Tone Mapping HDR Panoramas for Viewing in Head Mounted Displays

Miguel Melo[1], Kadi Bouatouch[2], Maximino Bessa[1,3], Hugo Coelho[3], Remi Cozot[2]
and Alan Chalmers[4]

[1]*INESC TEC, Porto, Portugal*
[2]*IRISA, University of Rennes I, Rennes, France*
[3]*UTAD, Vila Real, Portugal*
[4]*WMG, University of Warwick, Coventry, U.K.*

Keywords: Head Mounted Displays, HDR Panoramas, Tone Mapping, High Dynamic Range.

Abstract: Head-mounted displays enable a user to view a complete environment as if he/she was there; providing an immersive experience. However, the lighting in a full environment can vary significantly. Panoramic images captured with conventional, Low Dynamic Range (LDR), imaging of scenes with a large range of lighting conditions, can include areas of under- or over-exposed pixels. High Dynamic Range (HDR) imaging, on the other hand, is able to capture the full range of detail in a scene. However, HMDs are not currently HDR and thus the HDR panorama needs to be tone mapped before it can be displayed on the LDR HMD. While a large number of tone mapping operators have been proposed in the last 25 years, these were not designed for panoramic images, or for use with HMDs. This paper undertakes a two part subjective study to investigate which of the current, state-of-the-art tone mappers is most suitable for use with HMDs.

1   INTRODUCTION

Panoramic images capture the full 360° of detail in a scene (Figure 1). When shown within a Head-Mounted Display (HMD) the viewer can look around the captured environment as if he/she was there, providing a highly immersive experience. There can be a wide range of lighting in any scene, from the sun to dark shadows. Traditional panoramic imaging techniques, known as Low Dynamic Range (LDR) are only able to capture a dynamic range of up to 256 to 1 (8 stops) and thus will fail to capture all the detail, resulting in areas of the panoramic image with under- or over-exposed pixels. By using 32 bit IEEE floating point values to represent each colour channel, High Dynamic Range (HDR) imaging is able to capture the full range of lighting in a scene.

Worldwide sales of Head Mounted Displays (HMDs) are increasing rapidly with cumulative total sales expected to top 200 million from 2015 to 2020 (Tractica, 2017) and the revenue for VR related hardware is due to exceed $3.6 billion in 2017 (Nafarrete, 2017). The low cost availability of high quality HMDs, such as the Oculus Rift, HTC Vive, etc., means that these devices are now being used for a wide range of applications from computer games

to films. In particular 360° content is becoming more widely available helped by companies, such as YouTube which provides the ability to upload and view 360° videos on their website in March 2015.

Currently HMDs are not HDR, and thus an HDR image has to be tone mapped before it is displayed within the HMD. A large number of tone mapping operators (TMOs) have been presented in the last 25 years (Banterle et al., 2011). However, these TMOs were not designed for panoramic images and how the perception of a scene may change as the user looks around within the environment. This paper undertakes a detailed two part subjective evaluation of the use of TMOs for 360° HDR panoramic images. In the first study, participants were asked to rank according to their preference a set of tone mapped HDR panoramas without reference. In the second part of the study, the same participants were asked to rank non-panoramic images tone mapped with the same operators within the HMD against the reference HDR image displayed on an HDR display. As can been seen in Figure 1, a 360° HDR panoramic image displayed on an HDR display would be distorted and thus not be a fair comparison with the image as it would be viewed in the HMD. In addition, the study also compared the performance of two makes of HMD, the Oculus Rift

Figure 1: HDR panorama.

DK2 and the Oculus Rift CV1, to investigate whether the technical specifications of the HMD itself had an influence on the choice of TMO for the HMD.

