically linear floating-point values to address any HDR display (inverse tone mapping). In our study the HDR content is tone mapped to 8 and 10 bits. We call them SDR8 and SDR10 respectively throughout the paper

B. *Video Compression*

We encoded the HDR (perceptually quantized to 10 bits), SDR10 and SDR8 videos at four different QP levels using the latest HEVC encoder software HM 16.2 [18]. While for HDR and SDR10 videos, the main 10 profile [19] is used, the SDR 8 videos are coded based on the main 8 profile. The QPs used for each of the HDR videos are the ones recommended in [3]. The

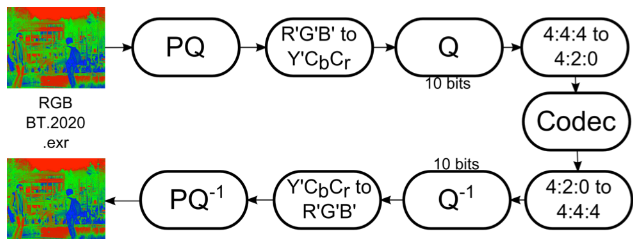

Fig. 1. MPEG workflow for encoding/decoding HDR streams

QPs used for SDR10 and SDR8 were adjusted for each video stream to achieve the same bit-rates as the corresponding HDR ones.

C. *Display*

In our subjective study we used an HDR display prototype based on the concept explained in [21]. This system consists of two main parts: 1) a 40 inch full HD LCD panel in the front, and 2) a projector with HD resolution at the back to provide the backside luminance. The contrast range of the projector is 2000:1. The original HDR video signal is split into two streams, which are sent to the projector and the LCD (see [21] for details). The input signal to the projector includes only the luminance information of the HDR content and the input signal to the LCD includes both luma and chroma information of the HDR video. Using this configuration, the light output of each pixel is effectively the result of two modulations with the two individual dynamic ranges multiplied, yielding an HDR signal. This HDR display system is capable of emitting light at a maximum brightness level of 2700 cd/m^2.

D. *Video Data Preparation for Display*

At the decoding stage, once the content is converted back to physical values, we need to consider the characteristics of our display. The HDR videos provided by MPEG [16] have been graded for a SIM2 display whose peak luminance is 4000 cd/m^2. Furthermore, they were encoded in YC_bC_r with the BT.2020 primaries [22], although the gamut of those sequences is not exceeding the BT.709 gamut [23]. As our display (please refer to Section *III.C* for more details on the display) can only achieve a peak luminance of 2700 cd/m^2 with a BT.709 gamut, both the original, decoded and reconstructed videos need to be converted into a BT.709 container. Note that this adaptation could have also been performed on the original resources and use the resulting videos as source for any comparison. While both solutions are valid, we chose the former one for its compliance with the MPEG workflow [3]. Furthermore, in practice the end-display is never known and hence adaptation at the decoding stage is mandatory.

The main issue with this method is that the Perceptual Quantizer (PQ) [24] assumes that the luminance value encoded will be displayed at the same luminance level. As we reproduce luminance values at a different level to which it was encoded, we decrease the efficiency of the PQ encoding and hence increase the chance of visible degradations in the decoded HDR sequences. Note that this remark does not affect the SDR decoded videos (Scheme II).

The display adaptation process consists of two steps: a gamut conversion and a display conversion. For the gamut conversion, we relied on the HDRTools software developed for MPEG [20]. Regarding the display conversion, we performed a pseudo-tone mapping similar to [25] before separating the back-light and the LCD displayed images using dual-modulation techniques [21]. Note that the adaptation process is the same for every sequence (HDR or SDR).