2 RELATED WORK

Tone mapping attempts to preserve the perception of a scene, captured with an HDR image, on an LDR display. Many TMOs have been proposed since the first one was published in 1993 (Chiu et al., 1993), For a detailed overview of these see (Banterle et al., 2011). TMOs are typically classifed as *global* or *local*. Global TMOs apply the same operation to all pixels in the image. They are typically able to run in real time. Local TMOs, on the other hand, apply apply the operator differently to each pixel, taking into account the pixels around the pixel being considered. Local TMOs tend to be much more complex than global ones but they are, unlike global TMOs, able to preserve both global and local contrast. An additional category *temporal* was recently introduced to take into account how a tone mapper may need to deal with how a scene changes over time, for example when tone mapping HDR video (Banterle et al., 2011; Boitard et al., 2014).

In this paper four state-of-the-art TMOs, which have previously been shown to be very effective on a number of different displays (Melo et al., 2014), were used:

Fer: Developed by Ferweda et al. (Ferwerda et al., 1996), this TMO includes a model of visual adaptation which takes into account aspects of the Human Visual System (HVS), including visibility, visual acuity, and illumination changes adaptation. Threshold-versus-intensity (TVI) functions are used to model photopic and scotopic vision, while a linear combination of both photopic and scotopic vision is used for mesopic vision.

Man: This display adaptive TMO uses a model of the HVS to minimise visible contrast distortions (Mantiuk et al., 2008). In addition, this TMO takes into account the ambient light levels and display characteristics.

Boi: Designed especially for HDR video, this is a temporal coherent TMO that preserves overall contrast (Boitard et al., 2012). This is achieved by considering the perceptual consistency of an object throughout the video. A frame is processed in two steps. In the first step the frame is tone mapped with (any) global tone mapper, while in the second step, the luminance of the frame is adjusted taking into account the lighting of the entire video.

Pat: Known as the time-dependent visual adaptation TMO (Pattanaik et al., 2000), this method takes into account that the HVS does not react instantly to large changes in luminance levels. The TMO simulates how the HVS responds to real world lighting levels.

2.1 Evaluating Tone Mappers

The different TMOs have been evaluated in detail over the years using both objective and subjective metrics. The most widely used objective metric has been VDP (Daly, 1993) and its HDR versions, HDR-VDP (Mantiuk et al., 2004) and HDR-VDP-2 (Mantiuk et al., 2011). Other objective metrics include VQM (Narwaria et al., 2015) and the dynamic range independent quality assessment metric (Aydin et al., 2010).

One of the first subjective evaluations of TMOs was by Drago et al. (Drago et al., 2002) in which

4 HDR scenes were tone mapped with 7 different TMOs. Ledda et al. (Ledda et al., 2005) were the first to use an HDR display as a reference. Tone mapped images displayed on two LDR displays on either side of the HDR display were compared with the HDR image on the HDR display. Other evaluations of TMOs include work by (Yoshida et al., 2005), (Kuang et al., 2007), (Čadík et al., 2008) and (Melo et al., 2014; Melo et al., 2015). Recently temporal TMOs have been subjectively evaluated by Eilertsen et al. (Eilertsen et al., 2013).

2.2 Tone Mapping for HMDs

Cutchin et al. presented an interactive view dependent tone mapping technique for HDR panoramas on HMDs via a view-adjusted mapping function stored in a separate texture file, which they termed Tone-Texture (Cutchin and Li, 2016). This technique was used to expand the perceived colours within the HMD in order to improve the panorama's visual appearance. The approach was shown to provide a better visual appearance compared to the panorama tone mapped with the Photographic TMO (Reinhard et al., 2002). They provided demonstration systems for WebGL and the Oculus Rift and GearVR HMDs.

Yu presented a dynamic TMO for HDR panoramas which takes into account the user's head position and thus the viewport (Yu, 2015). A simple model of eye adaptation is used to simulate light and dark adaptation. Yu's TMO was shown in a study of only 8 participants to be preferred over the global Photographic tone mapper (Reinhard et al., 2002).

Perrin et al. (Perrin et al., 2017) undertook a quality assessment of 360° HDR content when displayed in HMDs. Three global (Simple linear scaling, Photographic (Reinhard et al., 2002), Display adaptive (Mantiuk et al., 2008))) and two local TMOs (Detail preserving (Mantiuk et al., 2006) and Exposure fusing (Mertens et al., 2009)) were evaluated using a simple pairwise comparison. There results showed that the participants did not have a clear preference for any of the TMOs evaluated.

3 SUBJECTIVE EVALUATION

This subjective study evaluated four state-of-the-art TMOs across two consumer-grade HMDs, namely Oculus RIFT CV1 and Oculus RIFT DK2. The study was divided into two. The first experiment consisted of evaluating five HDR panorama images viewed on the HMD. These panoramas had each been tone mapped with the four TMOs. Participants were asked to record their preference. The second experiment asked the participants to rank the TMOs applied to set of HDR images viewed on an HMD, compared to the HDR images displayed on a SIM2 HDR display. A participant only performed the experiments on one of the two HMDs (between subjects) but every participant performed the rankings for both HDR panoramas and HDR images (within-subjects).

3.1 Sample

A total of 15 participants with ages between 20 and 44 (M = 27.62, SD = 7.38) were divided into two groups. One group performed the experiments using the Oculus RIFT CV1 (N=7) and the other group performed the experiments using the Oculus RIFT DK2 (N=8). All participants reported normal or corrected-to-normal vision.

3.2 Variables

There were three independent variables: the HMDs, the Content Type, and the selected TMOs. The dependent variable is the TMO ranking provided by participants. Ranking was preferred over rating as it avoids narrow distributions by guarantying that each ranked item has a unique value for all tone mapped stimuli. Ranking was also chosen over pair-wise comparisons as this would require a significantly larger number of comparisons to compare all stimuli.

3.3 Materials

The HMDs used for the experiments were the Oculus Rift CV 1 that has a resolution of 2160 × 1200, a refresh rate of 90Hz and a field of view of 110° and the Oculus Rift DK2 with a resolution of 980 × 1080, a refresh rate of 75Hz and a field of view of 100°. The stimuli consisted of 5 HDR panoramas, and 5 HDR images [1]. These were tone-mapped using the 4 different TMOs: Fer, Man, Boi and Pat, as described in section 2.

Custom VR software was developed to enable participants to perform the experiments with ease. This software consisted of different screens (one per each panorama) where participants could watch every tone-mapped panorama the number of times they needed to perform the rankings. For ranking the tone mapped panoramas, participants had to drag the corresponding tone mapped panorama thumbnails to a set of ranking boxes numbered from 1 to 4 where 1

---

[1]Please refer to http://www.miguelmelo.pt/hdrhmd/ for the images used in this study.

was the preferred (or closer to the reference in experiment 2) and 4 the less preferred (or less close to the reference in experiment 2). The experiments were conducted using a desktop PC with an Intel i7-5820K CPU, a NVIDIA GeForce GTX 980 Ti graphics card and 16GB of RAM.

3.4 Procedure

The experiments were conducted in a laboratory setting where all the environment variables such as sound, temperature and lighting levels were controlled. The first step was to present the experimental study to the participants without disclosing its purpose to avoid bias. A consent form was given to participants for them to agree to take part in the study. How the software worked was then explained and a demonstration scene instructed them how to perform the two-part experiments.

- Part 1: Participants were asked to rank 4 TMOs applied to the HDR panoramas according to their preference taking into consideration overall image quality, naturalness and detail across 5 scenes. In this experiment, there was no reference.

- Part 2: Participants were asked to rank 4 TMOs applied to HDR images according to a reference that was shown on an HDR display across 5 scenes.

The experimental team helped participants take the HMDs on and off and remained in the experimental room in case participants had any questions and to ensure that the data was collected properly. Participants were free to drop-out from the experiments at any given time. On average, the entire procedure took approximately 15 minutes.

3.5 Statistical Procedures

To study the impact of HMDs across Content Types over TMOs rankings, a 2 (HMD) $\times$ 2 (Content Type) $\times$ 4 (TMOs) mixed design factorial ANOVA was performed where TMO was a within-participant variable and HMD and Content Type were between-participant variables. A Boxs M test was used to test the multivariate homogeneity of the data. To investigate the performance of TMOs across the different scenarios, a Mann-Whitney U was performed. The Friedman test was carried out to verify if TMOs rankings were statistically significantly different in each scenario. Additionally, post hoc analysis with Wilcoxon signed-rank tests was conducted with a Bonferroni correction applied, to identify possible groupings of TMOs. Kendalls Coefficient of Concordance W was also used to give an estimate of

the agreement amongst participants. A Coefficient of Concordance of W = 1 signifies perfect agreement amongst the participants and W = 0 indicates complete disagreement.