E. *Subjective Tests*

To compare the performance of the two HDR encoding schemes discussed in Section II, we performed a subjective test. The objective here is to determine how human visual system perceives the compression impairments on decoded HDR, SDR10 and SDR8 video streams compared to the original uncompressed HDR video. The evaluation method was Side-by-Side presentation on a single HDR display based on Recommendation BT.500-13 DSIS [26]. One side was always the original HDR video (uncompressed) while on the other side the stimuli was decoded HDR, SDR10 and SDR8 content. As described in Section *III.B*, four QPs were used, i.e., four impairment levels for each content and input data type (HDR, SDR10 and SDR8) resulting in 4x3x3 = 36 comparisons. The scale for rating the impairments ranged from imperceptible, perceivable but not annoying, slightly annoying, annoying and very annoying, corresponding to discrete numbers of 1 to 5. The order of videos in each session of the test was randomized and extra care was taken for the same sequence not to be shown consecutively. The test took approximately 10 minutes.

We also investigated whether viewers prefer HDR over SDR streams in presence of compression distortions. In this test at each bitrate level, the viewers rate their preference of decoded HDR, SDR10 or SDR8. The evaluation method was a full-paired comparison with the video pairs were presented side-by-side. This method has been shown to be accurate and reliable to construct a scale of perceptual preference [10] while it minimizes working visual memory limitation. The two stimuli shown on each side were the input data type (HDR, SDR10, and SDR8) at similar bitrates for each video content. Overall, all possible combinations of the 3 data type (HDR vs SDR10, SDR10 vs SDR8, and HDR vs. SDR8), i.e., 3 pairs for each content and bitrate level (4 levels), resulting in 3x3x4 = 36 paired comparisons were tested. The test took about 10 minutes. The stimuli shown on left side of the display was labeled A, and the one on the right hand side was labeled B. Subjects were asked to rate which one of A and B they preferred in quality. The option 'same' was also included to avoid random preference selections.

Note that for both of the above-mentioned subjective test scenarios, the side-by-side video frames were cropped (instead of scaling) to fit on the display. The size of the cropped window was kept constant throughout the test. The side-by-side videos are preceded by a message on the screen announcing what will be shown next. Afterwards a 4-second message on a grey background was presented asking the viewers to vote.

F. *Viewers*

Eighteen adult subjects including 10 males and 8 females participated in our experiment. The subjects' age ranged from 24 to 30 years old with an average of 27. Prior to the tests, all subjects were screened for color blindness using the Ishihara chart and visual acuity using the Snellen charts. Subjects that failed the pre-screening test did not participate in the test. All subjects were considered naïve in the field and were not aware of the test objectives. An oral instruction of the test was presented to the subjects before each test began along with a training test, which consisted of 3 videos that were not in the actual test dataset, to allow subjects to adapt to the assessment procedure. All the tests were conducted with two subjects per session.

Table I. Accuracy and monotonicity indexes for the different metrics based on Subjective results

Index / Video	PLCC				SROCC				RMSE			
	mPSNR-Y	tPSNR-Y	HDR-VDP-2	VIF	mPSNR-Y	tPSNR-Y	HDR-VDP-2	VIF	mPSNR-Y	tPSNR-Y	HDR-VDP-2	VIF
FireEater2	0.9905	0.8085	0.9612	0.9969	1	1	1	1	0.4777	0.5989	0.1727	0.0381
Market3	0.4692	0.8421	0.9826	0.7908	0.8	0.8	0.8	0.8	0.5417	0.55	0.1571	0.2850
Tibul2	0.9304	0.9926	0.8086	0.8971	1	1	1	1	0.6554	0.5994	0.4862	0.3585
Overall	0.4801	0.4177	0.8744	0.8659	0.5359	0.3993	0.9212	0.9632	0.5631	0.5832	0.3114	0.3210

IV. RESULTS AND DISCUSSIONS

After collecting the subjective results, the outlier subjects were detected according to the ITU-R BT.500-13 recommendation in [26]. One outlier was detected in these tests and their rates were discarded from the results. The Mean Opinion Score (MOS) for each impaired video was calculated by averaging the scores over all the subjects with 95% confidence interval.

Fig. 2 shows the results of the subjective tests in MOS values versus the required bitrate for encoding the HDR, SDR10 and SDR8 videos. It can be observed from Fig. 2 that the MOS values of encoded HDR for FireEater2 and Tibul2 test sequences are higher compared to those of the SDR10 and SDR8 ones at the same bitrate. This difference in quality is a result of the superior quality of the HDR videos compared to SDR ones due to their capability in covering a wider color gamut and brightness range.

However, for Market3 test video, the subjective quality of HDR, SDR10 and SDR versions of the video are not quite distinctive. One reason may be because of the content of the Market3, which represents a natural scene with one guy walking further from the camera. The cropped version we used in our test does not include the cloth rack therefore it lacks texture. The walking guy, the most visually important object in the cropped scene, lacks high dynamic range information. Therefore, the contribution of HDR to the quality of the data is not apparent. Moreover, it appears that the QP step recommended by MPEG for this video is not large enough and the quality of compressed videos at the recommended QP stays somewhat unchanged.

Overall, the encoding of the HDR videos leads to an improvement of MOS by 2.5, 0.26, 2.26 compared to encoding SDR10 and by 2.17, 0.18, 1.45 compared to encoding SDR8 videos at the same bitrate level for videos FirEater2, Market3, and Tibul2 respectively. On the other hand, encoding SDR8 videos leads to an improvement of MOS by 0.37, 0.08, 0.80 compared to encoding SDR10 ones at the same bitrate level. The MOS improvements are computed by the method in [27].

In recent professional shows such as CES, NAB or IBC, it is claimed that encoding HDR content requires between 10% to 30% overhead in bit-rate, compared to its SDR counterpart. In this article, our subjective results show that HDR content requires less bit-rate compared to SDR, to achieve similar quality of reproduction. Although many reasons explain these results, the main one is perceptual encoding of HDR content before its encoding. This encoding effectively removes visual noise of the input signal, that is to say, information that when displayed on the HDR monitor will not be visible. In a nutshell, SDR content produced using a camera effectively adapt the bin distribution of the tonal value to the recorded content (through the exposure setting). Consequently, areas with the highest granularity do not correspond to where visual information should be kept but rather where the density of pixels is high (in the histogram domain).

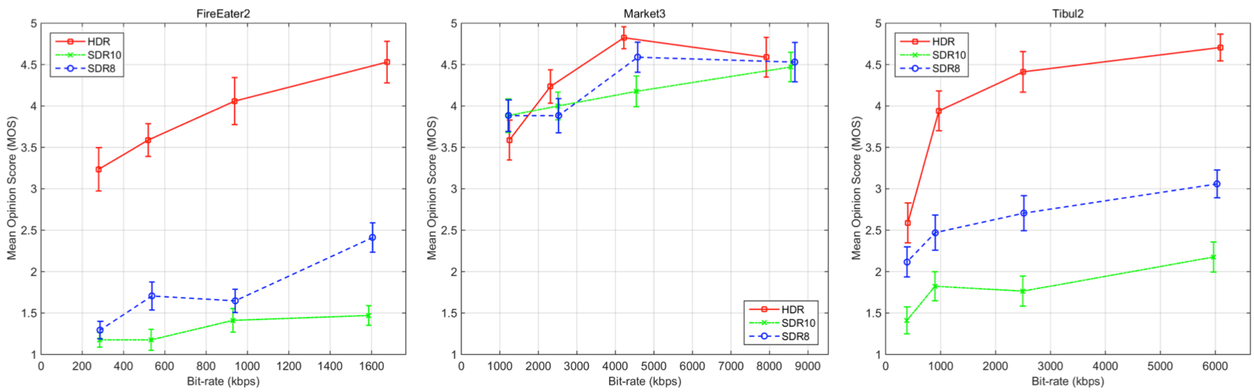

Fig. 2. MOS-rate comparison of the HDR, SDR10, and SDR8 versions of videos FireEater2, Market3, and Tibul2