4 RESULTS

For ease of understanding, the results are divided into subsections: Results across all HMDs across all content types; HMD impact on TMO performance for HDR panoramas; HMD impact on TMO performance for HDR images; Content type impact on TMOs performance on Oculus RIFT CV1; and ,Content type impact on TMOs performance on Oculus RIFT DK2.

4.1 Results Across All HMDs Across All Content Types

To test the multivariate homogeneity, a Box's M test was applied which revealed that the within-group covariance matrices are equal and that a MANOVA can be applied (Box's M = 1.511, $p > 0.05$). The MANOVA analysis showed that a statistically significant impact of the HMDs over TMOs ranking can be assumed but with some reservation as there is only a moderate effect F(4,143)= 2.369, $p = 0.055$, $Wilks'$ λ = .938, η_p^2 = .062, $O.P.$ =.673. For Content Type, statistically significant differences were found in TMOs rankings, F(4,143)= 5.668, $p = 0.000$, $Wilks'$ λ = .859, η_p^2 = .141, $O.P.$ = .981. For HMD $\times$ Content Type, no statistically significant differences were found, F(4,143)= .994, $p = 0.413$, $Wilks'$ λ = .973, η_p^2 = .027, $O.P.$ = .308.

4.2 HMD Impact on TMO Performance for HDR Panoramas

A Mann-Whitney U test showed that there was no statistically significant difference in the performance of TMOs across the two HMDs for the HDR panoramas:

- Man: U = 596.000, p = .227
- Fer: U = 641.500, p = .511
- Boi: U = 553.500, p = .105
- Pat: U = 558.500, p = .081

As there were no differences of performance across the two HMDs, the results for the two scenarios were grouped to perform a Friedman test to analyze if there were differences across the TMOs. The Friedman test revealed a statistically significant difference between the different TMOs ($\chi^2(3) = 70.041, p = 0.000$). Post hoc analysis with Wilcoxon signed-rank

Table 1: Wilcoxon Signed-Rank Test results across HMDs for HDR panoramas.

Paired Comparison	Z	p
Fer-Man	-4.849	0.000
Boi-Man	-2.662	0.008
Pat-Man	-6.338	0.000
Boi-Fer	-2.152	0.031
Pat-Fer	-3.979	0.000
Pat-Boi	-5.320	0.000

Table 2: Wilcoxon Signed-Rank Test results across HMDs for HDR images.

Paired Comparison	Z	p
Fer-Man	-0.146	0.884
Boi-Man	3.755	0.000
Pat-Man	-6.512	0.000
Boi-Fer	-3.314	0.001
Pat-Fer	-5.906	0.000
Pat-Boi	-3.701	0.000

tests was conducted with a Bonferroni correction applied, resulting in significance levels $p < .031$ for all the comparisons as shown in table 1. The TMOs ranking order was: Man (M=1.73), Boi (M=2.13), Fer (M=2.67), and Pat (M=3.41) with no groupings of TMOs. Kendalls coefficient of Concordance was W = 0.311 ($p < 0.05$).

4.3 HMD Impact on TMO Performance for HDR Images

A Mann-Whitney U test showed that there was no statistically significant difference in the performance of TMOs across the two HMDs for the HDR images:

- Man: U = 566.500, p = .113
- Fer: U = 624.500, p = .394
- Boi: U = 607.000, p = .295
- Pat: U = 635.500, p = .436

Again, as there were no differences of performance across the two HMDs for HDR images, the results for the two scenarios were grouped to perform a Friedman test to analyze if there were differences across the TMOs. The Friedman test revealed a statistically significant difference between the different TMOs ($\chi^2(3) = 70.041, p = 0.000$). Post hoc analysis with Wilcoxon signed-rank tests was conducted with a Bonferroni correction applied, resulting in significance levels across all TMOs except for the pair Fer - Man ($p = .884$) as shown in table 2. The TMOs ranking order was: Man (M=1.93), Fer (M=1.95), Boi (M=2.69), and Pat (M=3.44), with Man and Fer grouped together. Kendalls coefficient of Concordance was W = 0.312 ($p < 0.05$).