We also compared the quality of the decoded HDR, SDR10, and SDR8 video streams objectively with respect to the original HDR content. The objective metrics used are the ones publicly available in the latest version of the HDRTools [20]. We also use the HDR-VDP-2 [11] objective metric, as its efficiency has been proven for HDR data with compression artifacts [9]. Table I reports the accuracy and monotonicity indexes for mPSNR, t-PSNR, VIF and HDR-VDP-2 metrics results on the luma (Y) component of the content with those of the subjective evaluations on decoded HDR streams. To compute these indexes, first a regression was fitted to each data set (Video quality metric score, MOS) using logistic fitting. Then, the Pearson linear correlation coefficient (PCC), the root-mean-square error (RMSE), and Spearman rank order correlation coefficient (SROCC) were computed. As it is observed from Table I, the results of HDR-VDP-2 and VIF quality metrics resemble the subjective test results more than the other two metrics. Further more, HDR-VDP-2 shows slightly better results in terms of PLCC and RMSE compared to VIF, while VIF outperforms HDR-VDP-2 in terms of SROCC.

Fig. 3 illustrates the preference probabilities for each content at each bitrate level as well as for all the videos at each bitrate level. The R1 to R4 rates (in kbps) for FireEater2, Market3 and Tibul2 are {1674.832, 937.936, 519.781, 279.367}, {7914.499, 4219.563, 2312.979, 1248.112}, {6090.645, 2499.799, 970.131, 402.4} respectively. In HDR-SDR8 and HDR-SDR10 paired comparisons, the HDR videos were largely preferred while in SDR10-SDR8 paired comparison, half of the times SDR10 and SDR8 ones were not preferred over each other. In other words, SRR10 and SDR8 streams were rated as ‘same’ and for the rest half SDR8 was preferred over SDR10. However, these observations do not hold for Market3 videos for the reasons mentioned earlier.

V. CONLUSIONS

In this paper we investigated the performance of two coding schemes for HDR, one using a single HDR stream and one that involved a tone mapped SDR version and additional metadata. Subjective test evaluations showed that, for the same bit-rate, the single HDR videos were rated higher than the HDR streams reconstructed from the coded SDR and metadata while the TMO used for generating them was the camera TMO. In conclusion, this investigation suggests that in case we want to support backward compatibility for SDR displays, it is preferable to transmit the original HDR stream and generate SDR tone mapped version at the receiver end rather than tone mapping them with camera TMO before encoding.

VI. REFERENCES

[1] E. Reinhard, W. Heidrich, P. Debevec, S. Pattanaik, G. Ward, andK. Myszkowski, High Dynamic Range Imaging, 2nd Edition: Acquisition, Display, and Image-Based Lighting. Morgan Kaufmann, 2010.

[2] P. E. Debevec and J. Malik, “Recovering high dynamic range radiance maps from photographs,” in Proceedings of the 24th annual conference on Computer graphics and interactive techniques - SIGGRAPH ’97. New York, New York, USA: ACM Press, 1997, pp. 369–378.

[3] A. Luthra, E. Francois, and W. Husak, “Call for Evidence (CfE) for HDR and WCG Video Coding,” in ISO/IEC JTC1/SC29/WG11 MPEG2014/N15083. Geneva, Switzerland: Feb. 2015.

[4] M. Pourazad, C. Doutre, M. Azimi, and P. Nasiopoulos, “HEVC: The New Gold Standard for Video Compression: How Does HEVC Compare with H.264/AVC?” IEEE Consumer Electronics Magazine, vol. 1, no. June, pp. 36–46, 2012.

[5] G. J. Sullivan, J.-R. Ohm, W.-J. Han, and T. Wiegand, “Overview of the High Efficiency Video Coding (HEVC) Standard,” IEEE Transactions on Circuits and Systems for Video Technology, vol. 22, no. 12, pp. 1649–1668, Dec. 2012.

[6] A. Koz and F. Dufaux, “A comparative survey on high dynamic range video compression,” in SPIE 8499, Applications of Digital Image Processing XXXV, Oct. 2012, p. 84990E.

[7] R. Mantiuk, K. Myszkowski, and H.-p. Seidel, “Lossy Compression of High Dynamic Range Images and Video,” in SPIE Proceedings Vol.

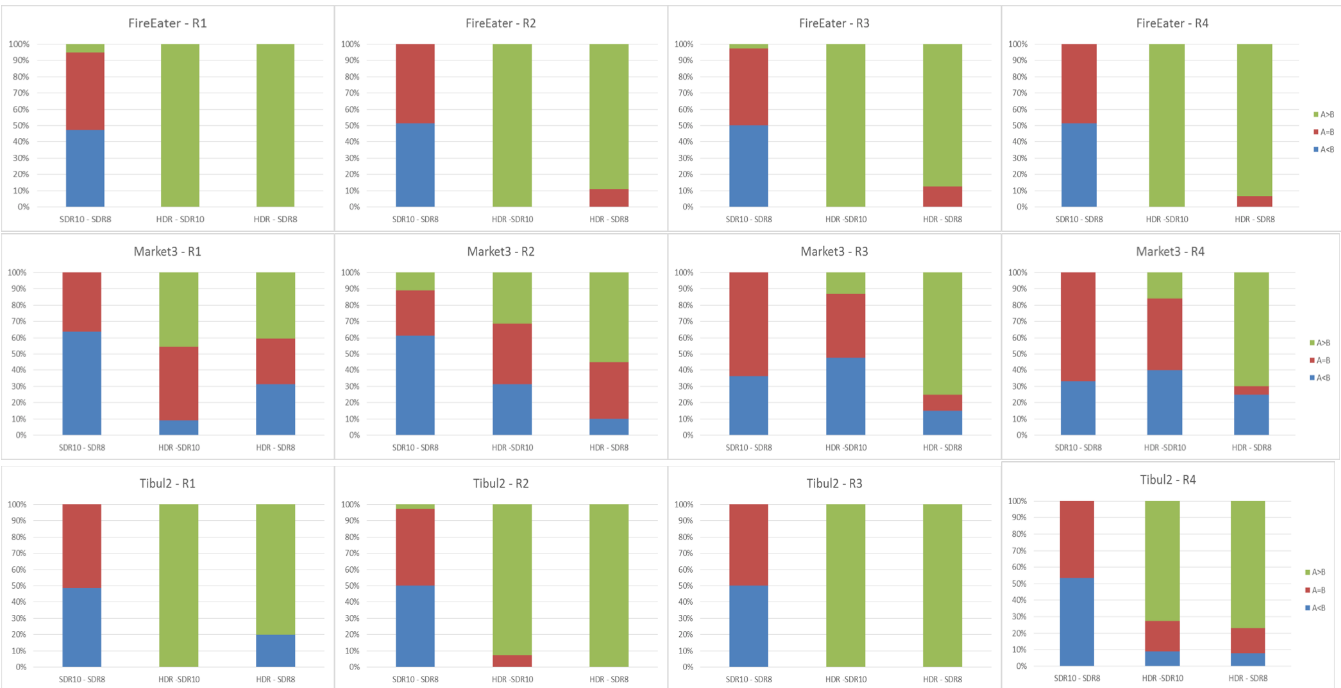

Fig. 3. Preference probabilities for FireEater2, Market3, and Tibul2 at different bit rate levels

6057: Human Vision and Electronic Imaging XI, Feb. 2006, pp. 60 570V–60 570V–10.

[8] S. Miller, M. Nezamabadi, and S. Daly, "Perceptual Signal Coding for More Efficient Usage of Bit Codes," SMPTE Motion Imaging Journal, vol. 122, no. 4, pp. 52–59, May 2013.

[9] M. Azimi, A. Banitalebi-Dehkordi, Y. Dong, M. Pourazad, and P. Nasiopoulos, "Evaluating the performance of existing full-reference quality metrics on High Dynamic Range (HDR) Video content," in International Conference on Multimedia Signal Processing (ICMSP), 2014.

[10] M. Rerabek, P. Korshunov, P. Hanhart and T. Ebrahimi "Correlation of subjective scores and objective metrics for HDR video quality assessment," in ISO/IEC JTC1/SC29/WG11 MPEG2014/ m35273. Strasbourg, France, October 2014.