4.4 Content Type Impact in TMOs Performance on Oculus RIFT CV1

A Mann-Whitney U test showed that there was no statistically significant difference in the performance of TMOs across the two content types:

- Man: U = 523.500, p = .258

Table 3: Wilcoxon Signed-Rank Test results across content types for Oculus Rift CV1.

Paired Comparison	Z	p
Fer-Man	-0.146	0.064
Boi-Man	3.755	0.016
Pat-Man	-6.512	0.000
Boi-Fer	-3.314	0.498
Pat-Fer	-5.906	0.000
Pat-Boi	-3.701	0.000

- Fer: U = 452.000, p = .051
- Boi: U = 564.000, p = .554
- Pat: U = 588.000, p = .750

As there were no differences of performance across the two content types when using the Oculus Rift CV 1, the results for the two scenarios were grouped to perform a Friedman test and verify if there were differences across the TMOs. The Friedman test revealed a statistically significant difference between the different TMOs ($\chi^2((3) = 35,584, p = 0.000$). Post hoc analysis with Wilcoxon signed-rank tests was conducted with a Bonferroni correction applied, resulting in significance levels across all TMOs except for the pair Fer - Man ($p = .064$) and Boi Fer ($p = 495$) as shown in table 3. The TMOs ranking order was: Man (M=1.99), Fer (M=2.31), Boi (M=2.45), and Pat (M=3.24), with Man and Fer grouped together as well as Fer and Boi. Kendalls coefficient of Concordance was W = 0.169 ($p < 0.05$).

4.5 Content Type Impact on TMOs Performance on Oculus RIFT DK2

The Mann-Whitney U test conducted to identify if there were differences of TMO performance across the two types of content have revealed that there are statistically significance differences for Fer and Boi:

- Man: U = 691,000, p = .241
- Fer: U = 386,000, p = .000
- Boi: U = 471,000, p = .001

Table 4: Wilcoxon Signed-Rank Test results for HDR panoramas for Oculus Rift DK2.

Paired Comparison	Z	p
Fer-Man	-4.560	0.000
Boi-Man	-1.810	0.070
Pat-Man	-5.463	0.000
Boi-Fer	-2.292	0.022
Pat-Fer	-3.771	0.000
Pat-Boi	-4.825	0.000

Table 5: Wilcoxon Signed-Rank Test results for HDR images for Oculus Rift DK2.

Paired Comparison	Z	p
Fer-Man	-4.560	0.948
Boi-Man	-1.810	0.000
Pat-Man	-5.463	0.000
Boi-Fer	-2.292	0.001
Pat-Fer	-3.771	0.000
Pat-Boi	-4.825	0.007

- Pat: U = 730,500, p = .425

As there were significant differences on TMOs performance across the two content types, the TMOs ranking was analyzed separately. For the HDR panoramas, a Friedman test revealed a statistically significant difference between the different TMOs ($\chi^2(3) = 57,316, p = 0.000$). Post hoc analysis with Wilcoxon signed-rank tests was conducted with a Bonferroni correction applied, resulting in significance levels across all as shown in table 4. The TMOs ranking order was: Man (M=1.60), Boi (M=2.01), Fer (M=2.76), and Pat (M=3.63), with Man and Boi grouped together. Kendalls coefficient of Concordance was W = 0.478 ($p < 0.05$).

Regarding HDR images, the Friedman test revealed that there are statistically significant differences between the different TMOs ($\chi^2(3) = 52,010., p = 0.000$). Post hoc analysis with Wilcoxon signed-rank tests was conducted with a Bonferroni correction applied, resulting in significance levels across all as shown in table 5. The TMOs ranking order was: Man (M=1.78), Fer (M=1.85), Boi (M=2.83), and Pat (M=3.55), being that Man and Fer were grouped together. Kendalls coefficient of Concordance was W = 0.433 ($p < 0.05$).

5 DISCUSSION

This paper has investigated three important research questions: 1) Is the choice of TMO for displaying an HDR image affected by the use of an HMD for the display?; 2) Does content type have an effect on the choice of TMO for use with an HMD?; and 3)Is the TMO performance consistent across HMDs across Content Types?

Table 6 summarizes the TMOs Rankings and groupings across scenarios from the analysed results.

From the statistics of the results we obtained for our user studies, the following is clear:

- **The choice of HMD seems to have an impact on TMO ranking:** Although MANOVA analysis revealed a $p = 0.055$, the null hypothesis is rejected with some reservation as there is only a moderate effect of the sample and an OP of near 67%. This latter result corroborates previous studies such as (IJsselsteijn et al., 2000), (IJsselsteijn et al., 2001), (Lin et al., 2002), (Baños et al., 2008), and (Young et al., 2014) which have identified technology is a factor that influences a VR experience. Our paper extends this previous work to now also include the influence of technology when displaying HDR content on HMDs.

- **The Content Type has impact on TMO ranking:** Depending on the HDR content to be shown, one should carefully select the appropriate TMO. For instance, for HDR panoramas, Boi performs better than Fer but the inverse happens when the content is HDR images. If one considers the features of both TMOs, the results seem to indicate that local TMOs perform better for HDR panoramas while global TMOs are more suited for HDR images. This might happen because of the distortion of the source panorama images; a local processing would consider that distortion and perform better tone mapping due use of the neighbouring pixels' information.

- **TMOs performance is consistent across scenarios** (except for Content Type when the Oculus Rift DK2): Despite the impact of the different HMDs and/or content types on the TMOs performance, the same TMO will have similar performance across conditions.

- **Concordance between participants on rankings are higher for Oculus Rift DK2** ($CV1 \approx .20$ *vs* $DK2 \approx .45$): Participants were able to identify more differences in CV1 between TMOs therefore there was less concordance. This suggests that less complex TMOs can be used for older devices as the difference between TMOs is not so noticeable.

From the statistics of the results we obtained for our user studies, the global MANOVA and the set of pairwise comparisons, the following is clear:

- The choice of HMD has impact on TMO ranking

Table 6: Overall results obtained for each scenario. Coloured groupings represent TMOs that were not found to be significantly different using pairwise comparisons to each other, via Bonferroni adjustment, at $p < 0.01$.

	Kendall's Co-efficient of Concordance*	1^{st}	2^{nd}	3^{rd}	4^{th}
Rankings across HMDs for HDR Panoramas**	0.311	Man	Boi	Fer	Pat
Rankings across HMDs for HDR images**	0.312	Man	Fer	Boi	Pat
Rankings across Content Type for CV1**	0.169	Man	Fer	Boi	Pat
Rankings for HDR panoramas using Oculus Rift DK2	0.478	Man	Boi	Fer	Pat
Rankings for HDR images using Oculus Rift DK2	0.433	Man	Fer	Boi	Pat

*for $(p < 0.01)$

**The independent variables were grouped as there was no statistically significant difference between conditions, $p < 0.005$.

- The content Type (panorama $\times$ image) has impact on TMO ranking

- Concordance between participants on rankings are much higher for DK2 (CV1 $\sim$ 20 vs DK2 $\sim$ 45). In particular, participants where able to identify more differences in CV1 between TMOs therefore there was less concordance. This suggests that less complex TMOs can be used for older devices as the difference between TMOs is not so noticeable.

6 CONCLUSIONS

With the rapid growth of VR technology and the increasing demand for more realism within VR experiences, more content is likely to be produced in HDR, as this has no areas of under- or over-exposed pixels. This is also very likely to be HDR panoramic content. Until HMDs include HDR displays, this HDR content will need to be tone mapped for enhanced display on the LDR displays of the HMD. This paper has shown that we cannot rely on previous evaluations of TMOs to chose the right TMO for use with HDR panoramic content for HMDs. A new choice of TMO specifically for HMDs is necessary to achieve a high level of perceived realism in the HMD, especially when considering panoramic images.