[11] R. Mantiuk, K. J. Kim, A. G. Rempel, and W. Heidrich, "HDR-VDP- 2: a calibrated visual metric for visibility and quality predictions in all luminance conditions," ACM Transactions on Graphics, vol. 30, p. 1, 2011.

[12] T. O. Aydin, R. Mantiuk, and H.-P. Seidel, "Extending quality metrics to full luminance range images," in Proc. SPIE 6806, Human Vision and Electronic Imaging XIII, vol. 6806. Spie, 2008, pp. 68 060B– 68 060B–10.

[13] R. Boitard, R. K. Mantiuk, and T. Pouli, "Evaluation of color encodings for high dynamic range pixels," in Proc. SPIE 9394, Human Vision and Electronic Imaging XX, 2015.

[14] T. Wiegand, G. Sullivan, G. Bjontegaard, and A. Luthra, "Overview of the H.264/AVC video coding standard," IEEE Transactions on Circuits and Systems for Video Technology, vol. 13, no. 7, pp. 560–576, Jul. 2003.

[15] R. Boitard, D. Thoreau, R. Cozot, and K. Bouatouch, "Impact of Temporal Coherence-Based Tone Mapping on Video Compression," in Proceedings of the 21st European Signal Processing Conference (EUSIPCO), Marrakech, Morocco, 2013.

[16] D. Touz´e and E. Francois, "Description of new version of HDR class A and A' sequences," in ISO/IEC JTC1/SC29/WG11 MPEG2014/M35477. Geneva, Switzerland: Feb. 2015.

[17] J. Petit and R. K. Mantiuk, "Assessment of video tone-mapping : Are cameras S-shaped tone-curves good enough?" Journal of Visual Communication and Image Representation, vol. 24, pp. 1020–1030, 2013.

[18] B. Bross, W.-J. Han, G. J. Sullivan, J.-R. Ohm, and T. Wiegand, "High efficiency video coding (HEVC) text specification draft 6," in ISO/IEC JTC1/SC29/WG11 JCTVC-H1003. IEEE, 2012

[19] https://hevc.hhi.fraunhofer.de/

[20] A. M. Tourapis and D. Singer, "HDRTools: Software updates," in ISO/IEC JTC1/SC29/WG11 MPEG2015/M35471, IEEE, Ed., Geneva, Switzerland, 2015.

[21] H. Seetzen, W. Heidrich, W. Stuerzlinger, G. Ward, L. A. Whitehead, M. Trentacoste, A. Ghosh, and A. Vorozcovs, "High dynamic range display systems," ACM Transactions on Graphics, vol. 23, no. 3, p. 760, Aug. 2004.

[22] ITU, "Recommendation ITU-R BT.2020: Parameter values for ultrahigh definition television systems for production and international programme exchange," International Telecommunications Union, 2012.

[23] ITU, "Recommendation ITU-R BT.709-3: Parameter values for the HDTV standards for production and international programme exchange," International Telecommunications Union, 1998.

[24] T. Kunkel, G. Ward, B. Lee, and S. Daly, "HDR and wide gamut appearance-based color encoding and its quantification," in 2013 Picture Coding Symposium (PCS). San Jose: IEEE, Dec. 2013, pp. 357–360.

[25] E. Reinhard, M. Stark, P. Shirley, and J. Ferwerda, "Photographic tone reproduction for digital images," ACM Trans. Graph., vol. 21, no. 3, pp. 267–276, Jul. 2002.

[26] International Telecommunication Union, "Methodology for the subjective assessment of the quality of television pictures BT Series Broadcasting service," in Recommendation ITU-R BT.500-13, vol. 13, 2012.

[27] P. Hanhart, T. Ebrahimi, "Calculation of average coding efficiency based on subjective quality scores", Journal of Visual Communication and Image Representation, Volume 25, Issue 3, April 2014, Pages 555-564, ISSN 1047-3203.