The subjective evaluation has considered only 5 HDR panoramas. To provide conclusive results more HDR panoramas need to be considered. In addition, these panoramas should be chosen from a mixture of scenes, such as outdoor daylight, outdoor night light, indoor lighting, mixture of indoor and outdoor lights, etc. We will investigate this in the future as we wish to investigate whether the final preference of a TMO may depend not only on the HMD but also on the content, for which we will need to capture this more complete HDR panorama database.

Future work will also investigate, including with the use of eyetrackers, whether more in-depth knowledge of how humans view content on HMD screens, could lead to even more effective TMOs for displaying HDR panoramic content on HMDs than the ones preferred in our experiments.

ACKNOWLEDGEMENTS

REFERENCES

Aydin, T. O., Cadik, M., Myszkowski, K., and Seidel, H.-P. (2010). Video quality assessment for computer graphics applications. In *SIGGRAPH Asia*. ACM.

Baños, R. M., Botella, C., Rubió, I., Quero, S., García-Palacios, A., and Alcañiz, M. (2008). Presence and emotions in virtual environments: The influence of stereoscopy. *CyberPsychology & Behavior*, 11(1):1–8.

Banterle, F., Artusi, A., Debattista, K., and Chalmers, A. (2011). *Advanced High Dynamic Range imaging*. CRC Press.

Boitard, R., Bouatouch, K., Cozot, R., Thoreau, D., and Gruson, A. (2012). Temporal coherency for video tone mapping. *SPIE Optical Engineering+ Applications*, pages 84990D–84990D.

Boitard, R., Cozot, R., Thoreau, D., and Bouatouch, K. (2014). Zonal brightness coherency for video tone mapping. *Signal Processing: Image Communication*, 29(2):229–246.

Čadík, M., Wimmer, M., Neumann, L., and Artusi, A. (2008). Evaluation of HDR tone mapping methods using essential perceptual attributes. *Computers & Graphics*, 32(3):330–349.

Chiu, K., Herf, M., Shirley, P., Swamy, S., Wang, C., and Zimmerman, K. (1993). Spatially Nonuniform Scaling Functions for High Contrast Images. In *Proceedings of Graphics Interface '93*, pages 245–253.

Cutchin, S. and Li, Y. (2016). View dependent tone mapping of HDR panoramas for head mounted displays. In *Proceedings of the 26th International Conference on Artificial Reality and Telexistence and the 21st Eurographics Symposium on Virtual Environments*, ICAT-EGVE '16, pages 29–36. Eurographics Association.

Daly, S. (1993). The visible differences predictor: an algorithm for the assessment of image fidelity. pages 189–206. Digital Image and Human Vision.

Drago, F., Martens, W., Myszkowski, K., and Seidel, H.-P. (2002). Perceptual evaluation of tone mapping operators with regard to similarity and preference. Research Report MPI-I-2002-4-002, Max-Planck-Institut für Informatik, Stuhlsatzenhausweg 85, 66123 Saarbrücken, Germany.

Eilertsen, G., Unger, J., Wanat, R., and Mantiuk, R. (2013). Survey and evaluation of tone mapping operators for HDR video. In *ACM SIGGRAPH 2013 Talks*.

Ferwerda, J. A., Pattanaik, S. N., Shirley, P., and Greenberg, D. P. (1996). A model of visual adaptation for realistic image synthesis. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 249–258. ACM.

IJsselsteijn, W., de Ridder, H., Freeman, J., Avons, S. E., and Bouwhuis, D. (2001). Effects of stereoscopic presentation, image motion, and screen size on subjective and objective corroborative measures of presence. *Presence*, 10(3):298–311.

IJsselsteijn, W. A., de Ridder, H., Freeman, J., and Avons, S. E. (2000). Presence: concept, determinants, and measurement. In *Electronic Imaging*, pages 520–529. International Society for Optics and Photonics.

Kuang, J., Yamaguchi, H., Liu, C., Johnson, G. M., and Fairchild, M. D. (2007). Evaluating HDR rendering algorithms. *ACM Trans. Appl. Percept.*, 4(2).

Ledda, P., Chalmers, A., Troscianko, T., and Seetzen, H. (2005). Evaluation of tone mapping operators using a high dynamic range display. In *ACM SIGGRAPH 2005 Papers*, SIGGRAPH '05, pages 640–648, New York, NY, USA. ACM.

Lin, J.-W., Duh, H. B.-L., Parker, D. E., Abi-Rached, H., and Furness, T. A. (2002). Effects of field of view on presence, enjoyment, memory, and simulator sickness in a virtual environment. In *Virtual Reality, 2002. Proceedings. IEEE*, pages 164–171. IEEE.

Mantiuk, R., Daly, S., and Kerofsky, L. (2008). Display adaptive tone mapping. In *ACM Transactions on Graphics (TOG)*, volume 27, page 68. ACM.

Mantiuk, R., Kim, K. J., Rempel, A. G., and Heidrich, W. (2011). HDR-VDP-2: A calibrated visual metric for visibility and quality predictions in all luminance conditions. *ACM Trans. Graph.*, 30(4):40:1–40:14.

Mantiuk, R., Myszkowski, K., and Seidel, H.-P. (2004). Visible difference predicator for high dynamic range images. In *Proceedings of IEEE International Conference on Systems, Man and Cybernetics*, pages 2763–2769.

Mantiuk, R., Myszkowski, K., and Seidel, H.-P. (2006). A perceptual framework for contrast processing of high dynamic range images. *ACM Transactions on Applied Perception*, 3(3):286–308.

Melo, M., Bessa, M., Debattista, K., and Chalmers, A. (2014). Evaluation of HDR video tone mapping for mobile devices. *Signal Processing: Image Communication, Special Issue on Advances in High Dynamic Range Video Research*, 29(2):247–256.

Melo, M., Bessa, M., Debattista, K., and Chalmers, A. (2015). Evaluation of Tone-Mapping Operators for HDR Video Under Different Ambient Luminance Levels. *Computer Graphics Forum*, 34(8):38–49.

Mertens, T., Kautz, J., and Van Reeth, F. (2009). Exposure fusion: A simple and practical alternative to high dynamic range photography. *Computer Graphics Forum*, 28(1):161–171.

Nafarrete, J. (2017). Report: Global VR hardware revenue to hit $3.6 billion in 2017.

Narwaria, M., Silva, D., and Le Callet, P. (2015). HDR-VQM: An objective quality measure for high dynamic range video. *Signal Processing: Image Communication*, 35:46–60.

Pattanaik, S. N., Tumblin, J., Yee, H., and Greenberg, D. P. (2000). Time-dependent visual adaptation for fast realistic image display. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 47–54. ACM Press/Addison-Wesley Publishing Co.

Perrin, A.-F. N. M., Bist, C., Cozot, R., and Ebrahimi, T. (2017). Measuring quality of omnidirectional high dynamic range content. In *Optics+Photonics Optical Engineering+Applications Applications of Digital Image Processing XL*. SPIE.

Reinhard, E., Stark, M., Shirley, P., and Ferwerda, J. (2002). Photographic tone reproduction for digital images. In *Proceedings of SIGGRAPH 2002*, pages 267–276.

Tractica (2017). More than 200 million virtual reality head-mounted displays to be sold by 2020.

Yoshida, A., Blanz, V., Myszkowski, K., and Seidel, H.-P. (2005). Perceptual evaluation of tone mapping operators with real-world scenes. In *17th Annual Symposium on Electronic Imaging*. SPIE.

Young, M. K., Gaylor, G. B., Andrus, S. M., and Bodenheimer, B. (2014). A comparison of two cost-differentiated virtual reality systems for perception and action tasks. In *Proceedings of the ACM Symposium on Applied Perception*, pages 83–90. ACM.

Yu, M. (2015). Dynamic tone mapping with head-mounted displays. Standford University Report.